Mitigating Ransomware's Impact

Investigative Cybercrime Series: Vol 1

A Collaboration Between

Arete Cyentia Institute

Table of Contents

Executive Summary

We're thrilled to partner with Arete in this first major installment focused on ransomware trends, and the financial exposure faced by past and future victims. This portion of the data contains nearly 1,300 completed ransomware engagements, including demands and payments totaling $797.7M and $218.4M, respectively. We explore specifics on ransom demands and payments, victims' industry and implemented controls, likelihood to pay, reasons for payment, and more! In the future, we will dive deep into attacker actions, intrusion methods, and malware families.

Victims of ransomware often find themselves in an unenviable position, not unlike that of someone who needs to buy a car, but has no knowledge of a vehicle's actual value, common sales and negotiation tactics, or their full range of alternatives. We seek to equip readers with exactly this information in the context of ransomware, and in so doing chip away at attackers' bottom lines while helping organizations better weather—or repel—the next attack.

If you have any stake at all in reducing risk posed by ransomware, you're probably getting excited about hard data from the frontlines. Spoiler alert: there's some good news ahead. What are we waiting for?!

Key Findings

Data associated to 1,288 completed ransomware investigations

Total demands: $797.7M

Total payments: $218.4M

7 of 10 payments less than demand

Typical dıfference between demand and payment is ~$100K

Typical % paid is 54% and trending down

<5% involve re-extortion

<1% fail to deliver after payment

1 in 5 victims chooses not to pay

18% of victims recover without attacker intervention

Even partial implementation of MFA reduced likelihood to pay by 12.5%

Victims who could successfully recover were 19.7% less likely to pay than those who couldn't

We explore specifics on ransom demands and payments, victims' industry and implemented controls, likelihood to pay, reasons for payment, and more!

Methods

Our dataset is composed of investigations and negotiations performed by Arete, a global cyber risk management company, over the course of about 2.5 years, beginning in May 2019. Engagements cover a range of incident response scenarios, such as business email compromise (BEC), forensic investigations, data breach analysis, and ransomware. Identifying information about the victims was sanitized prior to receipt.

How many ransomware events are we seeing?

Ransomware has come a long way since its first recorded manifestation in 1989, when it was distributed via infected 5.25” floppy disks sent through the mail. While its share of the overall malware landscape has grown gradually since, ransomware escaped the conversational orbit of threat researchers and risk managers in 2017, when WannaCry’s unprecedented scale made it a household name.

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 its current prominence (just over 60% of cases in the dataset are related to ransomware), and its continued—albeit modest—gains since January 2020, as we can see in the figure below. The apparent fall-off beginning around October 2021 is explained by unfinished engagements at the time of data collection, rather than criminals gaining some compassion before the holidays.

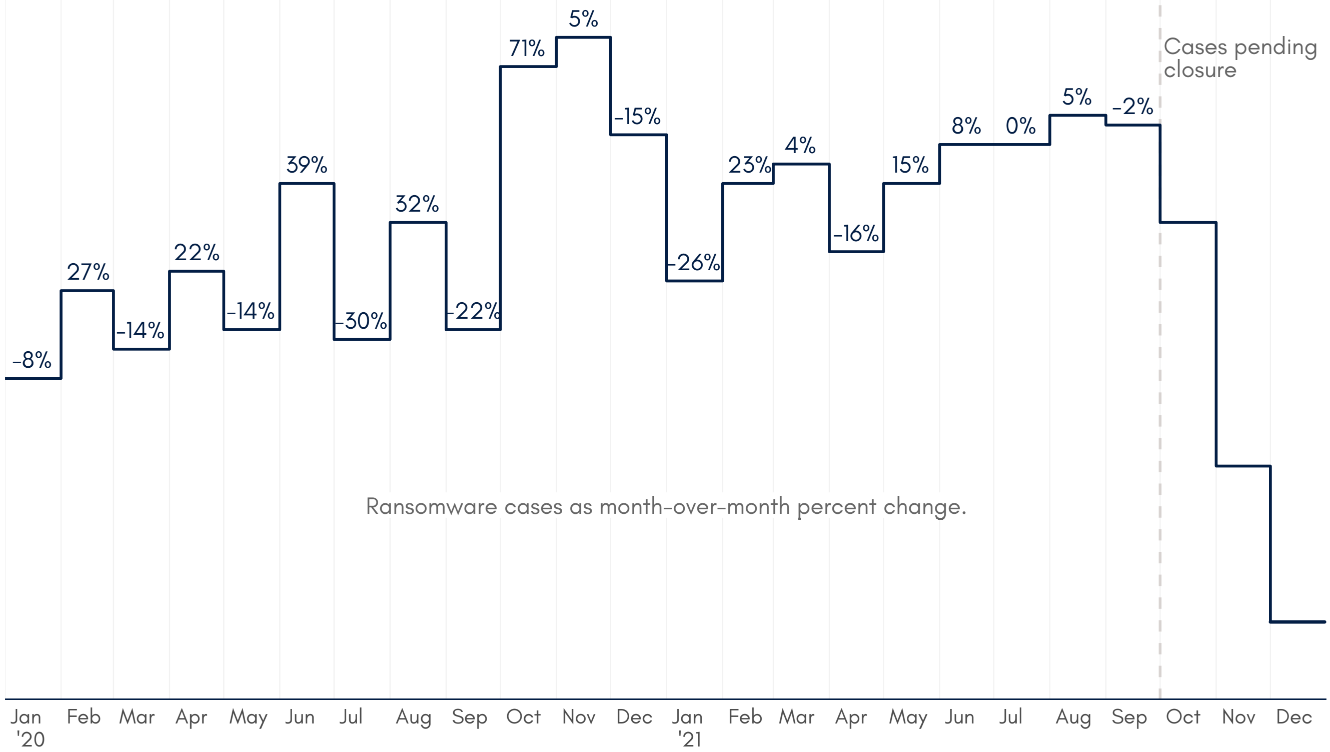

Figure 1 — Monthly ransomware cases

When it campaigns, it pours (sometimes)

Some may look at Figure 1 and wonder if the peaks and troughs exhibited in our ransomware caseload tie back to current events or attack campaigns. As with so many things, the answer is "kinda." We absolutely do see influxes of cases tied to specific threat actors or ransomware strains over time, and sometimes those are associated with events "ripped from the headlines." But we also see a steady drumbeat of diverse ransomware cases that aren't directly attributable to a particular campaign.

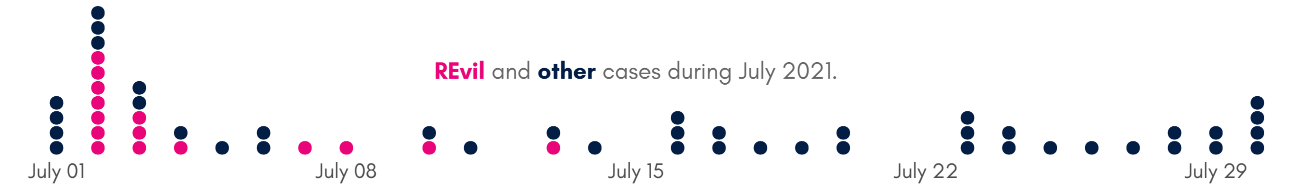

Figure 2 — Highlighting REvil cases, specifically around Kaseya attack

The Kaseya ransomware outbreak that hit hundreds of organizations in July 2021 is one such example where our caseload coincides with current events. The group behind the attacks, REvil, is a notorious ransomware–as–a–service provider that makes regular appearances in our investigations. Per Figure 2, those appearances clearly increased in the weeks surrounding the Kaseya incident. Yet it's also clear that other ransomware crime families didn't declare a ceasefire just because REvil grabbed the spotlight.

Ransomware victims

If attackers are coordinating to target the next "it" victim industry with ransomware, be it on chat boards or in coffee shops, it's not evident in the data. More likely, representation among cases stems from various other factors, including victims' general attack surface and need for external assistance when responding to incidents. Nonetheless, we hope you'll find it helpful to understand which industries are represented in the data, in the figure below.

