RELIAQUEST®

Annual Cyber-Threat Report:

ReliaQuest Annual Cyber-Threat Report

2023

Make Security Possible

Table of Contents

Executive Summary

This report offers an in-depth and inclusive view of the cyber-threat landscape, over what has been an exceptionally challenging 12 months for cyber security professionals (1 February 2022 to 1 February 2023). Over the course of this period, the ReliaQuest Threat Research Team has identified trends across several data sources and analyzed them to provide readers with insights into cyber-threat trends and observations. During a reporting period that saw Russia's invasion of Ukraine, a continued risk of ransomware attacks and data extortion, and an avalanche of high-risk vulnerabilities, we identified the following key events and patterns:

- The attempted exploitation of exposed remote services was the most commonly detected attack technique. Those services, including virtual private networks (VPNs) and remote desktop protocol (RDP), pose a very high risk. A wide range of attackers, including cybercriminals and nation-state–aligned groups, have exploited them to access a network or establish persistence on it.
- ReliaQuest identified wide use of defense evasion techniques, notably indicator removal, data destruction, and the sub-technique of clear command history. This emphasizes the significant efforts threat actors place on obfuscating their activity.
- ReliaQuest's GreyMatter Digital Risk Protection (GMDRP) service yielded data that identified especially vulnerable sectors. They are most susceptible to: fraud via impersonating retail web domains, significant risk from exposed credentials (particularly for financial services), and exploitation of open ports at utilities companies.
- CVE-2022-22965 (aka Spring4Shell) is regarded to pose the greatest risk of all high-risk vulnerabilities, for its available and easy exploits and its potential to cause a technical and business impact.
- The most common access type advertised by initial access brokers (IABs) was RDP, which accounted for 24.4% of all ReliaQuest "tippers" published in the reporting period. RDP access was also the costliest type being offered, with an average median price of approximately $1,000.
- Initial-access malware continued to be delivered mainly by phishing emails. Threat actors adapted their techniques to circumvent organizational controls, and in ReliaQuest customer environments, we detected many instances of the "Emotet," "SocGholish," "IcedID," "GootLoader," and "Bumblebee" malware.
- "LockBit" was, overwhelmingly, the most active ransomware group, and is increasingly using the SocGholish malware distribution framework to gain initial access to networks which is making their efforts more potent. We anticipate even more use of SocGholish by ransomware groups during 2023.
- NameCheap was the most common registrar of Cobalt Strike team servers, followed by Ename Technology and MarkMonitor. These registrars are primarily content delivery networks (CDNs) used for domain fronting. Domain fronting is used to conceal user traffic and is commonly used by threat actors for command and control (C2) purposes.

What's in the Report?

In this first ReliaQuest Annual Cyber-threat Report, we cover activity observed from 1 February 2022 to 1 February 2023:

- Trends in events recorded in GreyMatter
- The most commonly found risk types on GMDRP
- Trends related to initial access brokers (IABs): the cybercriminal gatekeepers that enable a raft of malicious activity
- An introduction to vulnerability intelligence, highlighting the "red-flag" flaws in platforms that ReliaQuest customers use most
- Trends related to ransomware, including a breakdown of the most commonly targeted sectors and regions
- Trends related to Cobalt Strike and command-and-control (C2) systems used by ransomware operators

The Evolving Cyber Threat Landscape: Understanding the Current Risks

The cyber threat landscape is increasingly complex and subject to consistent change. The rapid advancement of technology and reliance on digital systems leave individuals and organizations facing an array of cyber threats. While cybercriminals and nation-state aligned threats use varying techniques, there are consistent observable trends that are used by many of these groups. Taking a deep dive into these trends by collating and analyzing data sources across our customer environments, allowed us to identify several insights into the state of the current cyber threat landscape.

Our report starts by looking at observations on active times of the year, most commonly observed kill chain phases and attacker techniques, before moving onto insights into a typical attack lifecycle. This involves the identification and exploitation of security weaknesses by an IAB, vulnerability exploitation, before moving to a network compromise by a ransomware actor. We conclude the report by identifying trends related to Cobalt Strike usage, which remains one of the most common methods of facilitating C2 over a compromised network, often used by ransomware actors.

As our CEO Brian Murphy has previously stated, cybersecurity likely represents the greatest technical challenge of our generation. By taking a retrospective look at the data from the previous year, we aim to enable our customers to take the best stance on preventing the cyber threats of 2023.

What Our Data Revealed

The data for this analysis was extracted from customer incidents collated by ReliaQuest. The data set covers February 1, 2022, 00:00:00 UTC to February 1, 2023, 00:00:00 UTC (12 months). There were 35,024 true-positive incidents during that reporting period.

Active Months and Targeted Sectors

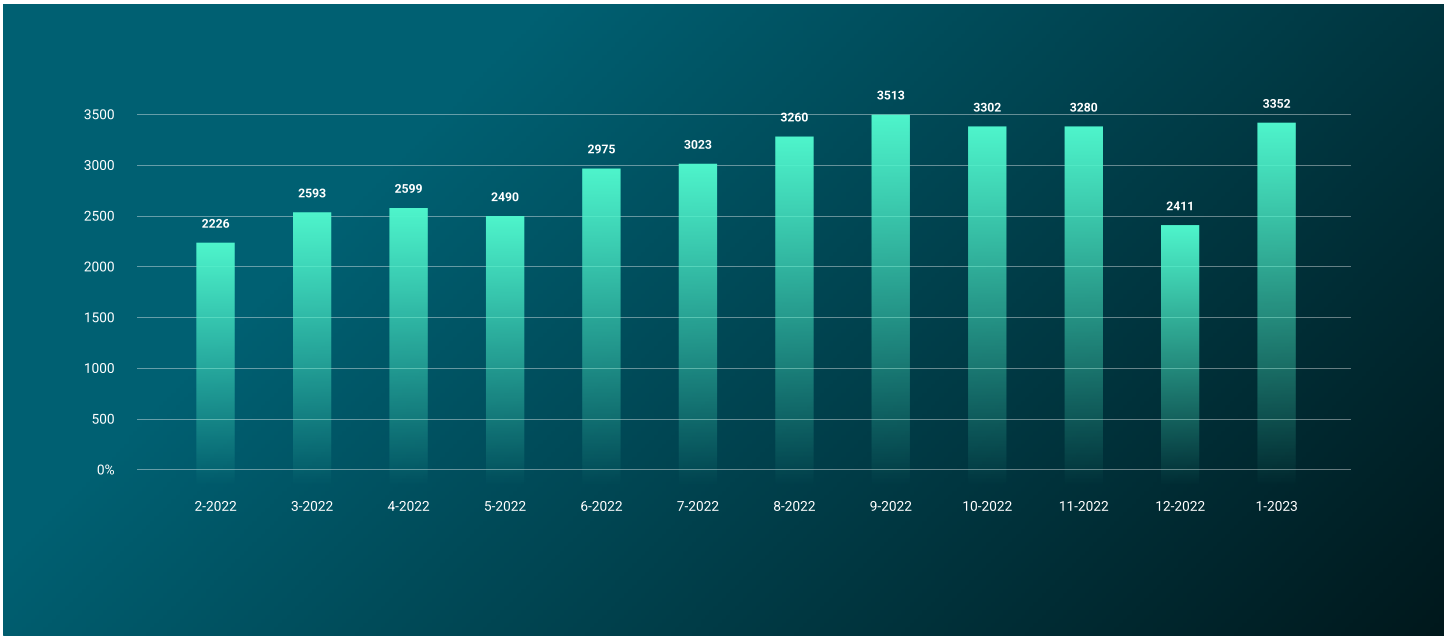

Figure 1: Number of true-positive tracked incidents per month

Figure 1 shows a steady increase in true positives[1], notably between August 2022 and January 2023, but with a noticeable drop-off in December 2023—potentially because threat actors were less active over the end-of-year festive holiday season.

[1] A True Positive or Confirmed Incident is an event or alert which identified malicious activity that resulted in either an attempt or successful unauthorized access, use, modification, or destruction of any information system or data.

