SOPHOS
Cybersecurity evolved.

SOPHOS 2021 THREAT REPORT

Navigating cybersecurity in an uncertain world

By **SophosLabs, Sophos Managed Threat Response, Sophos Rapid Response, Sophos AI, Cloud Security**

CONTENTS

THE POWER OF SHARING

Joe Levy, CTO, Sophos

"If you want to go quickly, go alone, but if you want to go far, go together."

This African proverb couldn't ring truer for the cybersecurity industry. By working collectively, with a strong sense of teamwork, we can achieve far more than fighting cybercrime as individual vendors.

But, only by improving our approach and sharing threat intelligence more comprehensively, and by expanding the pool of participants who contribute to (and benefit from) this sharing and collaboration, will cybersecurity vendors continue to drive up costs for attackers, and make lasting, impactful change.

In the spirit of that approach to working together, in 2017 Sophos joined the *Cyber Threat Alliance*, an organization dedicated to breaking down the barriers that, for years, stymied any chance for competitors in the information security industry to collaborate with one another. The CTA has succeeded far beyond its initial mandate to serve as a repository of shared threat intelligence and a place to resolve differences, and has become a sort of UN to the cybersecurity industry.

Through our partnership with the CTA, Sophos can better protect our customers, thanks to the early warnings and data exchange between vendors, made possible by the alliance. Sophos also shares the burden of protecting the other vendors' customers by contributing our own threat intel.

In March 2020, as lockdowns to contain the spread of COVID-19 were implemented rapidly across the world, Sophos chief scientist, Joshua Saxe put out a call on Twitter. Appalled that criminal groups were starting to incorporate references to COVID-19 into a range of crime campaigns, information security analysts – more than 4,000 of them – banded together in a collective show of defiance and formed the COVID-19 Cyber Threat Coalition (CCTC) in a Slack channel created that same day. This channel is building an enduring "commons" for the community to leverage in times of crisis, and is close to achieving not-for-profit status under the auspices of the CTA.

Ultimately, these stories about sharing threat intelligence tell us about more than just the organizations themselves. As another parable—that of the blind man and the elephant—teaches us, no one vendor can provide comprehensive or absolute truth through their subjective experiences alone. The true shape of complicated things emerges from the union of our experiences. These collaborative initiatives protected millions of people from becoming victims of cybercrime, but that alone wasn't *why* they were successful. They thrived because the core motivation of their members and founders has been to, first, protect anyone who might be in harm's way from harm. There was no profit motive, just a desire to defend those in need, while it seemed like the wolves were at the door.

This proves the model is correct, and bridges critical gaps in coverage none of us alone could generate, but we can do more with it. As an industry, we may in the future want to consider sharing machine learning models, or training datasets, just as we share block lists or Yara rules today. We could also strengthen and contribute to emerging standards like STIX and the ATT&CK framework. And we could participate in industry-specific ISACs and ISAOs.

The future will be more connected, and we'll all be better off (and better protected) for it.

EXECUTIVE SUMMARY

The Sophos 2021 Threat Report covers topic areas into which Sophos has gained insight from the work over the past 12 months by SophosLabs on malware and spam analysis, and by the Sophos Rapid Response, Cloud Security, and Data Science teams. These aspects of our daily work protecting customers provide insight into the threat landscape that can guide incident responders and IT security professionals on where they should direct their efforts to defend networks and endpoints in the coming year.

We've segmented the report into four main parts: Discussion of how ransomware has transformed itself, and where this threat is headed; analysis of the most common attacks large organizations face, and why these metaphorical canaries in the coal mine remain significant threats; how the emergence of a global pandemic affected information security in 2020; and a survey of the scope of attacks targeting platforms not traditionally considered part of an enterprise's attack surface.

To summarize the key takeaways from the report:

Ransomware

‣ Ransomware threat actors continue to innovate both their technology and their criminal modus operandi at an accelerating pace

‣ More ransomware groups now engage in data theft so they may threaten targets with extortion over the release of sensitive private data

‣ As ransom groups put more effort into active attacks against larger organizations, the ransoms they demand have risen precipitously

‣ Further, distinct threat actor groups that engage in ransomware attacks appear to be collaborating more closely with their peers in the criminal underground, behaving more like cybercrime cartels than independent groups

‣ Ransomware attacks that previously took weeks or days now may only require hours to complete

'Everyday' threats

‣ Server platforms running both Windows and Linux have been heavily targeted for attack, and leveraged to attack organizations from within

‣ Common services like RDP and VPN concentrators remain a focus for attack on the network perimeter, and threat actors also use RDP to move laterally within breached networks

‣ Even low-end "commodity" malware can lead to major breaches, as more malware families branch out into becoming "content distribution networks" for other malware

‣ A lack of attention to one or more aspects of basic security hygiene has been found to be at the root cause of many of the most damaging attacks we've investigated

COVID-19

‣ Working from home presents new challenges, expanding an organization's security perimeter to thousands of home networks protected by widely varying levels of security

‣ Cloud computing has successfully borne the brunt of a lot of enterprise needs for secure computing environments, yet still has its own challenges unique from those in a traditional enterprise network

‣ Threat actors have attempted to launder their reputations making promises not to target organizations involved in life-saving health operations, but later reneged on those promises

‣ Criminal enterprises have branched out into a service economy that eases new criminals into the fold

‣ Cybersecurity professionals from around the world self-organized in 2020 into a rapid reaction force to combat threats that leverage the social engineering potential of anything relating to the novel Coronavirus

Nontraditional platforms

‣ Attackers now routinely take advantage of the wealth of "red team" tools and utilities pioneered by penetration testers in live, active attacks

‣ Despite efforts on the part of operators of mobile platforms to monitor apps for malicious code, attackers continue to work around the edges, developing techniques to bypass these code scans

‣ Software classified in an earlier era as "potentially unwanted" because it delivered a plethora of advertisements (but was otherwise not malicious) has been engaging in tactics that are increasingly indistinguishable from overt malware

‣ Data scientists have applied approaches borrowed from the world of biological epidemiology to spam attacks and malware payloads, as a method to bridge gaps in detection

THE FUTURE OF RANSOMWARE

Ransomware attacks launched throughout 2020 magnified the suffering of an already wary population. As the pandemic ravaged lives and livelihoods, so did a host of ransomware families, whose efforts did not stop targeting the health and education sectors, even as hospitals became COVID-19 battlegrounds and schools struggled to invent an entirely new way to teach children through March and beyond.

You can't raise enough in a bake sale during a pandemic to pay a ransom, but some schools managed to recover from attacks that appeared targeted at the first day of school, by keeping secure backups.

Ransomware operators pioneered new ways to evade endpoint security products, spread rapidly, and even came up with a solution to the problem (from their perspective) of targeted individuals or companies having good backups, securely stored where the ransomware couldn't harm them.

But what appeared to be a wide variety of ransomware may not be as wide as it seems. As time went on, and we investigated an increasing number of attacks, Sophos analysts discovered that some ransomware code appeared to have been shared across families, and some of the ransomware groups appeared to work in collaboration more than in competition with one another.

Given all this, it's hard to make any kind of reliable prediction of what ransomware criminals will do next. Ransomware creators and operators have burned a lot of time working on defenses against endpoint security products. We counter their countermeasures. They show creativity and versatility in devising new tactics; we show tenacity in studying what they do and finding clever ways to stop them.

Data theft creates a secondary extortion market

Until this year, conventional wisdom among security companies that had any experience at all with ransomware ran quite uniformly: Lock down obvious ingress methods, such as internet-facing RDP ports; keep good offline backups; and deal with infections of small, innocuous malware such as Dridex or Emotet quickly, before they can deliver the killing payload.

Several high-profile ransom attacks, for instance, against school districts across the US, failed at least in part because the IT managers had maintained an unaffected backup of critical data.

As a countermeasure to their victims' preparedness, several ransomware families picked up on a side hustle designed to increase pressure on their victims to pay the ransom – even if every backup with essential data was safe. Not only would they hold the machines hostage, but they would steal the data on those machines and threaten to release it to the world if the targets fail to pay a bounty.

Over the past half year, Sophos analysts observed that ransomware adversaries have settled on a common (and slowly growing) toolset they use to exfiltrate data from a victim's network. This toolset of well-known, legitimate utilities anyone might have won't be detected by endpoint security products. The list of ransomware families that engage in this practice continues to grow, and now includes Doppelpaymer, REvil, Clop, DarkSide, Netwalker, Ragnar Locker, and Conti, among many others. The attackers operate "leaks" sites, where they publicize what data they've stolen; REvil allows anyone to buy the data from them right from its website.

The criminals use the toolset to copy sensitive internal information, compress it into an archive, and transfer it out of the network – and out of reach of the victim. These are some of the tools we've seen used, so far:

‣ Total Commander (file manager with built-in FTP Client)

‣ 7zip (Archive creation software)

‣ WinRAR (Archive creation software)

‣ psftp (PuTTY's SFTP client)

‣ Windows cURL

When it comes to data theft, the attackers are far less picky and exfiltrate entire folders, regardless of the file types that are contained within. (Ransomware typically prioritizes the encryption portion of the attack to key file types and excludes many others.)