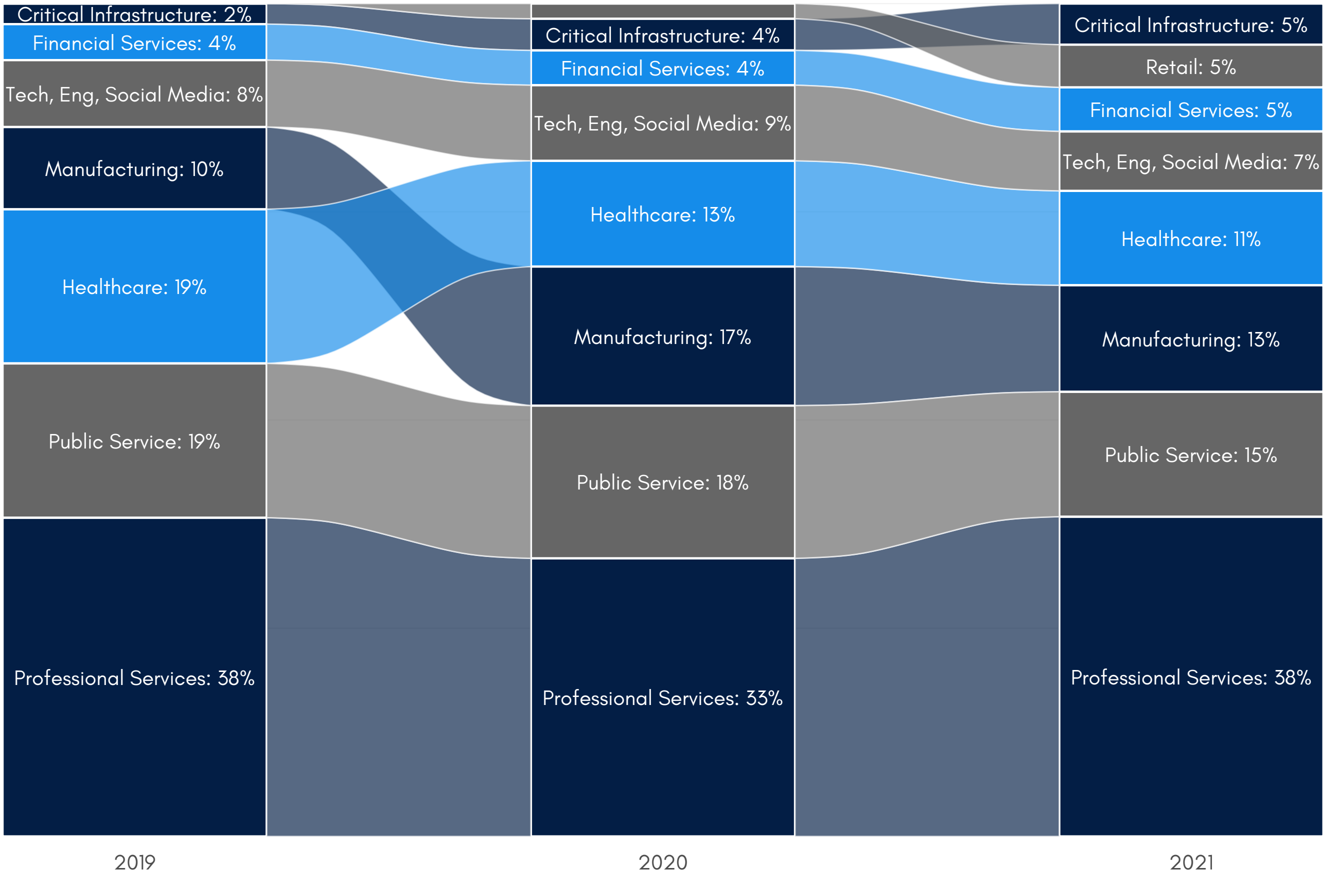

Figure 3 — Ransomware victim industries by year

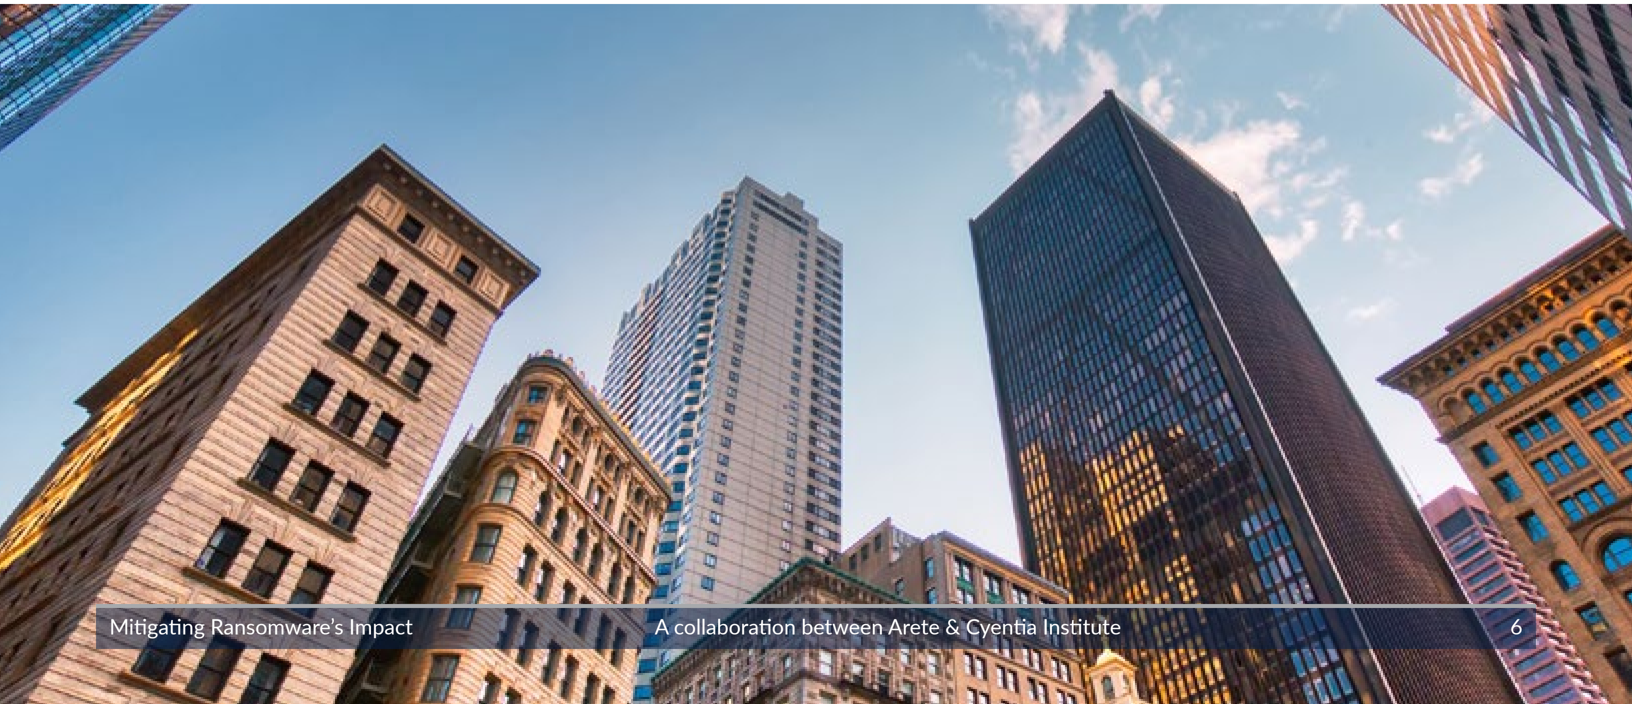

How big are typical demands and payments?

First, the good news is that seven out of 10 times a ransom is paid, it's less than the demand. For the remainder, victims are three times as likely to pay an amount equal to the demand, as one that exceeds it. So, while it's true that extortion attempts often come with eye-popping price tags attached (exceeding $100M in rare cases) in reality, that's usually an opening bid trying to price anchor subsequent negotiations (you knew you were negotiating, right?).

Figure 4 reveals the difference between demands and payments; note the higher peak among payments, suggesting some method to the madness when attackers size up victims' ability to pay. We'll learn more about what affects the percent of a demand paid further on.