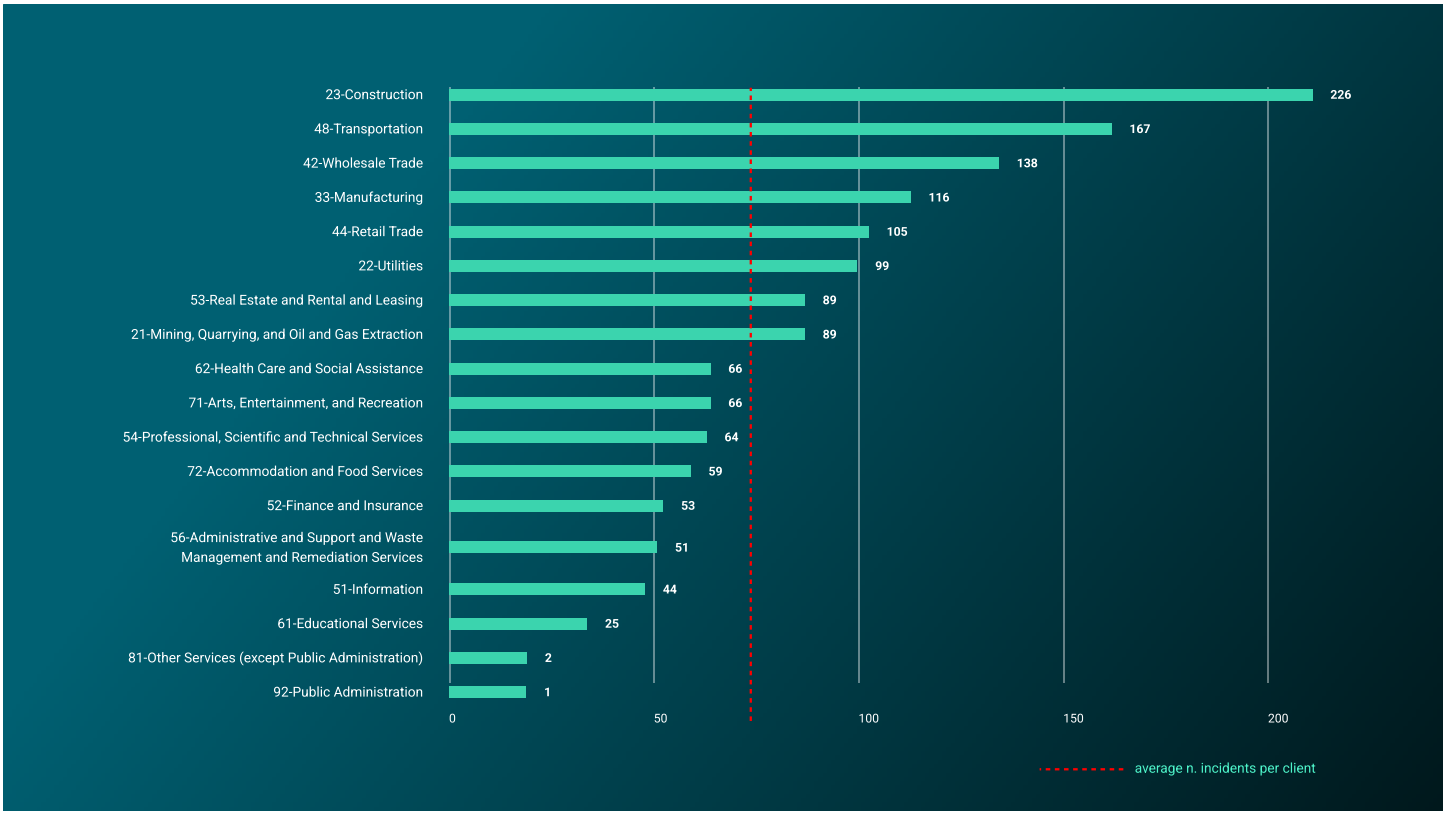

Figure 2: Average number of incidents per customer sector

The number in front of each Figure 2 sector shows how many incidents, on average, any given client in that sector can expect to see over the course of a year. We calculated it by dividing the number of incidents affecting that sector (the number at the end of each teal bar) by the number of GreyMatter customers operating in that sector. The red line indicates the average number of incidents all customers can expect to see over a year: 71. For that we divided the total number of GreyMatter incidents for all customers by the total number of customers. If a sector has the red line going through the grey bar, that sector has more incidents than average. Construction had the most during this reporting period, followed by transportation, then wholesale trade.

Each of these sectors are increasingly **relying on IT to drive efficiencies** making them susceptible to cyber attacks.

The perceived lack of cybersecurity maturity, controls, and tools paired with the significant impacts of outages is likely to have placed the Construction, Transportation, and Wholesale Trade sectors in the cross-hairs of threat actors.

Each of these sectors are increasingly relying on IT to drive efficiencies making them susceptible to cyber attacks.

Most Common Kill Chain Phases

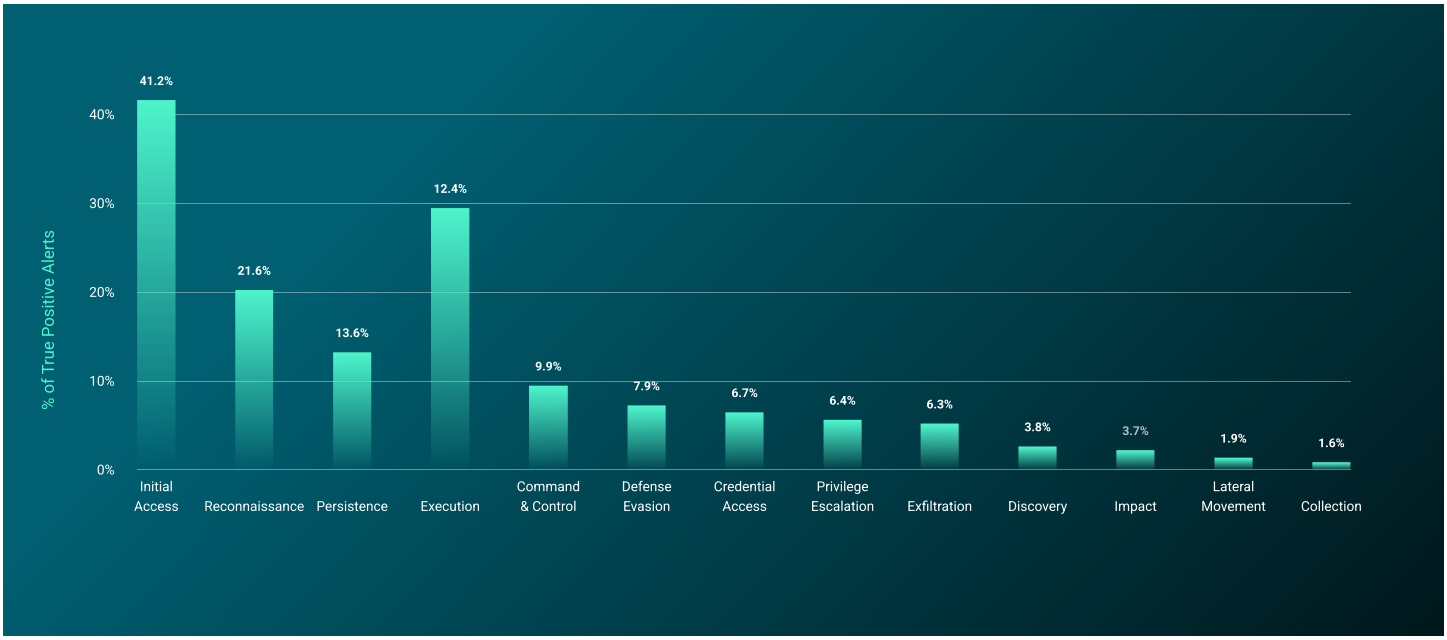

Figure 3: Kill chain phases detected in incidents, ranked by prevalence

Figure 3 highlights the most commonly detected kill-chain phases, identified through incident data. You can see the efficiency of GreyMatter in identifying malicious activity in the early stages of an attack lifecycle (e.g., initial access and reconnaissance), before an attacker has time to progress their intrusion or establish persistence on a targeted network. ReliaQuest focuses heavily on detection, and ensuring customers are in the best possible position to respond in the early stages of threat actor activity.

