SOPHOS

Everything Everywhere All At Once: The 2023 Active Adversary Report for Business Leaders

A deep dive into over 150 incident-response cases reveals both attackers and defenders picking up the pace

Written by John Shier

A Sophos Whitepaper. August 2023

The 2023 Active Adversary Report for Business Leaders presents what the Sophos X-Ops Incident Response (IR) team has learned about the current adversary landscape from tackling security crises around the world. Our report is based on data from over 150 cases selected from the 2022 workload of the IR team. We provide more detail on the demographics represented in this analysis at the end of the report.

In addition to the 2022 data, this year's report incorporates data from the previous two years of our Active Adversary reports (2021 data, 2020 data). As our IR program continues to expand, we have found various interesting year-to-year trends in that data – some expected, some less so.

As our dataset continues to expand, so does our ability to derive information useful for various constituencies – heads of business, infosecurity leaders, and hands-on security practitioners. Over the course of 2023, we'll release reports based on our data (and incorporating 2023 findings as appropriate) for each of those constituencies. This, our first report for the year, focuses on information useful for heads of business seeking an all-up understanding of what their enterprises are facing.

Key takeaways

- Ransomware is still a pervasive threat
- Detection-and-response tech and services are making measurable inroads against attackers
- Patch, patch, patch
- Enterprises put themselves at risk of repeated attack when they don't have, or don't retain, log data
- Attacker dwell time is shrinking – for better or worse
- Once a ransomware attacker is inside your network, the odds are dangerously high that your data will be exfiltrated

Where the data comes from

For this report, 81% of the dataset was derived from organizations with fewer than 1000 employees. As in previous years, 50% of organizations requiring our assistance have 250 employees or fewer.

However, larger enterprises are still very much part of our equation: While most enterprise organizations already have the required resources to undertake some incident response in-house, one-fifth of the organizations with which IR worked in 2023 were companies with over 1000 employees. (Sometimes even the best-staffed in-house IR teams can use an assist from experts with specialized skillsets and knowledge.)

And what do these organizations do? For a third consecutive year, the manufacturing sector (20%) was the most likely to request Sophos IR services, followed by healthcare (12%), education (9%), and retail (8%). In total, 22 different sectors are represented in this dataset. Overall, the numbers shouldn't be taken as a comprehensive statement on this aspect of the overall threat landscape – no industry is safe from attack, unfortunately – but that said, we have witnessed numerous sustained attacks against both healthcare and education institutions over the last few years, and it is not surprising to find these sectors near the top of our list yet again.

Attack Types: Ransomware runs the game

Organizations across all demographics experienced a sustained onslaught of ransomware attacks this past year. While there have been claims in the news about ransomware attacks having plateaued or even declined in 2022, over two-thirds (68%) of incidents recorded in this year's Active Adversary data were ransomware attacks, followed by non-ransomware network breaches (18%) as the second most common finding.

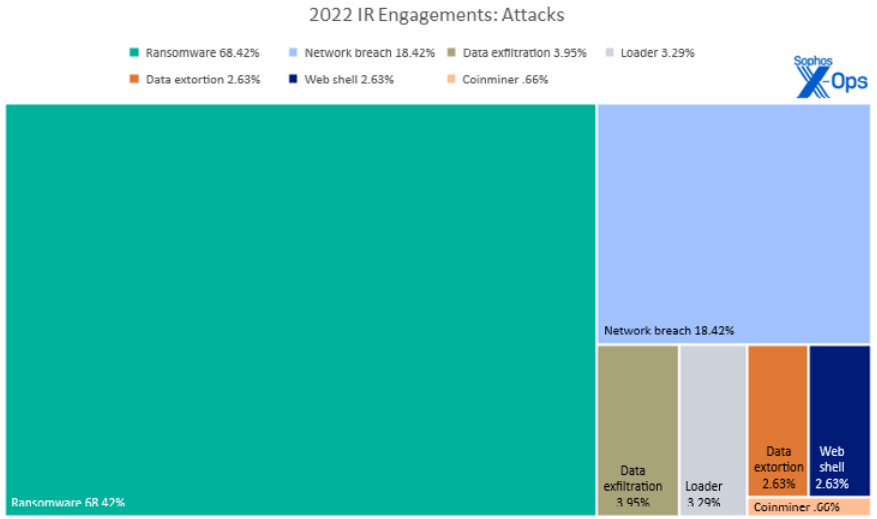

Figure 1: Two-thirds of the attacks covered in this year's Active Adversary dataset were ransomware attacks

Though ransomware may have stopped its exponential growth as attackers diversify their goals, it was still vastly more common than all other forms of attacks in 2022. Taking the longer view, we see ransomware unsurprisingly in the top spot for all three years of our Active Adversary reports, with nearly three-quarters (73%) of IR investigations involving ransomware attacks in that timespan.