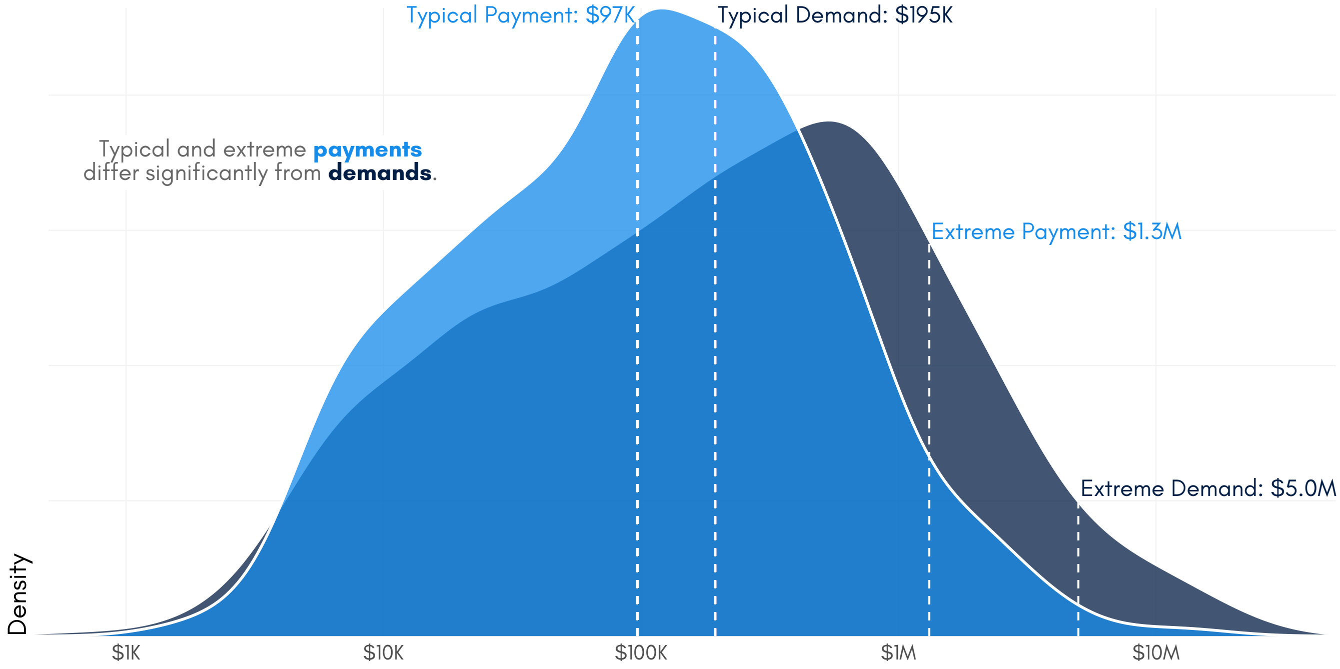

Figure 4 — Ransom demands vs. payments

For now, take a moment to look over the differences between demands and payments in Table 1.

> "Arete was involved in an engagement with a company targeted by a well-known threat actor. The initial demand was $4.7M, and the company's negotiation strategy focused on restoring business operations with the threat actor's decryptor under the best discount available. After several weeks of negotiations, Arete was able to negotiate the ransom down over 93% to an amount of $321K, with all the requested recovery tools and reports successfully delivered to the client."

We'll leave this as an exercise to the reader to consider the implications of selecting a poor central measure—e.g. an average payment of about $339K, versus a more typical payment around $97K when planning mitigations or counteroffers.

	Typical	Median	Average	Extreme
Demands	$194.6K	$226.8K	$1.1M	$5.0M
Payments	$96.9K	$100.0K	$338.7K	$1.3M

Table 1 — Important values for ransom demands and payments

Are demands and payments changing over time?

Yes and no. That is to say, demand amounts are changing over time, mostly increasing across 2020 and 2021, but payment amounts? In a surprising twist: flat, possibly even declining. Needless to say, we were surprised by this promising trend, observable in Figure 5. Each point represents the typical amount of a payment or demand for the previous 90 days.

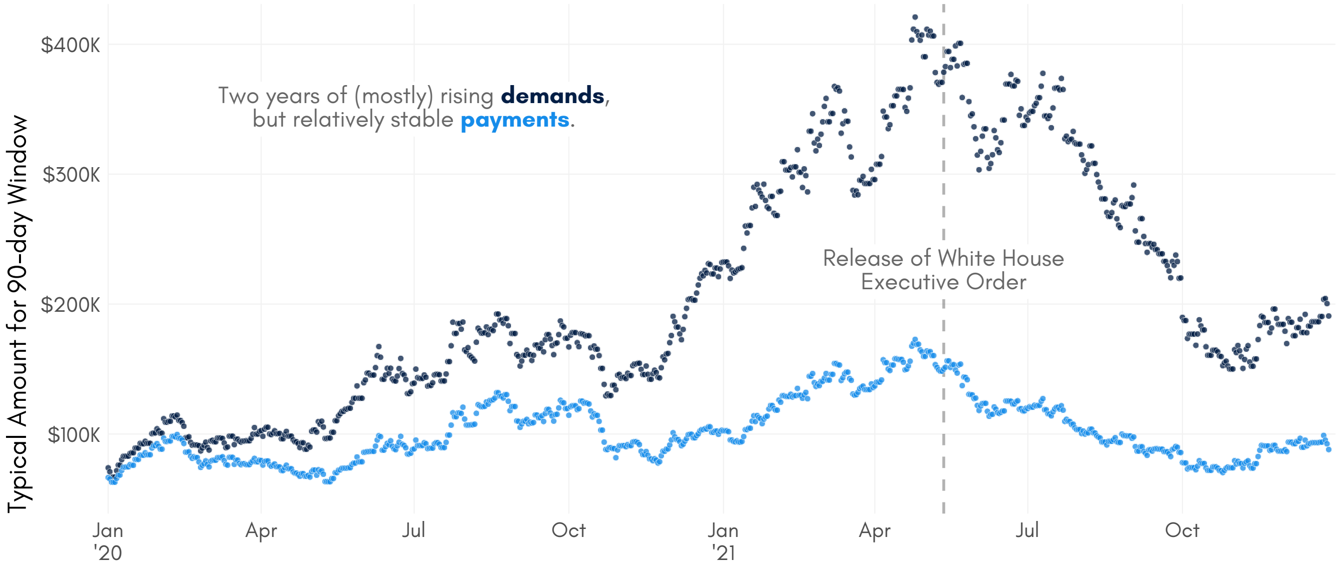

Figure 5 — Demands and payments over time

We can think of a few possible explanations: experienced negotiators getting better deals for victims, attackers jacking up demands to game the appearance of steeper discounts on payments, the changing price of Bitcoin, given its prominence among purveyors of ransomware. Another explanation could be federal actions following the release of the White House's Executive Order on Improving the Nation's Cybersecurity—or WHEOOITNC, if you will—on May 12, 2021. With so many moving parts, it's hard to be sure of any one cause; however, in terms of (good) news you can use, the recent trend in payments is steady and downward.

The good news, recent trends in payments are steady and downward.

How often do victims pay?

About one in five victims opts not to pay, and this seems to be a growing trend. The explanation is welcome news for advocates of defense–in–depth: four out of five times a victim didn't pay, they were able to recover without the attacker's "help", evidenced by Figure 6.

It's no secret that some ransomware attacks are carried out by actors involved in other nefarious activities, sometimes landing them on lists of sanctioned entities, blocking payment (the roughly 7% identified as "Sanctions" in Figure 6). We'd like to give particular attention to the small, but important, group of victims that ran the numbers and determined that a more–involved recovery process was preferable to paying a ransom. In some cases, victims simply refused to pay on principle. Needless to say, what constitutes an organization's bottom line can be as varied as the personalities involved, and it's probably best to know your leadership's perspective beforehand.

> This boils down to 18% of ransomware victims recovering without the attacker's support.

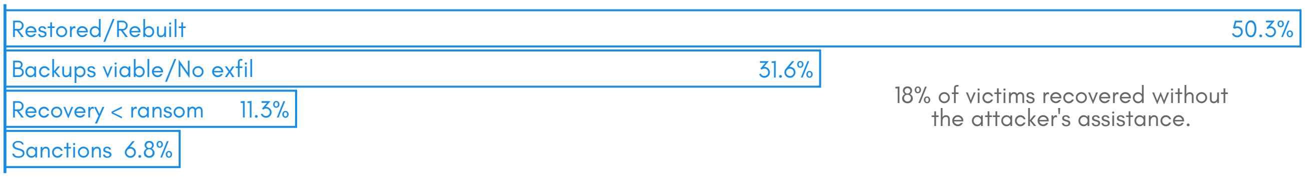

18% of victims recovered without the attacker's assistance.