Most Commonly Detected Techniques

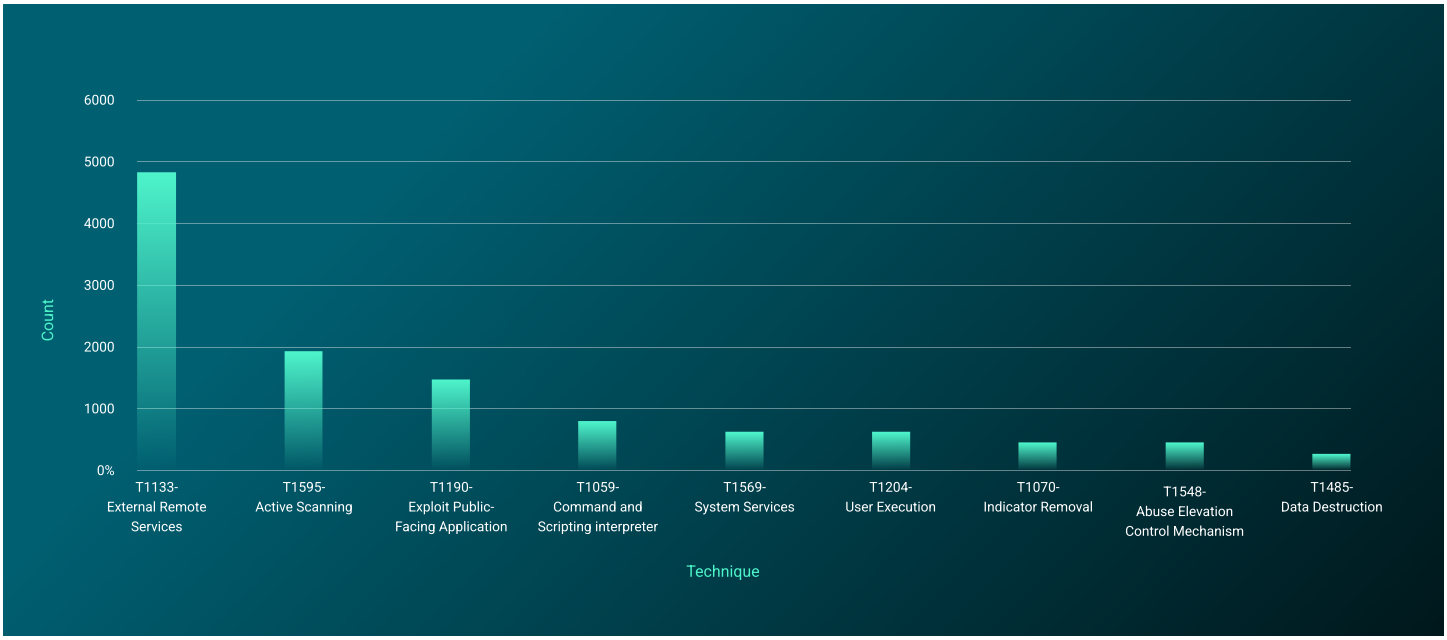

Figure 4: Top 10 most detected techniques

Taken from our collection of true positive incident data was the most commonly observed attacker techniques, detailing methods identified by ReliaQuest as part of investigations on customer environments. Figure 4 above highlights the top 10 of these techniques, including T1133 External Remotes Services which was overwhelmingly the most commonly seen technique. This comes as no surprise; exposed remote services, including VPN, Citrix, TeamViewer or RDP, represent one of the most common methods of enabling initial access onto a targeted network, or establishing persistence. We have observed significant threat actor interest in identifying exposed RDP servers, which has resulted in a flourishing ecosystem of cybercriminal activity in identifying, exploiting, then selling RDP accesses onto interested third parties; we go into more detail on these accesses in our section related to initial access broker (IAB) activity later in this paper.

Also represented within the top ten techniques was T1070 Indicator Removal and T1485 Data Destruction, which were the 8th and 10th most observed respectively. Indicator Removal refers to the attempt from adversaries to delete or modify artifacts generated within systems, to remove evidence of their actions on objective and hinder investigative efforts. Data destruction refers to adversary attempts to destroy data and files on specific systems or in large numbers on a network to interrupt availability to systems, services, and network resources. We found these techniques notable in particular due to the implications they have on defenders. To deny Defense Evasion and Impact techniques, organizations should prioritize technical controls and ensure sufficient monitoring is in place to identify suspicious activity. As we move on to the top sub-techniques identified in our data set, we also see sub-techniques that map back to the Defense Evasion tactic which again reiterates our previous stance. The top 10 most commonly observed can be seen in Figure 5.

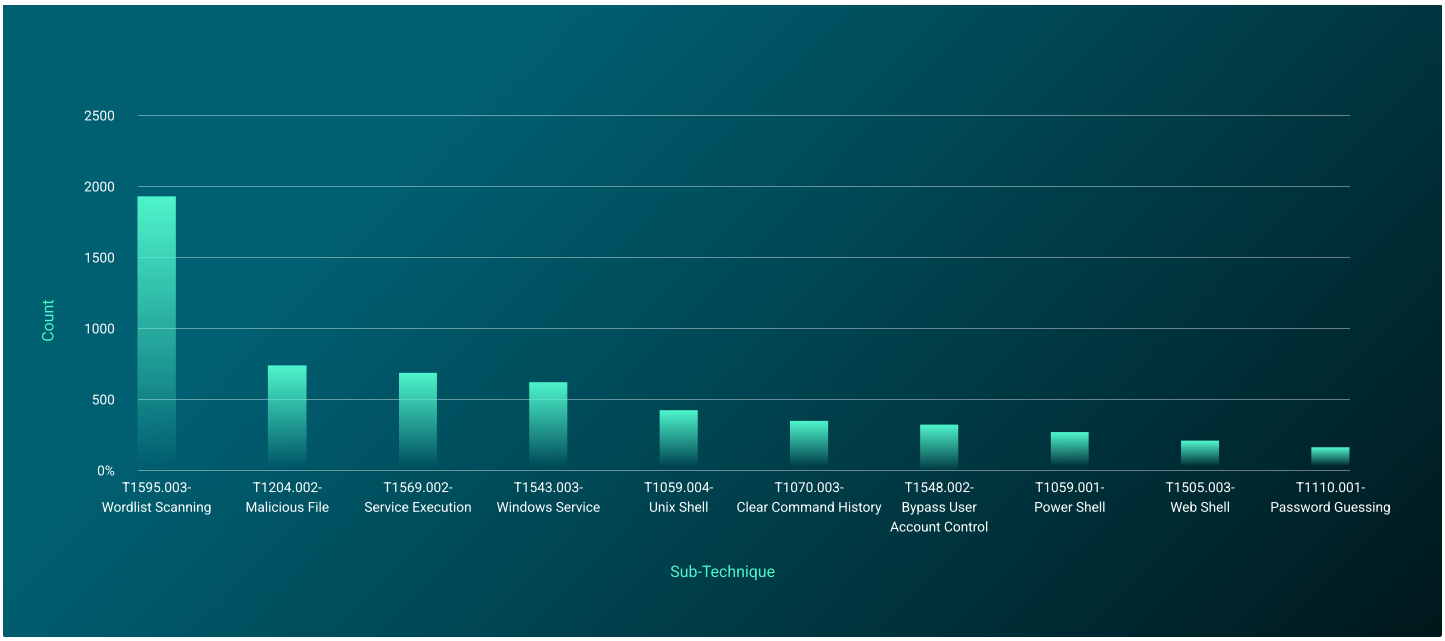

Figure 5: Top 10 most commonly observed sub-techniques

Within the top 10 sub-techniques in the 8th spot is T1070.003 Clear Command History, which involves an adversary attempting to clear the command history of a compromised account, in an attempt to conceal the actions undertaken during an intrusion. This emphasizes the importance of detecting suspicious activity early in the attack lifecycle before a threat actor has a chance to establish persistence on your network; with capable threat actors it's often extremely difficult to fully remove their presence if they have developed several methods of gaining access.

The most commonly observed sub technique was T1595.003 Word List Scanning. This refers to threat actor attempts to iteratively probe infrastructure using brute-forcing and crawling techniques. While this technique employs similar methods to brute force techniques, its goal is the identification of content and infrastructure rather than the discovery of valid credentials. Word lists used in these scans may contain generic, commonly used names and file extensions or terms, that are specific to a particular software. Adversaries may also create custom, target-specific wordlists using data gathered from other reconnaissance techniques. The best method of minimizing the risk from Wordlist scans is ensuring services are not unnecessarily exposed to the internet, i.e., make sure your systems, resources, and infrastructure are not exposed externally unless they have a specific business requirement.

GreyMatter Digital Risk Protection (GMDRP) Alert Trends:

Known as GreyMatter Digital Risk Protection (GMDRP) ReliaQuest customers can receive intelligence, risks, and alerts that enable security operations to make informed decisions about threats to their environments. This capability also enhances visibility of previously unknown threats, by enriching incidents with actionable intel, reducing the number of false positives, and driving faster business outcomes.

By continually monitoring open, deep, and dark web sources to isolate legitimate threats and provide quick and easy remediation, it provides a unique view of security threats outside an organization. If a risk is detected, GMDRP customers receive context-rich alerts with clear response steps via the GreyMatter alert process and workflow. This includes alert-assignment, automated-action, and mitigation recommendations. For more information on GMDRP, see our solutions brief.

