Cisco Umbrella

Volume 1, November 2020

The modern cybersecurity landscape:

Scaling for threats in motion

Introduction

For the majority of 2020, in the face of a global pandemic, the entire world has been grappling with massive change — in how we live, how we work, how we connect. But one area that's always been dynamic and rapidly evolving is the cyberthreat landscape.

Here at Cisco, we've seen firsthand how commonplace threats have quickly evolved into complex, multi-stage attacks that use tried-and-true malware methodology paired with innovative new tactics to cover their tracks.

In the face of these new threats, InfoSec teams are feeling increasingly overwhelmed. The right information, however, can prepare you for what's out there.

Cisco Umbrella has identified a number of major threat trends in the first three quarters of 2020 that will have serious implications for years to come:

- **Trend #1:** Trojans and droppers are getting a second life as new forms of malware delivery.
- **Trend #2:** Orchestrated, multi-staged, evasive attacks are becoming the norm.
- **Trend #3:** Cryptomining is opening the door to other types of cyberthreats.
- **Trend #4:** Attackers are taking advantage of pandemic-related content to propagate attacks.

Before we dive into those larger trends, let's take a look at the overall threat landscape and what's changed between 2019 and 2020.

Today's threat landscape

Phishing is on the rise as second most common threat.

Top cyberthreats found on the Cisco Umbrella Network

2019 ▸▸▸▸ 2020

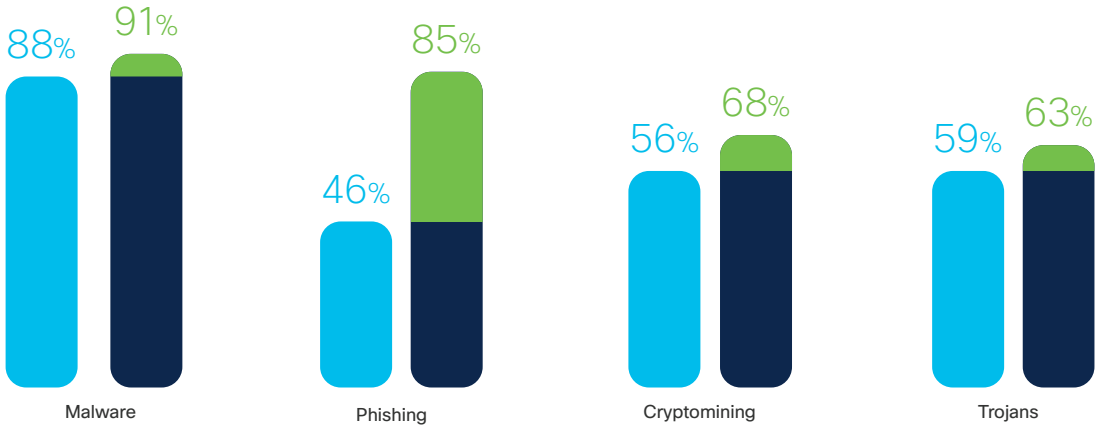

Looking at the broad threats Cisco Umbrella's customer base encountered in the first nine months of 2020:

91%
of customers saw a domain linked to **malware**.

85%
saw a domain linked to **phishing**.

68%
saw a domain linked to **cryptomining**.

63%
saw a domain linked to **trojans**.

Since 2019, trojans and phishing have swapped spots – in 2019, trojans were the number two threat at 59%, and phishing was in fourth with 46% impacted. Over the past year, phishing has risen by nearly 40%. Why the shift? One reason is related to the pandemic – a huge increase in malicious phishing sites preying on people's fears about the virus.

Malicious cryptomining is once again the top form of attack.

Top cyberattacks on the Cisco Umbrella Network

2019	2020
1. Cryptomining	1. Cryptomining
2. IcedID	2. Emotet
3. Emotet	3. Ursnif/Gozi
4. Locky	4. COVID-19
5. Necurs	5. IcedID
6. Andromeda	6. Andromeda

In terms of the specific attacks and infiltrations we've seen so far in 2020, it's no surprise that cryptomining is #1, having taken the top spot last year as well. Why does cryptomining continue to be the most prevalent form of attack? Key reasons include:

- The relative ease with which criminals can monetize their activities
- Cryptomining's relatively small footprint is easier to hide
- A misperception that cryptomining is less dangerous than other threats – more a nuisance than a real danger to your environment (More on this later in the report.)

