Reining in Ransomware

Investigative Cybercrime Series: Vol 2

A Collaboration Between

Arete Cyentia Institute

Table of Contents

The Investigative Cybercrime Series is an ongoing research effort to unmask insidious cyber threats and lessen their impact on insurers and the organizations they cover.

The data for this research comes directly from security incidents investigated by Arete and the intelligence operations that support those investigations.

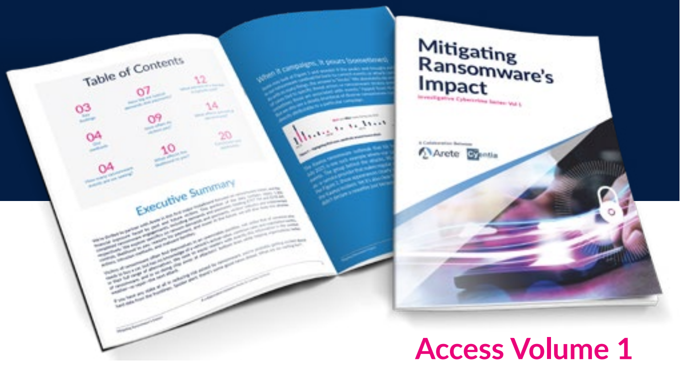

Access Volume 1

Executive Summary

What are the most common strains of ransomware over the last few years? Which strains are demanding (and receiving) the highest ransoms? How do the most "successful" ransomware groups infiltrate their victims? What techniques do they use to spread around the network? How (and how often) is data exfiltrated in addition to being encrypted? And most importantly—what can organizations do to prevent, contain, and otherwise rein in the most virulent ransomware?

If you want data-driven answers to any of those questions, this report is for you. By "data-driven," we mean that this isn't one of those fluff pieces in which we stand on our cyber expert soapbox and prattle off opinions about how you can live your best ransomware-free life in just five steps. Instead, we've deeply analyzed on-the-ground evidence collected while responding to nearly 1,500 ransomware events exceeding $1 billion in ransom demands. Through it all, we've helped our clients manage their response, minimize costs, and maintain business operations. And we hope this report helps many other organizations and insurers do the same.

In terms of agenda, we start with trends surrounding the most prevalent ransomware families and the industries they target. Then we examine the tactics, techniques, and procedures (TTPs) employed during three key phases of a ransomware attack: initial infection, post-compromise spread and impacts, and then methods of data exfiltration. For each phase, we prioritize recommended practices based on their scope of effectiveness against the TTPs involved. Enough talk; here's a few key findings and then let's get to it..

Key Findings

Seven of the top 10 ransomware strains so far in 2022 weren't in the top 10 last year. That indicates swiftly-changing dynamics among cybercriminals and their campaigns.

In 61% of attacks, ransomware initially infected victims by exploiting poorly-secured remote access services. That vector will only become more prominent as digital transformation continues.

The top 10 post-compromise techniques each factor into >50% of ransomware incidents. That's a lot of tools but also a lot of opportunities for detection and containment.

Ransom demands are 5X higher when data exfiltration is involved. And that's happening six times more often in 2022 than in 2019.

Ransomware Families

As mentioned in our **previous report**, ransomware has come a long way since it first arrived on the scene in 1989. Over the past 33 years, much has been written about ransomware's rise and possible explanations for it, like its ease of distribution, shortened path to monetization, parallel growth in cryptocurrency, etc. Our data reflects not only its prominence, but its growth as well.

Our data around ransomware shows a slow and steady climb through 2021, peaking in September. Curiously, however, we have since observed a steady dip back to early 2020 levels. This is a trend that we are starting to notice across the industry—**SonicWall's Cyber Threat Report** even reports a 23% decline in global ransomware attacks in the first half of 2022. There are a few things that we can look to as likely contributors in this new pattern, including but not limited to, the Russian-Ukraine war, inflation and the thought of an impending recession, and international sanctions against various crypto markets and countries. Along with these, we are also seeing increased regulations from governments on payment to threat actor groups. Only time will tell if this is truly a trend, or if it is merely a restabilization of a new normal.

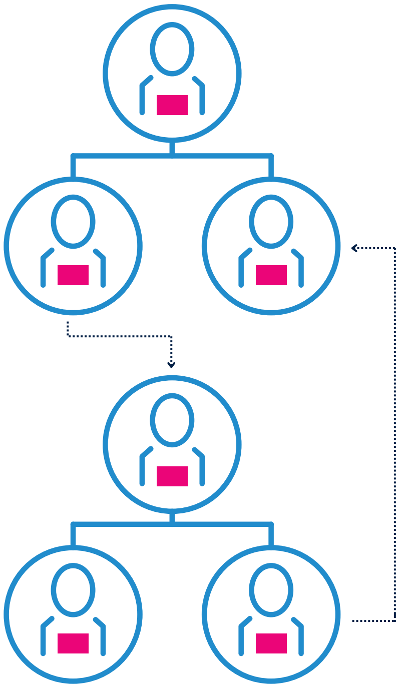

Takeaway for Insurers

It's no secret that ransomware events and claims have increased over the last few years. But multiple sources—including Arete's own caseload—point to at least a temporary downturn. That in no way means ransomware gangs have abandoned their schemes, but we are seeing a changing of the guard that will certainly influence future developments.

What are the most common ransomware families?

Let's begin to dive a little deeper into the dark world of ransomware families. With the proliferation and continuous development of Ransomware-as-a-Service (RaaS), we are seeing not only an increase in ransomware families, but also in the number of "family members" within each family.

In each case, investigators determined which family of ransomware was involved in the attack. In the next chart, we can see the top 20 families that were involved over this time period. What's interesting to note, is that these top 20 families account for 71% of all cases.

REvil, a Russian-based Ransomware-as-a-Service provider, is nearly twice as common as any other family on this top 20 list. As a RaaS provider, REvil arrived on the scene in 2019 as an evolution of GandCrab ransomware. In 2020, they quickly launched a few high-profile attacks, including one on the law firm Grubman Shire Meiselas & Sacks that represented then-U.S. President Donald Trump, Lady Gaga, and Madonna.

In July 2021, REvil returned to the public consciousness by exploiting zero-day vulnerabilities in Kaseya. Shortly after the media hype around Kaseya, REvil quietly disappeared and their websites were taken offline. They haven't been active since. Conti—another group that was recently shut down, in May 2022—and RYUK round out the top three ransomware families that we see responsible in overall cases.