GMDRP covers 40 unique risk types that affected most sectors during the reporting period. They include credential exposure, impersonating domains and phishing sites, leaked documents, and exposed code. The data for this section was extracted from all the alerts provided to ReliaQuest customers in GMDRP from February 1, 2022, 00:00:00 UTC to February 1, 2023, 00:00:00 UTC (12 months).

Most Common GMDRP Risk Alerts

Figure 6 shows the top 20 most common risk types escalated to ReliaQuest GMDRP customers. This intelligence can be used to consider how the types of risk are managed across your business

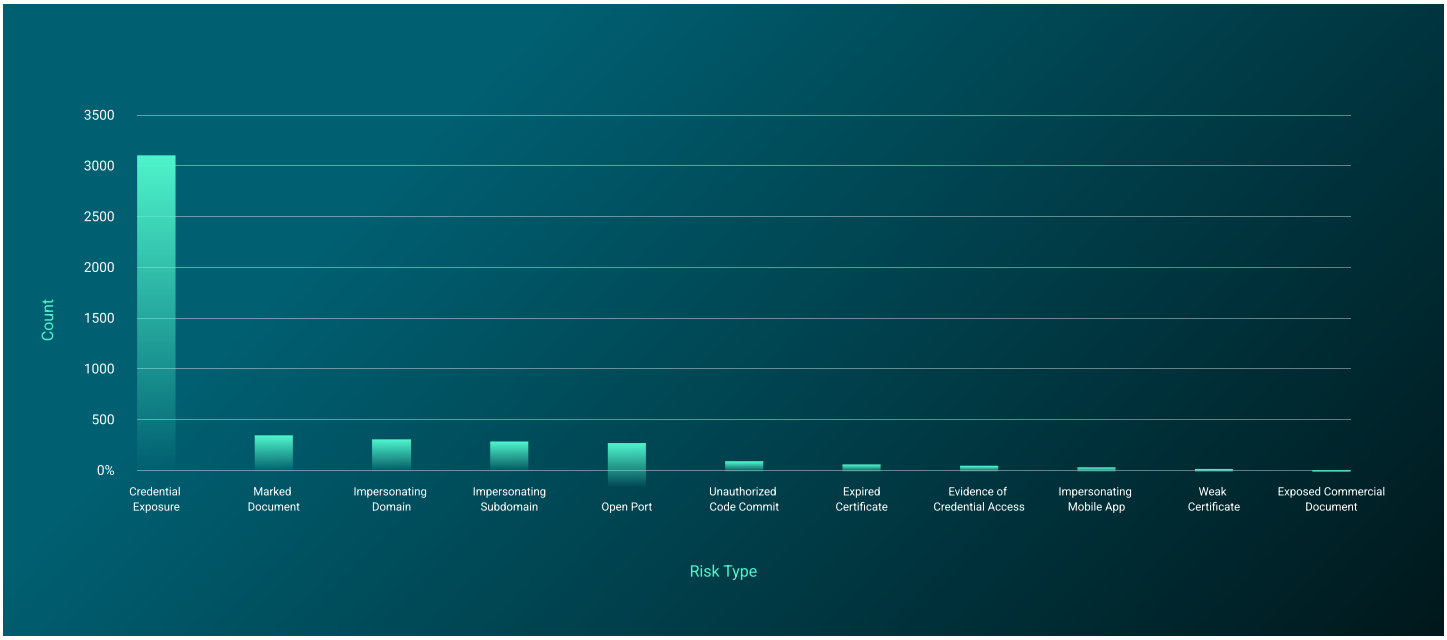

Figure 6: Most common risk types, by thousands of GMDRP alerts

The most common was, overwhelmingly, credential exposure—unsurprising, given the abundance of third-party breaches, risky user practices, and insufficient controls over managing credentials. In recent blogs we have commonly referred to the huge problem of credential exposure, which continues to fuel a host of cyber threats. Stolen credentials are the most common method of gaining initial access to networks. According to the 2022 Verizon Data Breach Investigations Report, stolen credentials enabled initial access in over 50% of 20,000 analyzed incidents.

The next most-common alert was for a detection of an exposed marked document. This can be documents with protective markings or data loss prevention (DLP) identifiers, but also technical and commercial documents (e.g., security assessments, product designs, legal documents, payroll data), which are often unmarked. Exactly why such sensitive documents get breached so often probably depends on the targeted organization, but most likely it is facilitated by a combination of risky user practices and insufficient data-loss prevention policies and controls.

The third and fourth most-commonly represented asset types were impersonating domains and impersonating subdomains. These are registered domains masquerading as a client's brand, company name, or domain, typically using typo-squatting and combosquatting. This is not a new issue, nor will it surprise anyone to see it represented so highly on our most commonly triggered alerts list. Domain impersonation occurs because it continues to work and many companies struggle to understand how they can detect such infringements and quickly take the sites down.

Other reasons for the explosion of fake domains include the sheer number of top-level domains (TLDs) that go beyond the usual .com, .net, and .org. Try to register a domain, and you'll be presented with probably dozens of options with various spellings and TLDs; it's this variety that is being exploited to trick internet users.

Another factor is the low barrier of entry to the world of cybercrime. The criminal underground has responded to novice threat actors' demand for registering and hosting domains, setting up phishing services or professionally-spoofed pages on a large scale, and lots of niche tools that support these enterprises; if you need it, someone will probably sell it to you, or even share it for free. There are also many tutorials and advice to aid fraud on the web—and they do not require substantial resources, whether technical or financial.

Alerts via Sector

As part of our analysis of the GMDRP data, we also determined which ReliaQuest customer sectors received the most alerts, divided by the total number of current customers. The output of this can be seen below in Figure 7, highlighting the top 10 with the highest number. (So, remember, the results reflect only specific companies/sectors making up ReliaQuest's customer base.)

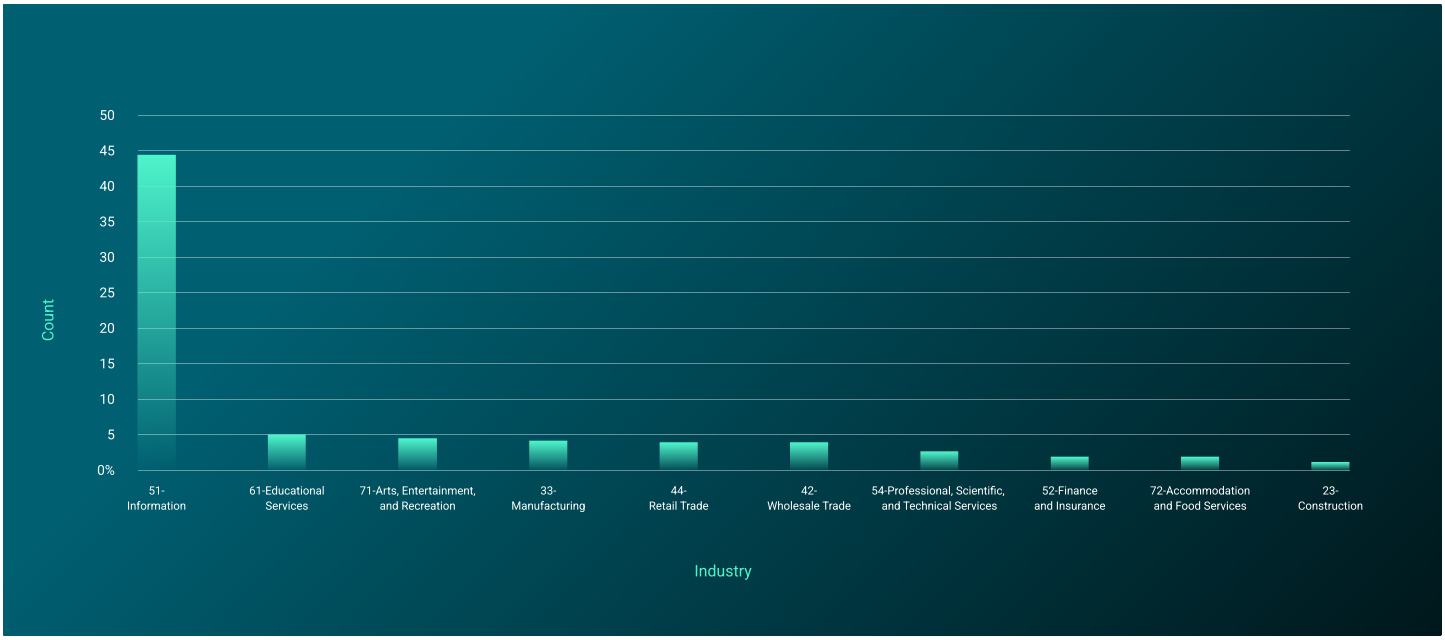

Figure 7: Ten ReliaQuest-customer sectors that received the most alerts, divided by total number of clients

During the reporting period (February 1, 2022 to February 1, 2023, **ReliaQuest identified 231,150 impersonating domains** and 162,895 impersonating subdomains.