Threat traffic by industry vertical for 2019 and 2020

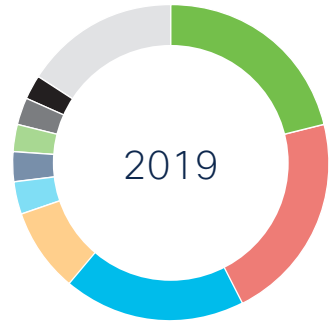

1. Higher Education - 21.38%
2. Financial Services - 21.37%
3. Managed Service Providers - 18.63%
4. Manufacturing - 8.36%
5. Professional Services - 3.42%
6. Technology: Soft/Hardware - 3.15%
7. Government - 2.81%
8. Energy/Utilities - 2.8%
9. Retail - 2.27%
10. Other - 15.81%

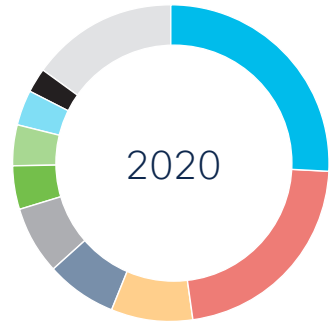

1. Managed Service Providers - 25.98%
2. Financial Services - 21.94%
3. Manufacturing - 8.40%
4. Healthcare - 7.12%
5. Technology: Soft/Hardware - 7.11%
6. Higher Education - 4.36%
7. Government - 4.17%
8. Professional Services - 3.55%
9. Retail - 2.39%
10. Other - 14.98%

Customers of managed service providers (MSPs) have become primary targets.

In terms of how threats are affecting business verticals, the distribution of threat traffic has changed slightly since 2019. In particular, managed service providers have now overtaken financial services for most impacted verticals. Indeed, there have been recent warnings by U.S. government agencies about the heightened risk of attacks by state actors on MSPs.

What's causing the jump in MSP threat traffic? MSPs are attractive targets because, if the MSP has not effectively secured their own environment, malicious actors can attack the MSP themselves first, then use hijacked remote monitoring management to go after the MSP's clients. As a result, the risk here is actually higher for customers, rather than the MSPs themselves.

Higher education traffic, meanwhile, has dropped considerably over the past year — from the #1 spot to #6 — most likely due to students being unable to attend classes in person.

The rise in malware using sophisticated hiding and evasion techniques — including steganography, LOLBins, macros, and data exfiltration tactics — has made cyber defense teams' jobs that much harder. These attacks orchestrate multi-staged end runs with callbacks to command-and-control (C2) servers designed to evade traditional endpoint defenses. Sweat equity is no longer a viable solution — today's cyber defenders must leverage machine speed and predictive intelligence to deliver scalable, adaptable protection.

So, where did the data come from? A closer look at Cisco Umbrella

At Cisco, we believe it's better to predict and prevent cyberattacks than to respond and remediate after they strike. And doing that means we need data.

Every day, Cisco Umbrella's 30+ data centers process more than 240 billion internet requests from across 190 countries. This real-time DNS data is further enriched with data from both private feeds and a handful of public ones.

With such a massive and diverse data set, our threat analysis can uncover patterns that signal malicious behavior. This analysis is based on aggregated DNS query logs paired with scrubbed and anonymized customer demographic information. Together, they give us a unique perspective on global DNS traffic, which helps us both see the trends and defend against them.

Cisco Umbrella protects against more than 7 million malicious domains and IPs, while discovering over 60,000 new malicious destinations (domains, IPs, and URLs) every day. Each node of attack infrastructure is an opportunity to identify and neutralize threat architecture before it can be used for new attacks.

Leveraging data from Cisco Talos, one of the largest commercial threat intelligence teams in the world, Umbrella uncovers and blocks a broad spectrum of malicious domains, IPs, URLs, and files being used in attacks. We also feed huge volumes of global internet activity into statistical and machine learning models to identify new attacks being staged. Tapping into anti-virus engines, Cisco Advanced Malware Protection (AMP), and sandboxing with Cisco Threat Grid, Umbrella takes advantage of intelligence from millions of daily malware samples to provide the most effective defense against malicious files.

Learn more about Cisco Umbrella, the intelligence that powers it, and the protection it provides.

Trend #1

Trojans and droppers are getting a second life as new forms of malware delivery.

Trojans themselves have never stopped being popular – they're a proven, tried-and-true attack methodology. Now, however, they're being reused in new ways, as part of an orchestrated effort in multi-staged attacks.

For instance, take Emotet – this year's #2 threat by query volume. It started as a successful banking trojan, but, with its sophisticated modular architecture, worm-like propagation, and ability to cast a wide net for victims, it quickly evolved into an even more successful delivery vehicle for malware. Or, there's Ursnif/Gozi, which is being spread as a standalone version and as a dropper. Leveraging email thread hijacking and abuse of trusted services such as Google Drive, its targeted use of varied delivery methods based on the victim has made it popular in a wide variety of attacks.