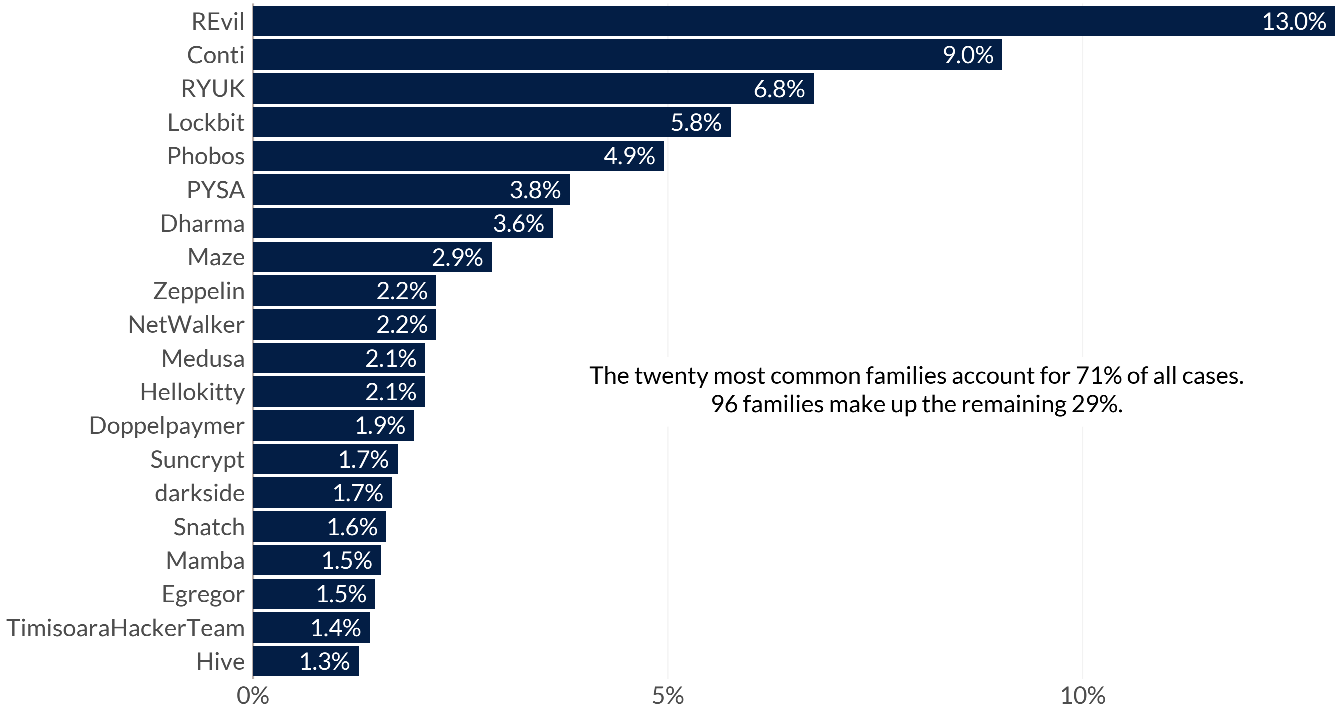

The twenty most common families account for 71% of all cases.
96 families make up the remaining 29%.

Figure 1 — Percent of cases among the overall top 20 ransomware families

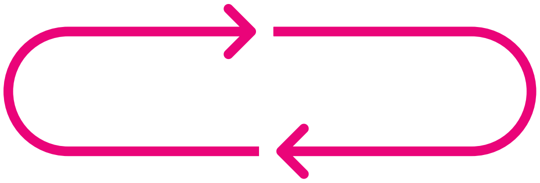

We see a lot of these family names over and over again, and more importantly, a lot of familiar code shared among them. As mentioned before, REvil was the evolution of GandCrab, and was able to quickly utilize already existing (and successful) code in order to make its attack.

RYUK, which came onto the scene in 2018, is a well-known variant of Hermes ransomware. RYUK fell off in recent years as Conti—a notorious group that appeared in 2020—climbed the charts. Conti ransomware worked by exploiting Microsoft Windows vulnerabilities, and no industry escaped its crosshairs. By using phishing campaigns to get access to the victim's system, it attacked the likes of Apple and Tesla, and even government resources like Ireland's Health Service and the city of Tulsa. In May 2022, the members of Conti disbanded shortly after trying to overthrow the Costa Rican government. However, that does not mean the attackers have retired. Many of Conti's members are suspected to have thrown in with other ransomware groups, including: BlackBasta, BlackByte, AvosLocker, Hive, and BlackCat.

Takeaway for Insurers

The top 20 ransomware families account for a large majority of incidents, meaning there's an opportunity to greatly reduce risk by encouraging organizations to focus defenses on those. The rest of this report will help do exactly that!

PYSA—or "Protect Your System, Amigo"—has also risen on the chart below. As the successor of Mespinoza, PYSA first arrived in December 2019, and seems to prefer targets in the education and government sectors in the United States and the UK. Its preferred method is known as "Big Game Hunting", meaning they specifically target high-value information and assets in victims that are sensitive to any type of downtime.

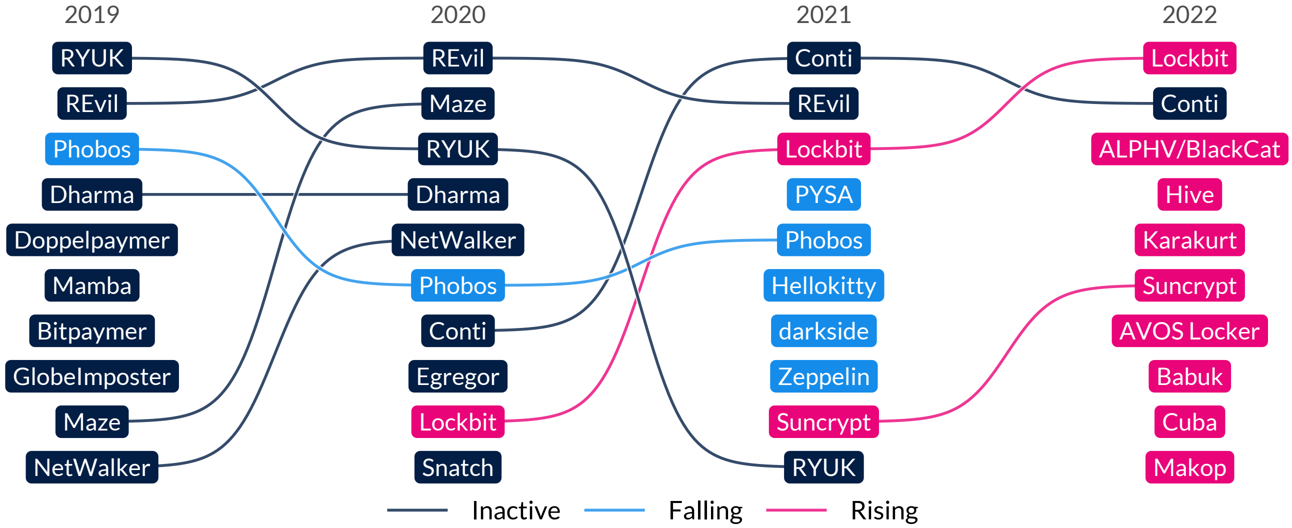

Figure 2 — Top 10 ransomware families by year

Looking at the top 10 ransomware families by year, it's not surprising to see the appearance of some of Conti's purported offshoots on this list—Hive, BlackCat, and AvosLocker. The spread of Ransomware-as-a-Service (RaaS) has allowed ransomware like Lockbit and Suncrypt to rise in prevalence. Makop will be one to keep an eye on—it is a trending piece of ransomware that asks users to contact the attacker via Tox (a P2P text messaging application). Once in contact, the malware then encrypts all the files until payment is received.

Takeaway for Insurers

Seven of the top 10 ransomware strains so far in 2022 didn't make this list last year. That indicates a high degree of churn among these groups, which makes it more critical to monitor trends to manage risk.