Table 1 cross-references those ten sectors with the ten most common risk types. We divided the total number of alerts by the total number of ReliaQuest customers in each sector, to give a more accurate representation of the exposure to each risk type. From this table we can see which risks are most pertinent for specific sectors.

Risk Type	Finance and Insurance	Professional, Scientific, and Technical Services	Health Care and Social Assistance	Retail Trade	Information	Accommodation and Food Services	Wholesale Trade	Manufacturing	Mining, Quarrying, and Oil and Gas	Public Administration
CREDENTIAL EXPOSURE	501.14	353.25	235.48	859.43	38949.24	108.19	590.47	883.08	88.64	178
MARKED DOCUMENT	331.37	789.32	8.48	456.48	77.67	17.12	1031.67	612.54	2.09	26.36
IMPERSONATING DOMAIN	220.53	192.61	72.1	795.1	385.62	216.38	322.93	364.31	152.73	61.73
IMPERSONATING SUBDOMAIN	71.63	271.12	17.56	99.05	445.52	31.19	214.2	500.38	24.91	71.18
EXPIRED CERTIFICATE	58.12	34.24	28.44	60.71	215.1	89.81	40.6	64.54	138.64	17.55
OPEN PORT	35.95	113.39	34.08	65.81	2812.43	272.31	125.6	291.38	81	13.64
EVIDENCE OF CREDENTIAL ACCESS	18.92	63.12	6.75	201	288.14	278.88	20.87	64.69	5.64	12.64
IMPERSONATING MOBILE APP	18.54	33.68	11.06	95.52	243.57	113.62	212.67	139.54	79.64	44.09
UNAUTHORIZED CODE COMMIT	13.9	84.17	5.73	56.33	84.9	31.38	50.87	92.92	20.45	39.27
WEAK CERTIFICATE	9.22	50.4	4.1	16	211.76	16.38	14.13	15.15	75.55	3.73

Table 1: Ten most common risk types per sector, by number of GMDRP alerts

The first data point that stands out is the abundant credentials of financial and insurance companies being exposed. The surface reason is obvious: access to credentials for financial services are a high priority for financially motivated cybercriminals. But the appeal is likely compounded by a lack of authentication practices; a recent study found that only 32% of financial service companies authenticate account logins with additional measures, such as two-factor authentication (2FA). For financial services, the level of risk is determined by the sufficiency of controls and the motive of a threat actor eyeing up accounts.

Open port exposures were also very common in the information sector, which includes telecommunication companies. This alert type relates to risky ports that have been left open on client domains or IP addresses. Ports are typically a security risk if the services running on them are misconfigured, unpatched, or otherwise vulnerable. Threat actors can easily identify unnecessarily exposed ports, which often can identify exploitable services.

For interested threat actors, there are several guides to scanning and exploiting exposed ports. The one quoted in Figure 8 highlights common mistakes companies make with ports, including insufficient authentication or misconfiguration of file transfer protocol (FTP) servers.

```
--[ 4 ]-- Scanning & Exploiting

Scan all the IP ranges you found with nmap to find all services running. Aside
from a standard port scan, scanning for SNMP is underrated.

Now for each service you find running:

1) Is it exposing something it shouldn't? Sometimes companies will have services
running that require no authentication and just assume it's safe because the url
or IP to access it isn't public. Maybe fierce found a git subdomain and you can
go to git.companyname.come/gitweb/ and browse their source code.

2) Is it horribly misconfigured? Maybe they have an ftp server that allows
anonymous read or write access to an important directory. Maybe they have a
database server with a blank admin password (lol stratfor). Maybe their embedded
devices (VOIP boxes, IP Cameras, routers etc) are using the manufacturer's
default password.

3) Is it running an old version of software vulnerable to a public exploit?


Webservers deserve their own category. For any webservers, including ones nmap
will often find running on nonstandard ports, I usually:

1) Browse them. Especially on subdomains that fierce finds which aren't intended
for public viewing like test.company.com or dev.company.com you'll often find
interesting stuff just by looking at them.
```

Figure 8: Excerpt of a port-scanning guide shared on a cybercriminal forum

Some ports should only be exposed internally; a good example is the infamous Port 445, which is used for Microsoft Directory Services for Active Directory (AD) and for the Server Message Block (SMB) protocol over TCP/IP. This port has been exploited over the years, notably with the use of the "WannaCry" ransomware and similar SMB exploits. It is not totally clear why telecommunications would be so susceptible to exposed ports, but it is realistically possible that many exposures are linked to the necessity for communication, plus the likelihood that telecommunications companies will have wider ranges of IP addresses to manage.

Domains targeting retail and trade sector companies can easily be created, often through dedicated phishing kits that are widely available on cybercriminal forums.

Another detail that won't shock most readers is the number of impersonating domains targeting retail and trade sector companies. These domains can easily be created, often through dedicated phishing kits that are widely available on cybercriminal forums. We found several examples, including the kit referenced in Figure 9, aimed at spoofing Google and available to rent for $2,000 per month. That kit was advertised on a high-profile Russian-language cybercriminal forum, and offers the ability to clone Google accounts, steal browser information, and reportedly bypass 2FA.

Retail domains might be spoofed to steal accounts or financial information—which then can be used to commit fraud, drop malware, or sell access on to third parties. Domain impersonation can have a devastating impact. If used as part of a social engineering campaign (such as involving business email compromise or BEC), financial losses could be huge. There are also fake domains that are used to trick users into downloading malware, or harvesting credentials; this is the most common reason.

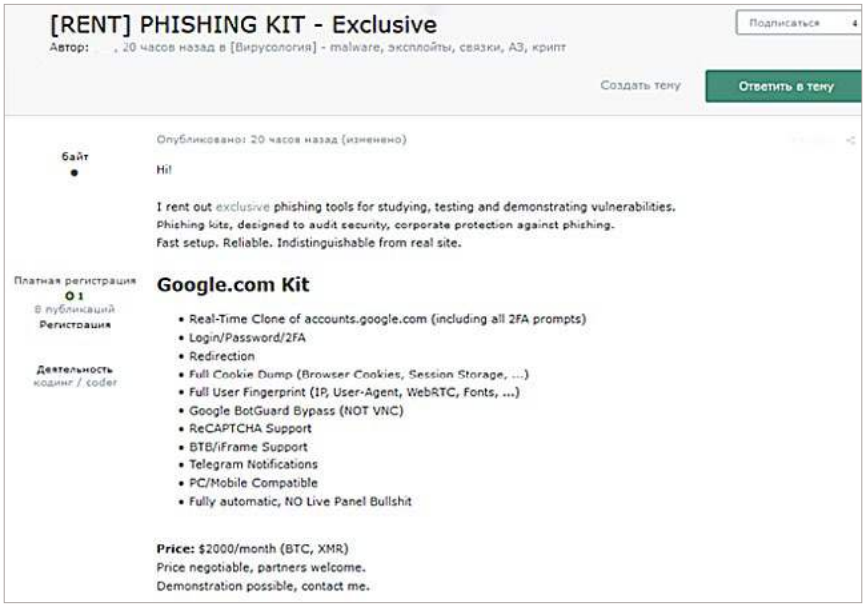

Figure 9: Forum advertisement of a phishing kit dedicated to spoofing Google

What Steps Should Defenders Take Now?

Based on the observations, we recommend taking certain steps to manage your external exposures.

- In order to minimize Defence Evasion and Impact techniques, organizations should ensure sufficient monitoring is in place to identify suspicious or abnormal activity. This activity could be unusual login attempts, changes to system configurations or permissions, or attempts to delete files.

- Credential exposure inevitably affects every company, to some extent. The risk can be managed best through a four-step program: Identify breached credentials, validate that the credentials are current, contain the credentials usage, and educate the credential's user. For more information, why not navigate to our blog on credential exposure?

- Exposed ports are a tricky business to manage; every organization is different, in terms of operational requirements and risk appetite. To manage the risk, first enumerate and document which services are publicly accessible from the internet. Perform a risk assessment; which services needs to be publicly accessible? (Close those that do not.)

- Limiting access to specific ports can also be an interim solution if they cannot be closed entirely. Restricting access to only the IP addresses or range used by your administrators and relevant systems (using the service on that port) will minimize the risk profile for the environment.

- Vulnerability management teams should prioritize vulnerabilities that have known exploits, if they are relevant to systems exposed to the network.