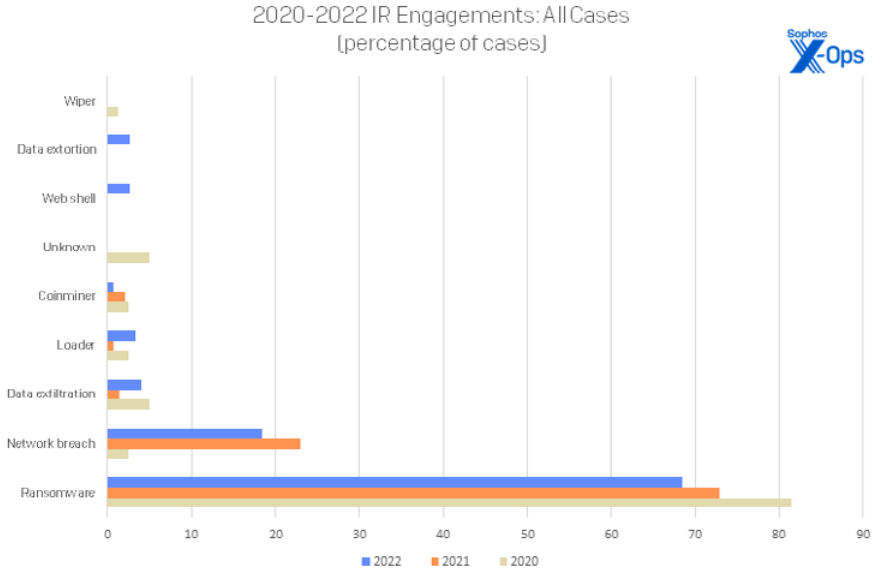

Figure 2: Ransomware has consistently dominated IR cases for all three years of the Active Adversary Report

We'll dig into the details of detected ransomware families later in this report. For now, it's most important to note that the names change, but the problem does not. Ransomware attacks will always be highly represented in IR datasets since those attacks are the most visible and most destructive, and often require the most expert help.

When we set aside ransomware in the results, network breaches dominate the rest of the field with 58% of the cases. In other words, more than half of all non-ransomware attacks consisted of an intrusion, but no clear motive was identified; data exfiltration could be neither confirmed nor excluded as the motive in these cases. This begs the question: ***How many of these were simply thwarted ransomware attacks?*** In fact, we were able to identify several attacks that were perpetrated by Cuba and Vice Society, both infamous ransomware purveyors, but crucially those attacks never reached the ransomware stage. The lesson here for business leadership is that prompt action can break even a tried-and-true attack chain such as that used by ransomware; in the case of a number of these incidents, that's likely what happened.

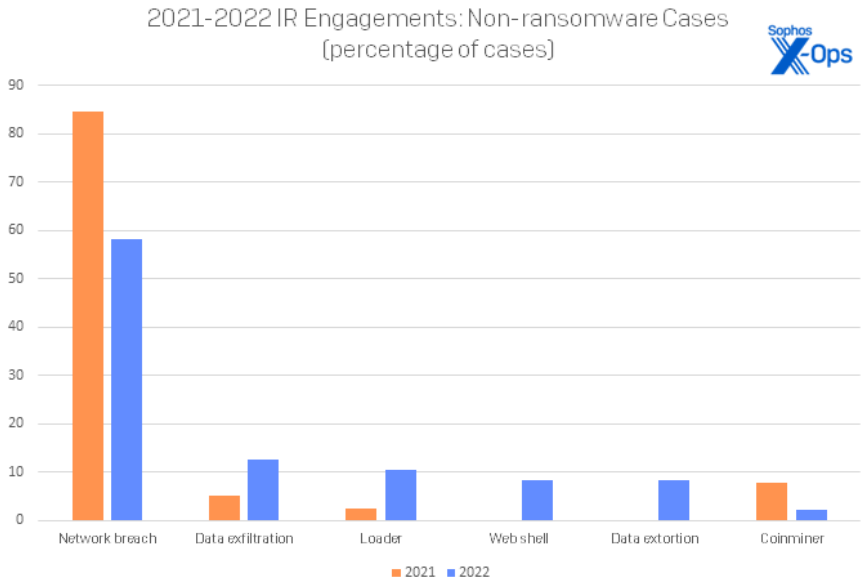

Figure 3: Beyond ransomware, network breaches have been the leading cause of calls to IR over the past two years. (We have excluded 2020 from this chart; in that year, just 18.52% of all cases handled involved something other than ransomware.)

We note a growing number of data exfiltration (13%) and data extortion (8%) attacks in the dataset. These types of attacks are defined by theft of data (exfiltration) in which payment may also be demanded (extortion) – also hallmarks of certain varieties of ransomware, but the other hallmarks of a ransomware attack (e.g., encryption of data in situ) were not present in these cases.

Beyond breaches and exfiltration, the variety of different attack types in this year's data rose slightly. It may be that this diversity is due to attackers not achieving their end objectives. More companies are adopting technologies like EDR (Endpoint Detection and Response), NDR (Network Detection and Response) and XDR (Xtended Detection & Response) or services like MDR (Managed Detection and Response), all of which allow them to spot trouble sooner. This in turn means they can stop an attack in progress and evict the intruders before the primary goal is achieved – or before another, more malignant intruder finds a protection gap first located by a lesser adversary. While a coinminer or a web shell on your network is still not acceptable, it is much better to detect and remediate threats such as these before they turn into full-blown ransomware attacks, or exfiltration, or extortion, or a reportable breach.

Root Causes: They're either breaking in or logging in

In our investigations, not only do we identify the initial access method (that is, how attackers got into the network) but we attempt to attribute their success to a root cause. For the second year running, exploited vulnerabilities (37%) contributed the most to the root causes of attacks. This is lower than last year's total (47%) but consistent with our three-year tally (35%).