Size doesn't matter. They don't seem to care about the amount of data targeted for exfiltration. Directory structures are unique to each business, and some file types can be compressed better than others. We have seen as little as 5 GB, and as much as 400 GB, of compressed data being stolen from a victim prior to deployment of the ransomware.
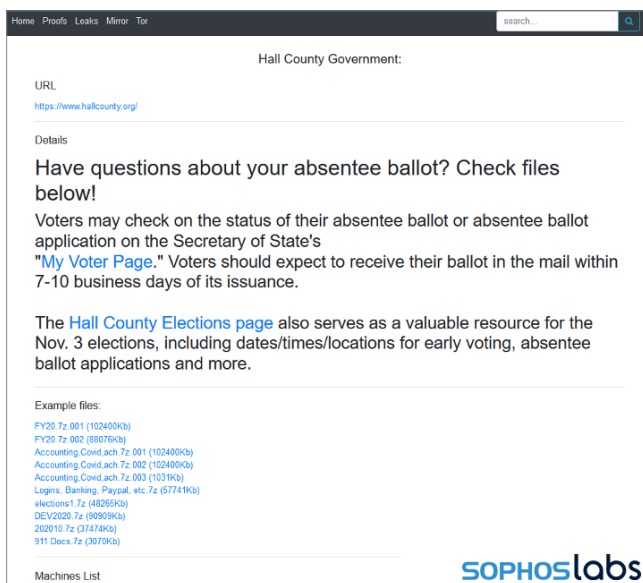
Fiig.1. In October, 2020, the Doppelpaymer ransomware leaks page revealed that the attackers had struck the networks of Hall County, Georgia. The leak included a reference to a file called "elections" which included sample ballot proofs for the state primary elections in 2020 and lists of poll workers and their phone numbers from the 2018 elections, among other sensitive files. The Associated Press reported that the ransomware encrypted the signature verification database the county uses to validate ballots. Source: SophosLabs.

The criminals typically send the exfiltrated data to legitimate cloud storage services, which make this activity harder to spot, since these are common, ordinary network traffic destinations. For attackers, the following three cloud storage services have been the most popular go-to for storing exfiltrated data:

- ‣ Google Drive
- ‣ Amazon S3 (Simple Storage Service)
- ‣ Mega.nz
- ‣ Private FTP servers

In a final act of destruction, ransomware attackers increasingly hunt for the local servers that contain backups of critical data; when found, they delete (or independently encrypt) these backups just before the network-wide encryption attack.

It's more important than ever to store a backup of key data offline. If they can find it, the ransomware criminals will destroy it.

Ransoms rise as attacks increase

It's hard to believe that just two years ago, Sophos analysts marveled at the $6 million haul brought in by the operators of the ransomware known as SamSam. In an attack Sophos responded to in 2020, the ransomware operators opened their negotiations at a dollar amount of more than twice what the SamSam gang earned in 32 months of operation.

Ransomware comes in weight classes, now: heavyweights that attack large enterprise networks, welterweights that target civil society (public safety and local government) and small-to-medium businesses, and featherweights that target individual computers and home users. While earning the dubious distinction of being the heaviest heavyweight sounds impressive, it isn't fair to compare high ransom demands to those that originate from the lower end of the ransomware spectrum.

Sophos has a dedicated team that investigates, and often works with the targets of, ransomware attacks. The team can forensically reconstruct the events of an attack after the fact, and sometimes disrupt attacks while they're still in progress. The Sophos Rapid Response team gets involved in cases when there's a chance to stop or limit the harm, but sometimes the attack happens so fast, there's nothing it can do, and the target must then decide whether or not to pay the ransom, at which point, Sophos is no longer involved.

That's where companies like Coveware come in. The company represents ransomware targets, as a high-stakes negotiator with their attackers. Coveware's CTO Alex Holdtman confirmed our suspicion, that ransomware heavyweights are the primary driving factor in the demand for sky-high ransoms.

Average ransom payouts, quarterly

Quarter	Amount
Q4 2019	$84,116.00
Q1 2020	$111,605.17
Q2 2020	$178,254.19
Q3 2020	$233,817.30

SOPHOS**labs**

Fig.2. The average ransom demand has risen 21% in the past quarter and has nearly tripled over the past year. Source: Coveware.

In just the past quarter, the average ransom payout has risen by 21%, but Coveware believes the averages can be skewed by just one or two very large ransom attacks. The average ransom payout in the just-completed quarter is now the equivalent of $233,817.30, payable in cryptocurrency. A year ago, the average payout was $84,116.

Ransomware threat actors understand how expensive downtime can be, and have been testing the upper limit of what they can extract in a ransom attack.

Several ransomware families have taken up extortion as a side-hustle to help close the deal. As mentioned earlier in our report, groups such as Netwalker and others are using this tactic. That way, even if the target of the attack has perfectly recoverable backups of their data, they may still be forced to pay in the hopes the ransomware criminals don't publish their internal information to the world.

At the lower end of the ransomware spectrum, demands have been increasing, but Holdtman says they're nowhere near the big fish. There are a lot of small businesses and individuals that get hit, but for them the ransom demands have remained relatively flat.

Days-in-the-life of a ransomware rapid responder

When an organization was targeted by the then still active Maze ransomware, it turned to the Sophos Rapid Response team. We investigated and actively countered the attack while it was still in progress. What follows is a day-by-day summary of the attack as it unfolded.

Before Day 1

At some point before the attack becomes active, the operators compromise a computer on the target's network.

This computer is then used as a 'beachhead' in the network. On multiple occasions, the attacker will connect from here to other computers using the Remote Desktop Protocol (RDP).

Day 1

The first evidence of malicious activity appears when a Cobalt Strike SMB beacon is installed as a service on an unprotected Domain Controller (DC). The attackers are able to control the DC from the previously compromised computer by exploiting a Domain Admin account with a weak password.

Day 2

The attackers create, execute, and then delete a series of scheduled tasks and batch scripts. From the evidence seen by investigators, the tasks were similar to a technique used later to deploy the ransomware attacks. It is possible that the attackers are testing the method they plan to use.

Using the compromised Domain Admin account and RDP access, the attackers move laterally across the network to other critical servers.

They use the legitimate network scanning tool, Advanced IP Scanner to start mapping out the network and make lists of IP addresses that would later have the ransomware deployed to them. The attackers create a separate list of IP addresses belonging to the computers used by the target's IT administrators.

Next, the attackers use the Microsoft tool ntdsutil to dump Active Directory's hashed credential database.

The attackers execute various WMI commands to collect information about compromised machines, and then their attention turns to the exfiltration of data: They identify a file server and, using the compromised Domain Admin account, access it remotely over RDP. They start compressing folders located on it.

The attackers move the archives to the domain controller, then try to install the cloud storage application Mega on the DC. This is blocked by security, so they switch to using the web-based version instead and upload the compressed files.

Day 3

Exfiltration of data to Mega continues throughout the day.

Day 4 and 5

No malicious activity is observed during this period. In previous incidents, we've observed ransomware attackers waiting to spring the attack over a weekend or holiday, when the IT security team isn't working or paying close attention to what is happening in the network.

Day 6

A Sunday. The first Maze ransomware attack is launched, using a compromised Domain Admin account and the lists of IP addresses that have been identified. Over 700 computers are targeted in the attack, which is promptly detected and blocked by security. Either the attackers don't realize the attack has been prevented, or they hope that having the stolen data to hold against the victim is enough, because this is the moment when they issue a ransom demand for $15 million.

Day 7

The security team installs additional security and engages 24/7 threat monitoring. The incident response investigation begins, quickly identifying the compromised admin account, discovering several malicious files, and blocking communication between the attacker and the infected machines.

Day 8

Further tools and techniques used by the attackers are discovered, as well as evidence relating to the exfiltration of data. More files and accounts are blocked.

Day 9

Despite the defensive activity, the attackers maintain their access to the network and a different compromised account, and launch a second attack. This attack is similar to the first one: execute commands on a DC, looping through the lists of IP addresses contained in txt files.

The attack is quickly identified. The ransomware is detected automatically and both the compromised account and the malware payload is disabled and deleted. No files are encrypted.

Cleary not wanting to give up, the attackers try again. The third attempt comes just a few hours after the second attack.

By now they seem to be growing desperate as this attack targets a single computer. This is the main file server that the exfiltrated data had been taken from.

The Maze attackers take a different approach, deploying a full copy of a virtual machine (VM) and a VirtualBox hypervisor installer, an attack detailed on SophosLabs Uncut in September, 2020.

The outcome of the third attempt is the same as before: the Sophos Rapid Response team detected and thwarted the attack, with no encryption of files. The team helped the customer lock out the criminal group, and the attackers ceased being able to press the attack further.

EVERYDAY THREATS TO ENTERPRISES – CANARIES IN THE COAL MINE

If everything you know about cyberattacks comes from news reports, you could be forgiven for thinking the sky was falling. Attacks that target large organizations happen every day, but they're not all the kind of black swan events, like a major data breach, that can send a company's fortunes (or stock price) tumbling and generate bad publicity. Many attacks are far more mundane, involving malware the SophosLabs team tracks in a sort of "Most Wanted" list of "The Usual Suspects".

But though these attacks, and some of the malware they deliver, are well understood and easily contained, every attack carries with it the potential to get far worse if it isn't dealt with speedily and effectively. To carry the bird metaphor forward, these routine, everyday attacks represent canaries in the coal mine, an early indication of a toxic presence that could quickly spiral out of control.

Attacks targeting Windows & Linux servers

While the vast majority of security incidents we responded to in 2020 involved desktop or laptop computers running variations of Windows, we saw a steady increase in attacks on both Windows and non-Windows servers. In general, servers have long been attractive attack targets for a variety of reasons: They often run for long periods unattended or unmonitored; servers often carry more CPU and memory capacity than individual laptops; and servers may occupy a privileged space on the network, often having access to the most sensitive and valuable data in an organization's operation. This makes them an attractive foothold for a persistent attacker. These characteristics won't change in 2021 and Sophos anticipates the volume of attacks targeting servers will continue to increase.