- For domain impersonation risk, visibility is key. Ensure that registered domains, branding information, IP address ranges, and other assets are all sufficiently monitored for suspicious activity. This will help ensure that any domains spoofing your brand—potentially through typo-squatting or use of an incorrect TLD—can be quickly identified and remediated based on the risk they pose. This risk can be tracked over time, enabling security teams to act quickly and proportionately to submit domain takedown requests or mitigate otherwise. Monitoring for domain impersonation can be achieved through ReliaQuest GMDRP.

Initial Access Broker (IAB) Trends

IABs give ransomware operators the tools to compromise a wealth of victims. A symbol of cybercrime professionalization, these brokers act as the "middlemen" in cyber threat operations: finding vulnerable organizations, exploiting them to access their systems, then selling that access to the highest bidder on dark-web forums. Their rise in popularity has aligned with the trend of ever-lower barriers to enter the world of cybercrime, which is also aided by the rise and spread of commodity malware and cybercriminal affiliate memberships.

Our monitoring of IABs dates back to 2014, when the sale of access to systems first began making ripples in the cybercriminal underground. IAB activity has been synonymous with the consistent threat posed by ransomware activity, and IABs often work directly with ransomware groups to identify susceptible networks for exploitation.

Our breakdown of IAB activity during the reporting period can be seen in Figure 10, taken from analysis of IAB "tipper" intelligence update reports provided to ReliaQuest customers. Each week, the Threat Research Team manually collects posts advertising initial access to compromised systems from select high-profile Russian-language forums, publishing the intelligence the threat actor provides about the victim as a tipper in GreyMatter Intel. We prioritize listings that contain the most intelligence about a victim (e.g., targeted sector or region) and that align with our priority intelligence requirements, client interests, and corporate strategy.

Most Common Access Types

The most common access type listed by IABs was RDP, which accounted for 24.4% of all tippers released during the reporting period. This was followed by VPN-RDP—representing VPN access to the RDP dedicated server of the victim—and VPN access. There was a significant difference among median prices of access types; the highest was $1,000, for RDP access. The greatest interquartile range was $2,700 for RDP access; the interquartile range refers to the middle 50% of the distribution. The average for VPN access was $500.

Listing Price (USD by Stated Access Type (p <0.05)

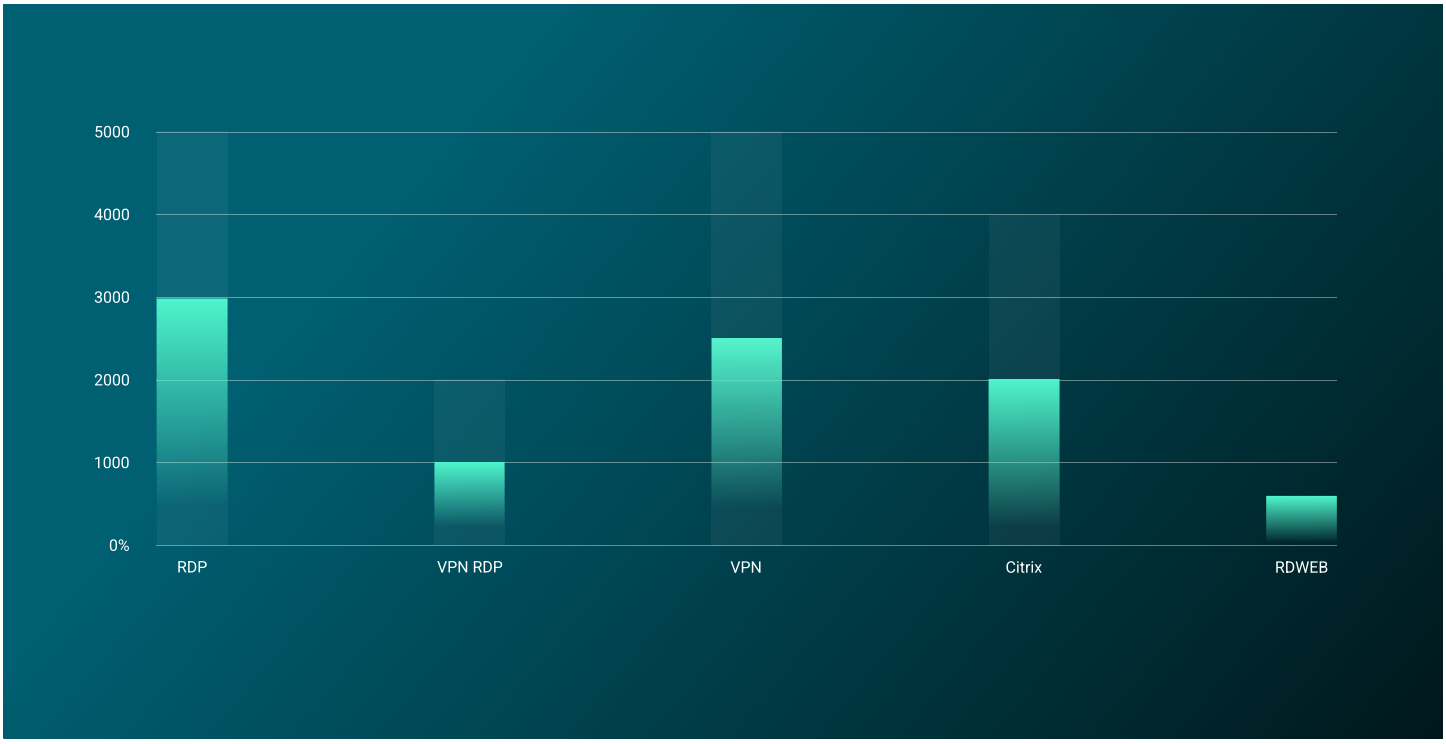

Figure 10: Most common IAB access types listed for sale, and price ranges from 01 Feb 2022 – 01 Feb 2023

RDP accesses being the type most sold and sought comes as no surprise. RDP applications remain favorites of IABs, given the relative ease with which they enable compromise through default or stolen passwords obtained via brute-force attacks. Despite significant research identifying unsecured RDP as one of the biggest threats facing business, organizations are failing to sufficiently secure their devices. Weak credentials often grant access, enabling threat actors to compromise these applications easily and drive more malicious activity.

Automated vending cart (AVC) websites, like Russian Market, have thrived in the ecosystem of stolen RDP credentials. Russian Market offers a one-stop shop for cybercriminals wishing to purchase RDP access and a range of other assets/services, including stolen card verification values (CVV), credit card dumps, account credentials extracted from malware stealer logs, and other illegally sourced items.

That AVC site offers easy-to-use search and filter functions to navigate the sections and tailor results to required needs. However, that functionality is only accessible to users who have deposited funds to their Russian Market account—an increasingly common tactic for AVC sites to encourage users to show intent to access and use the site. Russian Market is available on both the clear and dark web, and appears to be growing in popularity, although competing platforms are beginning to emerge.

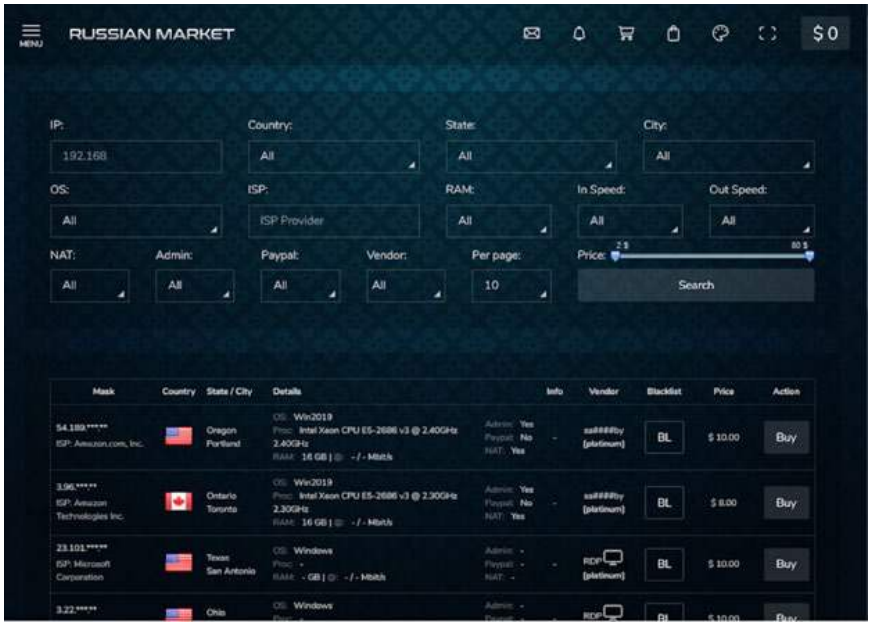

Figure 11: AVC site Russian Market sells compromised RDP credentials

Most Commonly Targeted Countries

The most commonly targeted country for IAB activity was overwhelmingly the US. This is unlikely to change in the long-term future (over one year), as cybercriminals frequently perceive US-based companies as offering large financial rewards, and spur other threat actors for political reasons. Some cybercriminals opportunistically target whatever entity can provide a profit, but for others, the US is still seen as the traditional enemy, and, for Russian-speaking cybercriminals, former Commonwealth of Independent States countries should not be targeted.

