

DEFENDING AGAINST CRITICAL THREATS

A 12 MONTH ROUNDUP

Election Security:

After spending over four years investigating election security, Matt Olney from Cisco Talos talks about the threats against democracy, and the spread of disinformation.

Big Game Hunting and the Evolution of Ransomware:
Talos investigates

Health Care Security in Focus:
Interview with CISO
Esmond Kane

Remote Working and the Attack Landscape



A YEAR LIKE NO OTHER



Towards the end of 2020, our team sat down to discuss how we would approach this year's "Threats of the Year" publication. I raised the question, "Does anybody want to hear about anything from 2020 now? Should we just mutually agree to let it evaporate from our memories and never speak about it again?"

Like many, I wanted to ride straight on through to 2021, in the vain hope that 2020 was simply an unlucky combination of numbers. However, the words of certain philosophers whom I thought I'd escaped after my History degree started echoing in my brain. If we never talk about our history, aren't we doomed to repeat the same mistakes?

Retrospection can of course provide greater insight into cyber threats and how they have evolved. It can also help inform strategic decision-making as organizations prepare for threats they may encounter in the future. And there's nothing more important to us than that.

So we got to work capturing the threat learnings from the most turbulent year in living memory. As difficult as this year was for us on a personal level, it was also a significant year in the threat landscape, and the two aren't entirely unrelated.

Take, for example, the situation surrounding the pandemic. As COVID-19 became the main headline across the world, an investigation by Cisco Talos found multiple threats capitalizing on the pandemic with Coronavirus lures and themes. Our researchers in Cisco Umbrella also discovered a surge in malicious sites that capitalized on the general populous' need for information.

In this publication, you'll find articles that address the ways cyber criminals sought to take advantage of the pandemic -- be it through phishing campaigns, leveraging the great migration to remote work, or even going after health care organizations


themselves. Our interview with Esmond Kane, CISO for Steward Health Care, shines a light on how COVID-19 impacted those on the security front line.

We have also seen a big evolution in ransomware over the past year. We have a terrific article written by Edmund Brumaghin, threat researcher from Cisco Talos, on Big Game Hunting attacks, whereby cyber criminals seek to monopolize a ransomware deployment by targeting backup systems, domain controllers, and other business-critical servers during a "post-compromise" phase. It's an incredible piece of research that illustrates what organizations can do to better prepare for and defend against the most unthinkable cyber attacks.

The year 2020 also saw one of the most momentous general elections in United States history. Cisco Talos spent four years conducting hands-on research into election security. In this publication, we have an interview with the leader of that research initiative, Matt Olney, to capture his thoughts post-election.

What this report is really all about is how cyber threats impact our lives on a human level -- from threats against our democracy, to our health care, to the organizations we work within.

I hope what we've pulled together is of interest and use to the cybersecurity community. Here's wishing all our readers a successful 2021 and beyond.

 Hazel Burton
Editor-in-Chief

INSIDE

4 [Remote working and how the threat landscape shifted](#)

7 [The sunburst supply chain attack](#)

8 [Election security: A conversation with Matt Olney from Cisco Talos](#)

14 [Stealing passwords with credential dumping](#)

16 [Big game hunting and the evolution of ransomware in 2020](#)

24 [RDP and the remote desktop](#)

28 [Health care security in focus Interview with Esmond Kane, CISO for Steward Health Care](#)

32 [Securing industrial IoT](#)

34 [News reel](#)

36 [How Cisco can help](#)

41 [Further resources](#)

Defending Against Critical Threats

Editor-in-Chief
Hazel Burton

Editorial Director
Ben Nahorney

Art Director
Sunny Powell

Advisory Editors
Angela Cannon
Cindy Valladares
Ben Munroe
Mitch Neff

Get in touch

 [@CiscoSecure](#)

 [facebook.com/ciscosecure/](#)

 [linkedin.com/showcase/cisco-secure](#)

 [blogs.cisco.com/security](#)

REMOTE WORKING AND HOW THE THREAT LANDSCAPE SHIFTED


As people transitioned from the office to remote work, we took a look at ways companies could improve their security posture. We discussed how you can shore up the security of older and personal devices now being used for work tasks, how to reduce your security footprint with company-sanctioned software, and ways to ensure that connections back into the company network are secure.

While we've all adapted to changes, the fact is that remote work introduces a number of security concerns that are different from working on premises. Here we'll examine some of the trends we've seen in a shifting threat landscape, including attackers who are adapting their techniques to take advantage of new opportunities.

The great migration

Before diving into what attackers are up to, let's take a look at just how significant the shift to remote work has been. To do this, we examined traffic running through Cisco Umbrella's DNS servers around the start of the pandemic to see where it was coming from, giving us a snapshot of remote internet activity. Cisco Umbrella is a cloud driven Secure Internet Gateway that provides protection from internet-based threats for users, wherever they go.

In mid-March of 2020, we saw a marked increase in remote connections. While it was interesting to see the office-based connections to Cisco Umbrella declining and remote connections increasing, even more interesting is analyzing how much the remote connections increased.



Volume of remote workers seen in Cisco Umbrella DNS traffic

Comparing the first and last weeks of March, the number of remote workers had effectively doubled. This means that IT teams have been dealing with setting up a lot of remote workers.

A topical evolution in spam

It's not news that spammers leverage the latest big stories in the news to help spread their wares. The pandemic has been no exception.

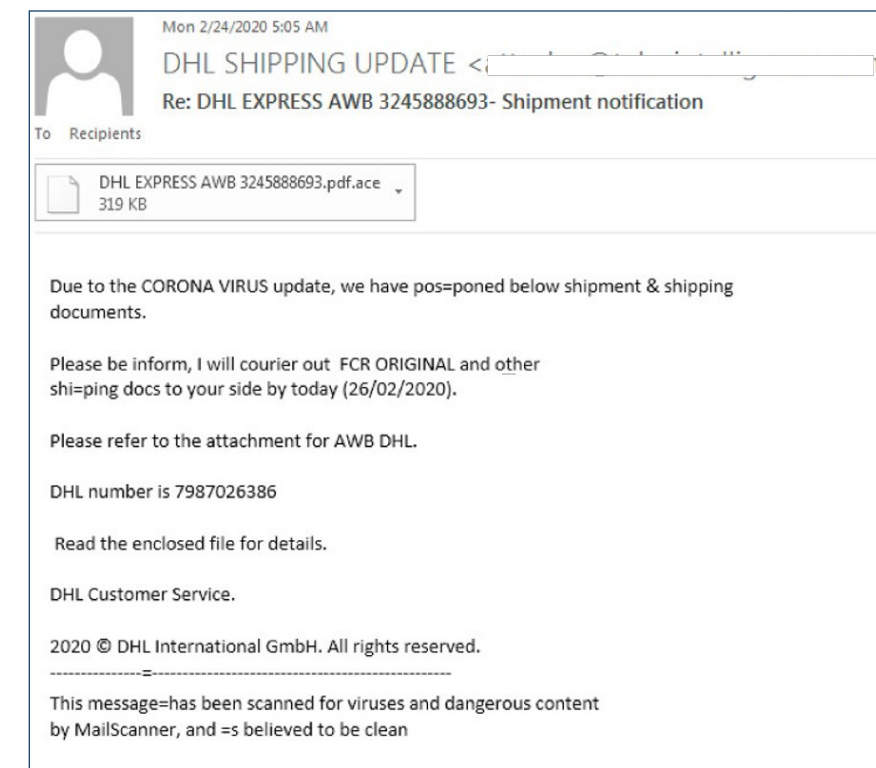
As reported by Talos on a number of occasions, threat actors have used it in a wide variety of malicious campaigns.

Some campaigns have sent out malicious emails that appear to share government information on the pandemic, while others claim to contain information regarding government stimulus payments.

This shift to pandemic-related campaigns is so pronounced that malicious spam campaigns focusing on package delivery have even pivoted to claim that deliveries have been postponed due to the pandemic.

The volume of pandemic-related spam campaigns is noteworthy. To determine just how much spam contained pandemic-based themes, Cisco Talos looked at distinct emails sent out earlier this year that contained the terms "pandemic," "COVID-19," and "corona."

While emails containing these keywords first began to grow in early February, there was a clear increase in mid-March when the pandemic was constantly in the headlines, and coinciding with the migration to remote work discussed above.



Package delivery spam with pandemic theme




Percent of observed emails tracked by Talos containing pandemic themes

Malicious domains

In early April, researchers from the Cisco Umbrella team looked at the increase in malicious domains that bad actors were leveraging to carry out attacks. According to the researchers, on March 19th, enterprise customers connected to 47,059 domains that contained “covid” or “corona” in the name. Of these, four percent were blocked as malicious.

By May 19th, this number had increased to 71,286 domains, where 34 percent of them were blocked as malicious.

Despite this being a marked increase from March, late April appeared to be the point where the most malicious activity took place. During this time the percentage of domains blocked as malicious frequently crossed 50 percent, even peaking as high as 75 percent.



Percentage of pandemic-related domains flagged as malicious

How to protect yourself

With the shifts in a corporation’s work processes, as well as the shift by attackers, advanced protection mechanisms must be implemented. Some ways to protect your organization include security audits, patching, provisioning for secure remote access, and promoting scam awareness. There’s no doubt that it’s been a scramble for many organizations to transition their workforce from an in-office to remote setup while keeping the organization safe. Fortunately, the following Cisco solutions can help.

Cisco AnyConnect Secure Mobility Client simplifies secure access to the company network and provides the security necessary to help keep your organization safe and protected.

Cisco Secure Access by Duo enables organizations to verify users’ identities and establish device trust before granting access to applications.

Cisco Umbrella’s cloud-based services can protect users from malicious Internet destinations.

Cisco Secure Email provides protection for email against phishing, business email compromise (BEC) and other cyber threats.

To protect those new devices that have been added to the network, as well as existing devices, there’s Cisco Secure Endpoint, which blocks malware as well as detects, contains, and remediates advanced threats.

Cisco SecureX is a built-in platform enabling coordinated defense against advanced threats to the business.

Learn more

For further information on how to sign up for these offerings, check out our Cisco Secure Remote Worker page.