The majority of attacks targeting servers fit one of three profiles – ransomware, cryptominers and data exfiltration – each of which has a corresponding, distinct set of tactics and techniques the attackers employ. Best practices for server admins is to avoid running conventional desktop apps, like email clients or a web browser, from the server as a safeguard against infections, so attacks targeting servers necessarily require a shift in tactics.

Internet-facing servers running Windows receive a never-ending barrage of RDP brute-forcing attempts, an attack tactic that, for at least the past three years, has been most often associated with (and predictive of) ransomware attacks. The Sophos Rapid Response team frequently finds that the root cause of ransomware attacks it investigates involve an initial access to the target's network by means of RDP, and then the use of those machines to gain a foothold within the network and take control of DC servers, from which they can mount the rest of the attack.

By contrast, cryptojacking attacks tend to target a wider range of vulnerabilities in Windows, and in applications that normally run on server hardware, such as database software.

For instance, one method used by the Lemon_Duck cryptominer involves a brute-force attack against internet-facing servers running Microsoft SQL Server. Once the attackers guess the correct database password, they use the database itself to reassemble the cryptojacker payload, write it out to the server's file system, and execute it. The infected machine then attempts to exploit the EternalBlue and/or SMBGhost vulnerabilities in a bid to spread the cryptojacker.

Lemon_Duck is an equal-opportunity attacker, and can infect Linux servers. The malware attempts to brute-force SSH passwords taken from a relatively small list. If successful, the attackers load malicious shellcode, which then establishes persistence by taking advantage of loopholes in a service called Redis. The cryptojacker can also conceal itself by executing the commands to start itself from within Hadoop clusters.
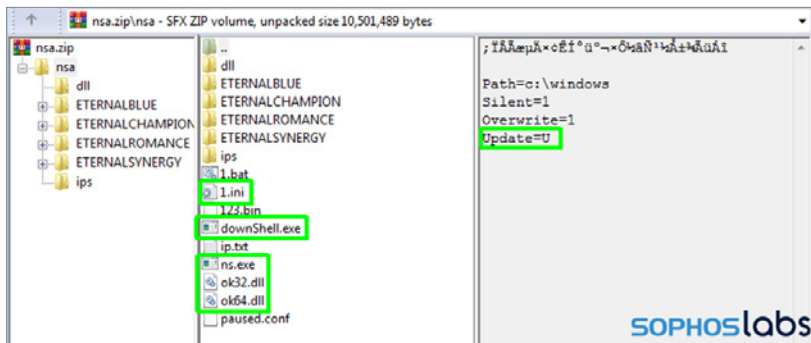

Fig.3. One of the more prolific cryptojackers, called MyKings, distributed the components responsible for installing the botnet (highlighted in green) inside a Zip archive along with several of the exploits leaked from the NSA by the Shadow Brokers. Source: SophosLabs.

Occasionally, attackers target servers because, rather than a quick payday or a steady trickle of cryptocurrency, they want to steal data of value stored on them. In 2020, Sophos discovered an attacker targeting Linux servers using malware we called Cloud Snooper. The servers in question were hosted in a cloud computing cluster, and evaded detection by inventing a clever message-relay system, piggybacking their command-and-control messages on routine HTTP connections.

Fig.4. A "wolf in sheep's clothing" metaphor illustration of how the Cloud Snooper APT malware concealed its commands and exfiltrated data in the guise of conventional HTTP requests and responses, with the help of a tool that monitored network traffic and rewrote TCP/IP packets in real time. Source: SophosLabs.

Server admins have not, historically, installed endpoint protection products on servers, but with the advent of these types of attacks, that conventional wisdom has shifted.

Underestimate "commodity" malware at your peril

Not everybody gets hit with a zero-day vulnerability by a nation-state-sponsored Advanced Persistent Threat (APT). Most attacks involve run-of-the-mill malware delivered by conventional means – which typically involve a spam email, a benign-looking attachment or link, and a lot of encouragement for the target to open that attachment. Sophos receives thousands of telemetry hits per month about such common malware, usually an indication that a computer protected by one of our products has blocked the attack.

In unprotected computers, where the malware can fully execute, it will profile the target's computer; extract any login credentials or saved passwords for websites that control something of value (usually, but not limited to, bank or financial services accounts); then send that information back to its operators and await further instructions, which may arrive in a few seconds...or several days later.

But don't let the fact that these malware families are *merely ordinary* lull you into a false sense of security. These malicious workhorses can cause huge problems if allowed to persist. As mentioned earlier in our report, the SophosLabs team maintains a "Most Wanted" malware list, with analysts dedicated to those families that remain stubbornly persistent. We've put together a short summary of some of them below.

Dridex and Zloader

One of the most common malware types is the *loader*. Loaders have features centered around delivering another malware payload on behalf of their operators or people who contract with their operators. The Dridex and Zloader malware families are both mature, established loader platforms. Attackers use both Dridex and Zloader to collect information about the target system and send it back to the criminals, who can decide at their leisure what components or payloads they will deliver, based on the information the bot sends back.

The Dridex loader's core function is to contact its command-and-control (C2) server, retrieve one or more encrypted payloads, and deploy them. It's very hard for analysts to get those payloads because the threat actors only distribute them on an as-needed basis, such as a hidden VNC (a remote-control application), or a SOCKS proxy. These payloads give attackers the ability to do things in the context of the user's device. They also allow the criminals to access resources on the victim system that are not directly reachable from their own system.

The server-side logic that determines what happens during an infection can be inscrutable, but we can infer some rules because the bots don't want to infect computers used by malware analysts. The bot sends its operators a list of installed programs; if there are analysis tools, or components of virtual machines, the bots don't deliver payloads to that machine. In Zloader's case, the bot's operators spread the malware via a spam message; if you take too long to infect your computer, within eight to 12 hours after the spam goes out, they stop sending payloads.

It also must be a really clean machine, but it can't be too clean, either. A plain-vanilla Windows installation won't trigger, but neither will a very full machine with a lot of tools.

Agent Tesla and RATicate, infostealers and RATs

Fig.5. We run all newly-discovered remote access trojan (RAT) malware samples through our internal sandbox system. This table illustrates how many new, unique samples we encountered over a seven-month period that we later classified to one of the 18 most common RAT families, broken out by the family names. * Partial-month data. Source: SophosLabs.

Remote Access Trojans, or RATs, and infostealers are among the oldest forms of malware. As the name suggests, RATs offer the attacker the ability to control the infected computer remotely. Infostealers also hew close to their name, engaging in theft and exfiltration of credentials, certificates, and other sensitive information. Two of the "Most Wanted" families we've dealt with over the past year are Agent Tesla (an infostealer) and RATicate (a RAT).

Like loaders, RATs also usually have a mechanism by which they can deliver additional payloads, including updated versions of themselves. We've observed RATicate distributing other malware – including Agent Tesla. We've also seen these RAT families being served up from, or communicating with, the same IP addresses or servers, which hints at something shared between otherwise unrelated groups.

Trickbot's takedown

Trickbot has been a persistent nuisance malware for at least four years. The infamous botnet pioneered a lot of what are now common behaviors and characteristics: for instance, it communicated with its C2 infrastructure using TLS. The bot is implicated in several high-profile ransomware attacks and is a competent credential thief in its own right.

```
    "type"  :  "TEXT",
    "size"  :  101
  },
  "controllers"  :  [   {
    "url"  :  "https://127.0.0.1.1"
  } ],
  "controllers"  :  {
```

Fig.6. Trickbot was taken down by a single line of code. Source: SophosLabs.

In October 2020, as we prepared this report, Microsoft and the US Department of Justice announced that they had seized several servers and sent a command via the botnet's command and control system that

caused some 90% of the botnet to stop communicating with the command-and-control infrastructure.

Investigators managed to upload a "poisoned" configuration into Trickbot's infrastructure that each bot downloaded. The configuration tricked the botnet into believing that its principal command-and-control server was the infected machine it was running on. The botnet then lost contact with the actual C2 servers and could no longer retrieve payloads or instructions.

The effort had a drastic impact on Trickbot's operator, but they're expected to slowly, eventually, return to normal operation.

Delivery mechanisms

There are a limited number of ways that malware or attackers can reach a targeted machine or penetrate a network. The methods of most malware attacks follow a well-worn path that may include the use of email containing links to or the attachment of a malicious file, or the attacker may take a more active role targeting RDP or some other vulnerable service hosted at the network perimeter, facing the public internet.

RDP, the #1 attack vector for ransomware

The Windows Remote Desktop Protocol, or RDP, is a standard service available in all current versions of Windows. With very little effort, RDP allows IT admins or computer users the ability to log in to a computer when not physically in the presence of that computer, which can be great in the case of a pandemic where everyone is suddenly forced to work from home. Unfortunately, for the past three years, ransomware threat actors have been abusing (at an accelerating rate) that same remote access platform as a way to gain a foothold and cause large-scale damage to enterprises, earning them an increasingly large payday from targeted organizations.

RDP login attempts per honeypot

Fig.7. We distributed honeypots to datacenters around the world and permitted attackers to try to brute-force their way in. The honeypot machines were discovered "organically," without being advertised in any way. Over the 1-month period of our tests, this map illustrates how many attacks each honeypot

location received.

The impact of the COVID-19 lockdown era has only exacerbated the problem, as increasing numbers of organizations and workers are forced by circumstances to rely upon RDP in order to remain operational. The main risk here is that RDP was never meant to withstand the kind of onslaught it can receive from the public-facing internet. If the RDP password is weak, easily guessed or brute-forced through automated login attempts, the attacker gains a foothold inside the network they can exploit at their leisure.