DO RANSOMWARE FAMILIES TARGET SPECIFIC INDUSTRIES?

A common question that is asked is "does certain ransomware only target certain industries?" As we can see in the next chart, it looks like the attacks are fairly consistent, which leads us to believe that there isn't any real targeting by industry. RYUK does seem to lean more heavily towards healthcare, public sector, and tech. RYUK, unlike REvil and Conti, doesn't make money off of offering RaaS, but operates more of a direct-to-victim extortion model.

	Critical Infrastructure	Financial Services	Healthcare	Manufacturing	Professional Services	Public Service	Retail	Tech, Eng, Social Media
REvil	24%	28%	30%	29%	23%	18%	30%	20%
Conti	18%	19%	11%	23%	17%	13%	30%	11%
Phobos	8%	8%	8%	10%	12%	7%	9%	4%
RYUK	8%	3%	22%	9%	11%	16%		17%
Lockbit	11%	14%	9%	8%	10%	13%	9%	11%
PYSA	8%	8%	5%	4%	7%	8%	9%	13%
Dharma	11%	6%	6%	6%	6%	9%	4%	4%
Maze		11%	3%	5%	5%	7%	4%	7%
Zeppelin	11%	3%	6%		5%	5%	4%	
NetWalker	3%		2%	6%	3%	5%		13%

Figure 3 — Overall top 10 ransomware families by industry

So, wait—if ransomware doesn't seem to be strongly targeted by industry, what does this mean for organizations trying to protect themselves? As we'll see, one may not be able to predict which specific family is dealt with, but there are times when family matters, and the group of popular families can (and should!) inform the prioritization of controls.

Arete does expect some geopolitical affiliations to eventually play out—e.g. targeting of financial, energy, and critical infrastructure sectors by pro-Russia threat actors in retaliation to sanctions. Furthermore, threat intelligence can really pay off in some scenarios, as some threat actors assert they will not extort organizations in specific sectors. Arete has occasionally secured decryption tools for free after advising an attacker of their victim's membership in an "off limits" industry.

Takeaway for Insurers

Each ransomware family represents a different set of geopolitical affiliations, guiding motives, common techniques, and demand strategies, etc. That means understanding who's on top, how they work, and what they want is critical to effectively managing risk across your portfolio of insured organizations.

FREQUENCY & IMPACT

Before we get too far into talking about ransomware demands and payment, it's important to note that the U.S. Treasury Department's Office of Foreign Asset Control (OFAC) still advises companies not to make any ransomware payments. Paying a ransom to an organization that is known to be located in a country that is under U.S. sanctions is illegal. OFAC encourages organizations that find themselves victims of a ransomware attack to report it to CISA and their local FBI office.

Let's take a moment to look at what portion of overall events and payments are represented by the most common families. Via this risk-like view, REvil clearly occupies the upper right corner—signaling high frequency among cases along with a large share of all payments. When we look at where the families land in respect to the dotted line, we can start to see if there are any unusually high or low frequency-to-cost instances. For example, darkside—along with others in fuchsia—are on the far-left side of the chart, representing relatively few cases but outlandish payments.

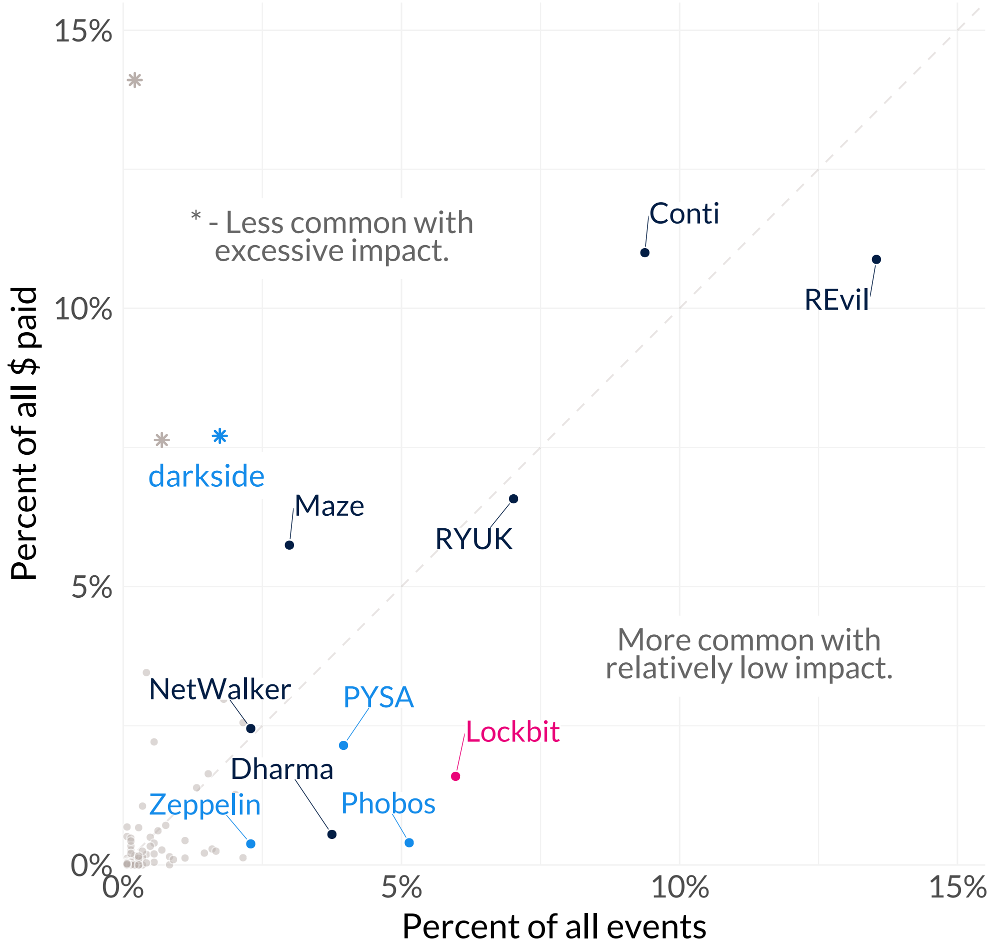

However, when victims receive extremely large demands, it raises a whole array of concerns, including whether or not they should even try to negotiate or pay—consistent with our findings in **volume 1**. This strategy doesn't typically last very long; after all, criminals want your money and if exceedingly high demands do not produce results, they tend to quickly pivot their strategy.

Figure 4 (Left) — Portion of all events and payments among highlighted families

Takeaways for Insurers

Like legit companies, ransomware groups have a business strategy. Some maximize per-unit cost (darkside), some go for low-cost volume (Phobos), and others seek a balanced approach (RYUK). It will help your clients to know who they're dealing with.

DEMANDS & PAYMENTS BY FAMILY

Besides their aggregated frequency and impact, families also vary widely across typical demands and payments. Each point in the next chart represents an intersection of a family's typical (read: geometric mean) demands and payments; the outlined area is the 50% highest-density region (read: the most likely 50% of values) which fades to 90% at the edges. Clearly, any negotiation strategy should take into account the family involved to minimize a victim's losses.