This article originally appeared as two ‘Threat of the Month’ articles which can be found at blogs.cisco.com/security/remote-work-threat-landscape



THE SUNBURST SUPPLY CHAIN ATTACK

Work in cybersecurity long enough and it seems that big security stories seem to disproportionately break on Fridays and in December. In other words, right before a break. Whether or not this observation is quantifiable, anyone in the security industry who spent the end of last year working or on call, while colleagues took a much-needed break, knows the feeling.

2020 was no exception, with news of a major supply chain attack emerging in early December. The attack campaign, which primarily leveraged compromised updates in SolarWinds’ Orion software, has compromised many high-profile targets, covering both the public and private sector. The details surrounding the attack, widely referred to as “Sunburst,” slowly came to light in the weeks following its discovery, on through the holiday break, and into the new year.

As with most supply chain attacks of this nature, not all compromised networks become targets. Ultimately, the attackers make a decision if

the compromised network is one they value or not, which is a highly subjective decision influenced by the bad actor’s particular modus operandi.

Since this is an ongoing investigation at the time of this writing, please refer to Talos’ threat advisory for further details concerning how the supply chain attack was carried out, some ways attackers have been seen pivoting once they gain access to a network, and the sophisticated methods they use to hide their communications. The blog also contains an updated list of indicators of compromise that you can use to check if your network has been impacted by this attack.

In the end, where the attackers go and what they do on a targeted network largely depends on the entity in question. This unfortunately pushes forensic analysis into a case-by-case basis, and it likely made for a lousy Christmas present and a shortened end-of-year break for security teams working on impacted networks.

Learn more

For more information about the Sunburst supply chain attack, be sure to check out the following resources:

- <https://blog.talosintelligence.com/2020/12/solarwinds-supplychain-coverage.html>
- https://tools.cisco.com/security/center/resources/solarwinds_orion_event_response
- <https://www.cisco.com/c/en/us/products/security/solarwinds-rapid-response.html>

ELECTION SECURITY: A CONVERSATION WITH MATT OLNEY

In July 2020, [Cisco Talos](#) published the first of three in-depth papers on election security. The paper, “[What to expect when you’re electing](#),” and the follow-on materials around disinformation campaigns are the result of four years of hands on research conducted by Talos. It was led by Talos Director of Threat Intelligence and Interdiction, Matt Olney.

The research process included building relationships with local, state, and national officials, performing independent research, and even watching one state plan an election in real time.

Hazel Burton, Editor-in-Chief, sat down with Matt to discuss his thoughts on election security in the aftermath of the United States general election in November 2020.

Post-election, what are your immediate thoughts on how things panned out from a security point of view?

Before the election, our assessment was that the goal of most foreign election interference was to attack the idea of democracy. Adversaries were, and are, essentially trying to inhibit the abilities of democratized nations to project power in the geopolitical space.

“A conversation with an election official in 2020 is fundamentally different than how it would be in 2016.”

They do that by targeting the faith of voters. So that if aspersions are cast about the way someone is elected to govern, they potentially then become a weakened governing power, because half of the population doesn’t believe that they were properly placed. It’s not really about who’s the President, it’s about making whoever is the President look as weak as possible.

Sadly, this is not just a foreign adversary problem. I’ve come to realize that disinformation is also a domestic problem. And we’ve seen domestic disinformation occurrences on both sides of political issues. We’ve seen it directly from campaigns, and also from a variety of difficult-to-trace actors outside of the campaigns.

My impression of 2020 from a disinformation point of view, unfortunately, is that our adversaries could have just sat back. Everything that they would have wanted to happen, happened, purely from domestic actors.

However, the good news is that when it comes to defending against election infrastructure interference, the whole thing has become an accelerated security maturation process. A standard state or local government agency, who might typically be under resourced and under prepared compared to a high-end enterprise

customer, understood in 2016 that they were the target of some of the most sophisticated actors on the planet. By 2020, they had the resources to defend themselves.

The vast amount of credit has to go to the state and local officials. Because ultimately, they’re the ones that took the learnings from 2016 and had to reimplement, upgrade, reconfigure, patch, and build procedures, etc.

As a result, a conversation with an election official in 2020 is fundamentally different than how it would be in 2016. Gone are the times where I would say, “Let me tell you about this threat,” because they’ve spent the last four years learning about those threats. I would have to do a lot more work to provide value today than I had to in 2016, which is what you want at the end of the day.

That’s why it’s so frustrating about the scale of disinformation.

Can you talk about the elements of the hands-on research you did, including the tabletop exercises with election officials in the State of Mississippi?

Yes, that was our second visit to Mississippi, so it was great to be invited back! I’m really proud of the relationship we have with the states that we worked with, and I think they get a lot out of the relationship as well.

We were lucky enough to have Liz Wadell from the Cisco Talos Incident Response Team come with us. She designed the tabletop scenario, having spent a few days getting to know the team and the systems they used.

“We had become certain that this was a battle of ideas.”

Part of this tabletop exercise was to presume some level of compromise, without proving that there was a vulnerability. We also gave them a scenario that pulled in all the various participants and personalities at the state’s security office.

For example, we made sure that the Head of Communications and the Head of Legal were part of the exercise. Because by this point, we had a real concern about the spread of disinformation. We had become certain that this was a battle of ideas, more than it was a battle over actual votes and the counting of those votes.

And so it was great that the people involved in responding to those claims and answering questions by reporters were also involved and knew every step. It was a very successful, highly customized exercise.

So what you’re saying is that the solution for dealing with election security attacks isn’t just technical, it’s also about communications and public relations? It’s about the words, the messaging, and how you communicate...

Yes, definitely. The most important moment of any potential election security incident is the moment when the Secretary of State steps up to the podium and begins to talk about what happened. This is when they can either handle or mishandle the narrative. As I mentioned earlier, it’s about voter perception, their faith in democracy, and protecting the institutions.

We talked with the states about leading with their values, stating their intentions, and being absolutely truthful about what’s going on.

If they stumble over the facts, even in good faith, if they make a mistake or leave something critical out, the adversaries can use that failure in further disinformation campaigns, and call into question the efficacy of the response.

“When a lesson is learned now, it’s propagated much more quickly than it would have been in 2016.”

Can you summarize the main things that were done to strengthen election security between 2016 and 2020?

As well as the outstanding work by the election officials to put in increased procedures and processes for every possible outcome, there were several steps that the federal government took to appropriately define its role in protecting democracy.

In 2018, the Elections Infrastructure Information Sharing and Analysis Center (EI-ISAC) was created. This allowed a set of network sensors and network flow monitors to be available at low cost to any state to protect election resources. Those are centrally monitored by the multi-state EI-ISAC.


There was also a successful designation of elections as critical infrastructure, which allowed for some increased focus.

And the Cybersecurity and Infrastructure Security Agency (CISA) was created, which became the point administration in the federal government for election security, led at the time by Director Chris Krebs.

His leadership was critical; he was definitely the right person at the right time in the right role. And it’s one of those situations where leadership really matters. He didn’t overextend, he understood the limited federal role. And he had the brilliant assistance of Matt Masterson, who is the key contact point between the Secretary of State’s offices and CISA for election security issues.

There’s now an entire infrastructure for moving information around from states to the federal government, and from the federal government to the states. Talos participated in some briefings through DHS where we talked to the entirety of the nation’s election community that was dialed into the phone call. When a lesson is learned now, it’s propagated much more quickly than it would have been in 2016.

To give you an example of the increased vigilance in the election community, one of our press contacts reached out to us after Governor DeSantis of Florida couldn’t vote because his voter registration had been altered, asking us to comment on it.



The inverse pyramid relationship between resources and threats faced by federal, state, and local agencies
From What to expect when you’re electing: What Talos learned after 4 years of research and hands-on experience

I read about what happened and realized that someone had impersonated the governor by having the correct information, and then made changes to the governor’s data. But it was spotted because of the IP address, and it was all turned around very quickly.

The fact that they had those controls and logging in place speaks to the increased vigilance that we as a nation have. I’m not saying for sure that this wouldn’t have happened in 2016. But what I saw in 2020 was the ability to see more of what’s happening on these networks.

“The Rumor Control website is an extension of the nimbleness and agility CISA has to create what’s needed.”

Another example is the [rapid response from CISA](#) to the Iranian actors that spread disinformation pretending to be the Proud Boys, an alt-right group. They were attempting to intimidate Democratic voters against voting for Joe Biden.

CISA saw the emails, and in about 24 to 36 hours, came back and was able to say who it actually was (i.e., Iranian actors), which is unprecedented in terms of speed of federal response.

That speaks to both the dedication of the people at CISA and the cooperation and support of other agencies like NSA and the FBI, as well as the preparation that has been done over the previous four years.

Director Krebs also helped to create the website ‘Rumor Control’ which seeks to disprove the aspersions that are being cast around the integrity of the election and election security...

Yes, and I think it’s fantastic. CISA is the agency that came out with the pineapple on pizza analogy to illustrate disinformation. It was very playful and creative, but also very much an apolitical example of disinformation.

The Rumor Control website is an extension of the nimbleness and agility CISA has to create what’s needed. One of our recommendations that we put in our disinformation paper, and which our PR team helped to put together, was to have a mechanism/website to do exactly this – to provide accurate, evidence-based information. We meant for those recommendations to be for state and county officials, but it works even more so at the federal level. The Rumor Control website is now a national repository for evidence against disinformation.

Director Krebs was also active on Twitter, speaking directly against disinformation such as

the “Hammer and Scorecard,” supposedly a CIA secret supercomputer that can hack election computers. It was completely bogus, extremist rhetoric, and they got called out for it.

What I think is also great about it is that other governments can build on programs like this. It allows for even the poorest county in the United States to have a resource that they can point their constituency to for information.

Is there anything that took you by surprise at all when you were researching disinformation campaigns?

First of all, I must give credit to Matt Valites, Kendall McKay, and Nick Biasini on my team, who did the bulk of the work on this.

Nick highlighted the emerging commercialization of disinformation campaigns. We've seen elements of this in the past decade, where specialists could be hired to pull some really dirty tricks, but nothing like what we've seen in the past few years.