The Sophos team that handles incident response for major incidents reports that RDP remains one of the top "root causes" of the ransomware events they handle. The advice to IT managers remains the same as it has all along: RDP should never face the public internet, but instead be placed behind a firewall that requires users to connect via a VPN or other zero-trust facility first; and admins should strengthen Windows password policies to require longer passwords and a multi-factor authentication token or app.

In research performed before the lockdown took effect, Sophos set up honeypots in 10 datacenters around the world to better understand just how bad the problem had become. Over a 30 day period, the honeypots received a median average of 467,000 RDP login attempts, or about 600 per hour at each location. The research revealed that each honeypot received a steadily increasing frequency and ferocity of login attempts until we finally pulled the plug.

Top 5 usernames used in all failed login attemps

USERNAME	FAILED LOGIN ATTEMPTS
administrator	2,647,428
admin	376,206
user	79,384
ssm-user	53,447
test	42,117

Fig.8. Remote Desktop brute-force attempts target the most common Windows usernames, including the default "administrator" account. Source: SophosLabs.

Business Email Compromise and Business Email Spoofing

Business email compromise (BEC) is the formal name for a specific kind of spam that centers around a fraudulent request for money. In a BEC attack, a spammer sends messages that have been crafted to look like they originate from a high-level executive within a company, directing someone at a lower level to perform some kind of financial transfer or complete a large purchase on behalf of that executive. Attackers may do this by spoofing the appearance of internal emails (sometimes called Business Email Spoofing) or they may try to take control of accounts on the organization's own mail server, and use that account to send the fraudulent request.
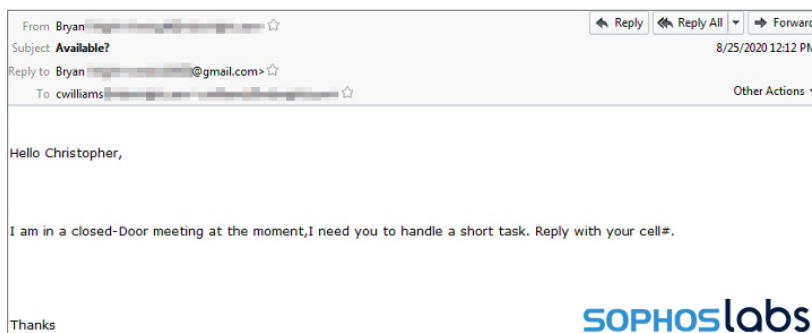

Fig.9. In this real-world example of a business-email compromise attempt, the fraudster poses as an executive asking an employee to respond to an urgent request. The email has a different Reply-To address (from a Gmail account) than the one in the From: header, a dead giveaway that something is awry – if the

target is paying attention to the mail headers. Source: SophosLabs.

BEC attackers, posing as an executive, may ask the targeted employee to buy expensive gift cards or expedite a financial transaction of some kind. The attacks are usually highly tailored to the targeted individuals and organizations. BEC email messages look nothing like malicious spam, because they fail to follow spam-like patterns; they (often) don't contain an attachment or malicious link, and they try to look like they originated from within the targeted organization, at times even incorporating the target organization's typical mail "signatures" or other elements that may be familiar to employees, to make them more convincing to the target than conventional malicious spam.
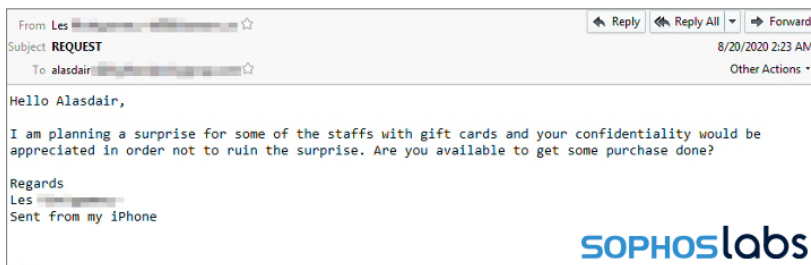

Fig. 10. After the target has acknowledged the initial request, the fraudster makes the "ask" - providing a pretext that appears plausible. Source: SophosLabs.

BEC scams rely on the target of the scam (the employee) being physically distant from the subject of the scam (the executive), and it also depends on the target acting quickly, before anyone can figure out what's going on and stop the target from buying gift cards or making bank transfers. BEC scammers may craft a message when they know the executive is out of the office on business.

These kinds of fraudulent requests often involve some back-and-forth between the attacker and target. The conversation may start with a simple request for the target to respond to the scammer and progress into a series of messages that eventually lead to an "ask" to make a purchase based on a plausible-sounding pretext.
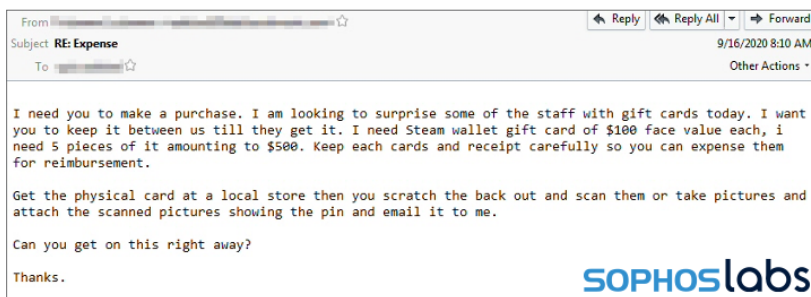
Fig.11. At some point during the attack, the BEC scammer will make a request that flies in the face of common sense, like a request to make a sudden, large wire transfer to an account unfamiliar to the scam's target. This provides another opportunity for a wary staffer to question the nature of the request: Why would the executive need a photograph of the back of a gift card with the PIN scratched off when they're going to be handed out as gifts? Source: SophosLabs.

Back when most of us were working in offices, physical proximity between the target and subject would have made the scam immediately apparent. But our current distributed work environment, where both the executive and employee are unlikely to be in the same physical proximity, reduces the opportunities for people to just walk over to someone's desk and ask them to confirm the request.

BEC scams existed prior to the COVID-19 era, but as more people are working remotely, BEC scammers are on the prowl. As an attack against the better nature of people who just want to be helpful and supportive, it is a particularly offensive type of scam. If you encounter emails like these, trust your gut and speak to the subject in question directly, if you can, or ask for guidance from someone else if you can't reach them. The more real employees get involved in handling these requests, the more likely it will be that the scam will be discovered before any damage is done.

Weird science: retro Office glitch strikes again

When it comes to malicious office documents (maldocs) and the exploits they attempt to deploy, what's old is often used again and again, goes away after Microsoft produces an update, and then (sometimes) resurfaces. For years, SophosLabs has tracked how attackers embed a wide and rapidly changing variety of exploits into maldocs. Newly-disclosed vulnerabilities often find favor among the criminals who use maldocs as a stepping stone to deliver a malware payload, because not everyone installs patches right away, and it sometimes takes a bit of time for security companies to craft an effective "safety net" defense, based on behavior or other characteristics of a novel vector.

Most of the maldocs we've seen throughout the past year have been constructed using tools called builders that give attackers a literal point-and-click menu system that lets them decide exactly what exploit(s) to craft into the malicious document. As endpoint protection tools get better at identifying these more modern exploits, which usually involve a script that has been embedded into the document, maldoc creators seem to have dug deep to find a very, very old bug that helps conceal the macros or other malicious content in the documents.

The bug is colloquially known as the **VelvetSweatshop** exploit, though it really isn't an exploit at all. In fact, VelvetSweatshop was introduced by Microsoft into Microsoft Office 2003, although we didn't see it abused until 2013, when Excel workbooks exploiting the CVE-2012-0158 vulnerability were cloaked with the help of the glitch. An Excel spreadsheet or Word .doc marked as "read-only" is just a password-protected document with a stock password of, you guessed it, VelvetSweatshop.

We've seen a lot of malicious Excel spreadsheets being delivered this year that use the technique as a way to evade advanced threat detection. Because of the encryption, the real malicious content is hidden behind strong crypto that scanners can't crack, and can't scan unless they support the latest algorithm used by attackers. Due to the use of the default-password, Excel opens the decoded content without prompting for the password, so from the execution point of view the encryption is transparent. Endpoint security programs added support for the encryption and the default password, but the criminals keep finding additional cryptographic algorithms that have the same feature and are not (yet) implemented by AV scanners.

It was quite a surprise to discover a bug old enough that, if it were human, would be in its last year of school. But it's no surprise that the makers of weaponized document builders would try to take advantage of it.

Information security: A 20-year retrospective

While an annual report gives us an opportunity to look back at significant events of the past year, we thought a look further back – at the past two decades – would provide context for how we arrived in our current threat landscape. The turn of the millennium marked a milestone, when information security became a professional discipline and a bona fide industry. This timeline of threats and events represent significant, representative moments in the evolution of threat behavior.

As both enterprises and individuals adopted the internet for both business and entertainment, large networks were ripe targets for the emergence of prolific worms – self-propagating malware. Cumulatively, worms infected tens of millions of systems worldwide and cost over $100 billion in damages and remediation costs.