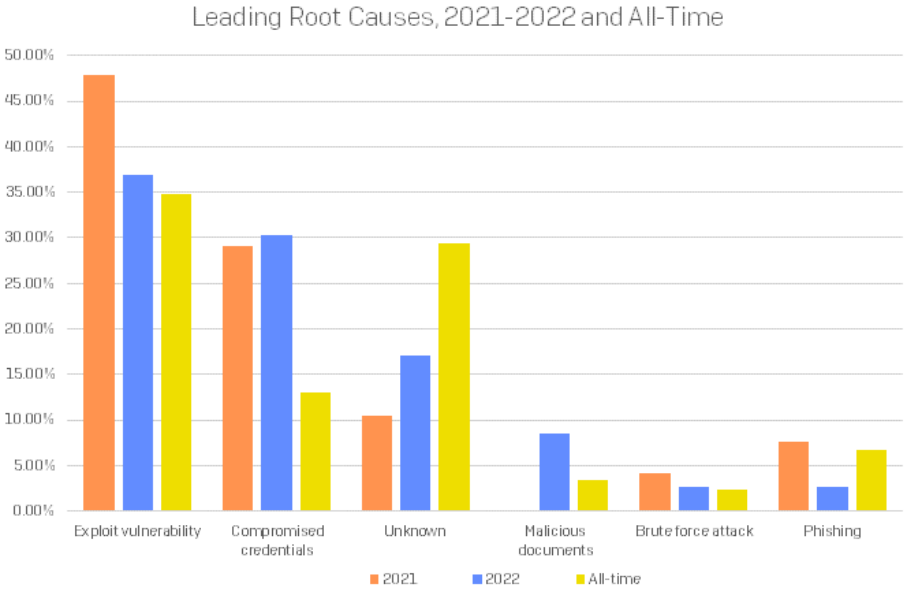

Figure 4: A 2022 uptick in compromised credentials may be related to the increase in Initial Access Broker (IAB) activity, but even that spike is dwarfed by the ongoing problem of vulnerabilities remaining unpatched and available to attackers. (Due to a change in how incidents were recorded after 2020, root-cause data from that era does not easily translate into this structure and is not broken out separately in these numbers; in addition, a handful of "corner case" incidents are not represented in this leading-causes chart. This chart was updated 26 May 2023.)

Many of these attacks could have been prevented if only the available patches had been implemented. To put numbers on it, in 55% of all investigations in which exploit vulnerability was the root cause, the exploitation of either the ProxyShell or the Log4Shell vulnerability was to blame. Patches for these vulnerabilities were made available in April/May 2021 and December 2021 respectively. There was precisely

one incident of a true zero-day attack in our dataset, where we saw the Log4Shell vulnerability used in July 2021, five months before the patch was released. "Zero-day" discoveries have always made a big splash in the public consciousness, and certainly your board of directors will wish to hear that the enterprise is aware of and addressing those high-profile threats, but it's the thankless monthly and quarterly grind of patching that truly makes the difference in your enterprise's risk profile.

The second most prevalent root cause was compromised credentials (30%). The provenance of these credentials is not usually known, but this root cause can often indicate the presence of initial access brokers (IAB) in the network. Where external remote services are involved, the first sign of an attack is often a successful logon with a valid account. If your services are not protected with multi-factor authentication (MFA), they should be. If MFA is unavailable for the service, it should be protected by something capable. Knowing where to draw the line between IAB activity and that of other criminals can be more art than science, but we have tried to isolate and identify trends associated with IABs in the Dwell Time section that follows below.

If there is one root cause we would love to see disappear from the next report, it is "Unknown," which is our notation for incidents in which the loss of forensic evidence on the network was so significant as to make it impossible to assign a root cause with a sufficiently high degree of confidence. Coming in third this year, "Unknown" accounts for 17% of all root causes, but is still the second most common all-time root cause at 29%. The problem with "Unknown" is that it prevents full remediation. If the organization does not know how the attackers get in, how will it fix the problem to prevent future attacks?

To be clear, loss of forensic evidence happens in many ways, some of them *not* due to the actions of attackers. Sometimes attackers wipe the data to erase their tracks, certainly, but other times the defenders will re-image systems prior to starting an investigation. Some systems are configured to overwrite their logs too quickly and/or frequently. Worst of all, some organizations do not collect the evidence in the first place. Whether the evidence was wiped by the attackers or the defenders, this loss of forensic data removes precious insight that might have been gained by its presence. Much like data backups, log backups are invaluable when faced with an incident response investigation.

As we parsed the 2022 cases for insight into root causes, a trio of related attacks caught our attention in the data. All three of the targets were healthcare organizations attacked by the Hive ransomware group, which was shut down by law enforcement in January 2023. The first attack leveraged compromised credentials to authenticate to the victim organization's VPN (Virtual Private Network), which had no MFA enabled. From this first foothold, the attackers used trusted relationships to access two additional targets using the credentials from the first one. This unfortunate chain of events illustrates how a compromise of one organization can lead to multiple attacks, all using the same or similar methods.

In a coincidentally inverse case, at one point we spotted three different threat actors attacking one target. All three attackers used the same initial access method — an exposed RDP (Remote Desktop Protocol) server — and leveraged compromised credentials. The first two attacks were separated by only a few hours, while the third occurred two weeks later. This example shows how opportunistic attackers can leverage the same root cause and initial access method to repeatedly victimize an organization. (As for RDP, a "star performer" in both of our past two Active Adversary reports, it's unfortunately just as prevalent in the tools mix as ever. We'll touch on RDP later in this report.)

Dwell Time: The good news, the bad news

Median dwell time decreased this year, which could signal both good and bad news depending on how you choose to interpret the data. The good news is that it *might* signal improvement in the detection of active attacks – a real improvement for defenders and their capabilities. The median dwell time for all attacks in 2022 was 10 days, down from 15 days in our last report. We saw the most improvement in non-ransomware dwell times, down more than 23 days from 34 days to 11 days. (There were, however, some outliers, with one victim hosting attackers in their network for more than 2.5 years.) Ransomware dwell times were also down in 2022, from 11 days to 9 days.