In short, when it comes to the backbone of their attacks, attackers are sticking with what works. But that isn't due to a lack of inventiveness on their part. These examples illustrate a trend wherein attackers use successful malware infrastructure as part of a larger multi-staged cyberattack chain. Why use them?

Reasons why attackers reuse malware

1. Their "Swiss Army knife" abilities allow them to deploy follow-up malware in a Loader-as-a-Service model that does further damage down the cyberattack chain.

2. Their highly distributed command-and-control (C2) infrastructure makes takedown much harder to implement.

A sample of follow-up malware that various loaders can deploy down the multi-staged cyberattack chain.

Loaders	Follow-up malware: Information stealers / Secondary loaders	Follow-up malware: Ransomware
Emotet	TrickBot, IcedID/BokBot, QuakBot, CobaltStrike	Ryuk, rEvil
Ursnif/Gozi/DreamBot	IcedID/BokBot, Valak, Dridex, CobaltStrike	Maze, LockBit, WastedLocker
Valak	IcedID/BokBot, CobaltStrike	
Zloader	QuakBot	WastedLocker, NetWalker
Hancitor	Ursnif/Gozi/DreamBot, CobaltStrike	

Emotet C2 infrastructure

Widely spread across the globe, Emotet uses compromised servers on every continent, making it much more difficult to stop than threats with just a few points of failure.

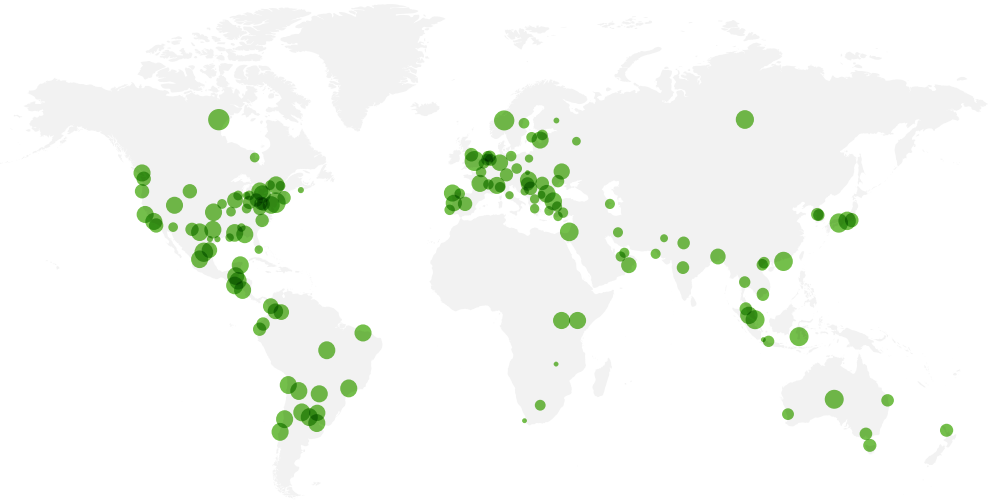

- = 1 Loader
- = 50 Loaders
- = 100 Loaders
- = 125 Loaders

As these extended kill chains become more complex, the attacker is able to use proven elements like Emotet and Ursnif/Gozi to reduce risk and coding, while focusing their efforts on the orchestrated movements that will hide their true intent. We'll cover this trend in more detail in the next section.

Trend #2

Orchestrated, multi-staged, evasive attacks are becoming the norm.

One of the biggest trends we've continued to see in the first nine months of 2020 has been the rise of complex, multi-staged cyberattacks. These attacks use new delivery mechanisms – like macros and other legitimate app functionality – to evade detection by antivirus software, hide data exfiltration actions (e.g., steganography), and coordinate multi-staged maneuvers through command-and-control (C2) infrastructure.

In the attack chain below, for example:

1. Malspam delivers an innocent-looking document that uses a macro or PowerShell (i.e., functionality embedded into the app used to open the file) leading to a loader.
2. The loader disables security controls, establishes persistence, and downloads follow-up malware.
3. When all of the targeted data has been exfiltrated, follow-up malware is launched.
4. Follow-up malware processes take control of the domain controller via SMB Exploit, resulting in network compromise.
5. This in turn leads to ransomware, which encrypts every affected component.

Sample attack chain

Malspam attacks have become increasingly complicated, with multiple levels of orchestrated actions – and multiple levels of risk as a result.