2000–2004
The Worm Era

- **2000** ILOVEYOU
- **July 2001** CodeRed
- **January 2003** SQL Slammer
- **January 2004** Bagle
- **August 2001** CodeRed II
- **August 2003** Blaster
- **January 2004** MyDoom
- **September 2001** Nimda
- **August 2003** Welchia
- **February 2004** Netsky
- **August 2003** Sobig
- **April 2004** Sasser
- **October 2003** Sober

2005–2012
The Malware Monetization Era

- **2006** Rx Spam
- **2007** Storm
- **2008** Conflicker
- **2009** Stuxnet
- **2010** Blackhole exploit kit
- **2011** Malvertising
- **2007** Zeus

2013–Present
The Ransomware Era

- **2013** Snowden leaks
- **2014** Point-of-sale (POS) malware
- **2016** Mirai
- **May 2017** WannaCry
- **2018** Magecart attacks
- **2019** Extortion ransomware
- **2020** APT tactics by threat actors
- **2013** CryptoLocker
- **June 2017** NotPetya

Fig.12. Source: Sophos

SOPHOS

2000–2004 - The Worm era

2000 - ILOVEYOU
The ILOVEYOU worm used a social engineering trick that persists even today: It arrived as a spam email attachment, eventually infecting about 10% of all internet-connected Windows computers.

July 2001 - CodeRed
Named after the flavor of Mountain Dew its discoverers were drinking at the time, CodeRed used a buffer overflow vulnerability in IIS to spread itself and deface websites. It was followed a month later by an upgraded version that installed a backdoor on networked computers.

August 2001 - CodeRed II

September 2001 - Nimda

January 2003 - SQL Slammer
At only 376 bytes, Slammer exploited a buffer overflow in Microsoft database applications. Doubling its infections every 8.5 seconds, Slammer took down large swaths of the internet in only 15 minutes.

August 2003 - Blaster
Blaster was created by reverse engineering a Microsoft patch a couple months ahead of the first Patch Tuesday. It exploited a buffer overflow vulnerability in the RPC service of Windows XP and 2000 systems and launched a DDoS attack against windowsupdate.com if the day of the month was greater than 15, or the month was September or later.

August 2003 - Welchia

August 2003 - Sobig

October 2003 - Sober

January 2004 - Bagle

January 2004 - MyDoom
It is estimated that 25% of all emails sent in 2004 originated with the MyDoom worm, which prolifically emailed itself to new victims and engaged in a denial-of-service (DDoS) attack.

February 2004 - Netsky

April 2004 - Sasser

2005-2012 - The Malware Monetization era
Until around 2005, malware incidents could be chalked up to curiosity or disruption. Botnet malware, designed for stealth and profit, dominated. This era also saw the start of so-called pharmacy spam. Exploits against software vulnerabilities became key components of malware, which enabled malvertising. Wherever there was the potential for financial gain, cybercriminals exploited those opportunities.

2006 – Rx Spam
What had been a mere annoyance (or a way to propagate worms), became a lucrative business selling mostly counterfeit prescription medicines advertised through spam. It's estimated that pharma spammers made billions of dollars selling medicines most people could get just by going to their doctors.

2007 - Storm

2007 - Zeus

2008 - Conficker
Conficker rapidly infected millions of computers worldwide but did not result in much damage. We still don't know the worm's true purpose, but thousands of hosts remain infected to this day, and Conficker scan traffic is routinely detected as part of the internet's "background radiation."

2009 - Stuxnet
Stuxnet was one of the first digital weapons to target a physical system: Nuclear refinement centrifuges used by Iran to enrich uranium. Stuxnet's enduring legacy is that it permanently opened the door to nation-states' use of malware as a tool of war.

2010 - Blackhole exploit kit
Exploit kits – toolkits targeting software vulnerabilities – bound different parts of the cybercrime ecosystem together. Crimeware-as-a-Service was born when the creators of the Blackhole Exploit Kit began offering their services.

2011 - Malvertising

2013–Present - The Ransomware era
Ransomware has had the most profound impact on this era. While worms, banking trojans, malvertising and spam persist, nothing has come close to rivaling ransomware's destructive force. Damage estimates from ransomware attacks over the past seven years are in the trillions of dollars. Ransomware is also most likely the first form of malware linked to a human death. Moreover, many of today's threats ultimately deliver ransomware and, like exploit kits, it has provided a nitro-fueled boost to an already thriving cybercrime ecosystem.

2013 - Snowden leaks

2013 - CryptoLocker
During its short existence, CryptoLocker provided future criminals with a winning formula by mating two existing technologies: encryption and cryptocurrencies. The threat landscape was forever changed by CryptoLocker and its aftershocks are still being felt today. Three months after launch, the bitcoin wallet used by CryptoLocker contained nearly $30 million.

2014 - Point-of-sale (POS) malware

2016 - Mirai

May 2017 - WannaCry
WannaCry, the most widespread ransomware-worm hybrid seen, demonstrated (again) how a lapse in patching can have dire consequences. It relied on exploits stolen from the NSA and publicly released by The Shadow Brokers. The attacks forced Microsoft to release out-of-band updates for unsupported products.

June 2017 - NotPetya
NotPetya crippled some of the world's largest shipping and logistics companies, reportedly causing over $10 billion in damages. Some of the affected companies have yet to fully recover.

2018 - Magecart attacks

2019 - Extortion ransomware
In an attack against the city of Johannesburg, South Africa, the criminals behind Maze ransomware pioneered the use of extortion. They not only encrypted and stole data, but also threatened to publish the stolen data if companies didn't pay. This tactic has been copied by many other ransomware crews as a hedge against the targets having good backups.

2020 - APT tactics by threat actors
The adoption of nation-state tools and tactics, which began in the past couple of years, went mainstream in 2020. Professional cybercrime gangs use sophisticated tools like Cobalt Strike to devastating effect, while some groups (Dharma) are baking it into point and shoot toolkits for novices to use.

COVID-19 AS A FORCE-MULTIPLIER IN ATTACKS

The novel coronavirus, COVID-19 dramatically affected every aspect of cybersecurity. Attackers felt emboldened to target the newly-homebound white collar workforce. An already high level of anxiety and fear permeating the public sphere was only exacerbated by waves of spam campaigns, ransomware that targeted weakened or broken institutions or civil society that was already under financial pressure, and every manner of rent-seeking fraud and profiteering from scarcity of everything from PPE to toilet paper.

Home is the new perimeter

Everything normal ended in March 2020, when workers who could work remotely, and students at almost all levels, were sent home in a mad dash to halt the spread of COVID-19 and relieve the pressure on overcrowded hospitals. Suddenly, we weren't so much working from home as living at work.

Many people struggled to find the new normal without a commute to the office. Demand for VPN access and multifactor authentication services surged. Chromebooks became rare commodities. Zoom went through about ten years' worth of evolutionary growth in two months. And through it all, Microsoft, Adobe, Apple, and Google were releasing updates and maintenance patch releases for a multitude of platforms.

COVID-19 and Coronavirus Email Scams on the Rise

Fig.13. Globally, a significant portion of spam email mentioned COVID-19 or Coronavirus in the weeks after the lockdown. Source: SophosLabs.

COVID-19 turned us all into our own IT departments, managing patches, security updates, and connectivity issues that kept us from getting into meetings or the kids from being able to attend a virtual classroom. Demand soared for headsets, microphones, better lighting, and security both on the network and the endpoint. And it meant giving even young kids a crash course in phishing, spam, online trolls, cyberbullying and malware disguised as a free copy of your favorite game, ready to play.

It hasn't been easy, and we're still not operating where we were in February 2020, but many people are finding the new normal might, in some ways, be an improvement. More offices have decided to continue allowing remote work even after lockdowns end and people could return to the workplace, which will have a significant benefit to both the environment and to people's quality of life.

As those workplace perimeters stretch and expand to encompass big portions of the workforce in their remote locations, circumstances have enhanced the seriousness with which we view home networks' role as the last line of defense. The modem in the hall closet is now the network perimeter. We need a complete rethink of how to provide that structure with defense in depth.

Crimeware as a service

It can be helpful to think of malware makers as a form of software startup. Scrappy at first, the successful makers eventually earn a loyal following. And there can be just as many business models for malicious software as there are for legitimate software.

The term "crimeware" is intentionally broad; some creators of malware, or of the tools that enable malware to be delivered with ease or to enhance it with new features, don't sell their product outright, but license it as you might buy a one-year license for the Adobe Creative Suite. We've called this class of business model "crimeware-as-a-service," (CaaS) and it seems poised to be the new normal.

One of the more notorious examples of a CaaS malware is Emotet. The spam-delivered trojan has been around for years, and seems to be centered around a smooth experience for the would-be criminal. Emotet is one of a class of malware collectively referred to by security researchers as loaders. Emotet exists primarily to deliver other malware to a target's computer. It accomplishes this job with a sophisticated network that distributes weaponized spam emails to large volumes of targets.


Fig.14. Source: SophosLabs.

Emotet, however, has gone through two dark periods so far this year. The malware remained in communication with its C2 servers over a nearly five-month period in which the spam emails that normally deliver attacks evaporated completely. Spam emails delivering Emotet mysteriously resumed in July.

Dharma ransomware is another CaaS malware of note. Unlike its pricier relations, Dharma maintains a fixed, small ransom. The reason comes down to Dharma's business model: It's the ransomware with training wheels on, for aspiring criminals who need to learn the ropes. Those criminals pay essentially a subscription fee to obtain payloads from the Dharma creators, and split the proceeds of any attacks with them.

As attackers branch out into specialties and sub-specialties, it seems the business model in which criminals work with independent contractors, freelancers and affiliates is one that doesn't seem to be going away anytime soon.

Spam, scams, and broken promises

The lockdowns across the world were accompanied by a flood of scams abetted by spam email. In the best of times, the most effective spam campaigns introduce a sense of urgency to demands that the recipient act on the message. This is a well-known psychological trick, because if you take a couple of moments to think about the contents of the spam message, you'll probably realize it's a fake. If the spammer triggers a fear reaction, you act before you think, and you get snared in the trap.

COVID-19 already had everyone on a hair trigger, so spammers didn't even have to try particularly hard.