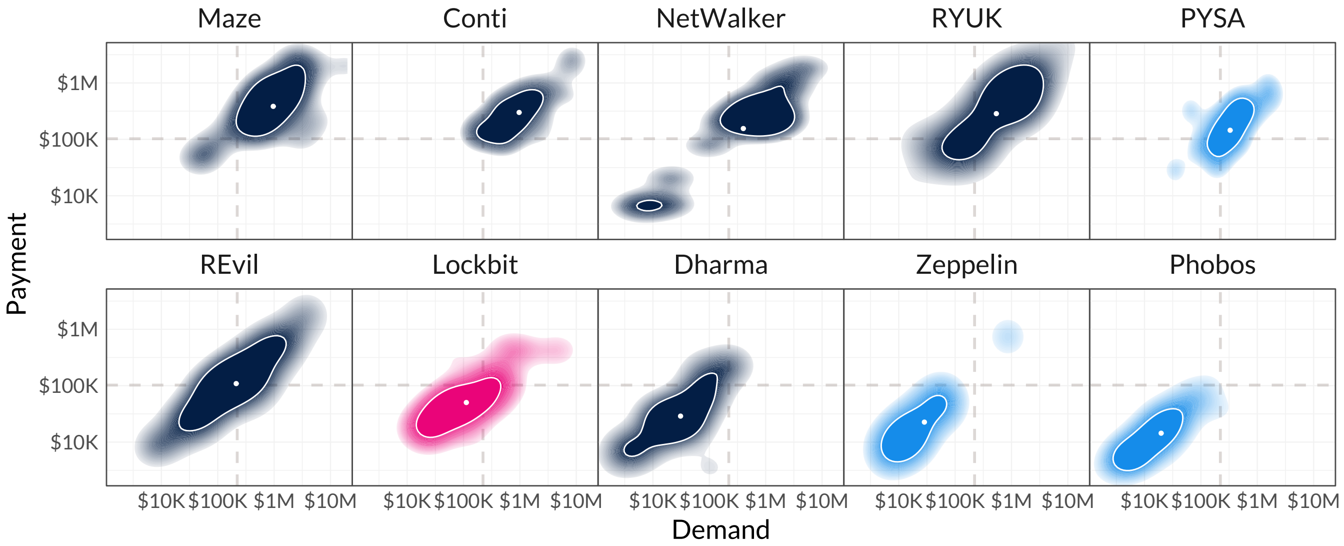

Figure 5 — Typical demands and payments among the most common ransomware strains

Maze and Conti consistently demanded and received above-typical amounts, denoted by the dashed lines for the overall values. NetWalker is a good example of shifting negotiating tactics: its two distinct regions represent a change around April 2020 toward significantly higher demands and payments. REvil's influence on the landscape—and our dataset—is observable in how much of the demand and payment range it spans. We also see that there is a strong correlation between demands and payments across each family—i.e., none has an inordinately high payment ratio.

Ransomware Capabilities & Defenses

Now that we know more about the trends and demand tactics of the top ransomware families, let's shift our attention to the modi operandi behind their schemes. We'll leverage evidence gathered by our investigators along with intel from cyber threat analysts to examine the top tactics and techniques employed by the most common ransomware. We've organized these into three key stages of a ransomware event: infection vectors, post-compromise, and data exfiltration.

But knowing what ransomware does is only half the battle. The other half is doing something about it. Thus, we'll assess the most effective defensive strategies for combating ransomware at each stage of a ransomware event based on common techniques.

*The methods and mitigations presented in this section are based on the **MITRE ATT&CK framework**. We've done this partly because ATT&CK is quickly becoming the common language of threat tactics and techniques used across the cybersecurity industry. But another benefit of using ATT&CK is that it enables readers to easily find definitions and examples of each technique referenced as well as explore a wealth of information on associated **threat groups**, **malicious software**, **mitigations**, **attack simulations**, etc.*

How does ransomware infect victims?

During a ransomware investigation, Arete responders take special care to determine the infection vector (or initial access technique in ATT&CK parlance). Everything else that happens afterwards hinges on attackers successfully introducing malware into the victim's environment, so preventing that from happening in the first place would be ideal. Understanding common infection vectors will help focus those preventive strategies, and identifying those is our goal here.

As per Figure 6, **external remote services** are the most common infection vector observed across ransomware cases. That's not terribly surprising if you consider prevailing trends that involve growing numbers of remote workers and third parties who need to access corporate systems. Unfortunately, remote access services provisioned for authorized users also present an opportunity for abuse. And ransomware gangs are clearly availing themselves of that opportunity in the majority of attacks.

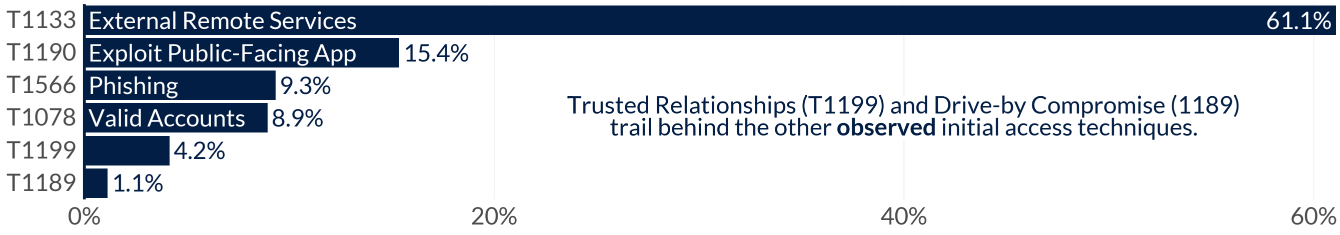

Trusted Relationships (T1199) and Drive-by Compromise (1189) trail behind the other **observed** initial access techniques.

Figure 6 — Infection vectors observed across ransomware cases

The **exploitation of public-facing applications** is next on the list, observed in 15% of ransomware cases. It's worth noting this technique is moving up the list over the last few years and currently sits in first place for 2022. Attackers typically take advantage of unpatched vulnerabilities that are easily identified through remote scanning. Openings introduced through administrative misconfigurations are also routinely taken advantage of to gain access to the environment.

The techniques in Figure 6 are those observed and deemed by investigators to be the primary infection vector. But most ransomware doesn't have just one trick up its sleeve to accomplish this. To emphasize that point, this table shows intrusion capabilities possessed by the 20 most common ransomware strains based on community intelligence collected by MITRE.

ID	Technique	Cases (%)
T1566	Phishing	36.8%
T1189	Drive-by Compromise	18.5%
T1078	Valid Accounts	17.2%
T1190	Exploit Public-Facing Application	15.1%
T1133	External Remote Services	14.7%
T1091	Replication Through Removable Media	12.2%

*Think of these techniques as a set of options that ransomware *could use* to gain access to your environment. With that in mind, it's worth noting that the top three differ from those in Figure 6. Thus, it's not a bad idea to add **phishing**, **drive-by compromise**, and **valid accounts** (stolen credentials) to your threat model if not included already.*

DO INFECTION VECTORS DIFFER BY FAMILY?

Short answer—Yes. But probably not as much as you might think. Earlier we talked about the similarities that many ransomware families have, thanks to shared code and even shared members. But the similarities don't stop there; most are content to mimic what works for others, leading to a convergence of go-to intrusion techniques among the top ransomware families as seen in Figure 7.