The bad news is that the attackers might be speeding up their efforts in response to improvements in detection capabilities. We'll be watching dwell-time statistics in particular throughout 2023 to see if we're observing a sea change in the ongoing back-and-forth between defenders and attackers.

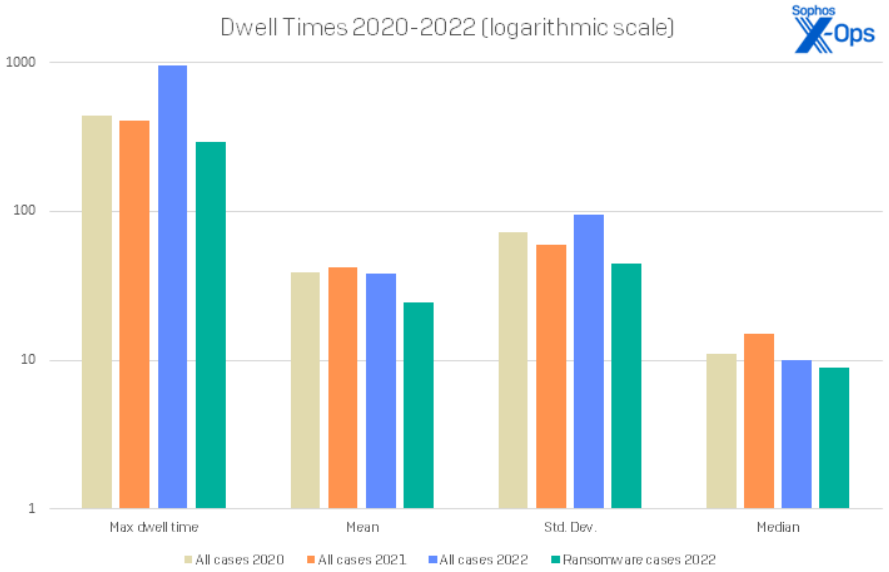

Figure 5: Even taking into account a single outlier that doubled the maximum dwell-time number not just for 2022 but for the entire history of the Active Adversary Report, the numbers for 2022 showed that median dwell times are dropping significantly – down to 10 days for the first time, and a day less for ransomware cases

It's possible that law-enforcement takedowns of IABs, and even of marketplaces for IABs such as Genesis, will drop dwell times even further; for now, though, it's simply too soon to be sure. In our previous Active Adversary report and Sophos Threat Reports we discussed the role that IABs play in the threat landscape. This year we attempted to tease out this factor by looking at the earliest signs of attack; in every investigation there is an attempt to identify when the attack started. Sometimes this is quite straightforward, since the first sign of attack is the start of the attack. However, there are many cases where additional suspicious or malicious activity pre-dates the accepted start date. This is the gray zone in which IABs operate. We know this information is more subjective, but we offer it as additional insight into attacker behavior.

This year saw no significant difference in dwell time among organization of different sizes or sectors. However, we did look at when these attacks were happening to understand if attackers showed a preference for a particular day of the week to either start the attack or launch their payload. The data showed no significant result for either, as their standard deviations were quite low. This reinforces the idea that most organizations are victims of opportunistic attacks, which can start or end any day of the week. With this kind of spread, having a team of trained analysts constantly monitoring the environment is of paramount importance.

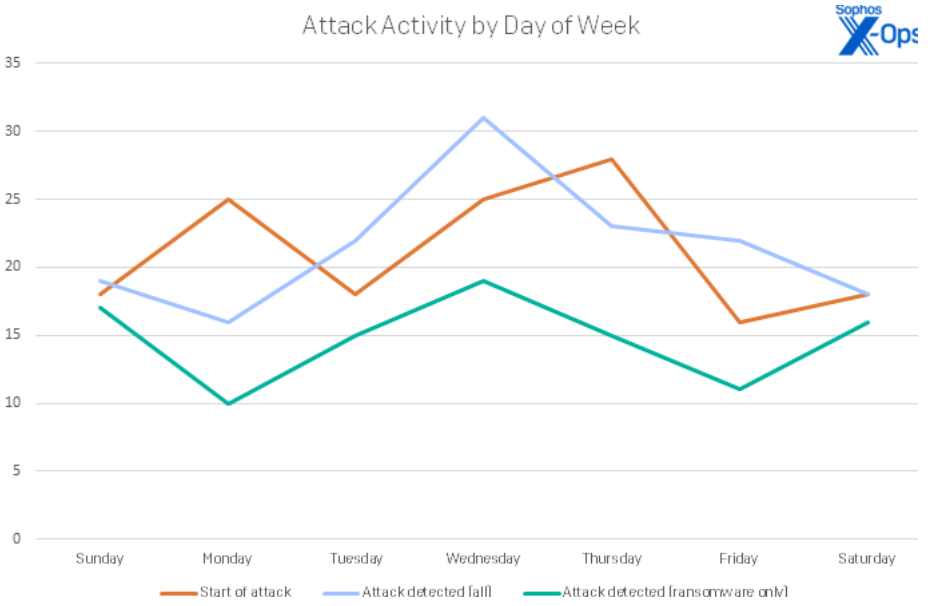

Figure 6: Attackers work a seven-day week, though it's arguable that everyone likes to take things a little easier on Fridays

There were some eye-watering revelations in what we call pre-attack dwell time. The first was the maximum dwell time. In one case, we saw evidence of attacker activity and previously undetected malware going back nearly nine years. The immense caveat to this is that we are unable to prove there was not any timestomping involved or other exacerbating factors contributing to this long dwell time. As such, we decided to focus on pre-attack dwell times where the maximum was less than two years. In this dataset we found 75 cases where the median dwell time was 91 days – better than 4000-odd days, of course, but not the hoped-for level of enterprise self-awareness concerning their own networks.