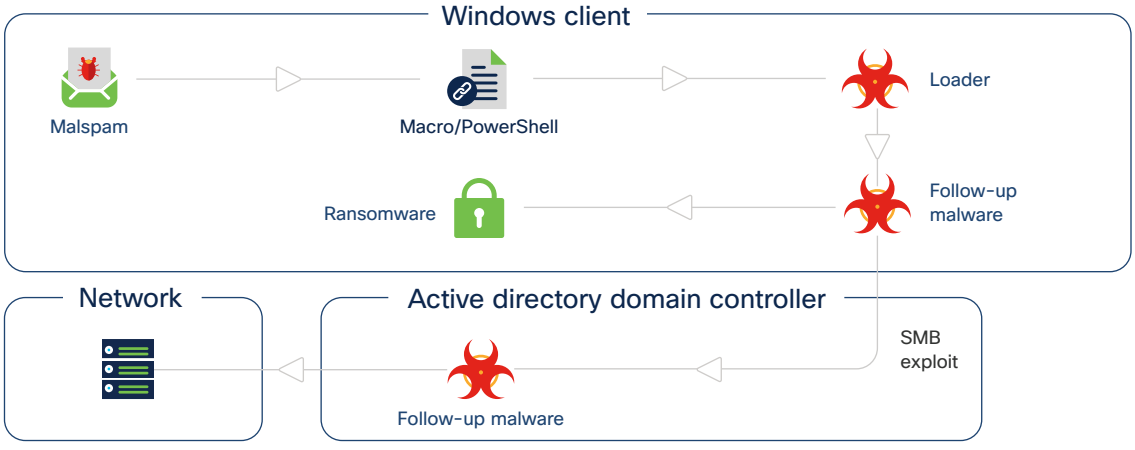

Similarly, in the phase where the file employs some sort of fileless automation – Macros 4.0, VBA, or PowerShell, for instance – the attack can make use of legitimate software automation to hide and then reveal commands. Below is an example of a Macros 4.0 exploit that uses a Binary Interchangeable File Format (BIFF) to hide an embedded Microsoft Excel file.

Example Macros 4.0 exploit

```
003c     19 CONTINUE : Continues Long Records
003c     16 CONTINUE : Continues Long Records
00ec   2182 MSODRAWING : Microsoft Office Drawing
005d     26 OBJ : Describes a Graphic Object
00ec      8 MSODRAWING : Microsoft Office Drawing
01b6     18 TXO : Text Object
003c    110 CONTINUE : Continues Long Records
003c     32 CONTINUE : Continues Long Records
023e     18 WINDOW2 : Sheet Window Information
088b     16 PLV : Page Layout View Settings in Excel 2007
001d     15 SELECTION : Current Selection
0867     23 FEATHEADR : Shared Feature Header
000a      0 EOF : End of File
remnux@remnux:~/JulyResearch$ oledump.py -p plugin_biff --pluginoptions "-x" 81
a66989b16d6b8005d23e80750031849cdfdd5beded1534b7f2d44cd4352f5 | grep Boundsheet
remnux@remnux:~/JulyResearch$ oledump.py -p plugin_biff --pluginoptions "-x" 81
a66989b16d6b8005d23e80750031849cdfdd5beded1534b7f2d44cd4352f5 | grep BOUNDSHEET
0085     10 BOUNDSHEET : Sheet Information
0085     14 BOUNDSHEET : Sheet Information
remnux@remnux:~/JulyResearch$ oledump.py -p plugin_biff --pluginoptions "-x" 81
a66989b16d6b8005d23e80750031849cdfdd5beded1534b7f2d44cd4352f5 | grep mh
remnux@remnux:~/JulyResearch$ oledump.py -p plugin_biff --pluginoptions "-x" 81
a66989b16d6b8005d23e80750031849cdfdd5beded1534b7f2d44cd4352f5 | grep yNVWfDxb
remnux@remnux:~/JulyResearch$
```

Meanwhile, attackers are continually running parallel campaigns targeting WordPress sites of small and midsize businesses, adding compromised sites to a cache of domains. These domains are loaded within the Excel files, so that when de-obfuscation occurs, it ends up downloading the actual .exe malware from compromised (yet seemingly safe) WordPress domains.

The last piece of the puzzle is exfiltration. Stolen data from the infected system needs to be sent to the attacker's infrastructure in some way – usually HTTP/HTTPS callouts to domains. However, rather than calling attention using these protocols, this new form of malware uses steganography – the practice of concealing information inside pictures – to hide the transfer of information while appearing to be part of legitimate traffic.

By eluding traditional endpoint defenses, these complex multi-staged attacks pose a much greater challenge for defenders. Raw labor alone is no longer enough to get the job done – instead, machine speed and predictive intelligence must be used to deliver proper protection.

An example of a recent steganography technique

One rare steganography technique we've seen recently hides the data in plain sight within the EXIF headers of the file's metadata – rather than in the image pixels or packet transmission.