It's now disinformation-as-a-service. If you want to spread fake news, there are organizations who will help you do that. That speaks to a very difficult future, but on the other hand, what's different in 2020 from 2016 is that generally we're a lot quicker to say, “That's a bot,” or “That's disinformation.”

People tend to be much quicker to pick up on red flags, and they help to invalidate things that don't necessarily line up with reality.

What's the biggest takeaway you have from all this research?

In every county in the United States, from the poorest to the richest, there are real, genuine, and talented people who run the elections.

It's the same people you pass at the grocery store, or stand in line with at the bank. There's no place in the United States where the elites run the elections. And these people have spent the last four years making things better.

I can't give them enough credit for what they've done, particularly when they've been treated as poorly as they have been in the days and weeks after the election. Election officials are incredible patriots and a fundamental part of democracy. They've done an outstanding job.



Here is an overview of all the election research that Talos published in 2020:

What to expect when you're electing: Talos' 2020 election security primer

The first research paper in our series, this outlines our major findings after spending months researching election security and talking to experts and election officials.

What to expect when you're electing: The building blocks of disinformation campaigns

Insights into how threat actors start, maintain, and grow disinformation and fake news campaigns.

What to expect when you're electing: Information hygiene and the human levers of disinformation

Why disinformation works on voters, and advice for detecting intentionally misleading blog posts, memes, and social media posts.

Roundtable video: Disinformation and election security

Easy-to-share tips to avoid disinformation campaigns and stay sane during the election cycle. These apply to anyone—from security experts to novices.

What to expect when you're electing: How election officials can counter disinformation

Advice for local, state, and national election officials on how to keep Americans' faith in the election process and counter disinformation efforts.

STEALING PASSWORDS WITH CREDENTIAL DUMPING

From a malicious standpoint, stealing and using legitimate credentials to gain access is more likely to go undetected as an attacker attempts to move through a network. Dropping a trojan or exploiting a vulnerability can certainly gain you initial access, but authorized credentials help you navigate laterally under the radar.

According to Verizon's 2020 Data Breach Investigations Report, using stolen credentials is the second most common activity conducted by attackers during a breach. Here we'll look at the practice of 'credential dumping.'

What is credential dumping?

Credential dumping is a technique whereby an attacker scours a compromised computer for credentials in order to carry out further attacks. The fact that this is an unfamiliar attack method only underscores the importance of understanding it better.

There are a variety of places within operating systems where credentials are stored. If an attacker can gain access to those parts of the system, they can attempt to copy and "dump" the credentials.

Credential dumping is possible because software and operating systems store passwords in memory, databases, or files. The operating system will initially request a password, but then use the cached password for successive logins,

```
.#####. mimikatz 2.0 alpha (x86) release "Kiwi en C" (Apr 6 2014 22:02:03)
.## ^ ##.
## / \ ## /* * *
## \ / ## Benjamin DELPY 'gentilkiwi' ( benjamin@gentilkiwi.com )
'## v #' http://blog.gentilkiwi.com/mimikatz (00.00)
'#####' with 13 modules * * */
```

An image of the header when Mimikatz is loaded

saving the user from having to enter it again.

Tools of the trade

Problems arise when an attacker gains low-level access to a computer. Credentials can then be harvested with various credential dumping tools. Although there are several tools an attacker can wield to steal credentials, Mimikatz, which was first released in 2007 by Benjamin Delpy, is the most popular. Its purpose was to highlight a flaw in the Windows LSASS process, which stores credentials to streamline access to system resources. The flaw in question was eventually fixed, yet Mimikatz has evolved to become a dual use tool, used by both security professionals and malicious actors.

Using the credentials

Once an attacker has gathered credentials, how do they use them? It's pretty straightforward when it comes to user names and passwords that have been stolen through phishing or keylogging, or have been stolen and successfully decrypted.

However, not all credentials can easily be decrypted. There's a group of attack techniques centered around using these credentials as-is.

For instance, consider that many user names and passwords are encrypted (a.k.a. "hashed") on the authenticating server. When you log in, they generally decrypt the password on the server and compare them. Another way to compare is to encrypt the password that arrives, then compare it to the encrypted password on file. Either way, if there's a match, access is granted.

“Mimikatz has evolved to become a dual use tool, used by both security professionals and malicious actors.”

If an attacker manages to steal user credentials, but can't decrypt them, they can attempt to pass them to the authentication server. If the server simply compares the two hashed passwords, and if they match, access is granted. This technique is often called "passing the hash."

There are a number of similar authentication attacks. For instance,

an attacker could also dump Kerberos tickets from a compromised system, then use them to attempt to log in. As a variation of the overall theme, this attack is called "pass the ticket."

“Not all credentials can easily be decrypted.”

There are plenty of variations out there. An attacker can "overpass the hash," by which they pass a hash to an NT LAN Manager in the hopes that it will pass them back a Kerberos ticket, which they can then use to log into network resources. There are also techniques that can grant them "golden" and "silver" Kerberos

tickets, which as the names suggest, offer elevated privileges and access throughout a network administered by Kerberos.

What to do

Fortunately, there are many ways to defend against credential dumping:

- Monitor access to LSASS services and SAM databases.
- Watch for command-line arguments used in credential dumping attacks.
- On domain controllers, monitor logs for unscheduled activity. Look for unexpected and unassigned connections from IP addresses to known domain controllers.

Cisco offers a broad solution set for monitoring and protecting your environment against credential dumping attacks. Products such as Cisco Secure Endpoint, Cisco Duo, and Cisco Identity Services Engine can help to keep your environment secure. To learn more, visit the Cisco Secure homepage.

This article originally appeared as a 'Threat of the Month' article which can be found at blogs.cisco.com/security/stealing-passwords-with-credential-dumping



BIG GAME HUNTING AND THE EVOLUTION OF RANSOMWARE IN 2020

By Edmund Brumaghin
Threat Researcher, Cisco Talos

Ransomware attacks have evolved from a threat that many organizations once considered a mere nuisance, to one that often leads to widespread business disruption, disclosure of sensitive information and reputational damages – all at the same time.

One of the biggest trends of 2020, which affected organizations around the world, was the widespread adoption of new tactics, techniques, and

procedures (TTPs) related to the deployment of ransomware on corporate networks, as referenced in [Talos' Quarterly Incident Response Trends](#) reports.

Over the past year, threat actors have increasingly adopted new strategies such as big game hunting and double extortion to maximize their attacks and force organizations to pay their ransom demands.

Let's take a closer look at these attacks and how they have evolved.



What is big game hunting?

Historically, most ransomware attacks focused on individual systems within organizational environments. Once compromised, data stored on the system is encrypted or otherwise made inaccessible, and the victim is delivered a ransom note with instructions for paying the ransom to regain access to their information.

Figure 1 is an example of one of these ransom notes, which is displayed to the victim following a successful ransomware infection.

The entire business model associated with ransomware relies on the attacker successfully coercing the victim into paying. If victims choose not to pay the ransom, attackers do not generate revenue, which is the primary motive in most ransomware attacks.

As the use of ransomware gained in popularity among threat actors, many attackers refined their approaches to maximize the effectiveness of their extortion attempts. An example of this can be seen with the introduction of countdown timers along with threats of the permanent destruction of data unless ransom payments are received quickly. Other ransomware experimented with the use of iconography designed to elicit fear in their victims.

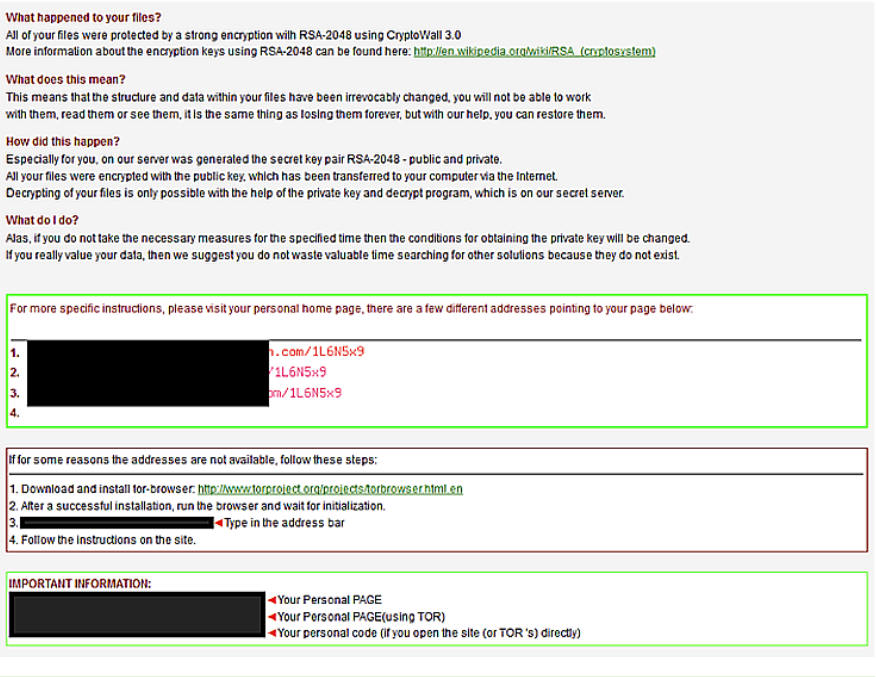


Figure 1 | An example of a ransom note

In most corporate environments, it takes minimal effort from IT staff to reimage an infected system, which limits the effectiveness of this type of attack. To overcome these limitations, attackers began to investigate alternative approaches.

In 2016, SamSam emerged as one of the earliest examples of ransomware targeting the health care industry. Rather than simply activating ransomware on the first successfully compromised system, the adversary behind SamSam

leveraged the compromised system as an initial access point into the network.