	Conti	Dharma	Lockbit	Maze	NetWalker	Phobos	PYSA	REvil	RYUK	Zeppelin
External Remote Services	36%	95%	67%	48%	93%	87%	82%	69%	46%	74%
Exploit Public-Facing Application	39%	2%	14%	4%		2%		7%	3%	4%
Phishing	19%		5%	30%				2%	30%	4%
Valid Accounts	7%		14%	4%		10%	14%	6%	19%	9%
Trusted Relationships		2%			7%	2%	5%	17%	3%	9%
Drive-by Compromise				15%						

Figure 7 — Infection vectors among overall top 10 ransomware families

Targeting remote services is a strongly favored technique used by all families. Conti also relies heavily on exploiting vulnerabilities in applications. Maze and RYUK apparently enjoy regular phishing expeditions, but occasionally swap over to drive-by compromises or abusing valid accounts, respectively. Our overall takeaway for the typical enterprise defender is to focus on locking down common infection vectors rather than getting caught up in all the family drama. The prevention strategies listed in the next section will help with that.

TOP RANSOMWARE PREVENTION PRACTICES

What can organizations do to prevent ransomware from infecting their systems? Figure 8 lists recommended practices based on **ATT&CK mitigations** associated with the initial access capabilities exhibited by the top malware families across Arete's caseload. The percentages are based on the proportion of incidents potentially thwarted by each practice.

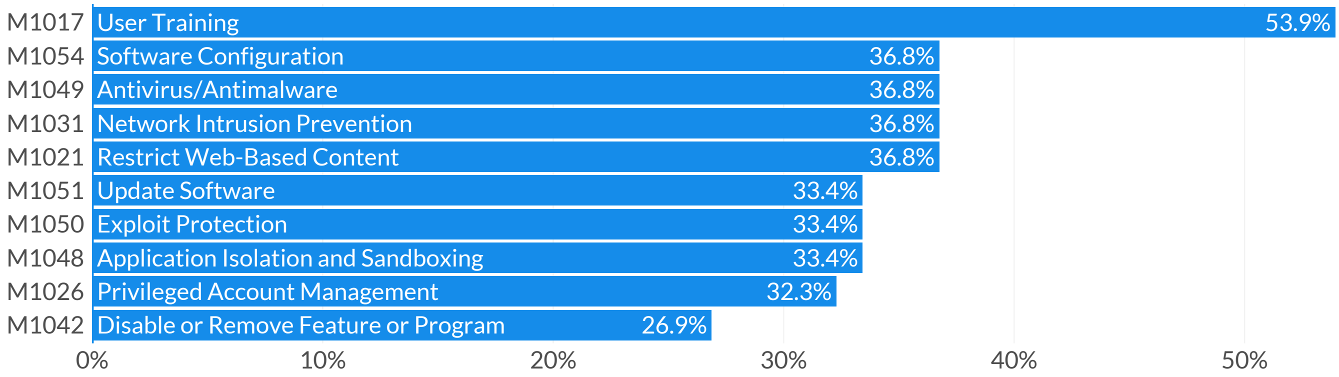

Figure 8 — Top practices for preventing common ransomware infection vectors

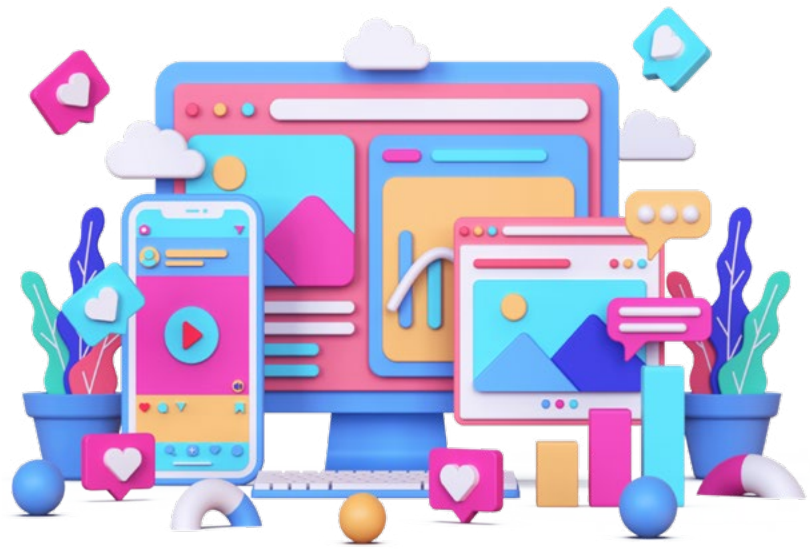

User training, specifically around common social engineering schemes and promoting norms of healthy skepticism might have helped in more than half of these cases. But we can't lay all the responsibility at the feet (fingertips?) of those who mistakenly take the bait. The four-way tie for second place among preventive controls goes to **software configuration**, **antivirus/antimalware solutions**, **network intrusion prevention**, and **restricting web-based content**. An honorable mention goes to multi-factor authentication (M1032) which, despite not making the top 10, still has plenty to offer defenders. Besides being associated with a lower likelihood of payment and lower percent of demand paid (**see volume 1**), it can help limit post-compromise spread (as we'll see in the next section).

Note that each of these defensive measures, along with all the others in Figure 8, are things that can neutralize ransomware despite the opening of a dangerous link or attachment. They'll also harden those oft-targeted remote services and public-facing applications to help organizations minimize their external attack surface. That means they're less likely to become an easy score for ransomware gangs probing for open doors into corporate networks.

Takeaway for Insurers

Many ransomware recommendations tend to focus "right of boom" once ransomware has already infected an environment (i.e., data backups). They also rarely prioritize defenses based on the scope of potential efficacy. Guiding your insured to implement preventive measures that thwart common infection vectors will reduce the likelihood they ever receive that dreaded ransom note.

What does ransomware do after initial infection?

Let's go beyond the initial infection vector and take a look at the fuller set of capabilities utilized by ransomware to carry out attacks. At the tactical level, this includes techniques to **maintain persistence** in the victim's environment, **escalate privileges** to gain more access, **discover** additional target systems and data, **move laterally** across the internal network, **evade security defenses**, establish **command and control channels**, **collect** and encrypt data, and other costly **impacts**.

Figure 9 ranks post-compromise techniques associated with the most common ransomware strains encountered by Arete. The percentages correspond to the proportion of cases involving ransomware possessing each capability. Since these techniques ostensibly contribute to the success of top campaigns, they offer a forewarning of what ransomware might attempt should an infection occur in your systems.