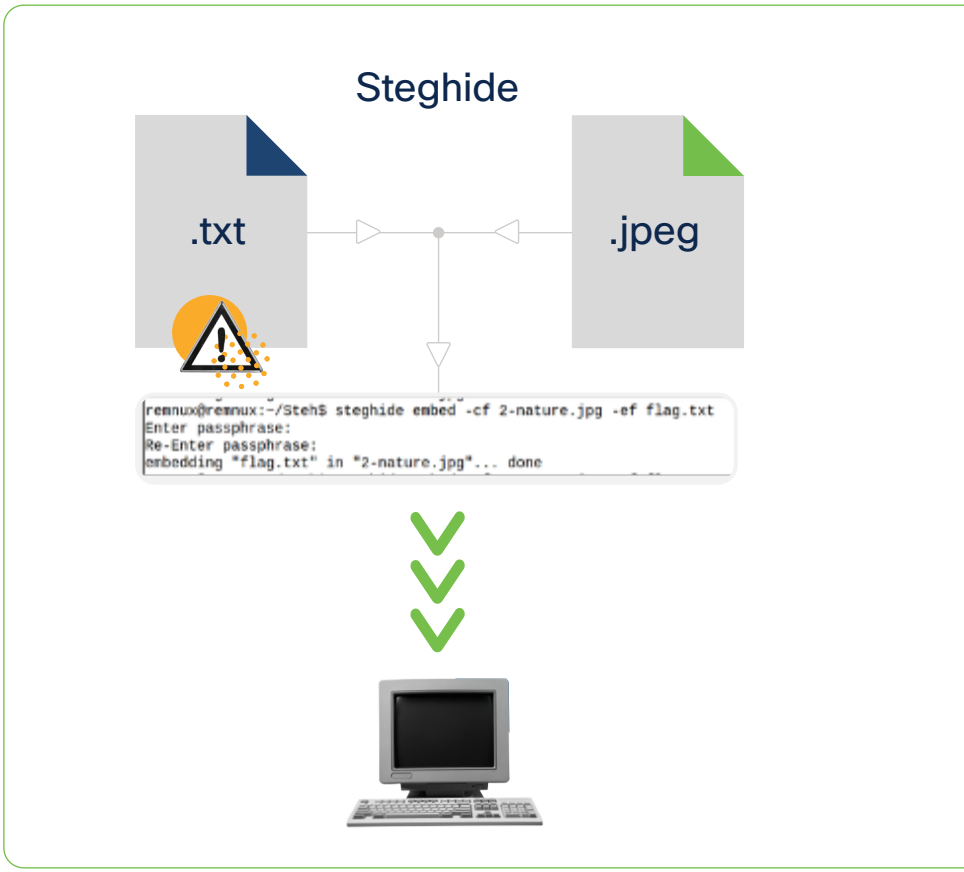

Further analysis of this data shows a search-and-replace function that exfiltrates phone numbers and system data via callbacks to a command-and-control (C2) server.

It's this multilayered approach to obfuscation that makes these new forms of attack so effective.

Combating complexity and evasion with entropy and machine learning

Analysis by Shyam S. Ramaswami

As we've seen, cyber actors are moving beyond simple and straightforward attacks into a realm of orchestrated, layered movements that are persistent, adept at hiding, and adaptive to shifting defenses. Cybersecurity professionals need new tricks to sort the hidden bad actors from the legitimate flow of data. One tool that could reveal these hidden threats is the use of entropy, assisted by machine learning.

The basic concept is that, the more complex malware gets, the less ordered the file hiding the malware becomes. This loss of order leads to entropy values that are much higher than would otherwise be expected – off-the-charts complexity that sticks out like a sore thumb. Using Shannon's entropy theory (which quantifies the amount of information contained in a variable), Python scripts, and some logic tweaking, a solution could be developed to take advantage of this phenomenon.

By applying machine learning to historical data records – system information, file URLs, passwords, etc. – automation could make it easier to flag the more anomalous files for further investigation. Working against a database of normal entropy values, threat researchers and incident response teams could then quickly identify those files where suspicious data transfer is occurring.

There have been several studies on how entropy calculation can be conducted to determine whether a file is packed or encrypted using entropy. The same could be applied to image files. One way to do so is to look for randomness or noisy data in EXIF header or image trailers. This is where the "word entropy" just described could help. Take a look at the image on the next page.