Fig.15. Source: SophosLabs.

A few weeks into the lockdown, we decided to take a closer look at another growing phenomenon: Domain registrations. Within weeks, people were registering thousands of new domain names per day that contained any combination of the strings *COVID-19*, *Corona*, or *virus*.

| Domain	First Seen	Nameserver	Ns Ip
coronavirusshaquilleoneal.com	2020-03-14 07:00:38	ns-cloud-b1.googledomains.com	216.239.32.107

Fig.16. Source: SophosLabs.

Some of the sites were obvious jokes, while others were confusingly similar to those used by legitimate, regional or national health authorities.

We also hunted for COVID-19-related domains and subdomains in TLS certificate transparency logs. Certificate transparency logs are useful for tracking subdomains that have their own TLS certificates – information that doesn't show up in raw domain registration data – and domain names.

## New COVID domain registrations, per day

February		
08		49
09		48
10		60
11		205
12		445
13		115
14		89
15		47
16		51
17		50
18		61
19		33
20		39
21		44
22		84
23		107
24		190
25		166
26		439
27		600
28		602
29		676

March		
01		425
02		521
03		491
04		613
05		759
06		866
07		607
08		623
09		750
10		820
11		1020
12		1890
13		2484
14		2585
15		2102
16		2732
17		2666
18		2555
19		3142
20		3011
21		2810
22		2647
23		2449
24		2385

Scale: 0, 1K, 2K, 3K, 4K

Total new COVID domain names, to date

February		
08		49
09		97
10		157
11		362
12		807
13		922
14		1011
15		1058
16		1109
17		1159
18		1220
19		1253
20		1292
21		1336
22		1420
23		1527
24		1717
25		1883
26		2322
27		2922
28		3524
29		4000

March		
01		4425
02		4946
03		5437
04		6050
05		6809
06		7675
07		8282
08		8905
09		9655
10		10475
11		11495
12		13385
13		15869
14		18454
15		20556
16		23288
17		25954
18		28509
19		31651
20		34662
21		37472
22		40119
23		42568
24		44953

Scale: 0, 12,5K, 25K, 37,5K, 50K

■ WHO declare COVID-19 a 'Public Health Emergency of International Concern'   ■ Lockdowns begin in North America

sophoslabs

Fig.17. Throughout the early months of the COVID-19 crisis, people registered thousands of domains and licensed at least that many TLS certificates per day, whose names contained the string "COVID-19" or "corona". Source: SophosLabs.

New TLS certificates per month with "COVID-19" or "Corona" hostnames

Fig.18. TLS certificate registrations that referred to the pandemic spiked at about the same time as the domain registrations. Source: SophosLabs.

We saw an average of over 200 certificate requests for COVID-19 domains per day in March, and the rate continued to climb in the subsequent months. By June, the average reached 625 a day. In October, that rate peaked at 951 new TLS certificates that were being requested, per day.

Most of these domains continue to be legitimate or benign, though many remain parked and have no content—a sign that the registrant may be "domain aging," and setting aside these domains for future reputation check-ups.

Fig.19. Even infamous pill mill scammers couldn't resist getting in on whatever miracle cure came across Twitter, and even posted tweets in the ads. Source: SophosLabs.

A small percentage (below one percent) have been identified as being associated with phishing or malware. Many are ephemeral, with hostnames that can no longer be resolved after as little as a day.

Remote work raises the importance of secure cloud computing

When the COVID-19 lockdowns began in March 2020, people and workplaces began a rapid and unprecedented transition that continues to this day. How we work, go to school, attend events and conferences, and entertain ourselves may have changed forever, and cloud computing was an essential element of that rapid evolution, but it faces a lot of challenges.

The overprovisioning of access permissions, limited visibility of assets and resources in the cloud and a lack of auditing, can all make cloud environments more vulnerable to cyberthreats and malware is about as bad in the cloud as it is everywhere else. For instance, cryptojacking is a growing problem in the cloud. The computing-cycles-heavy cryptominer processes are bad enough when they run on physical machines, and run up the electricity bill; they create an even more painful side effect when they run on cloud instances: The target gets billed by the cloud provider for the CPU cycles consumed by its virtual workstations performing the heavy math required to deliver a few pennies' worth of cryptocurrency.

Further, many dispersed, remote workforces have been hit by ransomware attacks, where criminals locked down the cloud infrastructure the same way they targeted physical machines. After all, ransomware can encrypt a virtual hard drive or object storage just as easily as physical storage. Organizations whose cloud infrastructure is attacked with ransomware can find themselves hit not only with a bill for the cycles spent encrypting the data, but for the ransom, too.

Organizations suffering security incidents in the last year


Fig.20. In our 2020 Cloud Security Report, Sophos surveyed more than 3,500 IT professionals about their experience using the cloud, and found that many of the security problems plaguing physical networks have translated over to the virtual ones. Source: SophosLabs.

On lockdown, IT departments needed a way to service a virtual helpdesk like they staffed a real one before many workplaces closed down. The big changes COVID-19 demanded came in three waves.

In the first few weeks after the lockdowns began, the first wave – an access wave – began to take shape. As millions of workers, suddenly unable to go to their workplace, needed to access resources inside their organization's environment, rapidly growing demand for virtual private network (VPN) or other zero-trust facility access overwhelmed existing resources. Along with VPNs, organizations found they needed to add new firewalls and other security appliances, deployments of modern unified threat management systems supplemented the rudimentary layer 3 firewalls provided by cloud services.

In the pre-COVID-19 world, VPNs only saw moderate use as employees in the workplace vastly outnumbered traveling or remote workers. As March turned into May and then June, for these workers, the VPN became an essential lifeline (if not the essential lifeline) that kept organizations in operation.

But those organizations also quickly realized that employees shouldn't use personal devices from home to access the VPN, and a dwindling supply of new laptops created a new challenge for organizations already struggling with the IT needs of a distributed workforce. Without enough physical machines, for the time being, organizations turned to the seemingly unlimited resource of virtual machines to fill the need for a secure computing workspace. That began the second wave – the virtual desktop wave.

As more employees transitioned into using a virtual corporate desktop, the move to hosting those desktops in the cloud made practical and cost sense, but they still required protection.

Suddenly, IT departments supported hundreds or thousands of employee VMs, and suddenly needed visibility tools to be able to inventory and securely configure the growing cloud estate of virtual servers, virtual desktops, and other cloud services – the cloud management wave.

Attack Timeline
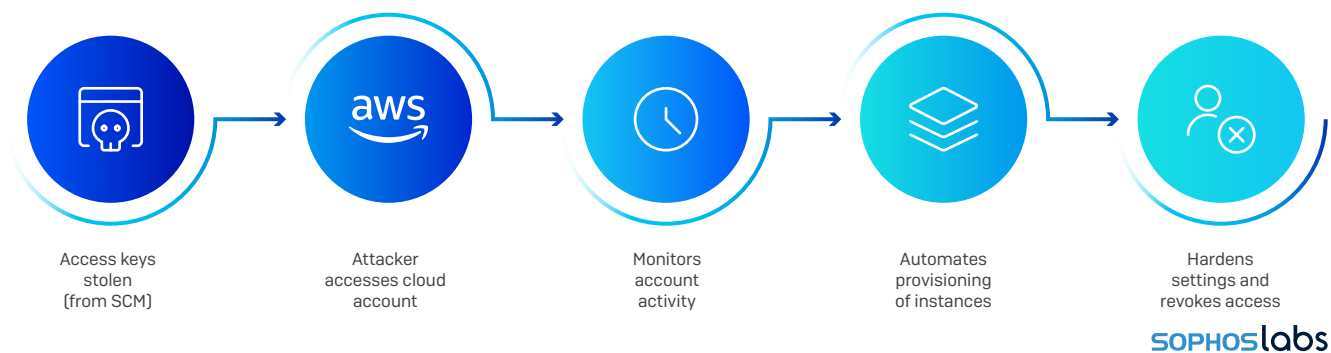
Fig.21. A cryptojacking attack we investigated began when a developer inadvertently embedded their cloud credentials into code on a public repository. The attacker discovered, then used those credentials to attack, using the native cloud provider's APIs – spinning up hundreds of VM instances to mine Bitcoin. At the same time, they automated features on those instances to make them harder to terminate. Later, they revoked access to other legitimate users. Source: SophosLabs.

The COVID-19 era has been marked by great transformation in every aspect of human life, including how many work. In a recent survey by Reuters, 97% of CEOs and CTOs surveyed said that the lockdowns sped up their transition to new technology. But in times of tight budgets and uncertainty, nearly one out of every three of those CTOs reported their mandate was to implement these changes in as cost-effective a way possible.

In Sophos' most recent Cloud Security Report, we found that the majority of security incidents involving cloud computing came down to two primary root causes: stolen or phished credentials, or misconfigurations that led to breaches. Seven out of ten of the more than 3,700 IT professionals surveyed for the report claimed that the cloud infrastructure they support had experienced a breach in the 12 months prior to the survey.

What the CCTC means for a rapid response to large scale threats

Fig.22. Source: Sophos.

About a week into the COVID-19 lockdown, Sophos chief scientist, Joshua Saxe put out a global call for volunteers. The virtual bucket brigade rapidly became the COVID-19 Cyber Threat Coalition (CCTC), an organization numbering more than 4,000 members in service of one goal: To put special effort into countering any type of threat or social engineering that attempted to leverage the public's fear about COVID-19, by name or inference.

"I'm not a firefighter, so I wouldn't know how to fight a building fire, but I can assist a team that's going to bolster the defenses of critical infrastructure, like hospitals," says CCTC's Nick Espinosa, a security analyst and podcaster based in Chicago.