Once this initial foothold was established, they then moved laterally throughout the network, gaining access to additional systems and escalating privileges. Once privileged access to critical network infrastructure and systems was successfully obtained, the ransomware could then be activated on all of these systems simultaneously, maximizing the

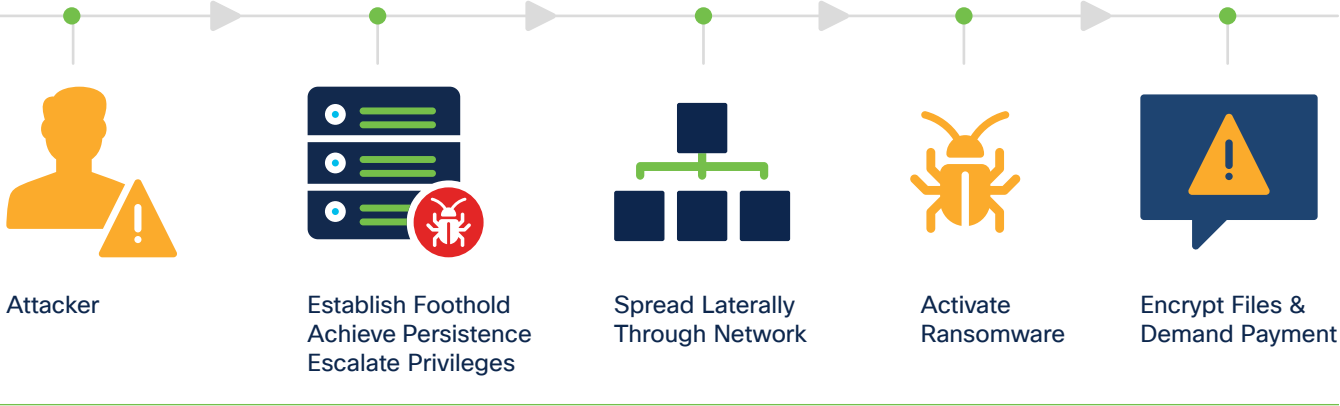


Figure 2 | A high-level diagram of ransomware activation

damage inflicted on the victim organization. Figure 2 is a high-level diagram illustrating this approach.

This new approach to ransomware deployment provides the following major benefits to attackers:

- Organizations are more likely to pay ransoms due to their inability to successfully recover from widespread business disruptions.
- The ransom being demanded can be set significantly higher, often ranging from hundreds of thousands to even millions of dollars in extreme cases, versus a few hundred for a single system.

As ransomware operators have recognized the viability of this approach to revenue generation, more threat actors have begun using this same methodology to attack organizations around the world. This approach to the attack lifecycle became known as “big game hunting” and has gained popularity across the threat landscape, with many adversaries actively targeting backup systems, domain controllers, and other business-critical servers during the post-compromise phase of their attacks.

As adversaries have gained more experience and refined their use of various attack frameworks and toolsets during the post-compromise phase of their attacks, their efficiency has increased as well.

In 2020, big game hunting emerged as one of the biggest threats that many organizations face, with new examples of successful attacks being publicized on a weekly basis.

Compromise-as-a-Service (CaaS) is here

As the crimeware ecosystem matured in recent years, new product offerings have emerged on hacking forums, darknet markets and in closed invite-only communities. These offerings cater to ransomware operators seeking to launch extortion attacks without having to first obtain initial access into the networks they are targeting.

The individuals selling these offerings have become referred to as “initial access brokers” who offer network access across a range of industries, geographic regions, etc.

“Big game hunting' has gained popularity across the threat landscape.”

Initial access brokers often seek to obtain an initial foothold in corporate environments, and may also perform the post-compromise activities necessary to escalate privileges. Once they have established this access capability, rather than attempt to deploy ransomware themselves, they then attempt to sell it to other threat actors so that it can be used in further attacks.

In recent years, online sales postings have become more frequent, with initial access brokers often listing multiple networks that they are selling privileged access to at any given point in time. For many ransomware operators, this is an attractive business model as they no longer have to worry about how to gain access. They can simply purchase it

which significantly lowers the barrier to entry and also streamlines the process of ransomware deployment.

In many cases, the price associated with this access ranges from hundreds to thousands of dollars, which is minimal considering that the ransom demands associated with big game hunting and double extortion attacks are often millions

Ransomware from the beginning

Early ransomware attacks targeted individual systems and featured ransom demands ranging anywhere from \$500 to \$1,000 to unlock the information the ransomware encrypts.

Many of the earliest examples were known as “screen lockers” and did not actually implement any encryption capabilities, simply relying on fear to convince victims to pay the ransoms demanded.

In late 2014, US-CERT estimated that approximately three percent of victims actually paid the ransom. As cryptocurrencies like Monero spiked in value, many malware distributors moved away from ransomware, instead delivering cryptomining malware to victims to attempt to generate revenue.

Over the past few years, ransomware has once again re-emerged as a payload of choice for a variety of threat actors. However, the nature of these ransomware attacks has changed.

of dollars. Figure 3 is an example of one threat actor advertising this type of network access to a variety of different types of organizations. In addition to procuring initial access on hacking forums and darknet markets, it can also be acquired from botnet operators.

Many big game hunting and double extortion attacks can be traced back to initial access gained through the use of commodity malware. Over the past several years, botnets like Emotet, Trickbot, ZLoader and others have amassed large numbers of infected systems that allow attackers remote access into the environments in which the systems are located. This access can then be used by adversaries seeking to deploy ransomware and extort their victims.

Post-compromise activities

Throughout 2020, one of the major trends we observed is an increased reliance on offensive security tools (OSTs) and “dual use tools” to conduct post-compromise activities.

Dual-use tools are applications that were primarily developed to assist legitimate users and/or system administrators, but are often co-opted by malicious attackers as well.

These tools have been featured in many high-profile intrusions over the past several years. Some examples include various components of the SysInternals Suite, PuTTY, VNC and TeamViewer.

“Rather than attempt to deploy ransomware themselves, they then attempt to sell it to other threat actors so that it can be used in further attacks.”

Attackers also often leverage the following types of tools in their attacks:

- **Native functionality** includes functionality that is natively provided by the operating systems and applications that are already present in network environments. Examples include PowerShell, WMI, WinRM/PS-Remoting, VBScript, JScript and others.
- **Living-off-the-land binaries (LoLBins)** are executables built into modern operating systems

that provide useful functionality to attackers such as retrieving additional malware, bypassing application controls, and more. Examples include MSBuild, Certutil, Bitsadmin and others.

- **Offensive security tools (OSTs)** are applications and frameworks that were originally developed to help penetration testers and red teams more effectively assess the security posture of organizations by emulating the activities of malicious threat actors.

OSTs often play a pivotal role in the post-compromise phase of the ransomware attack lifecycle. Cobalt Strike is one example of a framework that was initially developed for adversarial emulation activities that is now regularly abused by attackers as noted in Talos’ Quarterly Report: Incident Response trends in Summer 2020 which highlighted that it was observed in 66 percent of all ransomware attacks that Cisco Talos Incident Response responded to in that quarter.

Our recently published analysis of a WastedLocker ransomware attack also provides insight into how various open-source offensive security tools are used to facilitate these post-compromise activities.

Double extortion: The new hotness

The increased popularity of big game hunting attack methodologies illustrates threat actors’ continued dedication to ensuring they take whatever steps are necessary to put maximum pressure on organizations to give in to their ransom demands and make their extortion schemes more effective.

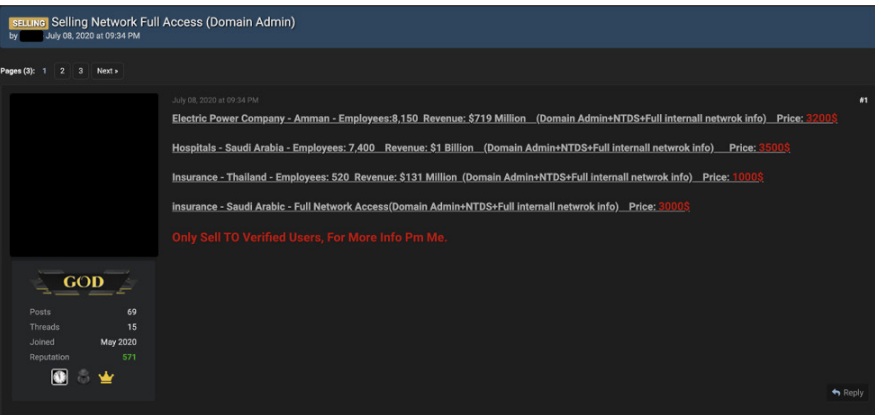


Figure 3 | An example of a threat actor advertising privileged network access

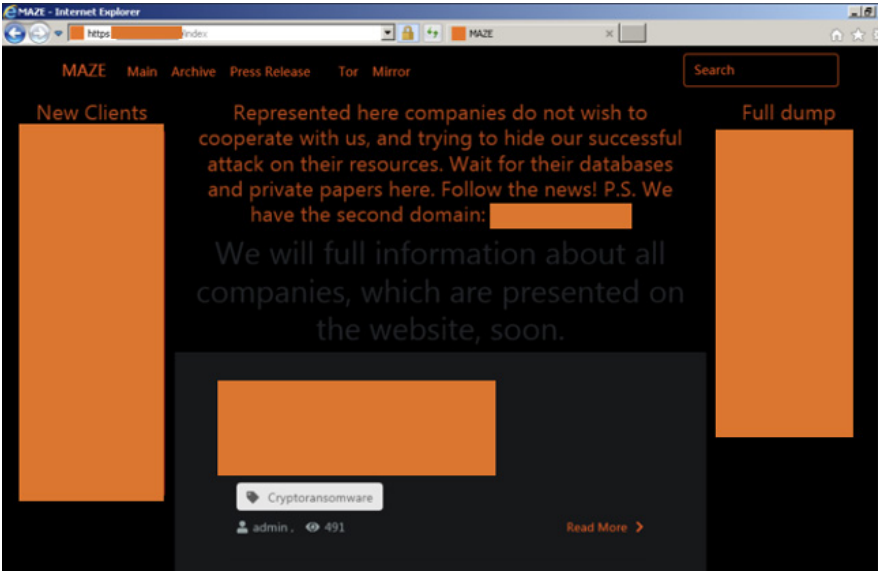


Figure 4 | An example of a Data Leak Site

Throughout 2020, we’ve observed that many threat actors are no longer content with simply causing widespread disruption to business operations. Many threat actors are now also exfiltrating large quantities of sensitive information from corporate networks prior to activating their ransomware and making their presence in the environment known to the victim. This enables them to conduct what has become commonly referred to as “double extortion attacks.”