Figure 6 — Reasons for non–payment

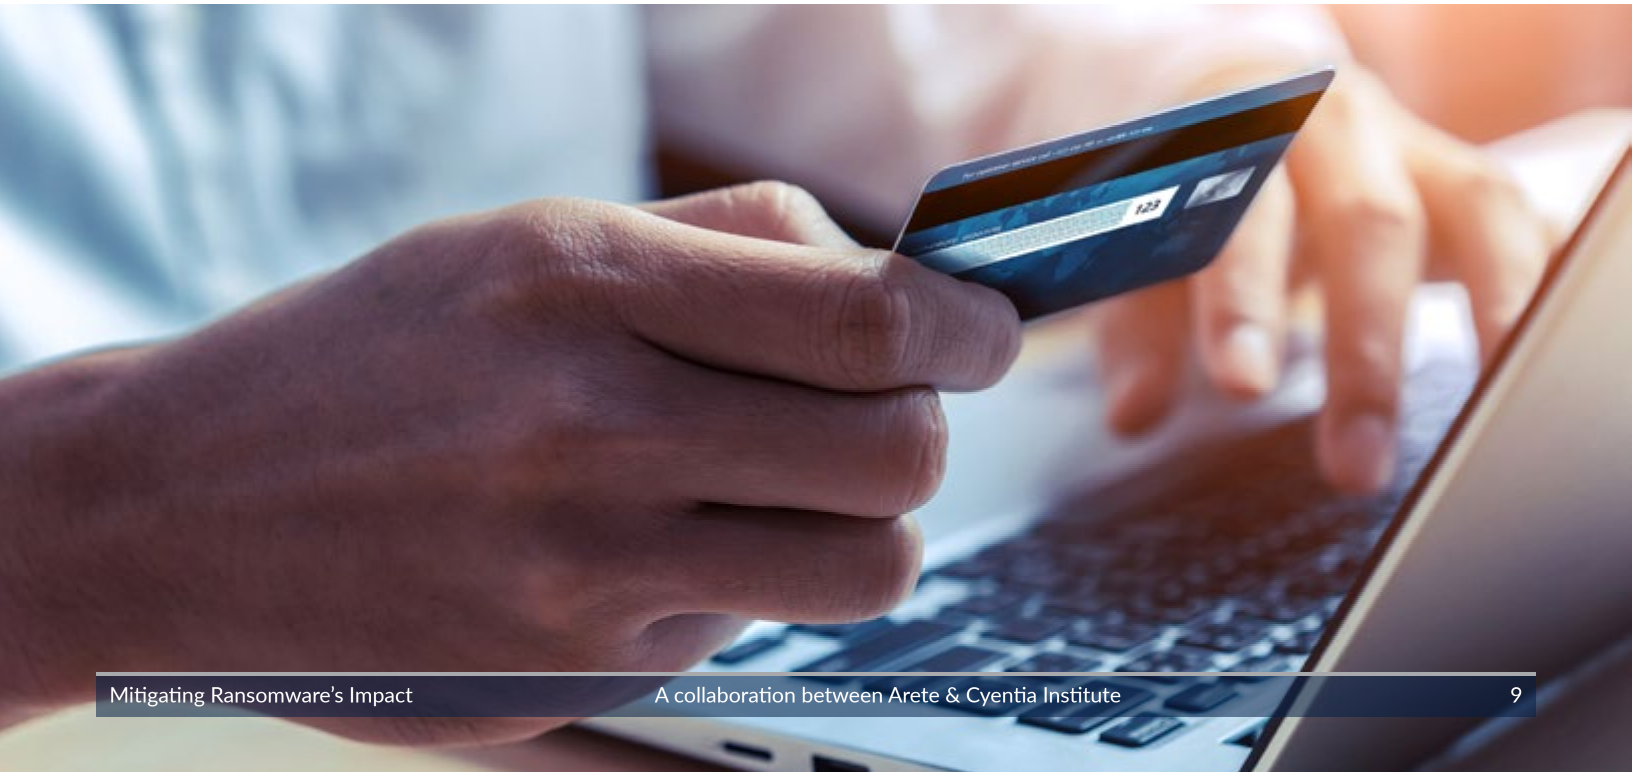

What affects likelihood to pay?

Controls

Oh yeah, we do backups.

~Literally every organization

In addition to victims' recovery outcomes, we were thrilled to find that many of the cases included information about implementation of backups and multi-factor authentication (MFA). A natural assumption about non-payment would be the role of backups. After all, if you don't need to go through the attacker to regain access to data or systems, why would you? Well, only having backups wasn't enough to reduce a victim's likelihood of payment (Figure 7), but the capability to successfully recover reduces the likelihood of payment by nearly 20% (Figure 8).

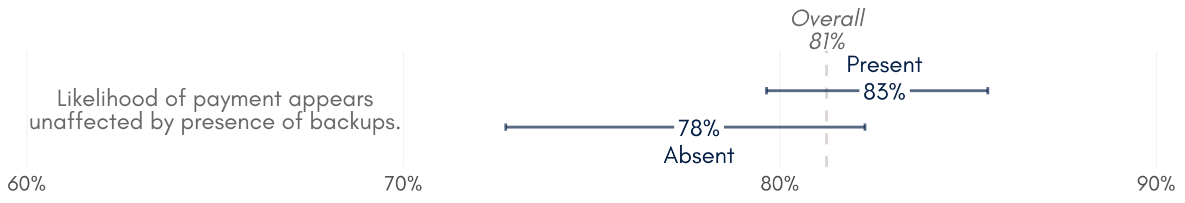

Figure 7 — Presence of backups vs. likelihood of payment

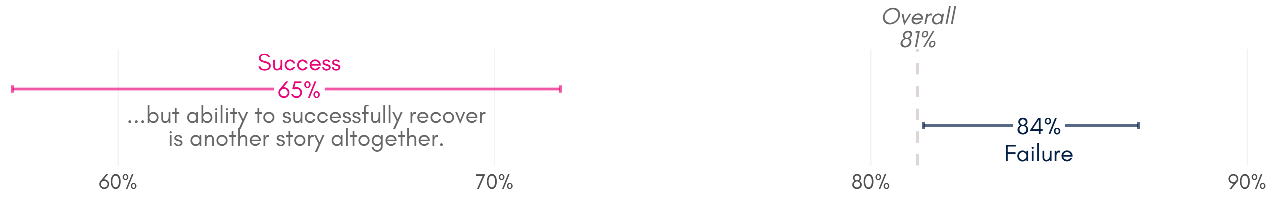

Figure 8 — Ability to recover vs. likelihood of payment

An interesting and perhaps less intuitive relationship exists among victims who implemented MFA on at least their VPN, Email, or Administrator accounts. These organizations were also less likely to end up paying (Figure 9).

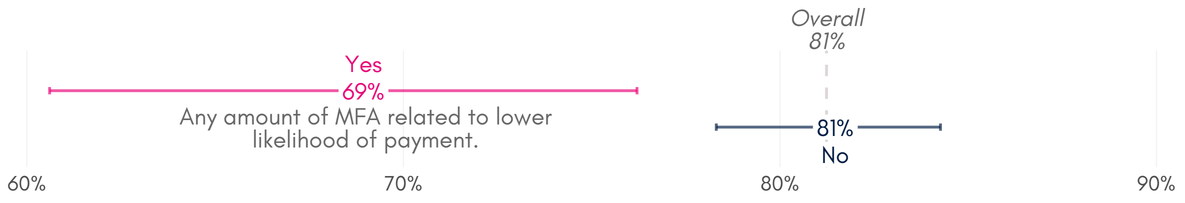

Figure 9 — MFA implementation vs. likelihood of payment

It's hard to pinpoint a causal relationship here. Are attackers being somewhat foiled by MFA in the victim's environment, despite having gained access? Is there something else about organizations implementing MFA which also makes them less likely to pay? Suffice it to say, given MFA's other benefits, this seems like icing on the cake.