Sample section of metadata in an image file, with entropy analysis

File Name	ce1670905817523d780e185002b3120.jpg
File Size	39 kB
File Type	JPEG
File Type Extension	.jpg
Mime Type	image/jpeg
Jiff Version	1.02
Exif Byte Order	Litte-endian (inter, ii)
Orientation	Horizontal (normal)
X Resolution	72
Y Resolution	72
Resolution Unit	cm
Software	Adobe Photoshop CS4 Windows
Modify Date	2013.05.28 16:32:49
Make	/*?
Model	eval(base64_decode('aWYkdfjowkdhhUDUej jf7823jfIDUFHZXZdfjIFU88FioLI89x798df9wD jfiILIifiDIZxmJikI1Nh3IFE30~"))
Color Space	Uncalibrated

The entropy analysis focuses on the abnormally long "Model" information, highlighted in green.

Here we see that the model number has Base64 malicious code embedded in it. Word entropy can be calculated for EXIF header values, image attribute values, and other key attributes; if we calculate word entropy for each of the attributes present in the metadata, we'll derive a base and a threshold score. Then, if we calculate the entropy score for the "Model" value, we'll see that it exceeds the threshold score and stands out as suspicious.

Word entropy analysis could be a promising piece of the puzzle in detecting threats, but it can't do the job alone. Entropy data must be used in conjunction with other attributes to effectively determine whether a file is good, bad, or ugly. Instruction-calling patterns, string similarities (when combined with attributes like mutex similarities), and variable naming patterns inside the code could provide a lot of information on which actor is engaging the malware and what similarities these samples have in common. Another trend in malware analysis is converting binary bitmap to grayscale image and then assigning classification models to the malware; this is touted as being able to tackle zero-day vulnerabilities with a high percentage of accuracy.

Regardless of technique, the rise of machine learning will make these innovations and advanced detection methods possible – helping us stay one step ahead of cyber actors, no matter how complex their methods.

Trend #3

Cryptomining is opening the door to other types of cyberthreats.

Of all the trends we've seen so far in 2020, the fact that cryptomining remains the top threat is the least surprising news – it's a well-documented trend across a variety of sources. However, because cryptomining is inherently chattier than other activities – i.e., it produces more DNS queries to properly sync with the blockchain network to successfully mine cryptocurrency – its strong lead in query volume over various forms of cyberattack is not as impressive as it appears. It's the consistently high volumes of DNS traffic over time that indicates a malicious party may be involved.

In addition, it's been argued that cryptomining isn't really an attack you actually need to worry about. This line of thinking, however, generally only considers *web-based* cryptomining, which only occurs while a user is on an infected web page; close the page and the threat is gone, so there's little risk of damaging hardware. Too often, though, organizations overlook *software-based* cryptomining, in which mining software installed on a machine operates any time the machine is on and connected to the Internet. Here, there is a much higher risk of damaging hardware. This could be considered an indicator of compromise (IOC).

Web-based vs. software-based malicious cryptomining

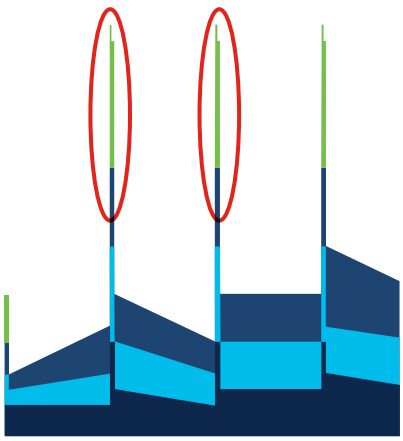

Web-Based Miners

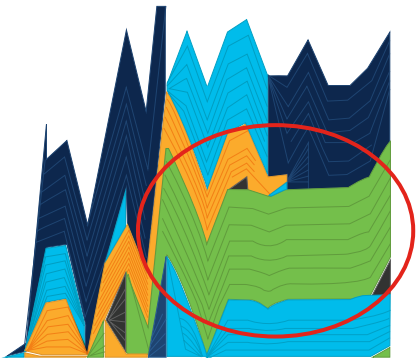

Software-Based Miners

The Prometei botnet

The Prometei botnet has more than 15 executable modules.

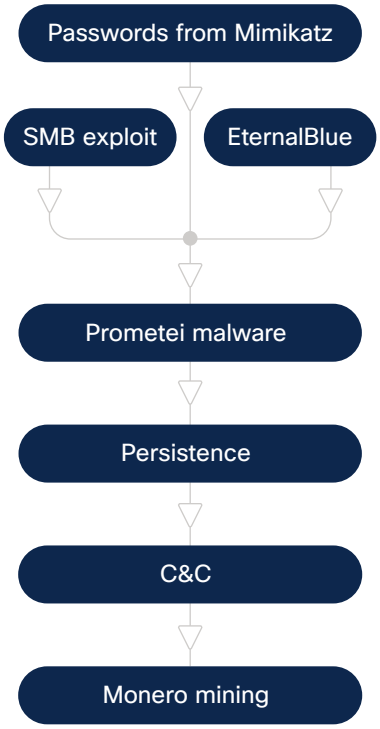

With cryptomining software running on your machines or in your public cloud, you're footing the bill for electricity, internet service, web services, and hardware replacement (if your machine wears out faster than it would otherwise). However, that might only be the beginning.