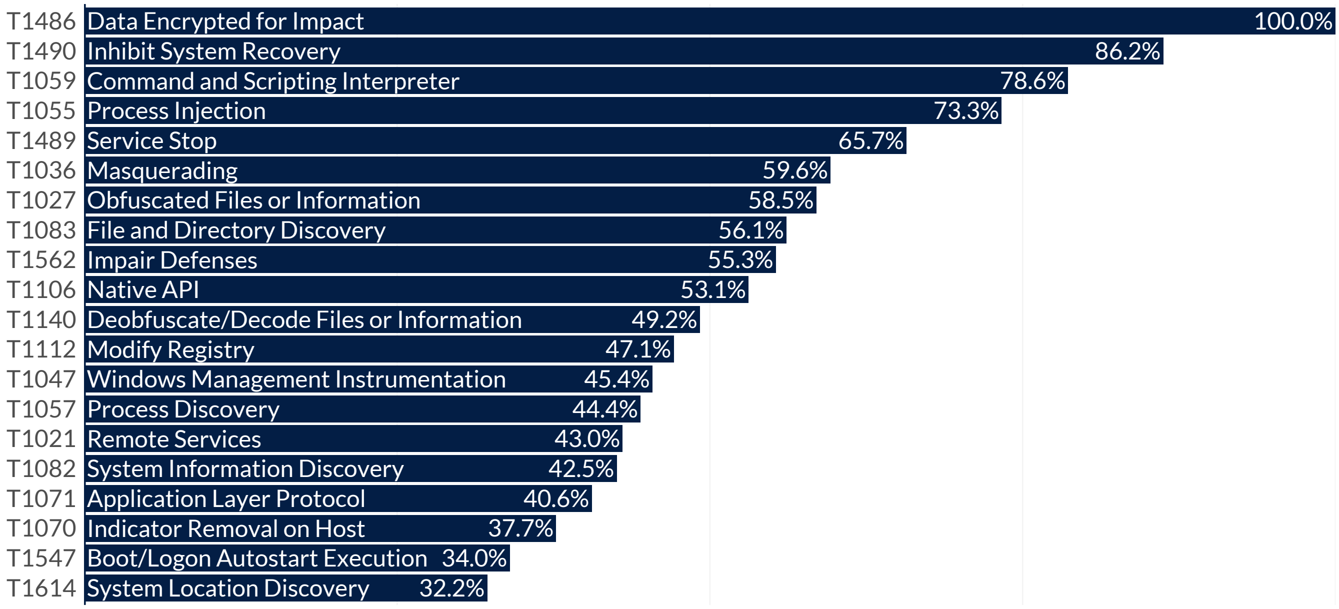

Figure 9 — Post-compromise techniques associated with the most common ransomware strains

And the award for least surprising ransomware technique goes to… **data encrypted for impact**, possible in 100% of cases involving the most common ransomware (although systems were encrypted in just over 96% of them). Kinda hard to demand payment for decryption if it's not encrypted, huh?

Next on the list of top capabilities is **inhibit system recovery** at 86% of incidents. Together with a slew of others in Figure 9 like **service stop**, **masquerading**, **obfuscated files or information**, **modify registry**, and **impair defenses**, the pervasiveness of these techniques demonstrate that ransomware gangs have a vested interest in remaining undetected and disabling their victims' ability to respond during an attack.

Along with the number three technique (**command and scripting interpreter**), **native API** and **Windows Management Instrumentation** constitute ways to execute malicious code on infected systems. Doing that successfully opens the door to a multitude of additional nefarious capabilities, worsening an already bad situation for the victim.

Fourth down the list is process injection, a technique still in play for nearly three-quarters of ransomware incidents. It's one of the primary ways ransomware escalates privileges in order to launch many of the other malicious techniques in Figure 9. Why is that important? Because the chance of victims having no recourse but to pay the ransom is substantially reduced for ransomware events that do not attain privilege escalation (see Figure 10). Interestingly, we've observed a small but measurable decline in privilege escalation in recent years (from 96% of cases in 2019 to 82% in 2022).

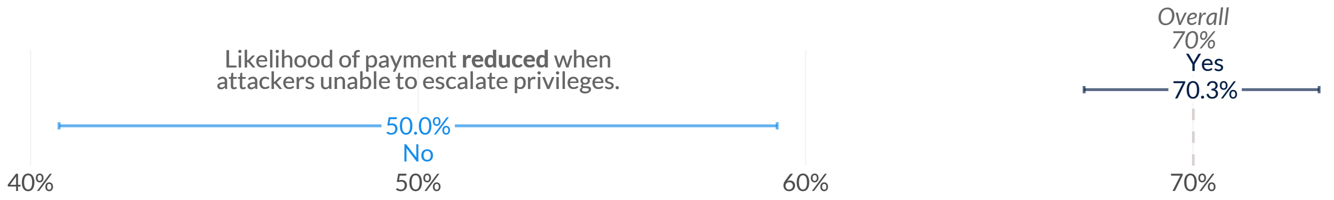

Figure 10 — Effect of privilege escalation on the likelihood of ransom payment

Before moving on from Figure 9, let the multitudinous nature of these post-compromise capabilities sink in. Even 20 deep, these techniques still factor into a third of ransomware incidents. That essentially means the top ransomware strains have a proverbial Swiss Army knife of tools at their disposal to accomplish their aims. We realize that's not terribly comforting to defenders, but the flip side of that fact is each of these techniques presents an opportunity for ransomware to be neutralized. Tips on doing just that come next.

TOP RANSOMWARE CONTAINMENT PRACTICES

What are the best approaches to containing the spread and damage of ransomware once it's gained access to your environment? To answer that, we pulled together recommended mitigations that address post-compromise techniques utilized by the most prevalent ransomware. These practices appear in Figure 11 and are sorted according to the percentage of ransomware incidents they potentially counter or contain.

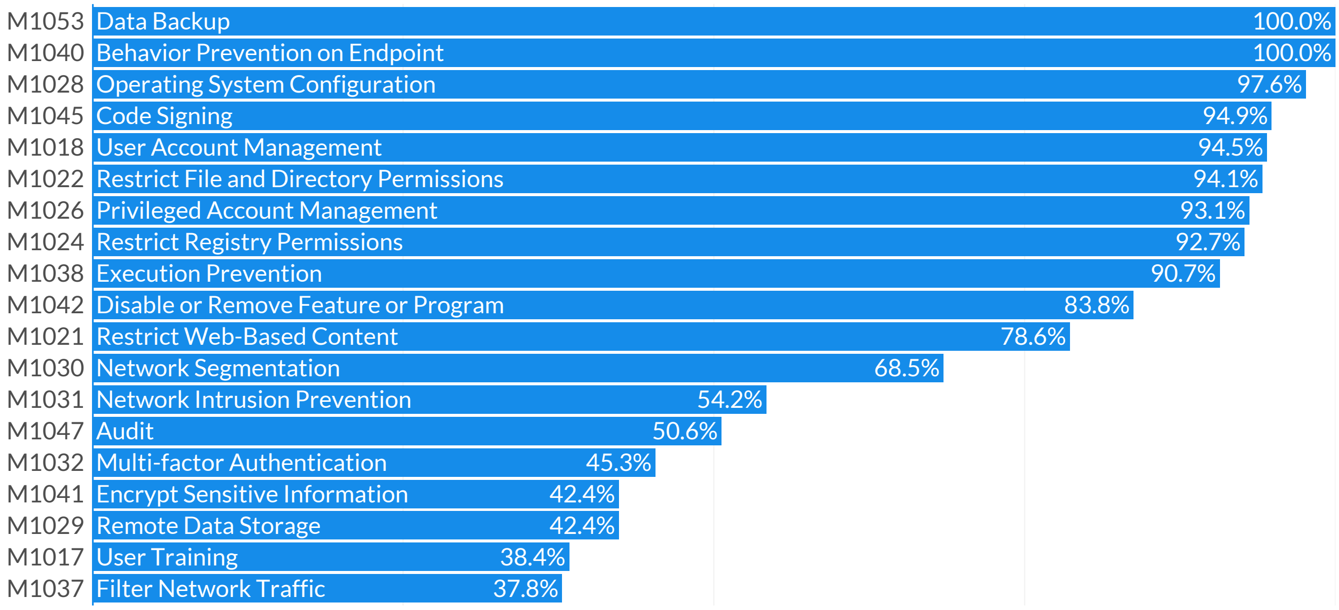

Figure 11 — Top practices for containing common ransomware post-compromise techniques

Regular data backups almost go without saying as a top practice for mitigating ransomware events, but we'll say it anyway. We'll also reiterate what we learned in our last report about the efficacy of backups. Namely, the mere act of having backups has no effect on the likelihood of payment, BUT the ability to fully recover data DOES reduce payouts by 20%. So, make sure those backups are ready to go when needed.