Initial demand

We were also curious about whether higher demands tended to scare off victims from ultimately paying up, and this appears to be the case, per the figure below. Apparently, sticker shock as a price anchoring technique occasionally backfires.

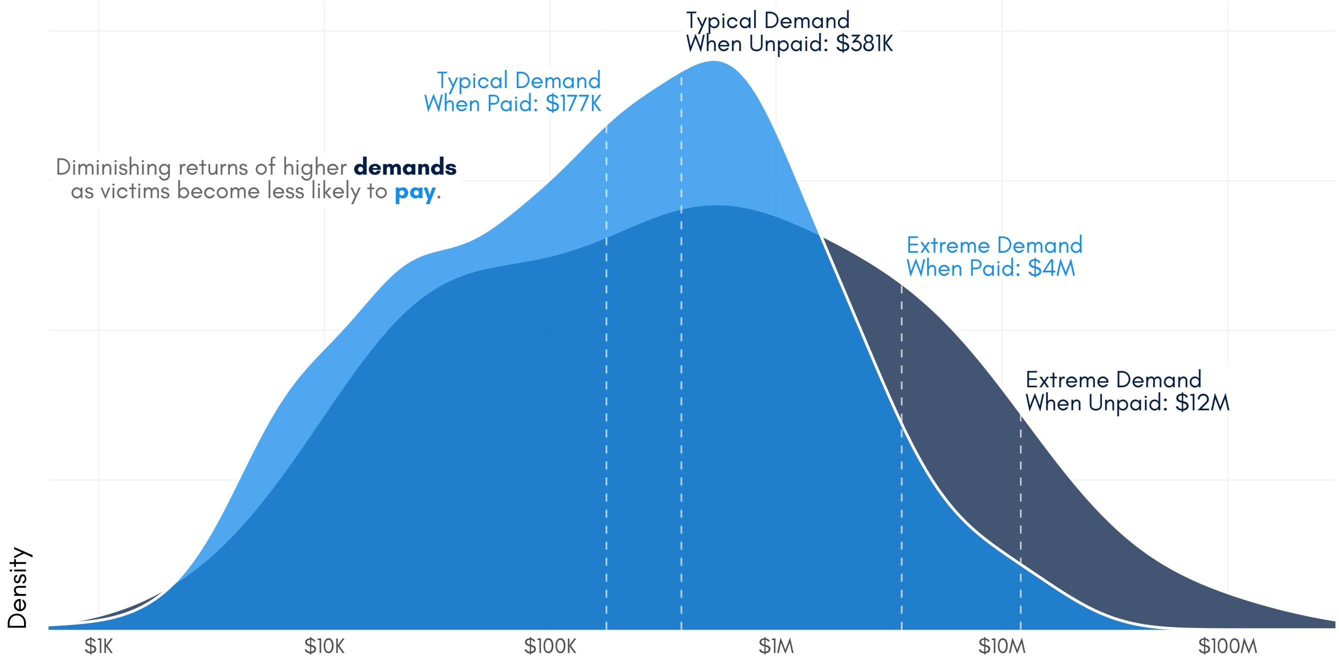

Figure 10 — Initial demand vs. likelihood of payment

What percent of a demand is typically paid?

We alluded to this before, but the majority of payments do not match the original demand. Percent of demand paid tends to decrease as the demand increases, but let us draw your attention to the highlighted points where the total ransom paid exceeded the original demand.

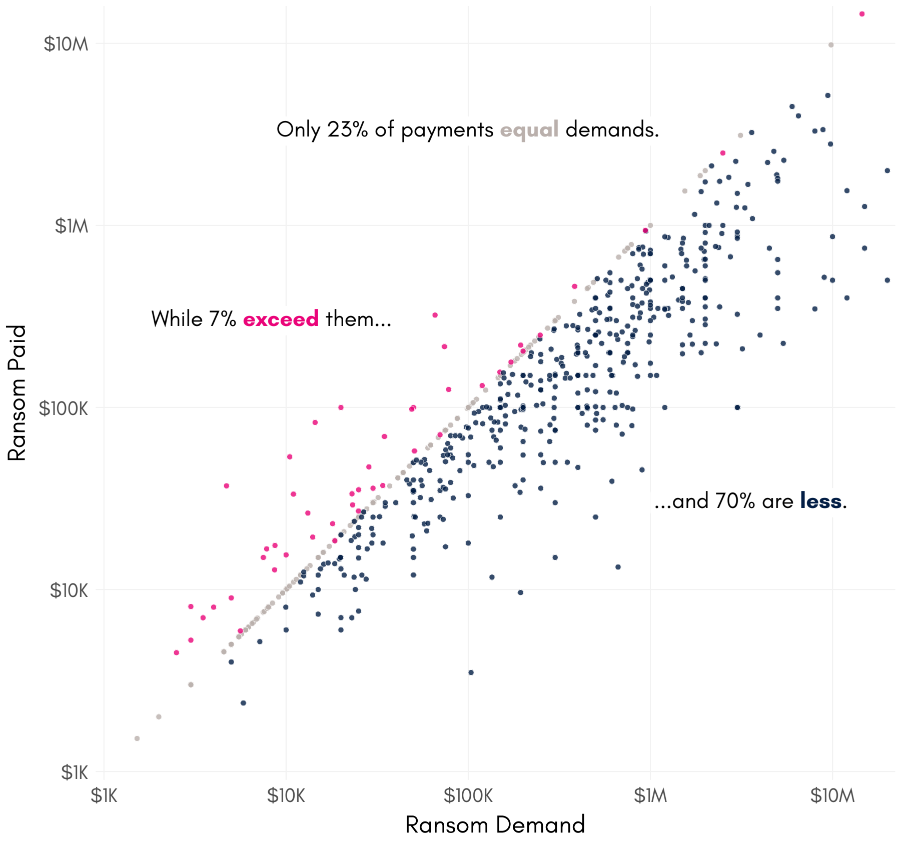

Figure 11 — Payments that do not match the demand

Most of these represent instances of re-extortion. These are cases where victims believe they are paying for a full resolution, only to be nickel-and-dimed for piecewise fixes—e.g. "Oh, sorry, that was the price per decryption key..." Others represent examples of attackers using time-pressure tactics, escalating demands while victims decide whether or not to pay.

"Hold on, did you say 'RE-extortion'? Tell me more"

Let's compare re-extortion to other cases in Table 2. We see a different business model, involving lowball demands followed by stringing victims along, likely motivated by sunk costs.

> "Arete had a case involving a company whose business operations were significantly affected by the attack. They initially did not want to negotiate, reducing our ability to properly drive the price of the ransom down. After some time, the company ended up needing decryption from the threat actor, and ended up paying 15% over the threat actor's initial demand to accommodate the threat actor's increased ransom once their negotiation timer expired. The business received all of the threat actor's deliverables shortly after payment was made."

> We were involved in a case where the client was re-extorted by a threat actor. We informed the client that there was a medium-high probability that the threat actor would attempt to re-extort them. During the engagement, the threat actor assured us that the cost of restoration for the entire system would be $2,000. After payment, the TA went back on their assurances and stated that it would cost $2,000 per decryption key. The client still wished to pay and proceeded to pay $2,000 per decryption key until their network was fully restored.

	Typical Demand	Typical Payment	Typical % Paid	Typical Response Cost
Standard (1,231)	$222.3K	$104.1K	50.2%	$42.3K
Re-extortion (57)	$41.4K	$45.7K	112.7%	$43.2K

Table 2 — Important values for re-extortion

There are a few upsides: re-extortion is relatively uncommon—less than five percent of cases—and they appear to have spiked in 2020 (30 of the 57 recorded instances). However, the "savings" of a lower total payment is partially offset by the headache and uncertainty from multiple rounds of negotiation resulting in response costs similar to standard cases.

What affects percent of a demand paid?

Initial demand

We noted in Figure 5 that percent of demand paid dropped from 2021 to 2020, mostly due to rising demands and stable payments. Next we look at how the initial demand itself relates to the percent which ultimately gets paid. Per Figure 12, we see that the highest demands come with the widest range of payments, and also the lowest payments as a percent of the demand.