In a more recent trend, cryptomining software running in your environment can also be just the first step in a multi-staged attack on your infrastructure. Malicious third parties can infiltrate your environment, then set up a miner to make passive income while they peruse your infrastructure to exfiltrate data or perform other malicious activities.

So, what can a multi-staged cryptomining attack look like? Meet Prometei, a cryptocurrency-mining, multi-modular botnet recently discovered by Cisco Talos. Employing multiple methods to spread across a network – including SMB with stolen credentials, PsExec, and WMI and SMB exploits – Prometei drops a payload focused on mining Monero cryptocurrency for the attacker. The software then uses a variety of crafted tools that help the botnet increase the number of systems that participate in the mining.

The infection starts with the main botnet file – which is copied from other infected systems by means of the SMB protocol, using passwords retrieved by a modified Mimikatz module and exploits like EternalBlue. From there, the botnet has more than 15 executable modules that are all downloaded and driven by the main module, which constantly communicates with the command-and-control (C2) server over HTTP. At the same time all of this is happening, Prometei also tries to recover administrator passwords; any discovered passwords are sent to the C2, then reused by other modules, which attempt to verify their validity on other systems using SMB and RDP protocols.

In short, cryptomining – already the most common threat – can also expand the potential attack surface, demanding a new type of security solution that can fight on multiple fronts.

Trend #4

Attackers are taking advantage of pandemic-related content to propagate threats.

The pandemic has us thirsty for information: Where to get free testing. The latest stats on cases in our area. Updates on vaccine development. A variety of emails and articles. Unfortunately, malicious actors have taken advantage of our need for news on the topic to set up numerous sites to phish for credentials and drop malware – often mimicking content from the CDC, ECDC, and other health and government authorities.