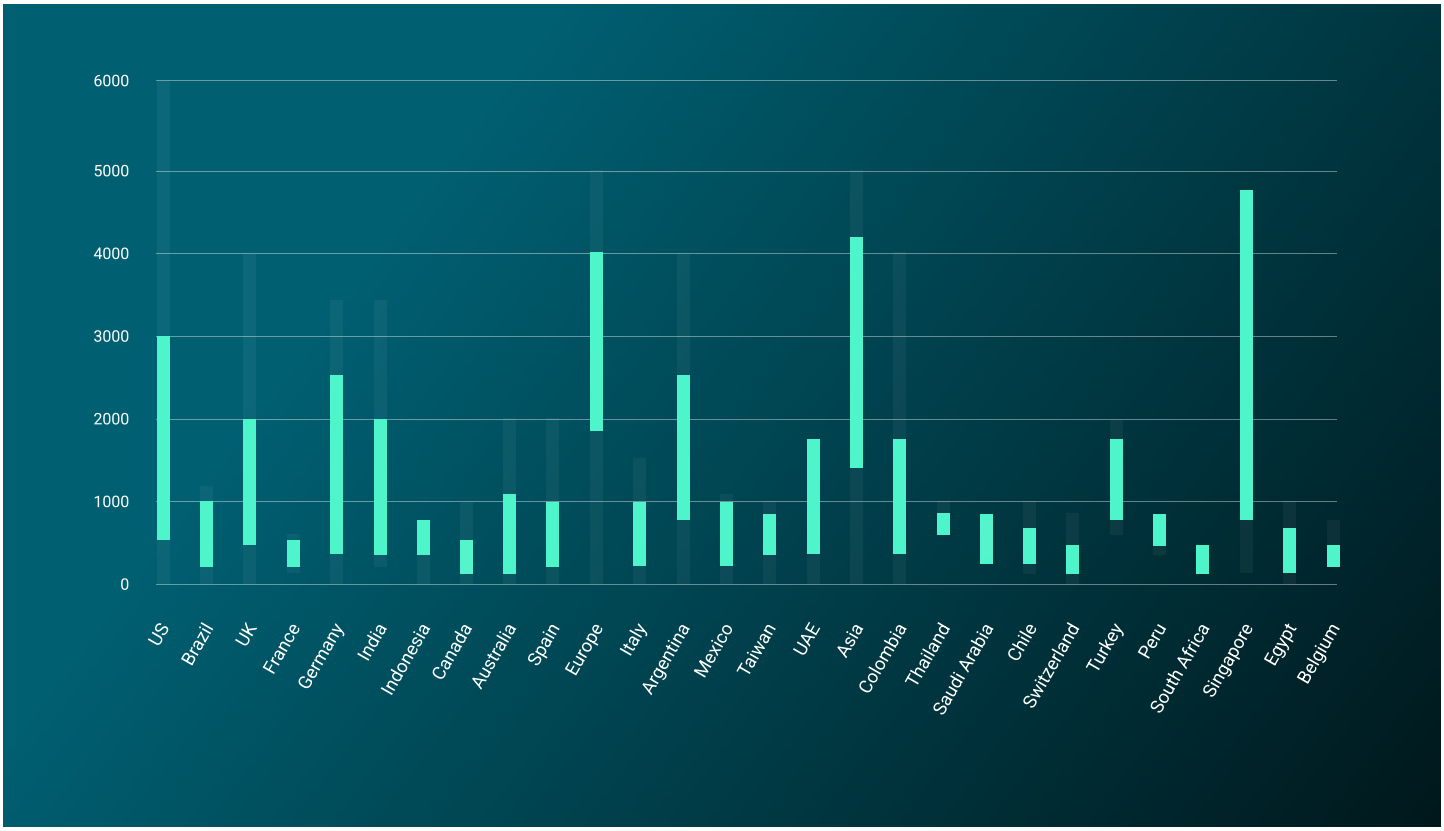

Figure 12: Countries most commonly targeted by IAB activity between 01 Feb 2022 – 01 Feb 2023

Cybercriminals frequently perceive **US-based companies as offering large financial rewards,** and spur other threat actors for political reasons.

Most Commonly Targeted Sectors

The most commonly targeted sectors list revealed interesting insights. Manufacturing took top place, prompting the publication of 142 tippers. Manufacturing is also the most commonly targeted sector for ransomware activity; this highlights the key role IABs play in identifying and supplying access to ransomware operators and other extortionists.

There was a significant sectoral difference among median prices of IAB access. The highest median price was $5,500, for a banking entity, and the greatest interquartile range was $23,000, also for a banking entity. Access to a banking entity is sold for the highest prices because it represents significant financial opportunity. But banks are likely to have substantial security budgets to avoid such risks—budgets likely not matched by other heavily targeted sectors. As a result, banking access is much less common—warranting only 14 tippers in 2022—and is more likely to generate interest from buyers; scarcity justifies the high price.

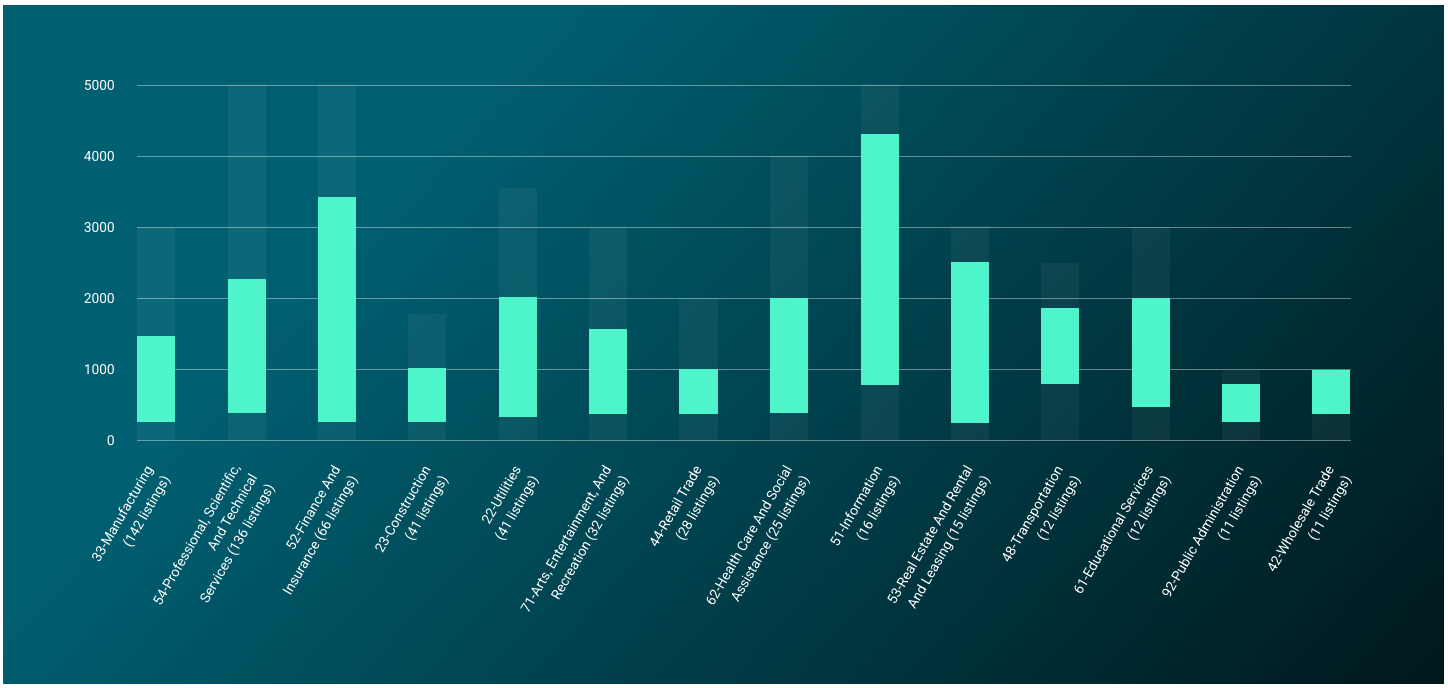

Figure 13: Sectors most commonly targeted by IAB activity

Case Study: Exotic Lily

The data behind IABs shows that all sectors have some sort of access being solicited on the dark web. How do these brokers get access to a company's environment? ReliaQuest uncovered a good example to hold up for examination. The "Exotic Lily" IAB began by sending elaborate phishing emails from what appeared to be a potential business prospect, to the inbox of a high-profile employee. This was accomplished by spoofing a legitimate domain, eaglemine[.]com, using a similar TLD: eaglemine[.]co.