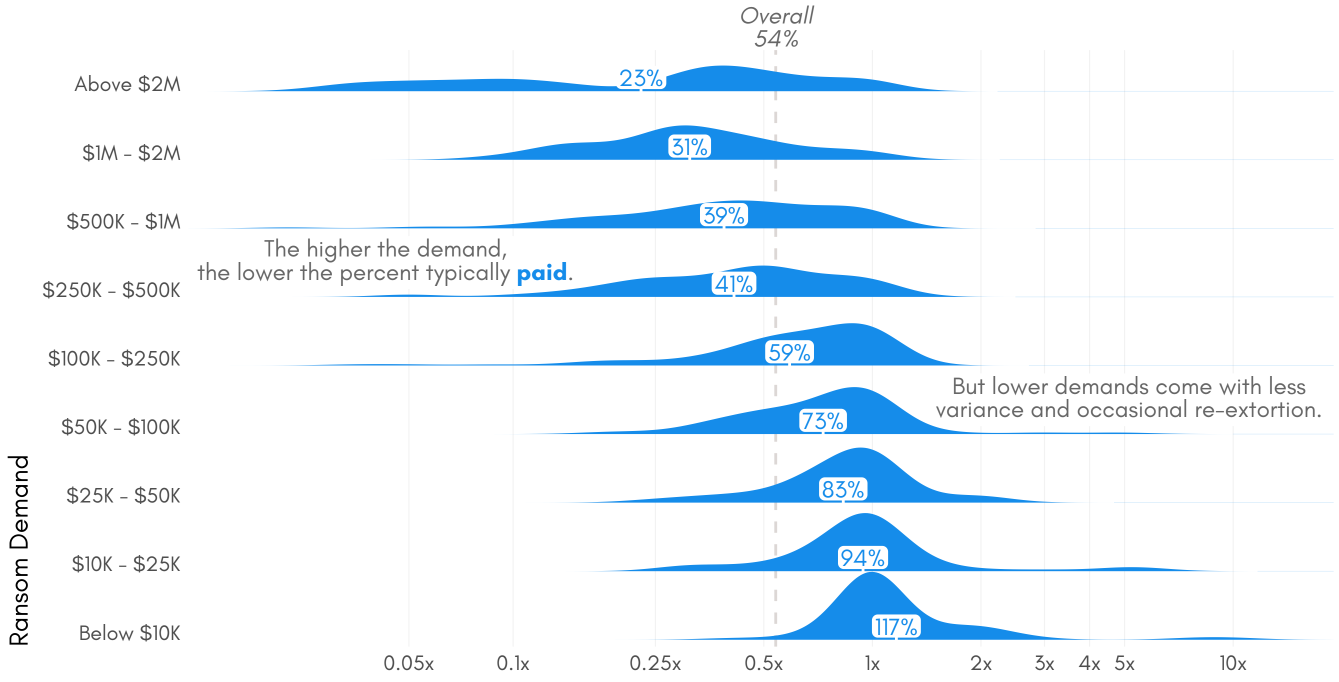

Figure 12 — Percent of demand paid by size of initial demand

Now, before you start thinking this is a lead-in to an Algebra II word problem (flashbacks to trains leaving stations at different times and different velocities…), let's get one thing straight: it certainly stings less to pay all or more of a $20K ransom than half of a $200K ransom. But, the consistent downward trend in typical percent paid as demands increase suggests some logic behind the scenes.

At the lower extreme, we see that the smallest demands tend to result in more concentrated—and often excessive—payments. This aligns with what we observed in Figure 11 where excessive payments tended to accompany smaller demands. So, if you're hit with a suspiciously low initial demand, be aware that there will probably be less wiggle room for negotiation, and an increased chance of re-extortion!

The smallest demands tend to result in more concentrated and excessive payments; in addition, low initial demands are signs there will be less negotiating room, and an increased chance of re-extortion!

Victim industry

> Though demands can vary depending on industry, payments are more driven by the size of initial demand.

There's plenty of conjecture about which sectors feel the most pain when it comes to ransom demands, and the data somewhat supports the idea that a victim's industry can play a role.

However, that's where it stops cooperating with the popular narrative as we see in Figure 13. We doubt many readers would have placed Healthcare and Critical Infrastructure's typical demands below the overall of ~$195K, let alone trailing all the other sectors in the dataset.

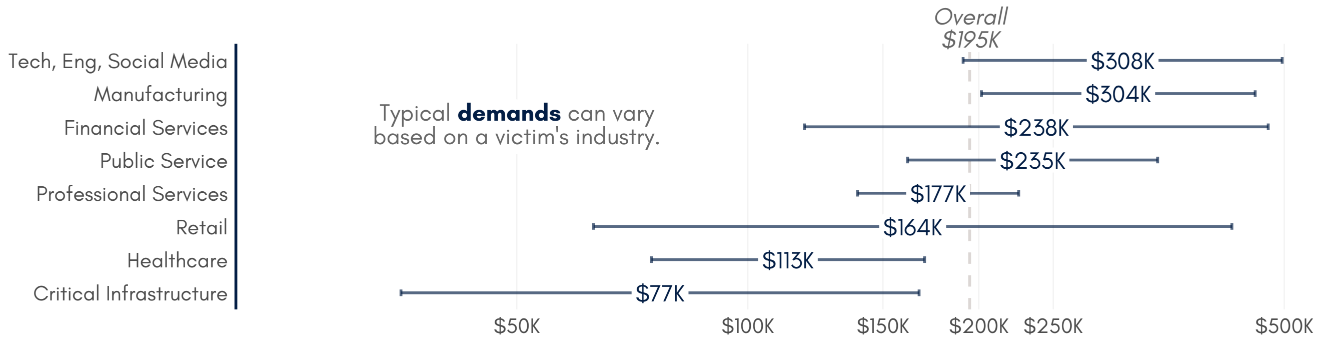

Figure 13 — Demands by industry vs. overall

Interestingly, this industry discount appears to be limited to demands. When it comes to payments, this effect is overpowered by that of the initial demand itself. You can see this in Figure 14 where the industries maintain their order as the range of payment amounts generally narrows.

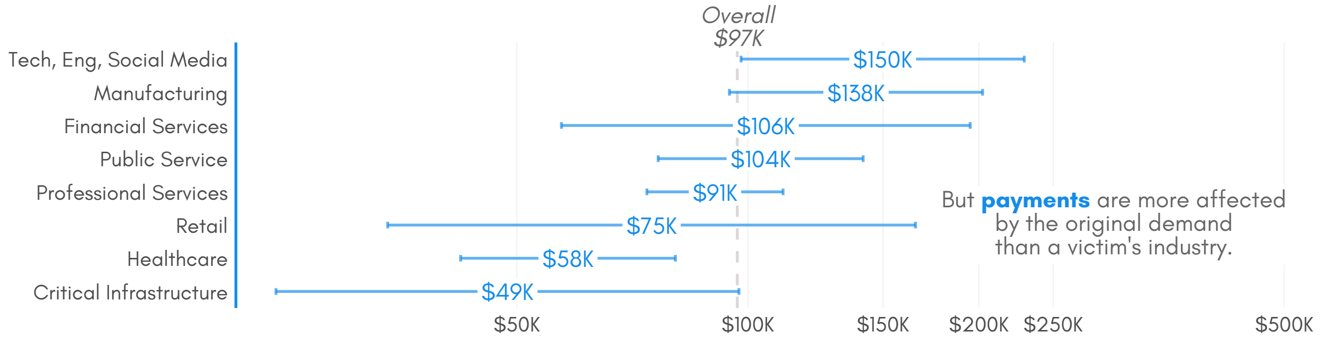

Figure 14 — Payments by industry vs. overall

Payment reasons

So if industry only indirectly affects percent paid, what about the amount of "service" required from the criminal? Victims pay in roughly four out of five cases, and they tend to do so for one of three main reasons: for a decryption key or return of stolen data, for both the key and data, or for proof of deletion.

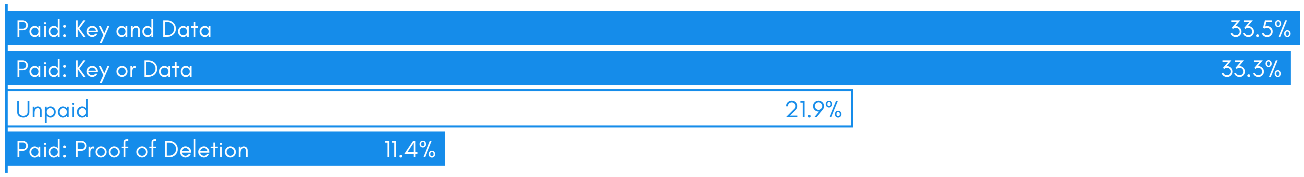

Figure 15 — Count of payment reasons

Victims tend to pay for one of three main reasons: for a decryption key or return of stolen data, for both the key and data, or for proof of deletion.