An example of a Dridex malspam email sent out during the pandemic

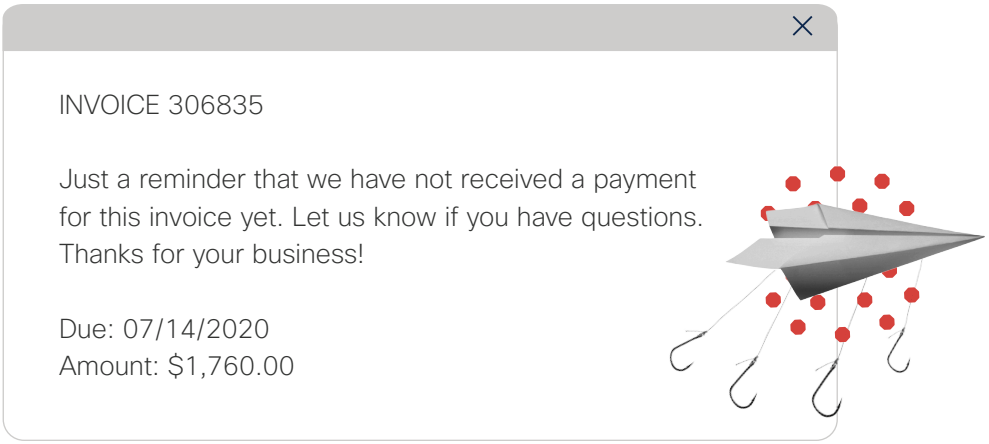

Pandemic-related traffic on the Cisco Umbrella Network by region

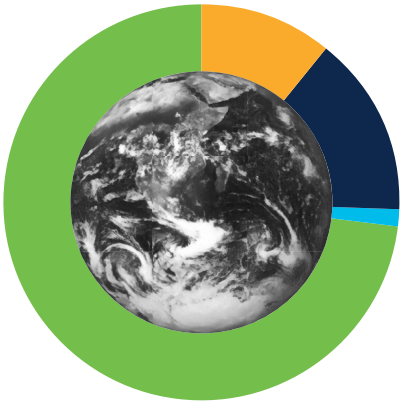

72.97%	North America
10.84%	APJC
14.89%	EMEA
1.30%	LATAM

Jumps in malicious pandemic query traffic by geographic region

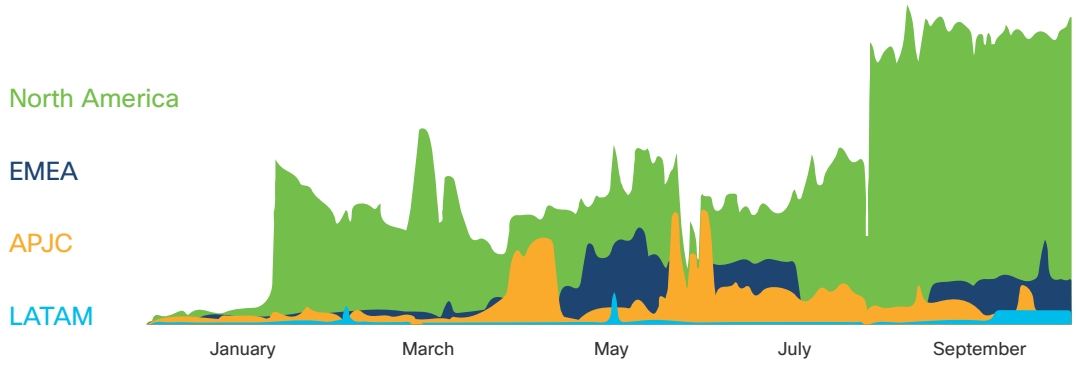

Traffic distributed by region

72.97%
North America

14.89%
EMEA

10.84%
APJC

1.30%
LATAM

- In its largest jump, traffic in North America increased 7.2x from the beginning of February through the end of August.
- The largest jump in traffic from all regions outside of North America was a 25.3x increase from the beginning of March through the end of August.

In the face of these new and rising pandemic-centered threats, it's more important than ever to ensure that all web traffic is inspected and protected, so people can get the information they need – and not the malware they don't.

Summary

The world is changing, threats are changing, and you should be changing too.

With the rise of the pandemic – and the accompanying move to remote work for most of the global workforce – the first nine months of 2020 have been a time of dramatic upheaval in work, education, governance, and more. The risk profile of individuals working at home is considerably different than it is with the security infrastructure provided by organizations in an office setting. The threat landscape has evolved – and will continue to evolve – but the bottom line is that malicious actors are still working hard to infiltrate your environment in a variety of new and changing ways. The pandemic may slow them down, but it won't stop them, and it shouldn't stop you either.

As attackers increasingly use complex, orchestrated attacks that leverage proven tactics, techniques, and procedures (TTPs), the cyber defender's job only grows more challenging. Throwing more time and more bodies at the problem no longer works – if it ever did. It's time to leverage novel solutions – like machine speed and predictive intelligence – to scale and adapt defenses to meet rising threats.

Here are a few recommendations for defending against these new threats:

- Implement scalable first-line-of-defense tools, like cloud security platforms and Secure Access Service Edge (SASE) solutions.
- Ensure your information security policies adhere to an internationally recognized information security management system (ISMS) like ISO 27001 or NIST.
- Employ network segmentation to help reduce outbreak exposures.
- Leverage timely, accurate threat intelligence that allows for that data to be incorporated into security monitoring.
- Fully embrace automated event sequencing and intelligent machine-generated analysis through machine learning.
- Existing Cisco Umbrella customers should enable the optional cryptomining policy in security settings.

About our experts

The data and analysis for this report were brought to you by the Cisco Umbrella Security Analytics and Research Teams. Recognized experts in the threat landscape, they identify new threats to help add to our intelligence, identify the latest attacker trends, and develop new processes and systems for locating malicious destinations. The research team takes a different approach to addressing threats; rather than waiting for them to hit, they proactively use Umbrella's predictive intelligence to spot the attackers.

Contributors to this report include:
Austin McBride | Shyam S. Ramaswami | Artsiom Holub

About our research methods

Using our massive and diverse dataset — collected across the more than 100 million users on our global network — our world-class team of engineers, mathematicians, and security researchers work to apply statistical and machine learning models that can predict what threats are coming next.

Together, the team is able to statistically score the "guilt" of domains and IPs to determine if they're part of an attacker's infrastructure. More than a reputation score (which merely focuses on the past), we analyze both historical and live data — using statistical models to automatically score and classify that data so we can detect anomalies and uncover known and emerging threats. (In addition to automated classifiers, our security analyst team also adds malicious domains to block lists.) In essence, our research is able to predict the likelihood of whether a domain, IP address, or entire ASN is going to originate an attack or pose a security threat before they can actually do so.

Cisco Umbrella

About Cisco Umbrella

Cisco Umbrella delivers the most secure, most reliable, and fastest internet experience to more than 100 million users daily. By unifying multiple security solutions into a single service, Cisco Umbrella helps businesses embrace direct internet access, secure cloud applications, and extend protection to roaming users and branch offices.

Learn more about how to protect yourself from threats:

View a live demo