The other option for potentially mitigating all the top ransomware strains is behavior prevention on the endpoint. According to MITRE, **behavior prevention at the endpoint** essentially suppresses any processes, files, API calls, etc. that raise a red flag. Several others on the list like **code signing**, restricting **file and directory** or **registry permissions**, and **execution prevention** offer additional layers of suppression functionality. That's good because as we observed in Figure 9, ransomware attempts all kinds of sketchy stuff on infected systems.

We also note mitigation strategies focused on both **highly-privileged** and **user-level** account management. While this may seem like a no-brainer, there are plenty of newsworthy examples where user account creation defaults to top privileges, which are then exploited. Organizations should make sure that their account management programs and subsequent permissions are continuously reviewed.

Network segmentation is worth calling out with regard to containing ransomware infections. Implementing physical and logical segmentation goes a long way to curtailing the ability of ransomware to wreak havoc across the environment.

It's often said that humans are the "weakest link in cyber security," and while there's an element of truth to that claim, the top containment strategies skew heavily toward process and technology controls. That doesn't mean, however, that you can't loop end users into the fight to rein in ransomware. **Multi-factor authentication** is often prescribed to prevent unauthorized access into the network, but it's also associated with lower likelihood to pay and percent of demand paid (**see volume 1**). This is potentially due to its ability to combat privilege escalation and lateral movement after the initial infection. Furthermore, **training users** to follow good security practices will help those process and technology controls be more effective.

Takeaway for Insurers

The top ransomware strains have no shortage of tools at their disposal once inside a victim's network. That's not great news for your clients, but the flip side of that fact is there are over 20 practices that can still mitigate substantial amounts of risk following the initial compromise.

How does ransomware exfiltrate data?

We all know that ransomware encrypts data and holds it for ransom. But it's becoming increasingly popular among ransomware gangs to also steal sensitive data from their victims and threaten to release it unless they pay up—see our **previous report**'s section on payment reasons over time for more info. In this section, we'll see how often data exfiltration occurs and what techniques are used to accomplish it.

HOW COMMON AND COSTLY IS DATA EXFILTRATION?

The tactic of exfiltrating data in addition to encrypting it is often referred to as "double extortion." But our analysis suggests that "quintuple extortion" would be more accurate.

"How's that?", you ask. Well, take a look at Figure 12. We've observed that the typical (geometric mean) ransom when data is exfiltrated from the victim's environment is five times higher than when data is just encrypted in place! What does this mean? Basically, criminals know that outing your data would be embarrassing (or even highly damaging) and they're banking on that fear enticing victims to pony up more cash.

Arete was involved in a case where the threat actor posted a multi-million-dollar ransom and provided proof of data exfiltration via a list of over a hundred thousand files over thousands of pages. In another instance, Arete was involved in a case where the threat actor posted a ransom of about $2,000. Though the threat actor claimed to have taken data, they failed to provide any proof that any data was exfiltrated. The Arete Forensics team was able to confirm that there was no evidence of data exfiltration after an in-depth investigation.