Maze was one of the first ransomware operators that adopted this approach and was at the forefront of this evolution in ransomware. While Maze has recently announced that they are shutting down their ransomware operation, other threat actors who were inspired by Maze’s approach are still going strong.

Double extortion attacks are a one-two punch for many organizations that are successfully compromised. Not only do they have to deal with the widespread disruption caused by

the ransomware deployed throughout the network, they now also have to consider the threat of the public release of sensitive information such as intellectual property and trade secrets, financial information, or other confidential information.

“We’ve observed that many threat actors are no longer content with simply causing widespread disruption to business operations.”

Organizations targeted by these attacks also risk reputational damage and decreased customer confidence due to the public disclosure of their successful compromise. All of these factors directly influence the decision of whether or not to pay the ransom being demanded by the attacker, which is why the popularity of this approach has grown significantly over the past year.

As ransomware operators began leveraging double extortion tactics, they also began creating websites commonly hosted on servers accessible via Tor. Tor is a technology designed to anonymize communications, that are used to announce new successful compromises and leak sensitive information collected in situations where the victim either chooses not to pay the ransom, or does not pay the ransom in a timely manner. These websites are commonly referred to as “Data Leak Sites,” an example of which is shown in Figure 4.

Throughout the course of 2020, this methodology became increasingly popular with new ransomware operators adopting it as part of their extortion attempts. Even in cases where organizations successfully recover from network-wide ransomware deployment, they may still choose to pay the ransom to avoid their data being publicly released.

It is important to note that there is no guarantee that threat actors will actually delete the exfiltrated information even if the organization pays the ransom, and the same threat of information disclosure could be used in the future.

The number of ransomware operators taking advantage of double extortion continues to grow at an alarming rate with new data leak sites being established. Figure 5 shows some examples of these sites that have been set up by various threat actors.

The list of ransomware threats that are known to leverage big game hunting and double extortion as part of their attack lifecycles continues to grow. Below is a list of some of the ransomware threats that have been

observed leveraging these tactics, though this list will likely continue to expand moving forward.

- Ako

Avaddon

Clop

Conti

Darkside

DoppelPaymer

LockBit

Maze

Mespinoza

MountLocker

Nefilim
- Nemty

NetWalker

Pay2Key

RagnarLocker

Ryuk

Snake

Sekhmet

Snatch

Suncrypt

Sodinokibi

As more threat actors begin to adopt these new approaches to ransomware deployment and extortion, organizations should be aware of the multi-faceted nature of these threats and implement comprehensive security controls to ensure that they remain protected. It is vital that organizations are proactive in their approach to defending against these attacks. Once the ransom notes begin to appear, the damage has likely already been done.

What can organizations do?

The good news is that while the nature of ransomware attacks has evolved over the past year, many of the security controls that have been recommended for combatting other types of attacks in recent years remain effective against big game hunting and double extortion attacks.

In 2017, following several high-profile ransomware attacks such as [WannaCry](#) and [NotPetya](#), we published recommendations for ways that organizations can defend against these types of threats. Organizations should seek to employ a comprehensive, defense-in-depth approach to securing their environments.

Emphasis should be placed on ensuring that the organization

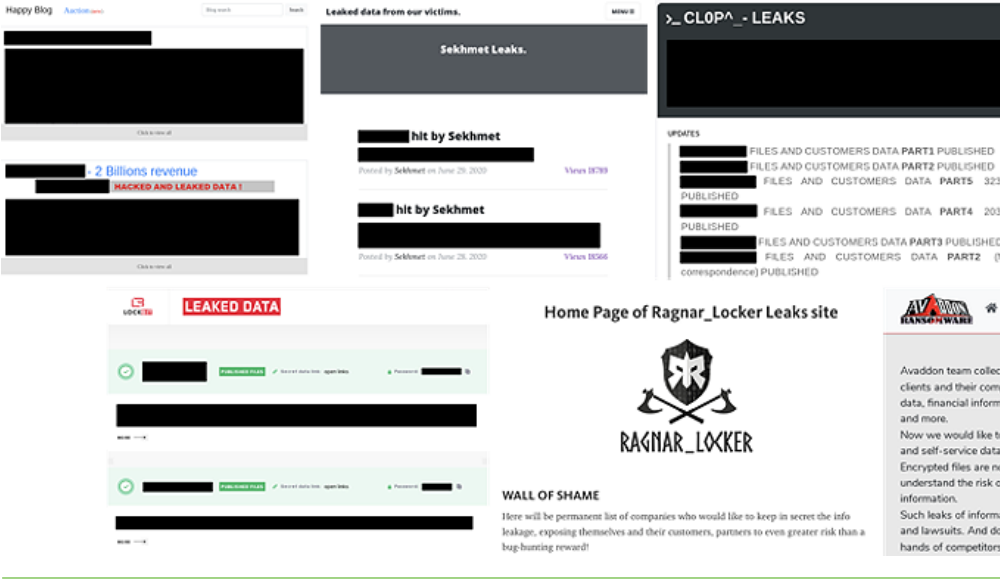


Figure 5 | Examples of sites that have been set up by various threat actors

focuses on all three phases of network defense: prevention, detection and response — such that the failure of one type of security control does not result in a complete inability to react and respond to attacks that may be encountered.

The National Institute of Standards and Technology (NIST) has published Special Publication 800-53 “Security and Privacy Controls for Information Systems and Organizations” which provides comprehensive guidance regarding recommended best practices and the selection of security controls that can be implemented to establish a sound defensive architecture within networked

environments. Revision 5 of NIST SP 800-53 was recently released and can be reviewed on the [NIST website](#).

Written by Edmund Brumaghin
Threat Researcher, Cisco Talos

Some of the basic types of security controls that organizations should consider include:

- Email security

Patch management

Least functionality

Least privilege

Systems and network monitoring
- Network segmentation

Backup and recovery

Policies and procedures

Security awareness training

How do they get in?

In the case of big game hunting and double extortion attacks, initial access to the environment can be obtained in a variety of ways. In many cases, the vector(s) used to achieve initial network access can be broken down into the following broad categories:

Web

Web-based attacks typically include the use of malvertising, drive-by downloads, or other malicious web-based content being delivered to victims via attacker-controlled websites.

Email

Email-based attacks may include malicious spam, or “malspam” campaigns that seek to entice victims to open an attachment or click on a malicious link to initiate the infection process.

Throughout 2020, email has been observed as one of the top initial vectors used in ransomware attacks, as referenced in [Talos’ Quarterly Incident Response Trends](#) reports.

Network

Network-based attacks are often conducted by identifying systems exposed to the internet that may be running software containing vulnerabilities that can be exploited by attackers to take control of the affected system. In most cases, the vulnerabilities being exploited are well-known, publicly-disclosed vulnerabilities for which exploit code is readily available online.

These threat vectors are typically used to deliver malware payloads to victims that can be used by adversaries to gain an initial foothold into the environment, establish C2 capabilities, achieve long-term persistent access, and facilitate the delivery of additional malware payloads that increase the functionality available to the attacker(s).

In some cases, network access can be purchased via online platforms like hacking forums or darknet markets (DNMs). In others, initial access may be obtained from botnet operators who have already successfully infected systems inside of the target organization using commodity malware.

RDP AND THE REMOTE DESKTOP

There are two sides to the shift to remote work. On one side, it’s about ensuring that your employees have the equipment that allows them to perform their daily tasks. On the other, it’s about providing a way for them to connect back to company resources in order to complete those tasks.

One solution to these challenges is deploying remote desktop technologies. These technologies allow a user to log into a computer remotely and operate it just as though they were sitting in front of it. There are a number of available remote desktop technologies, ranging from platform-independent implementations such as VNC, to third-party services that make remote desktop connections simple to set up and use.

One remote desktop implementation that stands out is Windows Remote Desktop Services and specifically Microsoft’s proprietary Remote Desktop Protocol (RDP). Enabling this feature and connecting to it is very easy.

RDP as a target

As convenient as it may be, RDP in particular poses some security concerns. Over the years, it has been targeted in a variety of ways. Brute-force attacks and logins using stolen credentials are a natural concern. The protocol has also suffered its fair share of vulnerabilities, allowing for man-in-the-middle attacks and remote code execution exploitation.

Then, in May 2019, a vulnerability in the way Remote Desktop Services handles RDP requests was disclosed. Dubbed BlueKeep, Microsoft noted that the vulnerability was likely “wormable” and could be used by malware to spread from one unpatched system to another, much like WannaCry had done two years earlier. In the subsequent months, two more remote code execution vulnerabilities were discovered in RDP. While no worm has materialized, RDP has gained significant attention as an avenue into a network.

Frequency of attack

Just how frequently is RDP targeted? It’s a difficult question to answer, simply because of how hard it can be to distinguish malicious from legitimate traffic. Since stolen credentials and brute-force login attempts are often used, unauthorized access will often look like legitimate log-in events.

There are a few ways we can qualify the level of RDP attacks:

Nmap is a popular network scanning utility that’s often used to check for open ports. The utility includes a list of commonly scanned ports that weights them by how likely they are found to be open. Here are the top 10 TCP ports most likely to be found open, based on Nmap’s current weightings:

Rank	Port	Service
1	80	HTTP
2	23	Telnet
3	443	HTTPS
4	21	FTP
5	22	SSH
6	25	SMTP
7	3389	RDP
8	110	POP3
9	445	SMB
10	139	SMB

RDP ranks seventh overall and is the highest-ranked proprietary port likely to be found open. What’s interesting is that only system ports used by well-known services precede RDP. The large number of RDP port-enabled systems that are exposed directly to the Internet further highlights concerns about RDP attacks.

According to data collected from shodan.io, there are over four million systems with RDP port 3389 open and exposed directly to the internet. Of course, just because a port is open doesn’t mean that attackers are targeting it.

“Unauthorized access will often look like legitimate log-in events.”

In order to investigate further, let’s take a look at data from Cisco Secure Endpoint. Here, we’ll specifically look at the exploit prevention technology included, which includes two signatures for detecting BlueKeep attacks.