Attribution: The more things change...

Every year we publish a leaderboard of the most active ransomware groups in our dataset. This year, of the 104 ransomware cases investigated, LockBit took the top spot with 15.24% of the cases handled, followed closely by BlackCat (13%), Hive (12%), and Phobos (11%). When looking at unique active groups, we found 34 active ransomware gangs in 2022, versus 2021's 38. Comparing the last two years' active groups, 13 remained active across 2021 and 2022, whereas the two years differed by 25 and 21 unique groups, respectively – reinforcing the observation above that the players change, but the game remains the same.

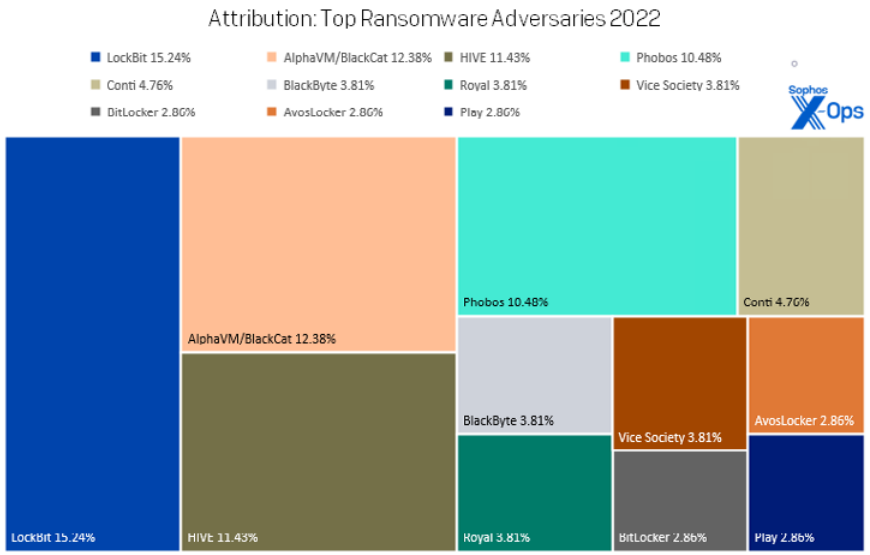

Figure 7: 2022's Top Ten leaderboard is a Top Eleven, with three different families in two separate three-way ties for position

One pattern that has emerged over time is that no group listing is permanent and any group from the past can rise to the top in any given year — all it takes is opportunity (whether that comes from intra-group conflicts or external disruption efforts by law enforcement or other means) and some room at the top. Given that these groups are nearly all ransomware-as-a-service (RaaS) operations, it's no surprise that affiliates will aggregate around certain well-known groups. Success and notoriety beget more of the same.

Some groups quietly persist year after year; for example, groups such as Cl0p, Cuba, and LockBit have been present in the ransomware landscape for many years and continue to attack organizations well into 2023. And, while LockBit has been steadily climbing our leaderboard, it only found its way to the top by sheer volume ***andthe*** disappearance of last year's leader, Conti. Conti, which shut down in early 2022 in the wake of the Russian ground invasion of Ukraine, still accounted for 5% of ransomware cases in 2022. We saw a similar situation in last year's report, where REvil operated at full power for only the first six months of 2021 before being taken offline but nevertheless attained the number two spot.

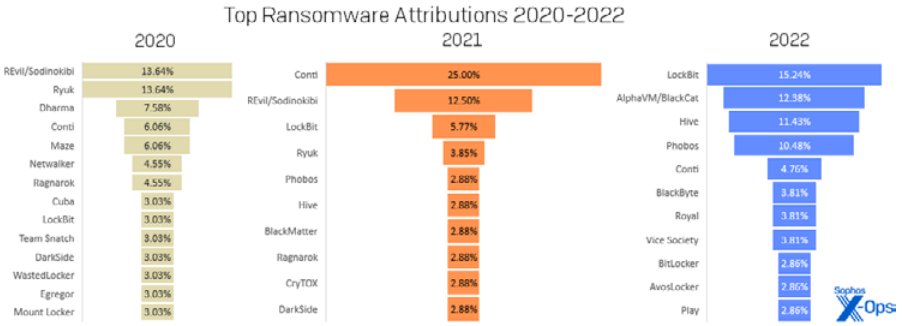

Figure 8: In 2021, Conti literally overshadowed the also-ran ransomware families. A year later, Conti was defunct, leaving other ransomware families to expand their share of the 2022 chart, and total incident numbers virtually unchanged. (The varied number of "top ten" families year-to-year represents a number of ties for infection prevalence, especially in 2020.)

Of the 70 different ransomware groups we have investigated in the first three years of this report, Conti is number one overall, followed by LockBit and REvil. Such was Conti's ubiquity in 2021 that they maintained their top spot despite their demise in 2022.

Despite the appeal of having a name to put with the mayhem, business leaders are cautioned to not overthink the rise or fall of any single ransomware group. We know that aside from Royal, which operates as a closed group, affiliates are frequently part of multiple ransomware-as-a-service (RaaS) programs – and, as will be clear when we look and tools and techniques, they all tend to look the same in the end.

Artifacts: Offensive weaponry, LOLbins, and random objects

We divide the artifacts (tools and techniques) we track into three categories. The first are tools that can be either legitimate offensive-security kit or bespoke hacking tools. The second are living-off-the-land binaries (LOLBins) found on most Windows operating systems. The third category is a catch-all of techniques, malware, cloud storage providers, and other difficult-to-categorize artifacts that we identified in our investigations.