This was a very necessary effort. From the earliest points of the lockdown period, attackers spread spam, malware, and a variety of other threats that made reference, in one form or another, to the frightening, new pandemic lingo. As mentioned in the main report, at one point, people were registering thousands of new domains with the words COVID-19, corona, or CoV in their names, each day. Sophos traced domains connected to TLS certificates with those same text strings in the certificate data and found thousands more.

COVID-19 Cyber Threat Coalition member growth

Fig.23. Source: Sophos.

Because of the unique nature of the threat posed by COVID-19, malicious spam that abuses the global crisis represents something particularly egregious and offensive. "We saw an explosion of criminal hacking using COVID-19 as a lure," said Espinosa. Spam campaigns spun up, in which the spammers dressed the messages like official communiques from the World Health Organization, the CDC in the US, the UK's NHS, drug companies, or national health authorities in countries outside the US and UK.

Analysts also saw references to COVID-19 in strings inside of binaries and used as variables in so-called living-off-the-land, LOLscripts.

CCTC participants shared samples and intelligence about all manner of incidents through a hastily set-up Slack. Chaotic at first, the organization formed a rudimentary structure quickly. "So many people came together and were fire-hosing so much information," Espinosa says.

The CCTC's product, the collected output, is its intelligence feed listing newly-gathered indicators of compromise. The feed is free to use, by anyone. These IoCs complement the defensive technologies already in place, in a vendor-neutral way. When the CCTC formed a partnership with the Cyber Threat Alliance, the security vendors that participate in the CTA amplified the protective effect of CCTC's threat intelligence by ingesting and defending against those threats.

The rapid coalescence of security professionals, sharing a common goal, was heartwarming, says Espinosa. "We probably were a hot mess out of the gate," he says, but the group self-organized quickly. The completion of the CCTC sharing platform means anyone who might need to respond to a COVID-19-like pandemic in the future won't have to reinvent the wheel, and can respond more readily to threats – a healthy metaphor for, and analogue of, the immune system itself.

NOT LETTING YOUR GUARD DOWN: THREATS VIA NONTRADITIONAL PLATFORMS

We live in a world surrounded by computing devices that don't look like a computer or server: routers, mobile phones, firewalls, smart TVs, streaming boxes, VoIP boxes, cameras and camera doorbells, network attached storage, some brands of kitchen and laundry appliances, and on and on.

But just because they don't look like traditional computers, it doesn't mean they can't be misused or abused in the same way.

Android Joker malware growing in volume

Android users find themselves in the middle of an arms race between Google (which owns the Android platform and its primary Google Play Store) and malware creators who want their malware listed for download on the Google Play Store. Google has spent years on a system designed to inspect source code of Android apps submitted for inclusion in the Google Play Store, in search of chunks of code that indicate a malicious intent or an undesirable outcome for an Android user. Malware app developers have had to work hard to evade the Google Play Store code checks.

Joker, aka Bread, is a premium SMS and billing fraud app, one of the more successful examples of a malware family that has evolved to evade these code checks. Google has removed thousands of these Joker-modified malicious apps from the Google Play Store since last year, when researchers first discovered it. Despite the amount of effort put in to getting rid of the malware, Joker keeps bouncing back.

Joker appears in the guise of a wide range of different apps: utilities and tools, wallpapers, translators, messaging services – just many clones of popular apps. Remember, Joker may actually be embedded in an app that looks and works exactly like the real version of almost any app you use. The Joker apps just have a little extra malware code buried deep, in one of the third-party libraries app makers routinely compile into their apps for a variety of legitimate reasons.

There are a few reasons why Joker manages a successful evasion of Google Play Store security code checks time and time again:

1. The malicious apps use obfuscation, from simple string substitution to complex commercial packers, to slow analysis and fool the Google Play Store.

2. When the Joker "developer" launches the app, it contains absolutely no malicious code. This establishes a history where the app that comes into Google Play Store is clean. Only later does the malicious code appear in the app, following an update.

3. The app either decrypts its payload at runtime or downloads it dynamically, later.

Joker malware uses native code (JNI) instead of the more common DEX. Native code uses C for programming, which slows down the analysis of malicious code. By comparison, DEX, being a variation on Java code, is much easier to decompile into something human-readable. The malware uses this JNI code for sending SMS messages, to make money and as one way of contacting its command-and-control network. The use of JNI and out-of-band signalling over the phone network instead of the internet may help Joker evade automated dex scanners that don't speak JNI.

Joker clearly has developed an edge in the battle against Google's automated code review on new apps, and we see no sign that Joker will slow down in 2021 and may be joined by competitors before long.

Ads & PUAs increasingly indistinguishable from malware

Malicious advertisements (malvertising) remain a major source of threats to a range of devices. We recently delved into two current trends in malvertising threats that fall outside the realm of malware attacks— technical support scams using "browser locking" web pages, and ads targeting mobile devices that are linked to fraudulent or "fleeceware" apps. Sophos classifies these as "fake alert" attacks—malvertisements that attempt to scare their targets into taking an action that will enrich the scammers behind them.
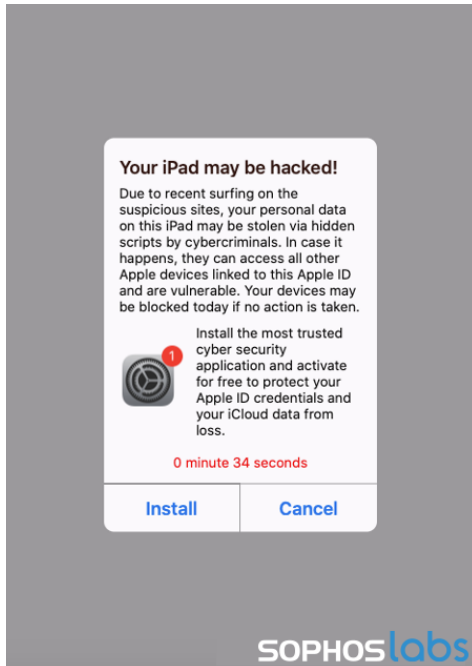

Fig. 24. Source: SophosLabs

Technical support scams typically attempt to steer targets into providing remote access to their computers and then convincing them to either purchase exorbitantly priced technical support software and services or obtain targets' credit card data for fraudulent purposes. While many of these scams have relied on direct telemarketing calls in the past, many scam operators have moved to a "pull" model—using malicious web advertisements that attempt to convince the user that their computers have been locked for security reasons, and directing them to call the scammers themselves.

To achieve this, the scammers deploy website kits containing scripts designed to make it difficult to navigate away from a page— including variations on the "evil cursor" (making the mouse pointer appear to be pointing somewhere it isn't, or rendering it invisible) and "infinite download" attacks that overwhelm the browser—while trying to look like an alert from Microsoft or Apple. Some of the kits we found exploited a bug SophosLabs' offensive security team discovered in Firefox earlier this year, while others executed similar attacks on other browsers—all of them being spread through malicious "pop-under" web ads.

Fig. 25. Source: SophosLabs

The same ad network infrastructure supporting these attacks on PC and Mac browsers also serves up tech support scams and fake alerts that link to potentially unwanted mobile applications—including apps claiming to be virtual private network services and "cleaner" tools advertised as removing malware, with built-in subscription fees (and in some cases, actual Android malware). Sophos found a collection of ad campaign servers delivering these ads, using commercial software from a Russian developer specifically built for running such campaigns.

Using your own strengths against you: Criminal abuse of security tools

Some attacks don't involve malware at all, or wait to deliver malware until the very end of the attack, instead using just the tools already on the operating systems running on computers across a network. Other criminals may leverage the power of a range of tools used by two large segments of the information security industry: incident responders and penetration testers.

The information security community has defined the style of attacks that involve very little or no malware, but instead harness the existing components of the operating system or popular software packages, pular software packages, as living-off-the-land (LOL). These attacks usually involve one or more forms of automation in the form of native scripting such as PowerShell, batch files, or VBScript scripts, collectively referred to as LOLscripts. The attackers use these LOLscripts to execute sequences of commands using living-off-the-land binaries (applications), colloquially called LOLbins.

Software originally designed for the "red team" segment of the industry comprises the bring your own attack method. Attackers, in this case, deploy and use off-the-shelf security tools that are commonly used by network administrators and penetration testers. These include tools such as Cobalt Strike and elements of the Metasploit framework, designed for use in security assessments and technical tests.

Netwalker threat actor toolset on the ATT&CK matrix

INITIAL ACCESS	EXECUTION	PRIVILEGE ESCALATION	DEFENSE EVASION	CREDENTIAL ACCESS	DISCOVERY	LATERAL MOVEMENT	IMPACT
Exploit Tomcat	PowerShell scripts	CVE-2020-0796	Fileless loading	mimikatz	SoftPerfect Network Scanner	psexec	Netwalker ransomware
Exploit Weblogic	psexec	CVE-2019-1458	Eset AV remover	Mimidogz	NLBrute	Teamviewer	Zeppelin ransomware
Phishing email		CVE-2017-0213	Gordon's Eset password recovery	Mimikittenz		Anydesk	Smaug ransomware
		CVE-2015-1701	Trend Micro's Security Agent Uninstall Tool	Windows Credentials Editor			Data exfiltration
			Microsoft Security Client uninstall	pwdump			
				NLBrute			
				LaZagne			
				WinPwn			

SOPHOSlabs

Fig. 26. The toolset used by a threat actor involved in Netwalker ransomware attacks used a panoply of open source, freeware and commercial utilities at different points during the attack. Source: SophosLabs.