The first of these signatures will alert when someone scans TCP

port 3389 to see if a system is vulnerable to BlueKeep (“CVE-2019-0708 scanning attempt detected”). The second signature alerts when someone attempts to exploit it (“CVE-2019-0708 detected”). We’re using the same methodology in the following chart to arrive at these numbers as we did in the recent Threat Landscape Trends blogs, examining the number of organizations that encountered these signatures each month.

It’s interesting to note that the type of alert shifted as the year progressed. In January, almost two-thirds of RDP alerts were exploit attempts; by June, more than two-thirds were scanning attempts.

On the surface, this appears to indicate that while attackers may have found initial success in exploiting the port, they may find it more efficient to test if a system was exploitable first.

One caveat is that within these numbers, some alerts were triggered by legitimate vulnerability scanners.

Having said that, there are a number of attack frameworks and dual-use tools, such as Metasploit, that have incorporated the BlueKeep exploit. This has significantly lowered the bar for entry when it comes to carrying out successful attacks against RDP enabled systems.

To use or not to use

Perhaps the simplest way to guard against RDP attacks is to not use the protocol. RDP is disabled in Windows by default, meaning that unless it’s manually enabled, systems are not susceptible to attack.

There’s a wide variety of less targeted remote desktop options available. If cost is an issue, there are


even open-source alternatives such as VNC that offer similar feature sets.

Nevertheless, any remote desktop solution, if compromised, grants an attacker entry into an organization through which malicious tools can be installed or privileges escalated. Considering these problems, remote desktop access may not be the best solution for everyone.

VPN solutions, on the other hand, allow remote working employees to access the resources they need, freeing up the network resources that would be used to mirror a desktop environment remotely, while also allowing for additional layers of security.

If RDP is necessary, we'd recommend connecting to a VPN first which means that clients can be authenticated to ensure security compliance. Plus, users can be authenticated by multi-factor authentication before connecting to RDP.

Understandably, in some organizations, moving away from RDP may not be an option. In these cases, RDP communication can be fortified in the following ways:



Example of BlueKeep vulnerability data from Cisco Secure Endpoint

- Do not connect RDP-enabled systems directly to the internet. Instead, use a VPN and/or proxy connection through a Remote Desktop Gateway.
- Block failed login attempts when they exceed a reasonable number. Apply policies that block certain IPs or disable suspicious user accounts.
- Use strong passwords and multi-factor authentication, such as Cisco Secure Acecss by Duo.

This article originally appeared as a "Threat of the Month" article which can be found at blogs.cisco.com/security/rdp-and-the-remote-desktop

How Cisco can help

Cisco's SecureX platform provides a number of touch points for detecting and blocking RDP-based attacks, which can be viewed within one easy-to-digest dashboard. Cisco Secure also offers a full suite of products to protect your network, including VPN and multi-factor authentication technology.

HEALTH CARE SECURITY IN FOCUS

WITH ESMOND KANE CISO OF STEWARD HEALTH CARE

Health care has long been a key target for cyber criminals. But in 2020, the COVID-19 crisis put an additional, incredible strain on the industry.

Hazel Burton sat down with Esmond Kane, CISO of Steward Health Care, to talk about the impact COVID-19 has had, and continues to have, on health care security.

What are the main cyber threats that target the health care industry?

The criticality of health care IT, the infrastructure supporting modern patient care, has never been more stark. COVID-19 has caused a staggering loss of life and we are only beginning to understand the long term impacts on population health, the economy and society at large. Now and to plan for the future, we must keep health care IT systems up and running because people's lives depend on it.

Unfortunately, the complexity and challenge to secure health care is daunting. We face all the threats of other industries, things like phishing, malware propagation, identity theft, insider threats and more. To that, we can add the targeting of health devices and the systems providing patient care, a large percentage of which is legacy and outdated technology. To compound the difficulty, cyber criminal adversaries appear unwilling to give us a break; they are more than willing to risk harming patients and to profit from misery.

Wow! You are clearly facing unique challenges. What would you say is the biggest challenge at the moment for health care security?

Steward focuses on community care hospitals and practices. Our employees are part of the local community and are impacted by COVID-19 both in work and at home. Our focus currently, is on keeping our patients and ourselves as safe and healthy as possible as we enter the winter months. Please do wear a mask, and get a flu shot!

“No one wants to slow down patient care, but at the same time, no one wants to harm a patient with immature and insecure medical devices/IoT.”

If I was to highlight one challenge, it would be around innovation and IoT in particular. We continually balance the risk from introducing new technology and all the benefits that come with that, with needing to update the old technology. No one wants to slow down patient care, but at the same time, no one wants to harm a patient with immature and insecure medical devices/IoT. It's a fine line, but the Industry must strive to embed security and risk into the culture, and that's what we certainly try to do at Steward Health Care.



Steward Health Care is the largest private physician-led health care network in the United States. The organization is headquartered in Dallas, Texas and cares for millions of people across the United States and beyond.

What has been the impact on health care security this year?

There's been a massive uptick in cyber attacks through the pandemic. It's clear that health care cybersecurity is in the crosshairs of sophisticated adversaries. We have had to continually educate our community on all the Coronavirus-related phishing scams; soon we will need to educate on the vaccine related ones!

As an industry, we've had to prioritize the business lines that are frontline in dealing with COVID-19. We have also worked hard to support our remote workforce and have done lots of related work to scale out VPNs and accelerate cloud and collaboration tools. There's also been some financial impact, necessitating some planning and budgeting exercises. Throughout, we have also needed to manage the impact on our staff, their families and communities.

For my role as CISO, I have tried to focus on picking the right strategies. It's my job as a leader to get accurate information across in the quickest and most pragmatic way possible so that Steward can make fast, informed decisions.

How did Steward Health Care react to the COVID-19 crisis?

It was amazing to work with a leadership team like Steward's. We were pioneering some of these approaches across the industry and in the United States in particular. It was a real team effort.

Steward was the first health care organization in the United States to create a facility dedicated to COVID-19, based in Carney Hospital in the northeast. We consolidated a lot of the initial efforts here, including our equipment and frontline staff. Massachusetts was a hotspot, and the work we did to ringfence the response was extraordinarily successful. It became a model for the industry and for other large healthcare delivery organizations and systems. We wanted to prepare not just the northeast, but other states and facilities across the country too.

I like to say that no business continuity or emergency management plan survives first contact. But we had a plan. And we were able to optimize it in-flight. COVID-19 hotspots can flare far and wide, so it's important that we have a framework to work from and adapt as the pandemic changes.

Do you find it hard to personally distance yourself from the context in which you're working?

I think there's definitely an element of perseverance you need when you work in this industry. It's far too easy just to see the worst of humanity, and that ransomware and phishing has had such a significant uptick in this pandemic. You can burnout and become cynical. It's true that cyber criminals are targeting the most vulnerable as well as the most prepared. They're doing big game hunting too which is a major threat towards the healthcare industry because of how even the

“Information security is “mission-adjacent.” Like Andy Ellis says, we’re a sidekick!”

well-prepared and the best-funded are sensitive to downtime and patient impact.

However, it's the job of a security professional to be measured, pragmatic, and to often be the coolest person in the room. If all you're doing is setting your hair on fire and running around and identifying threats rather than helping to mitigate, or not focusing on how you could be a business lead, well, to be frank, you're not helping. Information security is “mission-adjacent.” Like Andy Ellis says, we’re a sidekick!

So you very much focus on the positive?

Definitely. As much as you're telling people what to stop, you also need to be telling people what to start and what to continue. If you're the person who shows up and just says “no” you're not the person who says “go.”

The first rule of the Health Insurance Portability and Accountability Act (HIPAA) is to do no harm. Our colleagues in the clinical space are very good at keeping us on point to make sure that our controls are reasonable and measured, and most importantly, get the job done.

What's the biggest lesson security practitioners should take away from 2020?

The transition of security throughout this pandemic has taught us that cultural change leads to digital disruption. During these challenging times, you need to be able to communicate the value of security. And if you did take shortcuts, hopefully you can identify those for follow-up after the fact. Because shortcuts lead to sharp edges.

Sometimes it's far too easy to say that the job of a cybersecurity professional is just walking into the china shop after the bull has run through and cleaning up the mess.

It's not a mess, this is business. Health care is not about cybersecurity. It's about patients.

And it's our role to get in there and help them make sure that security doesn't enter the bedside of the patients.

Absolutely, I totally agree. Everything we do is for and about people at the end of the

day. Lastly, what would be your advice to anyone in the security industry reading this?

Cybersecurity isn't a weapon to beat people up with because they're not doing things right. Nobody does things right all the time. Educate, inspire, and motivate – that's the role of a cybersecurity professional.

I also think that we need to take time periodically to do strategic thinking. If you don't, it's far too easy to just become a firefighter and not focus on catching arsonists or building tolerance into systems so that fire doesn't spread if it does catch.

I try to spend 20 percent of the time being strategic, and 80 percent of the time dealing with the operational delivery issues and technical issues. We don't want to wait for another major catastrophe to drive these initiatives. The time is now.

Resources

[Read more about how the health care industry is under threat of trojans and ransomware with an investigation by Cisco Umbrella.](#)

[Learn about the top concerns and 2021 recommendations for health care organizations in Cisco's Security Outcomes Study.](#)



SECURING INDUSTRIAL IOT

It's hard to ignore the ubiquity of the internet of things (IoT). Even if you don't own any personal IoT devices, industrial IoT (IIoT) devices already play a part in your daily life. From the delivery of water and electricity, to manufacturing, IIoT devices are now part of many industries.

A big issue for many operational technology (OT) environments hosting IIoT assets is dealing with older industrial control systems (ICS) that have sometimes been in operation as long as 30 years. Many of these older assets make OT environments susceptible to attacks.

These legacy devices were often deployed at a time when the need for security was eclipsed by other priorities, such as high availability and performance. Patches cannot be easily applied to these devices, as the unwanted forced downtime in these environments is too disruptive not only for the organizations, but to human life. This leaves attackers with a large swath of vulnerabilities to exploit.