This year we saw 524 unique tools and techniques used by attackers — 204 offensive/hacking tools; 118 LOLBins; and 202 other unique artifacts, which includes various tactics recognized in MITRE's ATT&CK taxonomy. With such a wide diversity of options in play, focusing detection efforts on any single tool or technique is a futile effort. Instead, organizations should limit the tools that are allowed to be present on systems, limit the scope of what these tools can do, and audit all use of approved tools.

As has become the norm for offensive/hacking tools, Cobalt Strike (42.76%) led the way in 2022, followed by AnyDesk (30.26%), mimikatz (28.29%), Advanced IP Scanner (21.71%), and Netscan (19.74%). (A notable non-factor: the Brute Ratel toolkit about which so fuss much was made last summer. Our investigators saw evidence of it just twice.) Certain tool categories are prominent as well; notably, tools that allow remote control of computers make up 7 of the top 15 tools. Some of the items in this category, such as Cobalt Strike, should always be blocked, while others, such as TeamViewer (14.47%), should be strictly controlled and (potentially) their use audited.

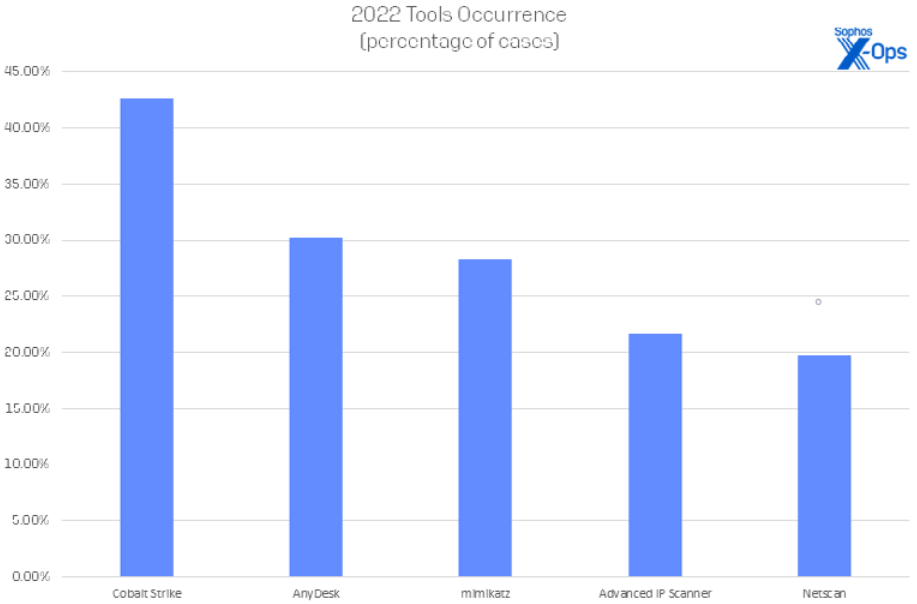

Figure 9: In 2022, we saw five tools used most frequently in attacks; these five occurred in at least 30 cases apiece. However, this is literally the tip of an iceberg for enterprises; an additional 200 tools were spotted at least once in the course of our investigations during the year.

In similar fashion, PowerShell (74.34%) leads the way in LOLBins in 2022 – sort of. (We traditionally exclude Remote Desktop Protocol (RDP) from these results due to its utter ubiquity in attacks; our next Active Adversary report, out in August, will delve into how infosec professionals can usefully address RDP-related risks.) PowerShell is trailed by cmd.exe (50.00%), PsExec (44.08%), Task Scheduler (28.29%), and net.exe (27.63%) to round out the top five. These were followed by rundll32.exe (25.66%) and WMI (19.74%) with smaller but still noteworthy presence.

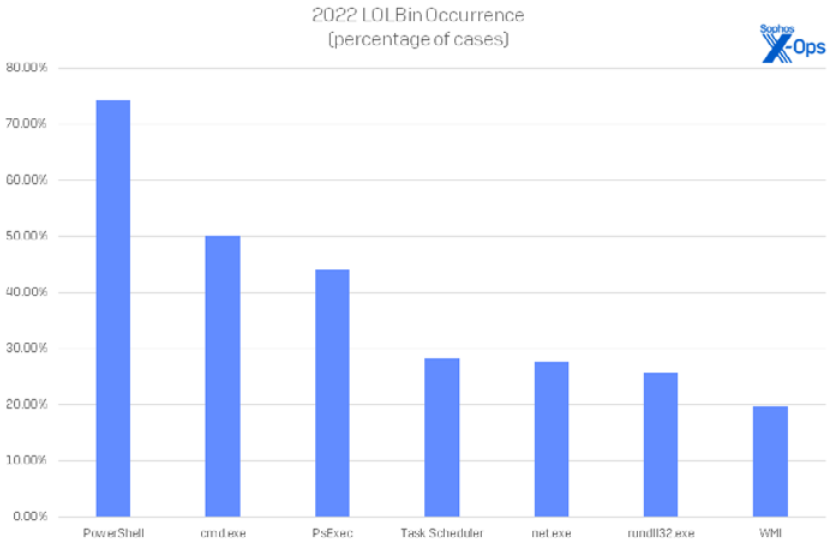

Figure 10: Seven Windows binaries (not including RDP) were significantly abused in 2022's caseload, meaning that we saw them in at least 30 cases. It should be noted, however, that an additional 111 LOLBins were seen in use at least once over the course of the year, demonstrating that attackers will use whatever's handy in the course of an attack.