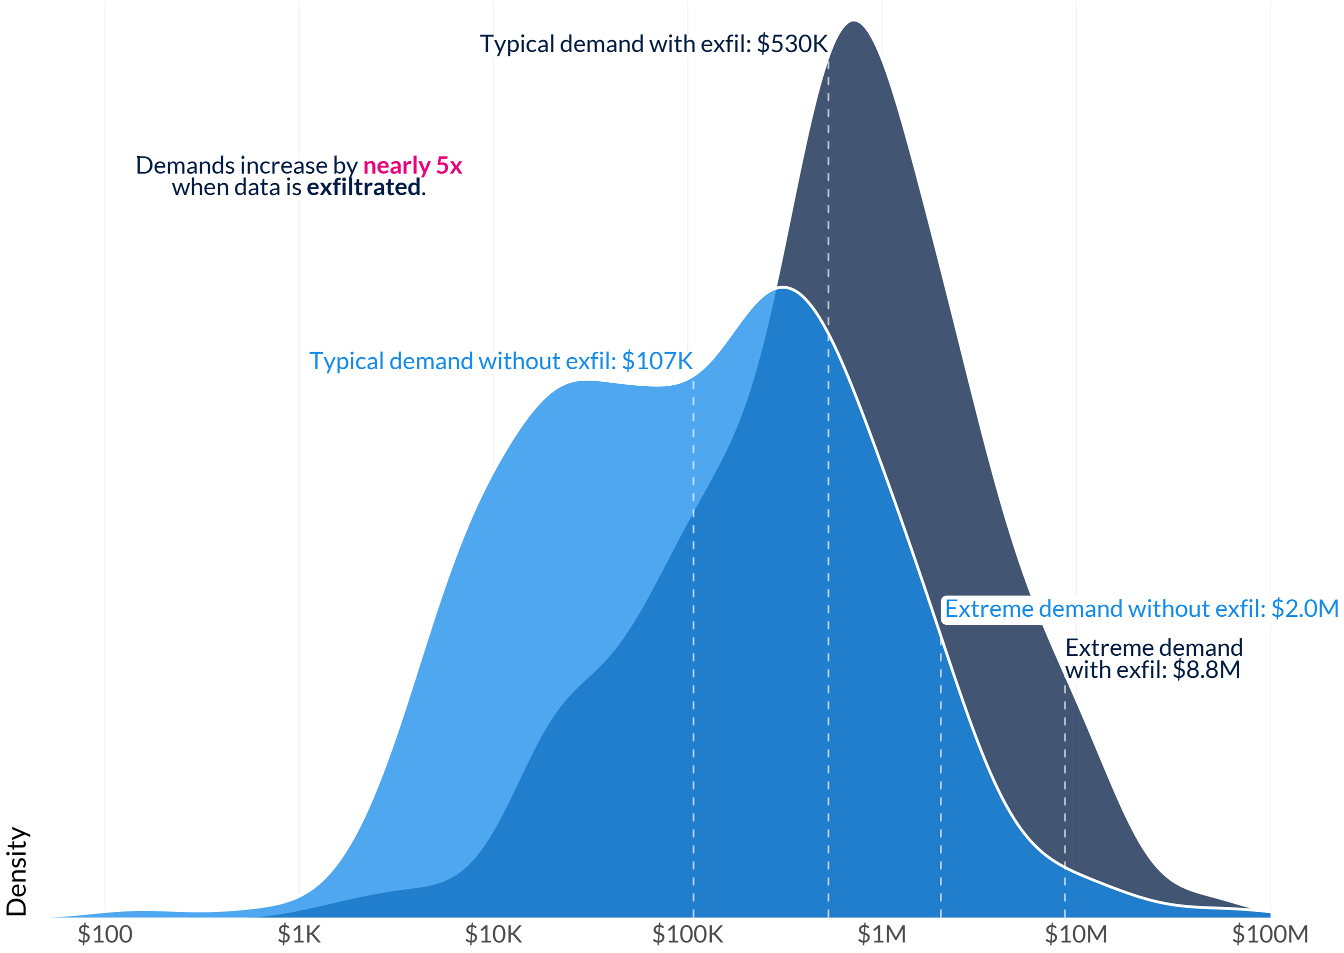

Figure 12 — Initial demand vs. data exfiltration

As if inflating the ransom weren't bad enough, our caseload also suggests that data exfiltration is occurring more and more often (see Figure 13). Back in 2019, a mere 9% of our investigations uncovered evidence that data was taken from the victim's environment. Over the next few years, that's risen nearly six-fold to 53% of our cases so far in 2022!

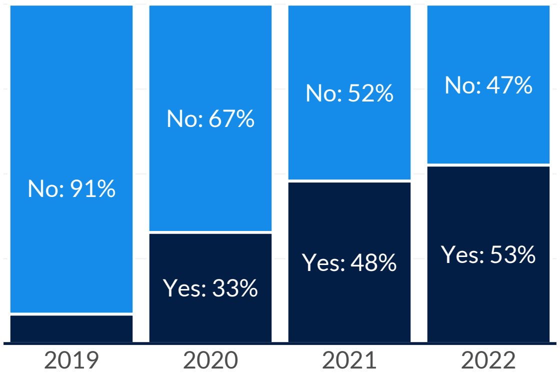

If you're wondering whether ransomware groups differ in their use of data exfiltration as a tactic, the answer is yes, and widely. We've observed data exfiltration in over 90% of Maze-related investigations and two-thirds of those tied to Conti. Lockbit, Netwalker, PYSA, and Zeppelin all hover around the 50% mark, while Dharma, Phobos, and RYUK fall below 10% for data exfiltration. As ransomware attacks pivot into more conventional data theft and extortion, how are criminals making their getaway?

Figure 13 (Left) — Exfiltration over time

COMMON DATA EXFILTRATION TECHNIQUES

MITRE ATT&CK includes **nine primary techniques** used by threat actors to exfiltrate data. It's a tad surprising then, that the top ransomware strains encountered across our caseload are known to utilize just five of them. Those techniques are featured in Figure 14 based on the proportion of incidents that involved ransomware known to use them.

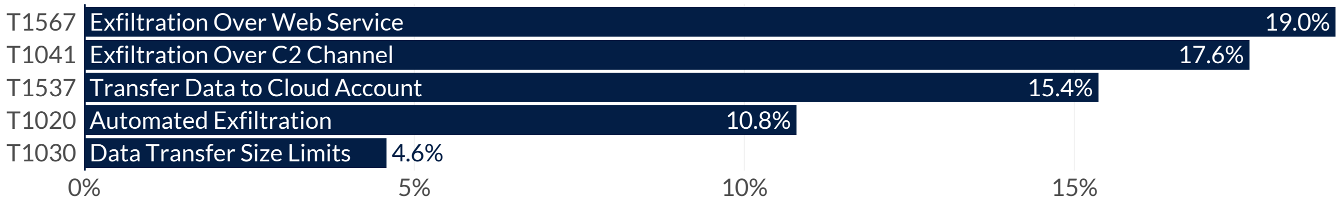

Figure 14 — Data exfiltration techniques associated with the most common ransomware strains

The most common capability for getting data out of the victim's environment is piggybacking **over web services**, hiding exfiltration in plain sight. Since such services are heavily used for legitimate data transfer, it's less likely that a few extra packets here and there would draw suspicion.

The next-most common option for getting data out of the victim's environment is **via command and control** (C2) channel. Some ransomware, and most other forms of modern malware, regularly establish C2 channels to phone home to attackers for additional updates and instructions. Encoding captured data into one of those channels often presents an easy way to send it out of the network.

Transferring data to a cloud account owned by the attacker on the same cloud service can also provide cover for exfiltration. **Automated exfiltration** is less of a standalone technique and more of an option in combination with others like C2 channels. **The same goes for data transfer size limits**, whereby data is sent piecemeal in small sizes using one of the prior techniques to avoid tripping detection thresholds that may be active.

TOP PRACTICES FOR THWARTING DATA EXFILTRATION

By now, you know how this works. We took the data exfiltration techniques listed for common ransomware strains and cross-referenced those with recommended mitigations from the good folks at MITRE. The result, weighted by the percentage of relevant cases, is found in Figure 15.

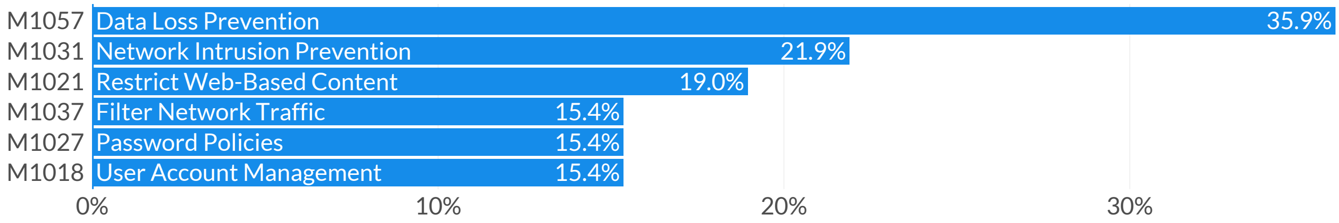

Figure 15 — Top practices for detecting and stopping common data exfiltration techniques

Of all the practices discussed so far, we'd wager these to be among the most familiar. **Data loss prevention** (DLP) and **network intrusion prevention** solutions have been around for a long time. And this is where we need to be careful because not all DLP or IPS are created equal. Old school DLP, for example, might detect social security numbers accidentally emailed in the clear, but wouldn't grant any visibility at all into encoded C2 traffic. When evaluating such solutions, it wouldn't be a bad idea to ask probing questions about how the exfiltration techniques in Figure 14 are handled.

Restricting web-based content and **filtering network traffic** both harken back to attackers using legitimate web and cloud services to smuggle data out of the network. It's nigh impossible to block all the bad things, but being choosy about what's allowed outbound via those services is a good place to start.

And finally, enforcing reasonable **password policies** and tight **user account management** will help ensure that attackers can't take the easy road to exfiltration by assuming the accounts and privileges of legitimate users.

Takeaway for Insurers

Exfiltrating data is fast becoming a favored tactic among ransomware gangs. We observe that ransom demands are 5X higher when that occurs, but there are a lot of proven practices your clients can implement to reduce risk even at this late stage of a ransomware attack.

A collaboration between

Arete®

119
Cyentia
INSTITUTE