These tools are valuable to attackers for a number of reasons: Because they're often used in a legitimate capacity (to audit or otherwise improve system security) it can also be difficult for anti-virus or security solutions to detect such tools or activity outright. As such, Sophos must rely more heavily on the study of the LOLscripted behavior to identify potential malicious activity. And of course, it's easier to use something that's already been created than to build your own tools from scratch.

While the use of LOLscripts and reverse shells wasn't new the past year, in 2020 they became a ubiquitous element of complex, manually operated ransomware break-in attacks. In fact, both the quantity and variety of the attack tools we observed during attacks seemed to increase.

Dharma RaaS Attack Tools Killchain

INITIAL ACCESS	EXECUTION	PRIVILEGE ESCALATION	DEFENSE EVASION	CREDENTIAL ACCESS	DISCOVERY	LATERAL MOVEMENT	EXFILTRATION	IMPACT
RDP credential spraying	PowerShell	CVE-2019-1388	Disables malware protection	mimikatz	PCHunter	Group Policy Objects	PowerShell screenshot emailer	Dharma Ransomware
Stolen RDP credentials	WMI	CVE-2018-8120	Revo Uninstaller	Remote Desktop Passview	Process Hacker	Remote Desktop	TOR	
	AutoIT	CVE-2017-0213	IOBit Uninstaller	LaZagne	GMER	WinRM Remote Management	dropmefiles[.]com	
	Command line/RDP			NLBrute	Advanced IP Scanner			
				Hash Suite Tools	NS2.EXE			

sophoslabs

Fig.27. Source: SophosLabs.

The wide variety of attack tools range from commercially available applications to open source GitHub repositories, with functionalities that may include:

‣ Botnet-like command-and-control frameworks

‣ Shellcode generation and obfuscation

‣ Anti-virus evasion and sandbox detection

‣ Password or credential extraction

‣ Kerberoasting (maintaining persistence of Domain Admin privileges)

‣ The ability to brute force passwords used by a variety of services

‣ System data exfiltration

Most of these types of tools contain benign payloads or no payload at all in their "out of the box" state, but in the past, we have been able to detect many of these tools engaging in malicious activity based on contextual information acquired via our behavioral detection technologies.

According to our telemetry, the ten attack tools we have seen most commonly in use are (in order of frequency of use) Metasploit, BloodHound, mimikatz, PowerShell Empire, Cobalt Strike, Veil Evasion, Hydra THC, Enigma, Nishang, and Shellter. Metasploit is far and away the most commonly seen tool, appearing about twice as often as the next most-common attack tool, BloodHound.

Sophos currently tracks the use of 99 different attack tools; it seems unlikely that we'll see a reprieve from attackers continuing to take advantage of these well-written tools throughout 2021.

Digital epidemiology

What percentage of computing devices are infected with undetected malware? What percentage of command line executions are executed by undetected adversaries? What percentage of targeted phishing emails go undetected? How do all these rates change as a function of industry, geolocation, and network posture?

Asking such questions is similar to asking, "what percentage of people are infected with COVID-19?" in a context in which many people might never get tested for the virus, those tests that are performed can have significant false positive and false negative rates.

In other words: it's hard.

Despite these challenges, epidemiologists answer critical questions about COVID-19 daily. Unfortunately, cybersecurity researchers fail to do the same for cyberattacks. We lag behind epidemiologists in the tools, techniques and procedures we've built to reason under uncertainty. There's no excuse, and it's time we built our own tools to understand the nature of the threat we face, accurately report risks to those we defend, and make decisions about where to direct our efforts.

To help in this mission, Sophos AI has embarked on the project of building a set of epidemiology-inspired statistical models for estimating the prevalence of malware infections in total. We combine a robust data collection pipeline that collects data from 100 million endpoints with a set of Bayesian statistical methods that let us tackle these difficult questions to build a complete picture of the performance of our models "in the field".

For example, consider the question: "How much malware is actually affecting our customers week over week, and how much of that are we detecting?"

If we already knew which files were malware and which were benign for all files, we'd already be done! Unfortunately, we have two problems.

1. We don't actually know the ground truth behind any given file – any endpoint product will miss at least some malware, and the occasional false positive (a normal file that gets flagged as malware) is inevitable

2. The balance between benign and malicious files is overwhelmingly tipped towards benign ones, so we probably can't figure it out with manual analysis – we'd have to do an in-depth analysis of thousands of files that our endpoint product labeled as benign to find a single malicious one

To address these problems, we turn to Bayesian statistics. In extremely simple terms, we build a "generative" model of the data: a mathematical program that can take guesses about parameters ("how much malware is there really?"), and turn those guesses into simulations of how many endpoint detections we might see. Then we try out different guesses, see which simulations match observed reality, and work backwards to find plausible values of the parameter we're interested in.

For example, imagine we have 2,000 endpoint detections and a good estimate of model true and false positive rates for a particular week. We can simulate worlds with malware rates of 0%, 2%, 5%, and so on, and see what the simulation predicts for endpoint detections; if we see close to 2,000 detections for some malware rates, then that's (perhaps) a plausible value.

Fig. 28. Propose a malware rate, sample, see if the simulation matches observed reality, tally the rates that do, and repeat. Source: SophosAI.

You can iterate this process millions of times to build a distribution of plausible values for malware rate, and because we use a Bayesian approach, error bars are "built in" to the estimate. In our example, the model thinks that the most likely value for "what percentage of files are malware?" is just over 3%, but anything from about 2.75% to about 3.35% is quite plausible.

And once we have a good idea of this number – how many files per hundred are likely to be malware across customer endpoints – missed detections and false positives become fairly simple to estimate. If we look at data from our deep learning ML-based malware detection system for a week in May (without any signature-based, behavioral, or heuristic options turned on) we can flesh out a complete matrix of true and false positives and negatives, and complete our picture of model performance. In this case we see that while we do have some false negatives, the number of false positives is low and skewed towards zero, while the number of true positives is high and skewed towards 161,000 (the total number of positive results in the sample). Looking at the scale, we can see that all three quantities are dwarfed by the number of true negatives – benign files that our ML labeled as benign.

Our epidemiology-inspired tool has allowed us to estimate, if not find, the needles in our PE file haystack.

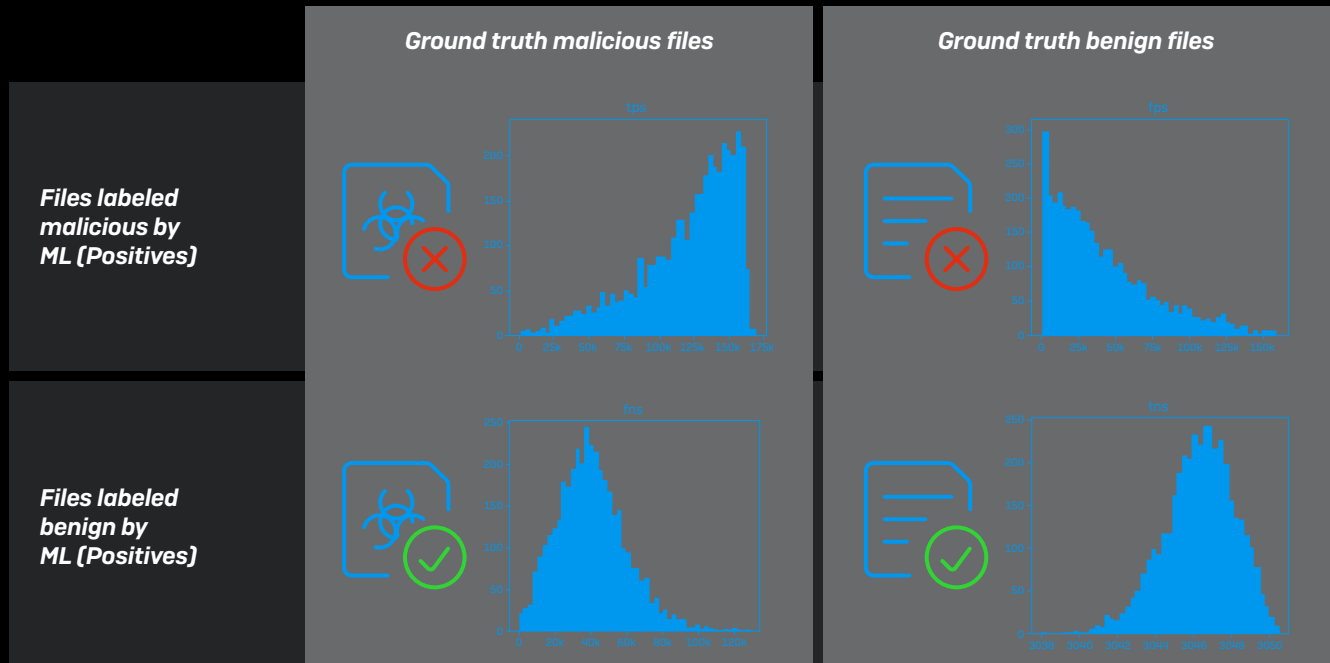|	**Ground truth malicious files**	**Ground truth benign files**
Files labeled malicious by ML (Positives)		
Files labeled benign by ML (Positives)		

Fig. 29. Analysis of ML model true and false positives in early May 2020. Source: SophosAI.

United Kingdom and Worldwide Sales
Tel: +44 (0)8447 671131
Email: sales@sophos.com

North American Sales
Toll Free: 1-866-866-2802
Email: nasales@sophos.com

Australia and New Zealand Sales
Tel: +61 2 9409 9100
Email: sales@sophos.com.au

Asia Sales
Tel: +65 62244168
Email: salesasia@sophos.com

20-11-01 EN (DD)

SOPHOS