The knee-jerk reaction to block or decommission these LOLBins is not useful risk management. Tools such as net.exe and rundll32.exe are versatile (and sometimes essential) and can be used for many different purposes; while they cannot practically speaking be blocked, it's good practice for your tech leadership to develop triggers for your XDR tools to catch activity involving these legitimate-but-often-abused binaries. Other, less-common binaries such as whoami.exe (16%) should always raise an alert, even in benign cases.

Finally, the catchall "other" category was topped by the Valid Accounts technique (71.05%), categorized by MITRE as T1078. In second place was System Services: Service Execution (aka installing a service, T1569.002, 63.82%); Command and Scripting Interpreter (aka executing a malicious script, T1059, 53.29%), File and Directory Discovery (aka browsing the network, T1083, 43.42%), and Impair Defenses (aka disabling protections, T1562, 36.18%) round out the top five. Techniques such as Indicator Removal (aka clearing logs, T1070, 34.87%), Modify Registry (T1112, 28.29%), and Create Account (T1136, 27.63%) often go together with disabling protections to evade defenses. Web shells (19.74%) were used as a persistence mechanism, while exposed RDP (19.74%) permitted easy access to victim networks.

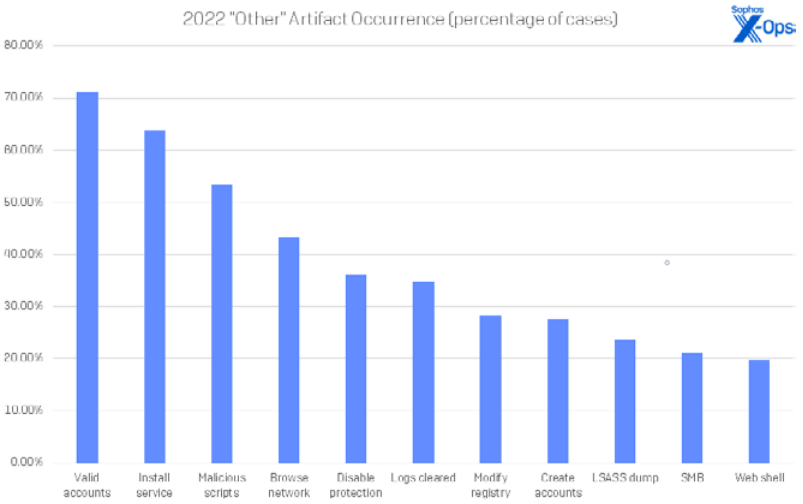

Figure 11: Eleven "other" artifacts occurred in 30 or more cases in 2022. Again, however, there is a very long tail of less-common "other" artifacts, with 191 unique items recorded in cases over the course of the year.

We give an overview of the top artifacts spotted at each stage of the MITRE ATT&CK sequence in Figure 12.

STAGE OF ATTACK	PRESENCE OF ARTIFACT IN INVESTIGATIONS					
Initial Access	External remote service	44.74%	Exploit public-facing application	33.55%	Valid accounts	32.89%
Execution	PowerShell	69.74%	cmd.exe	48.03%	PsExec	25.66%
Persistence	Valid accounts	64.47%	Install service	57.24%	Cobalt Strike	30.26%
Privilege Escalation	Valid accounts	68.42%	Modify local groups	10.53%	Create accounts	5.92%
Defense Evasion	Valid accounts	71.05%	Disable protection	34.87%	Logs cleared	33.55%
Credential Access	mimikatz	24.34%	LSASS dump	23.03%	Brute force attack	7.24%
Discovery	Browse network	41.45%	net.exe	21.05%	Netscan	18.42%
Lateral Movement	RDP	82.24%	SMB	11.18%	PsExec	7.24%
Collection	Browse Network	21.71%	Rclone	11.18%	WinRAR	8.55%
Command and Control	Cobalt Strike	37.50%	PowerShell	19.74%	AnyDesk	7.24%
Exfiltration	Mega (all TLDs)	12.50%	Rclone	11.84%	Megasync	6.58%
Impact	Data encrypted for impact	67.76%	"No impact" *	32.89%	Resource hijacking / inhibit system recovery (tie)	9.13%

Figure 12: The top three investigation artifacts noted in each of 12 of the 14 fields of the MITRE ATT&CK Matrix. The first two fields of that matrix, Reconnaissance and Resource Development, are out of scope for incident-response investigations of this sort and are thus not represented in our data. Since cases may exhibit more than one of certain types of artifact (eg., Execution), percentages may add up to over 100 percent

* "No impact" indicates that the impact of the incident, ranging from network breach to data extortion to Web shells to stolen data and beyond, didn't fit into MITRE's 13 subcategories of recognized impact. MITRE also doesn't cover such non-IT concerns as time spent to remediate, reputation damage, lost productivity, legal and compliance costs, etc.

Exfiltration, leakage, and theft: Goodbye, data

While it may be that exfiltration and extortion attacks are only slightly on the rise, the amount of data theft and leakage from all sources, including ransomware attacks, is still staggering. There were 65 confirmed data exfiltration events in our dataset in 2022. At nearly half (42.76%) of investigated cases, this statistic suggests your data stands a nearly even chance of being stolen during an attack, as shown in Figure 13.

EXFILTRATION OCCURRENCE, 2022 CASES	
No conclusive evidence	47.37%
Yes	42.76%
Possible	5.26%
Staging detected	3.29%
No exfiltration	1.32%

Figure 13: Data exfiltration occurrences in the 2022 dataset. Note that exfiltration could be conclusively disproven in less than 2% of all cases

In ransomware-only attacks, over half (55%) involved confirmed exfiltration, while another 12% of cases showed signs of possible exfiltration or data staging. Of those cases in which data was exfiltrated, half (49%) provably resulted in confirmed leaks, meaning that we were able to locate evidence of the exfiltrated data on one of the leak sites.