Getting in

The good news is that most IIoT assets aren't directly exposed to the internet, meaning attackers must rely on other methods to access them. In essence, the same techniques used in other attacks are used to get to IIoT assets.

The most common infection vector is email. A successful phishing attack against a person who has access to IIoT systems can be the most direct path to compromising those assets. Weaknesses through unpatched

systems, poor device passwords, and relaxed remote access policies also offer attackers avenues of approach.

The reality is that IIoT-specific threats are not that common of an occurrence. There are threats that have attacked general IoT devices en masse, such as Mirai and VPNFilter. And there are threats like Stuxnet, which specifically targeted PLCs. Of course such highly targeted threats are cause for concern. But it's far more likely that an IIoT device will be compromised and reconfigured by an attacker than be compromised by a trojan or a worm.

“Once in, the attacker performs reconnaissance, flagging the IIoT assets present. The attacker identifies vulnerable services, exploits them, and knocks them offline.”

Scorching the earth

Let's say an attacker sets their sights on bringing a particular business to its knees. He or she begins by crafting an enticing phishing email with a malicious PDF and sends it to HR under the guise of a job application. The employee responsible for monitoring job inquiries opens the PDF, effectively compromising the computer.

The attacker works his or her way laterally through the network, monitoring network traffic and scanning compromised systems, looking for logins and authentication tokens. Without multi-factor authentication enabled for access, they encounter few issues in doing so. The attacker eventually manages to compromise a domain controller, where they deploy malware using a Group Policy Object (GPO), successfully compromising the entire IT network.

Due to poor segmentation, the attacker manages to eventually work his or her way to the OT network. Once in, the attacker performs reconnaissance, flagging the IIoT assets present. The attacker identifies vulnerable services in the assets, exploits them, and knocks them offline. Production grinds to a halt and the business is effectively shut down.

Defense with an arm behind your back

How can you defend your IIoT assets and the OT network as a whole against attacks, especially for assets requiring high availability?

Monitoring: Network monitoring is often the most effective step you can take. However, passive IIoT traffic monitoring is also important. Active monitoring (generating traffic on the network specifically to observe its behavior), can result in an increased load on the network, causing disruptions to device performance and possible failure. In contrast, passive scanning listens to the traffic, fingerprinting what it sees rather than introducing new traffic into the OT environment.

Visibility: Keeping a current inventory of assets on the network is also a very important protection tool. This can help to identify rogue devices. And with a comprehensive list of devices, you can create policies for specific asset groups.

Segmentation: It's also very important to segment your networks. The complete asset inventory and policy organization mentioned above will help when figuring out how to segment your IIoT assets and the OT network. While this may not prevent a determined attacker from crossing the boundaries between different areas of the network, it can slow them down enough to respond in a timely fashion when this behavior is detected. Implementing zones and conduits for segmentation is discussed in ISA99 and IEC 62443.

Many IIoT assets leverage broadcast and multicast network communications. This can pose

a challenge when aggressively segmenting a network. To address this, strong dataflow mapping is also helpful towards knowing which assets are talking to each other and how they interact.

If it isn't possible to take a device offline to patch, then visibility becomes critical. An accurate network map can allow prioritization of what absolutely must be patched. IIoT redundancy can also allow you to take one device offline and bring a replacement online during maintenance cycles.

Detection: Detecting IIoT traffic anomalies is also very helpful. Look for unexpected behavior, such as two IIoT assets talking to each other that shouldn't be, unplanned firmware updates, and unauthorized configuration changes.

Threat Hunting: Finally, threat hunting is a great way to reveal threats within

your environment. Proactively looking for bad actors will go a long way in improving your security posture.

Easing the burden

Protecting IIoT assets is arguably one of the more difficult tasks in security. There are so many devices that operate in a very tailored manner and don't respond well to the disruptions caused by some security techniques.

Fortunately, there are many Cisco products that can help make protection of your critical infrastructure easier.

This article originally appeared as a 'Threat of the Month' article which can be found at blogs.cisco.com/security/securing-industrial-iiot



NEWS REEL

Throughout the year there were cybersecurity stories in the news involving the topics covered in this publication. While we talk about a number of these stories in the articles, here are a few other headline grabbers worth mentioning.

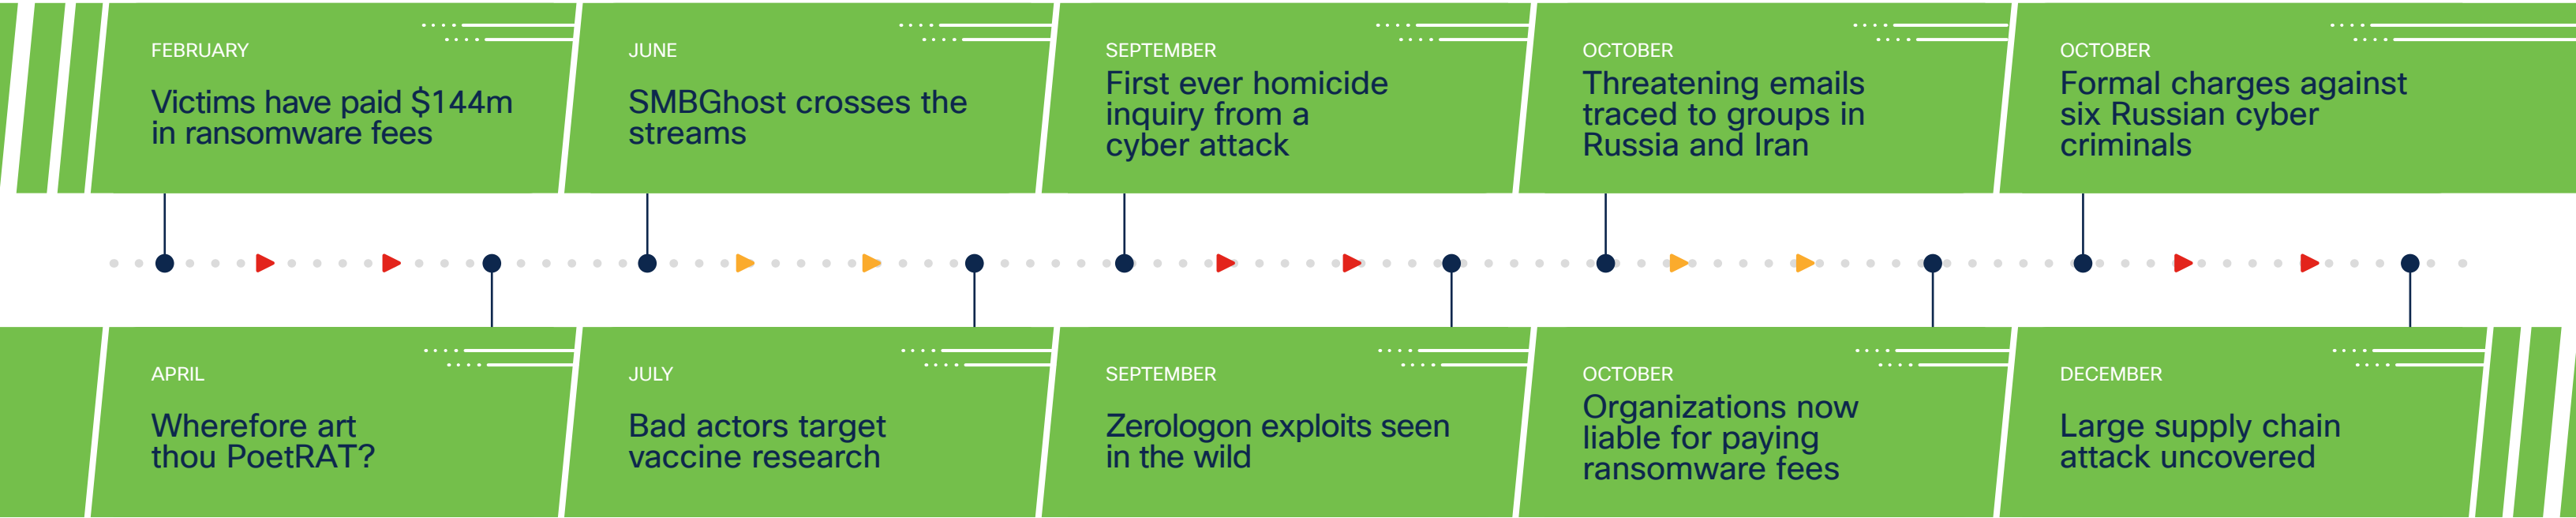
In February, the FBI reported that victims have paid US \$144 million over the last six years to malicious actors leveraging ransomware. Speaking at the RSA Conference, Special Agent Joel DeCapua also discussed how RDP was the primary vector for compromise in 70-80% of cases tracked by the FBI.

In June, The U.S. Department of Homeland Security's Cybersecurity and Infrastructure Security Agency issued a warning to all Windows 10 users to patch immediately, after a new proof-of-concept code was released that could exploit a flaw known as SMBGhost, or External Darkness. The disclosure of the vulnerability stoked fears that bad actors could create a worm similar to the WannaCry outbreak in 2017.

In September, the attacks on health care stepped up a notch when the first ever homicide inquiry as a result of cybercriminal activity was launched in Germany. Cyber criminals disabled computer systems at Dusseldorf University Hospital by exploiting a flaw in a VPN client, and a patient died when her doctors were attempting to transfer her to another facility. If the homicide inquiry finds that the woman's life was taken as a result of the cyber attack, it will be the first known case of this happening.

In October in the U.S., Democrat voters in key election battleground states were targeted with threatening emails, warning them not to vote against President Donald Trump. The emails were falsely believed to be from the alt-right group, 'Proud Boys,' but were actually traced to groups in Russia and Iran, according to the FBI. "These emails are meant to intimidate and undermine American voters' confidence in our elections," said Chris Krebs, the top election security official in the Department of Homeland Security at the time.

Also in October, the U.S. Department of Justice (DOJ) formally indicted six Russian cyber criminals who were behind some of the world's most devastating cyber attacks, including NotPetya and Olympic Destroyer. While officials have blamed Russian cyber criminals for these attacks for several years, the indictment puts charges on specific people. The DOJ worked with several technology companies, including Cisco, during the investigation.