We see a nearly even split between victims that need decryption and data retrieval vs. those that need only one of the two. Interestingly, about one in 10 payments wasn't to restore access, but rather to obtain proof of deletion, sometimes done to reassure stakeholders in the wake of a breach notification. This is something to keep in mind when accounting for the overall costs of prevention, response, and recovery.

Figure 16 focuses on payment amounts organized by payment reason. Payments trend upward as victims require more from attackers. It seems there's no escaping the upcharge business model, even when dealing with ransomware.

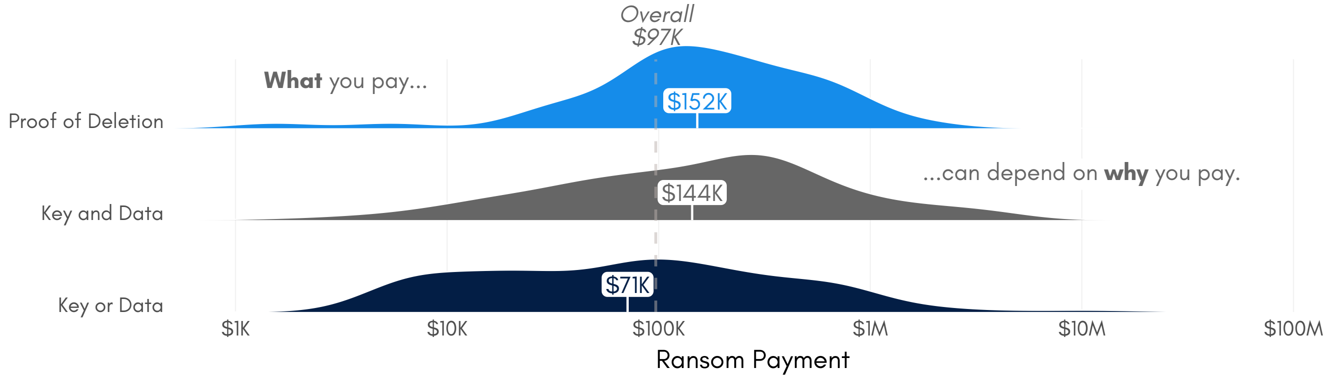

Figure 16 — Payment by payment reason

Are payment reasons changing over time?

Like us, you've probably heard rumblings of evolving tactics among ransomware groups—i.e. transitioning from extorting money for access to systems to then blackmailing organizations that seek to avoid public release of pilfered data. Our information on payment reasons both supports and contradicts this generalization, as seen in Figure 17.

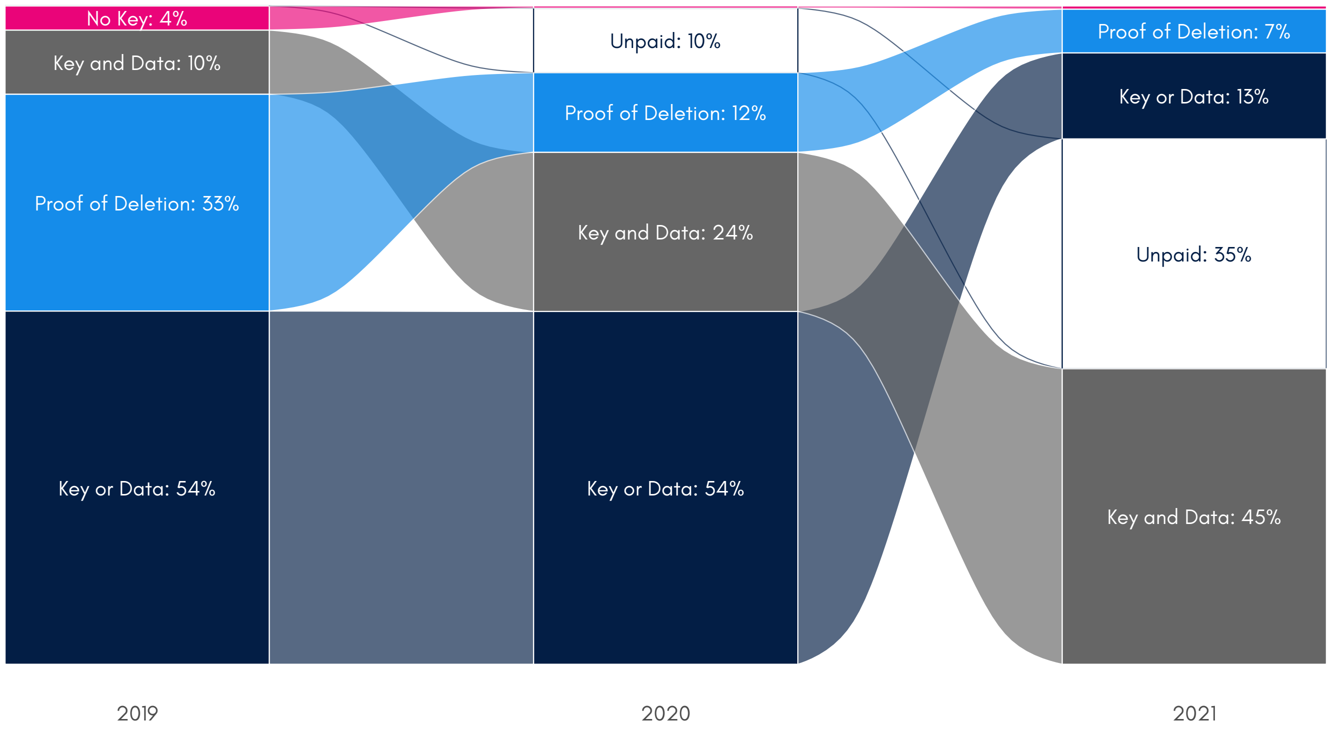

Figure 17 — Payment reasons over time

The shift from merely encrypting data to outright theft is evident from the growth in payments for both key and data. However, the parallel decline in payments for proof of deletion and rise in those unpaid is cause for some optimism. It's explained in part by the 18% of victims able to recover without attacker intervention, but also by organizations—and those negotiating on their behalf—adapting their response strategies. I.e.—if operations can more or less carry on as normal, and notification letters have to go out either way, why bother tipping the criminal?

> “Arete was notified that a client was completely knocked down by PYSA. PYSA deleted the client's backups, and encrypted all the servers and several workstations. Miraculously, the technical team found a two-day old snapshot for all servers stuck in a replicator application. They were able to use these snapshots to recover all their servers. The data was extremely sensitive so Arete worked with the Client and Counsel to negotiate the threat actor's demand to a reasonable amount and pay for proof of deletion.”

How common is failure to deliver?

It's worth taking a moment to acknowledge the boardroom naysayer that may haunt payment deliberations, murmuring “but what if they don't deliver…?” It's true that victims have paid, but never received a decryption key (if you squint at Figure 17, you can see them, mostly in 2019).

> “Arete handled a case with a company that was attacked by a threat actor who insisted that they would deliver a decryptor upon payment.
>
> After several weeks of negotiations and after payment was made, the threat actor advised that they could not provide the decryptor due to investigations by law enforcement agencies.
>
> The threat actor stopped replying to any further inquiries.”

Here’s the thing: in 2.5 years of cases where payment reason is captured, criminals failed to deliver <1% of the time (6 out of 632, to be exact.) We’re not taking sides, but that’s a hit rate that can’t be ignored. It’s not all that surprising when you consider how important credibility is, even in the downside–up world of ransomware: if you have a reputation for ghosting victims after receipt of funds, you’ll struggle the next time around. Criminals on dark web marketplaces have been known to show off their credibility with badges and even reviews from past victims.

Controls

Earlier, we saw the effects of backups and MFA on victims’ likelihood to pay, but how might they influence the percent of a demand paid? Brace yourself: there was no measurable difference in percent of demand paid between organizations that “had backups” and those that didn’t (Figure 18).