Sometimes when data is stolen and held for ransom, we can also identify the time between the exfiltration event and the ransom notification. This is an important window, because it provides another clear signal that you are under attack — but you need to be watching. The median time between the beginning of exfiltration and the ransom event was 2.14 days. This is slightly longer than last year's measure of 1.84 days.

Conversely, we noted that just over 47% of all attacks showed no conclusive evidence of data exfiltration. That is, however, not the good news it would appear to be at first glance. What is most worrying is that in many cases, it is not that the logs simply showed no evidence, but rather that they were incomplete or missing. Therefore, it is safe to conclude that much more data ***may*** have been stolen, since many of the ransomware groups involved in these attacks are known to exfiltrate data. This makes a compelling case for comprehensive and continuous network traffic logging and analysis.

Unfortunately, it is not possible to know in these cases if more data was leaked, even if no ransom demand is received. There remains the unfortunate possibility that the stolen data was sold privately to another criminal. In addition, multiple ransomware groups provide the option for bidders – not just the victim — to pay for access to stolen data. The bottom line is that half the time we do not know if your data was stolen or if that stolen data has been leaked, and crucially ***neither do you.***

Conclusion

Whatever the size of the organization, country, or industry, no one is perfectly safe from attack – even if years of exposure have numbed many enterprises to the basics of patching, keeping useful log data, and planning for ransomware. Industry advances in detection and response technology and services hold promise as a way to both repel attackers and to thwart their purposes if they manage to evade initial layers of defense. As a result, attacker dwell time has declined significantly in the past year. However, complacency is dangerous: Once attackers have established a foothold on your network, it is as likely as not that your data will be exfiltrated, whether to be sold back to you or to the highest bidder. Business leaders are advised to take these possible outcomes into consideration as they plan how to allocate their continued efforts to secure their networks in 2023 and beyond.

Appendix: Demographics and methodology

As we put together this report, we chose to narrow our focus to 152 cases that could be meaningfully parsed for useful information on the state of the adversary landscape as of the end of 2022. Protecting the confidential relationship between Sophos and our customers is of course our first priority, and the data you see here has been vetted at multiple stages during this process to ensure that no single customer is identifiable through this data – and that no single customer's data skews the aggregate inappropriately. When in doubt about a specific case, we excluded that customer's data from the dataset.

Nations

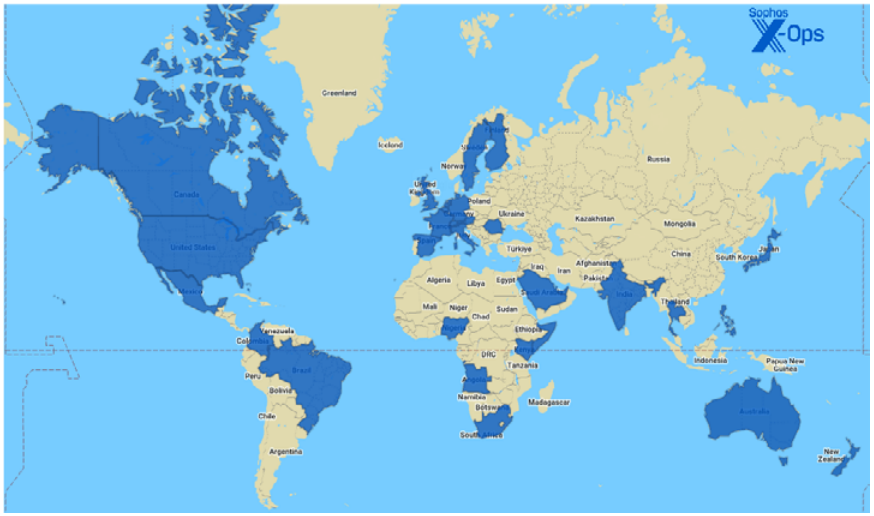

Figure 14: Sophos' incident-response travels

The full list of nations represented in this report is as follows:

Angola	Hong Kong	Saudi Arabia
Australia	India	Singapore
Austria	Italy	Somalia
Bahamas	Japan	South Africa
Bahrain	Kenya	Spain
Belgium	Mexico	Sweden
Brazil	Netherlands	Switzerland
Canada	New Zealand	Thailand
Colombia	Nigeria	United Arab Emirates
Finland	Philippines	United Kingdom
France	Qatar	United States
Germany	Romania	

Industries

The full list of industries represented in this report is as follows:

Agriculture	Food	News media
Architecture	Government	Non-profit
Communication	Healthcare	Pharmaceutical
Construction	Hospitality	Real estate
Education	Information technology	Retail
Energy	Legal	Services
Entertainment	Logistics	Transportation
Finance services	Manufacturing	Utilities
Financial	MSP/Hosting	

Methodology

The data in this report was captured over the course of individual investigations undertaken by Sophos' X-Ops Incident Response team. For this initial report of 2023, we gathered case information on all investigations undertaken by the team in 2022 and normalized it across 37 fields, examining each case to ensure that the data available was appropriate in detail and scope for aggregate reporting as defined by the focus of the proposed report.

When data was unclear or unavailable, the author worked with individual IR case leads to clear up questions or confusion. Incidents that could not be clarified sufficiently for the purpose of the report, or about which we concluded that inclusion risked exposure or other potential harm to the Sophos-client relationship, were set aside. We then examined each remaining case's timeline to gain further clarity on such matters as initial ingress, dwell time, exfiltration, and so forth. We retained 152 cases, and those are the foundation of the report.

2023-08-29 (WP-NP)

SOPHOS