In April, Cisco Talos discovered a new malware campaign based on a previously unknown family which we called "PoetRAT." Research showed the malware was distributed using URLs that mimicked some Azerbaijan government domains, targeting private companies in the SCADA sector like wind turbine systems. We named this malware PoetRAT due to the various references to William Shakespeare in the code. Talos observed multiple new campaigns from these threat actors over the course of the year, indicating a change in the actor's capabilities and showing their maturity toward better operational security.

In July, authorities in the UK, U.S., and Canada accused the Russian-based group APT-29 (or "Cozy Bear") of hacking into drug companies that were working on vaccine research for COVID-19. The campaign was alleged to include phishing emails sent to drug company employees, and attempted to exploit vulnerabilities in VPN products. APT-29 is the same group that made headlines in 2016 for breaching the U.S. Democratic National Committee.

Also in September, reports of cyber criminals exploiting a Netlogon vulnerability, dubbed 'ZeroLogon,' appeared, after attacks that exploited the vulnerability were seen in the wild. Netlogon is the protocol used by Windows systems to authenticate against a Windows Server running as a domain controller. Once exploited, the flaw could allow cyber criminals to take over the domain controller, and therefore a company's internal network.

We also saw an evolution in the way organizations could be liable if they paid a ransomware fee. In October, the U.S. Department of the Treasury warned companies that they could be fined for paying ransom payments to cyber criminals. This move could encourage other regulators around the world to take a similar stance.

Finally, in December a major supply chain attack was uncovered. SolarWinds, a company that produces infrastructure management applications, had been compromised and malicious code was introduced into product updates. A number of organizations that use the software, and had patched with the malicious updates, subsequently reported that they had been breached, including security companies, government agencies, and others.

HOW CISCO CAN HELP

This report covers a wide variety of security issues. From credential dumping to industrial IoT security, it can often seem as though each unique challenge requires a custom solution. The good news is that the approaches for protecting against these seemingly disparate issues aren't as varied as you may think.


Cisco Secure offers a number of security solutions that not only address the security concerns raised in this report, but can also be tailored to meet your organization's specific security needs. Our products work together through the integrated Cisco SecureX platform, a built-in experience that connects our portfolio and your infrastructure.

Let's take a quick look at the products in the Cisco Secure portfolio and see how they address the security issues raised in this report, as well as how they can work together to better protect your organization.



Cisco SecureX

Bring it all together in one place with [Cisco SecureX](#), a cloud-native, built-in platform experience within our portfolio. It can even connect with existing third-party solutions in your infrastructure, whether you want to ingest threat intelligence from Microsoft Graph Security or search Cisco SecureX threat response data in Splunk. Cisco SecureX is integrated and open for simplicity, unified in one location for visibility, and maximizes operational efficiency to secure your network, endpoints, cloud, and applications.




Cisco Secure Network Analytics

Many threats these days attempt to move laterally through the network. From big game hunting campaigns implanting ransomware, to attacks compromising industrial IoT devices, to threat actors taking over endpoints by exploiting RDP, these attacks must quietly move through your network to be successful.

Outsmart these threats with industry-leading machine learning and behavioral modeling provided by [Cisco Secure Network Analytics](#) and [Cisco Secure Cloud Analytics](#).


Know who is connected and what they are doing using telemetry from your infrastructure. You can even identify malicious intent with machine learning-based analytics across managed endpoints and unmanaged network or cloud entities.



Cisco Secure Endpoint

Every threat we've discussed in this report relies on compromising endpoints in one way or another. Attackers siphon usernames and passwords off of them with credential dumping. Ransomware locks up endpoints when activated. And it goes without saying how necessary an endpoint is for conducting work remotely.

[Cisco Secure Endpoint](#) can protect your organization's endpoints from attack, whether they're on-premises or remote. But it's so much more than that. Cisco Secure Endpoint integrates prevention, detection, threat hunting, and response capabilities in a single solution, leveraging the power of cloud-based analytics. It can protect all your endpoints, be it a Windows, Mac, Linux, Android, or iOS device.



Cisco Secure Email

Due to its ubiquity, email is the number one vector attackers use to distribute threats. Whether they're targeting remote workers, attempting to gain a foothold within an organization to deploy ransomware, or targeting a plant worker with direct access to IIoT devices, email is a quick and easy method for kicking off an attack.

[Cisco Secure Email](#) is your best defense against phishing, business email compromise (BEC), and other email-borne threats. It protects against stealthy malware in attachments and combats malicious links with industry-leading threat intelligence.



Cisco Secure Firewall

Attacking the perimeter of a network is another popular method for gaining access. Bad actors regularly look for systems connected directly to the internet with RDP enabled. Attackers often jump from the IT to the OT network in order to compromise IIoT devices.

Whatever the technique, you need someone watching these perimeters, keeping such attacks at bay. [Cisco Secure Firewall](#) has the power and

flexibility that you need to stay one step ahead of threats, detecting attempted attacks before they breach the walls of your network.



Cisco Umbrella

Many attacks rely on persuading an individual to visit a malicious webpage, then exploiting and installing malicious code on the device. Remote workers in particular, when connecting to the organization’s network from their home, are frequently targeted with emails containing malicious links.

DNS-layer security provided by [Cisco Umbrella](#) can stop such attacks dead in their tracks by blocking access to malicious sites. The fact is, if your users can’t get to malicious webpages, they can’t get infected by them.




Cisco Secure Access by Duo

Many of the attacks discussed in this report rely on gaining access to machines or networks simply by logging in. This is at the heart of credential dumping practices, whereby an attacker scrapes login

details from compromised computers in order to take over more systems.

This is where multi-factor authentication comes in. [Cisco Secure Access by Duo](#) protects your systems by using a second source of validation, like a phone or token, to verify user identity before granting access. This second layer of protection can even double as an alert for users if their credentials are stolen and unauthorized login attempts are made.



Cisco Identity Services Engine

Whether big game hunting, targeting the OT network, or exploiting RDP vulnerabilities, bad actors attempt to traverse networks in order to dig their way deeper into an organization. One effective way to slow attackers down is by implementing network segmentation. This breaks the network down into smaller pieces and makes it more difficult for attackers to move laterally, enabling rapid threat containment.


The [Cisco Identity Services Engine](#) (ISE) provides visibility-driven segmentation to extend zero trust in the workplace. With Cisco ISE, not only can you control access and limit the lateral movement of threats, but also simplify bring-your-own-device (BYOD) and guest access in zero trust environments.



Cisco AnyConnect

It almost goes without saying how important it is to have secure connections to company resources. Whether it’s remote workers connecting from home, or adding a level of protection to more securely connect to systems using RDP, the virtual private network (VPN) has become the de facto solution.

[Cisco AnyConnect](#) provides the visibility and control you need to identify who and which devices are accessing the extended enterprise. Cisco AnyConnect is more than a VPN, but rather a unified security endpoint agent that delivers multiple security services to protect the enterprise, including posture enforcement, web security features, and roaming protection.



Cisco Cyber Vision

Protecting Industrial IoT assets is arguably one of the more difficult tasks in security. There’s a wide variety of IIoT devices, many of which operate in a very tailored manner and don’t respond well to the disruption that could be caused by many security processes and procedures.


[Cisco Cyber Vision](#) gives OT teams and network managers full visibility into their industrial assets and application flows. Embedded in Cisco industrial network equipment, it decodes industrial protocols to map your OT network and detect process anomalies or unwanted asset modifications.



Cisco Secure Workload

In a number of scenarios, attackers take advantage of applications and services that reside on the endpoints and networks they’re targeting. For instance, in big game hunting attacks, bad actors often leverage living-off-the-land binaries (LoLBins) built into modern operating systems, which can provide useful functionality to attackers.

Micro-segmentation can secure applications by expressly allowing particular application traffic and, by default, denying all other traffic. Leverage [Cisco Secure Workload](#) to automate micro-segmentation through customized recommendations based on your environment and applications. Secure Workload can also provide visibility into who is accessing certain applications and when.



Combining forces

It’s often said that there is power in numbers. This certainly holds true when you combine Cisco Secure products to protect your organization, since the integration of multiple Cisco Secure products brings forth even more benefits. For example:

- Configure Cisco Secure Endpoint to notify Cisco Secure Access by Duo when an endpoint has potentially been compromised, allowing it to automatically block that endpoint from accessing critical applications.
- Protect your remote workers with Cisco AnyConnect even when they’re not connected to the VPN by integrating Cisco Umbrella into the client.
- Use Cisco ISE to leverage the asset inventory built by Cisco Cyber Vision to create dynamic security groups and automatically enforce segmentation.
- Protect your Cisco Secure Email by leveraging multi-factor authentication provided by Cisco Secure Access by Duo.



Cisco Talos Incident Response

Looking for help assessing your security posture? Need assistance in the event of a security incident? [Cisco Talos Incident Response \(CTIR\)](#) provides a full suite of proactive and emergency services to help you prepare, respond and recover from a breach. CTIR enables 24-hour emergency response capabilities and direct access to Cisco Talos, the world’s largest threat intelligence and research group. Let our experts work with you to bolster your defenses and provide rapid assistance when you need it most.

Cisco Secure

You can learn more about the full Cisco Secure portfolio of products and services at <https://cisco.com/go/secure>.



FURTHER RESOURCES



A Roadmap for Success:
Read Cisco's 2021 Security Outcomes Study, a report that empowers security leaders to manage risk, operate efficiently and drive business growth. This is a global double-blinded survey of over 4800 respondents from 25 countries.

Cisco's most popular security blog series
Our cybersecurity experts select the most notable cybersecurity threat to feature. With in-depth analysis and reporting, the Threat of the Month series provides clear explanations and recommendations for defenders. From phishing to ransomware to banking Trojans, we address the threats you need to know about.



Security Stories
Discover the unique, strange, and often raw stories behind what it takes to lead cybersecurity efforts in an organization. Security Stories is an interview-based podcast full of insights from those who are carving a path in this weird and wonderful industry.

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA), Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV
Amsterdam, The Netherlands

Published January 2021

RPT_01_2021

© 2021 Cisco and/or its affiliates. All rights reserved.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. 2231377 | 01/21