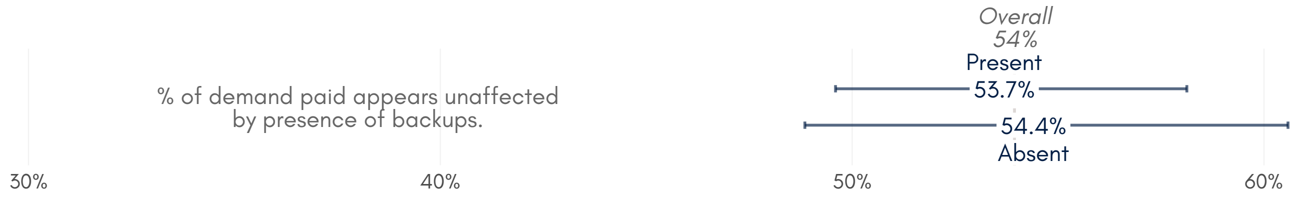

Figure 18 — Presence of backups vs. percent of demand paid

This, and the previous result about likelihood to pay, probably come as no surprise to anyone who has worked with disaster recovery plans charitably described as “binderware.” The data is unambiguous on this point: when it comes to reducing the percent of a demand paid, provable ability to recover is what counts (Figure 19).

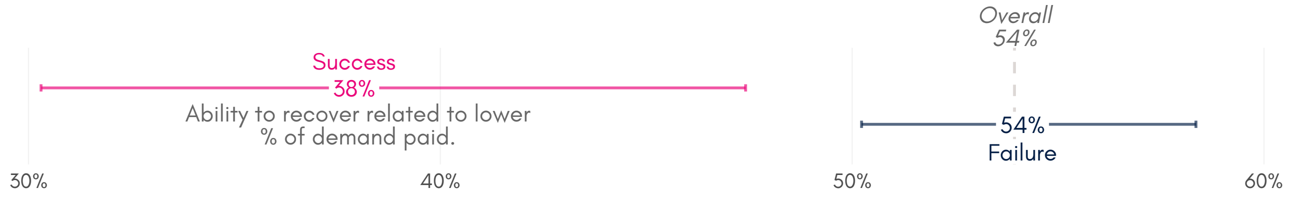

Figure 19 — Ability to recover vs. percent of demand paid

It's worth noting that a ransomware event is relatively tame on the spectrum of disaster recovery scenarios. Yes, it is more adversarial than a natural disaster, but personnel typically retain access to their primary facility and aren't juggling simultaneous hazards at home and the workplace. If you have a vested interest in reducing the impact of ransomware—whether in part by reducing percent paid, or in whole by eliminating the need to pay—get proof that your systems can successfully recover.

> "Recently, Arete dealt with a threat actor we had not seen in years. In the past, this TA was known to access victim's systems through open RDP to the internet. Before Arete was engaged, the technical person, who was backing up the client's systems to an external hard drive, attempted to restore from the hard drive, as they had done during previous ransomware incidents. Unfortunately, the TA was waiting for the hard drive to be connected, and immediately deleted the backup images."

Now, about MFA. You might be wondering why we've included it, given that its main value proposition is in preventing these incidents in the first place. Yet, in terms of percent of demand paid, those victims who implemented MFA on at least their VPN, Email, or Administrator accounts paid less (Figure 20). We can already hear the distant sound of a thousand business cases being updated.

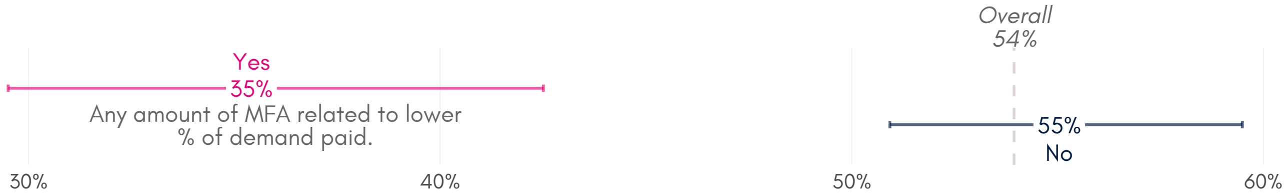

Figure 20 — MFA implementation vs. percent of demand paid

Conclusion and Application

How might defenders apply lessons from the data to protect their organizations?

This dataset can't speak to preventing incidents, so we'll settle for avoiding or reducing payment in the wake of a ransomware attack:

Attackers will use an array of slimy tactics to bully you into payment: price anchoring, time pressure, lowballing, false dilemma, sunk cost, and bandwagoning, to name a few. Your best bet is to already know how much your data and daily operations are worth to your organization, who exactly would authorize this kind of payment, and how much it would cost to fully restore operations via a third party.

Payments are almost always below sticker price. The typical percent of demand paid may be 54%, but it slides lower the more a demand exceeds $250K, and higher the more a demand undercuts $250K.

Know that some executives will refuse to pay on principle, and that the U.S. Government sanctions some threat actors, prohibiting payment.

Those with a demonstrated capacity to recover were less likely to pay and paid a lower percent of demands compared to those who simply claimed to perform backups.

Those with MFA implemented on at least their VPN, Email, or Administrator accounts were less likely to pay and paid a lower percent of demands compared to those without MFA.

What's next?

Are you curious about the rise and fall of major ransomware families?

Or perhaps the implications of different intrusion methods and attacker actions on likelihood and amount of payment? You're in luck!

We've only scratched the surface here, and the next installment will be threat–focused, covering those and more. And while we think there is plenty to digest in this report, here is a taste of what's to come. Figure 21 shows the prevalence of identified intrusion methods within each industry.

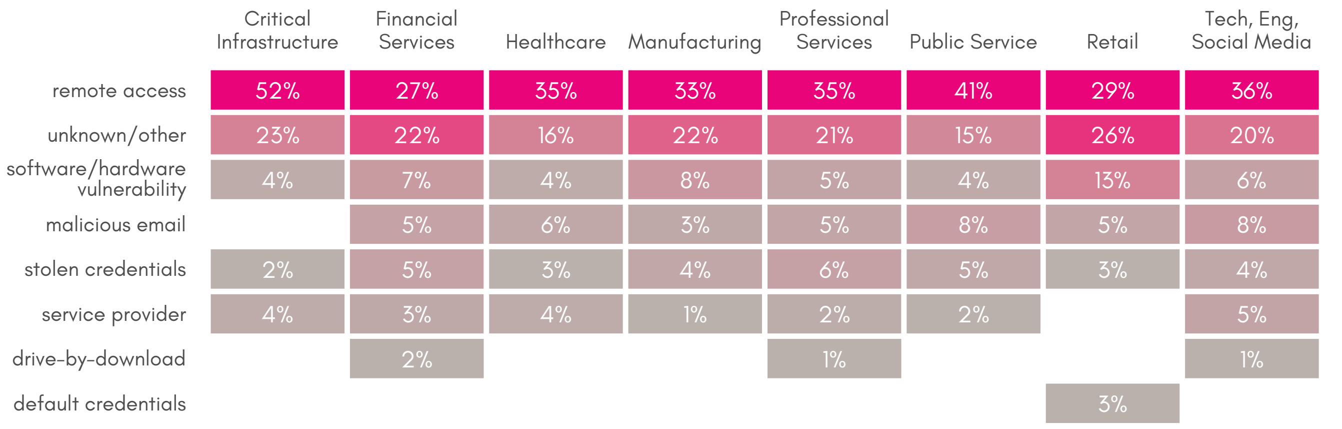

	Critical Infrastructure	Financial Services	Healthcare	Manufacturing	Professional Services	Public Service	Retail	Tech, Eng, Social Media
remote access	52%	27%	35%	33%	35%	41%	29%	36%
unknown/other	23%	22%	16%	22%	21%	15%	26%	20%
software/hardware vulnerability	4%	7%	4%	8%	5%	4%	13%	6%
malicious email		5%	6%	3%	5%	8%	5%	8%
stolen credentials	2%	5%	3%	4%	6%	5%	3%	4%
service provider	4%	3%	4%	1%	2%	2%		5%
drive-by-download		2%			1%			1%
default credentials							3%	

Figure 21 — Prevalence of intrusion methods by industry

A collaboration between

Arete SM

119
Cyentia
INSTITUTE