Because the purported organization was a potential sales lead for the recipient, Exotic Lily succeeded in establishing communication. They next sent the recipient a ZIP file, named for_company.zip, from the legitimate file-transfer website WeTransfer[.]com, to bypass email security gateways. Unzipped, the file contained an IMG file with the name project_requirments.img; when mounted, it likely contained a LNK file that the user clicked. This resulted in a Python interpreter being loading into memory and executing a Python script, which downloaded a Cobalt Strike beacon. ReliaQuest attributed the malicious activity to Exotic Lily by continuously monitoring the known infrastructure of that IAB.

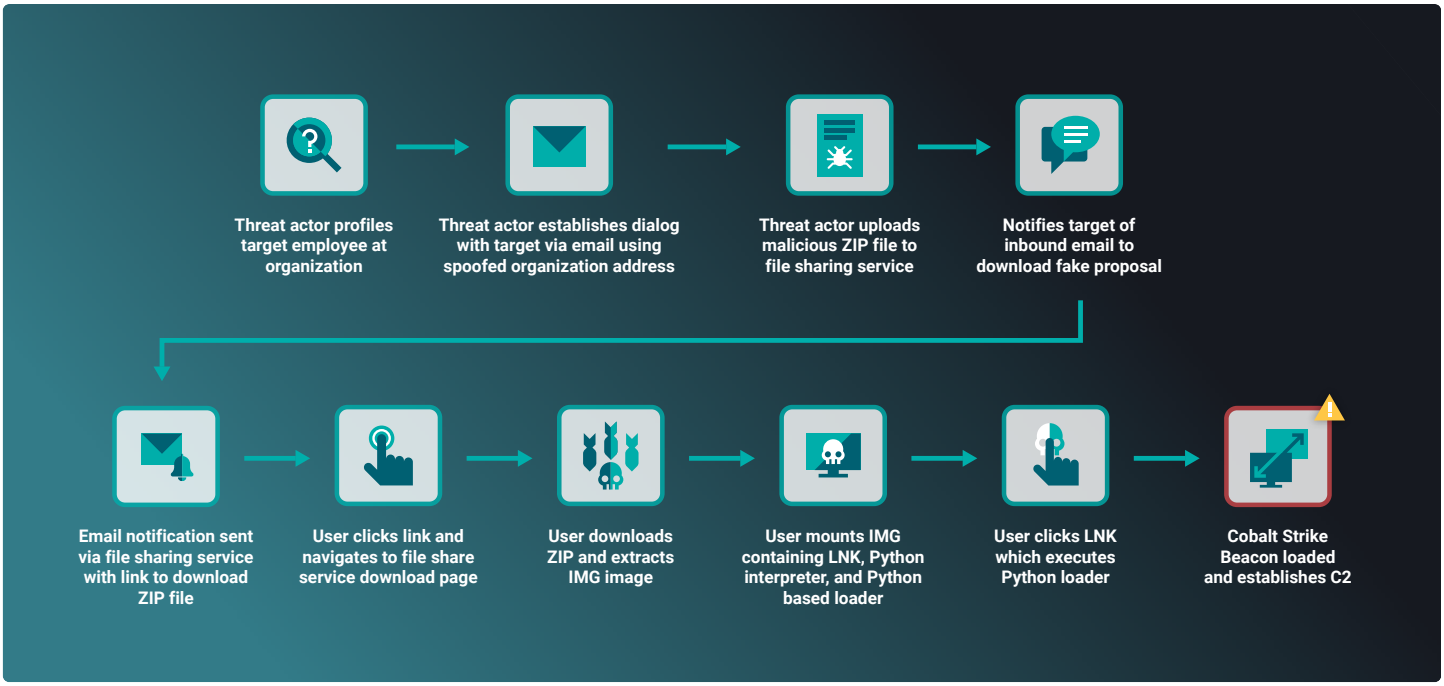

Figure 14: Typical IAB lifecycle, as observed with Exotic Lily

Security operations need internal- and external-event visibility. Securing an environment against IABs is accomplished through a robust security program, using defense in depth (DiD) strategies for internal and external locations. IABs like Exotic Lily are highly skilled at profiling users and developing phishing campaigns tailored to them. A robust security awareness program allows users to spot indicators of suspicious activity and act accordingly.

Initial Access Malware

In addition to identifying and stealing credentials for RDP, VPN, and other remote-access software, IABs are prolific at using malware to access a network.

QakBot

IAB-run "QakBot" spam campaigns in 2023 have frequently used malicious OneNote files to deliver malware that grants initial access to a system. Our research across the cybercriminal underground revealed users of multiple dark-web forums sharing information and articles about QakBot's operators. They said the operators used HTML smuggling techniques to deliver malware using SVG images embedded in HTML email attachments, toward the end of 2022.

HTML smuggling is an evasive malware-delivery mechanism, by which threat actors use legitimate HTML5 and JavaScript features to smuggle malware, remote-access trojans (RATs), or other payloads into targeted mailboxes. The technique is used by a wide variety of threat actors, including cybercriminals and nation-state actors, but it is not new. HTML smuggling has become more commonplace since Microsoft's decision to block macros in Office documents by default.

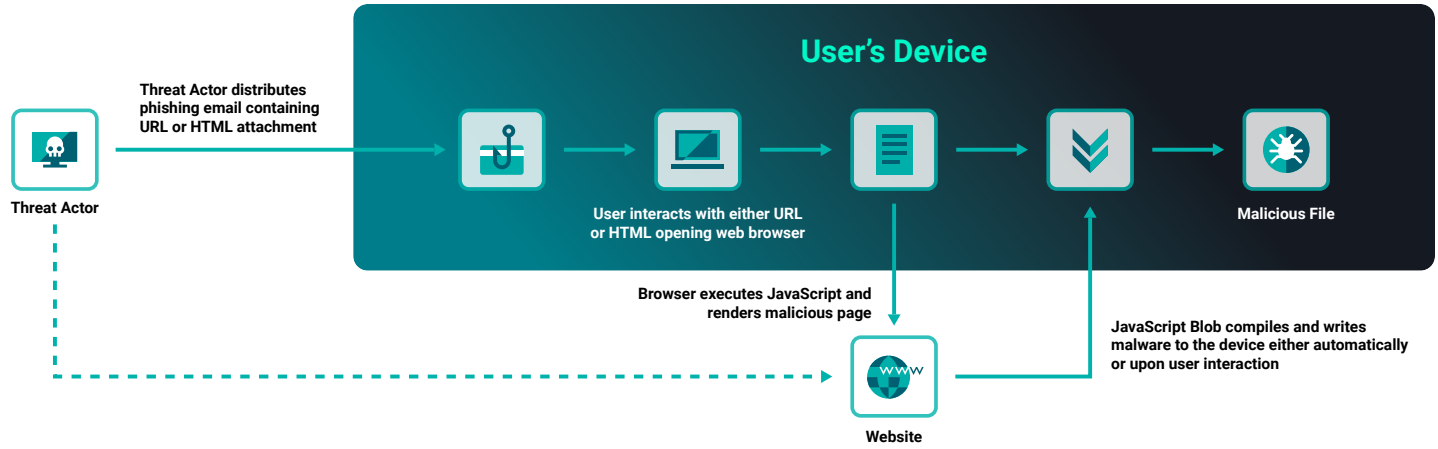

Figure 15: HTML Smuggling attack lifecycle

Emotet

The Emotet malware has had a checkered history. It was, in effect, disabled in January 2021 after a law-enforcement operation, only to be resurrected through the assistance of the operators of the "TrickBot" malware and operators of the now-defunct "Conti" ransomware. Recently Emotet has been distributed through malicious Word files containing macros that, if enabled, start the infection chain and execute the file Emotet.dll—a significant approach, given Microsoft's recent changes to control macro enablement.

GootLoader

Also popular for initial access is the stealthy GootLoader malware, which was classified as a first-stage downloader designed to attack Windows-based systems. GootLoader's earliest second-stage payload, and the source of its name, is GootKit: a banking trojan and stealer in use since 2014. It includes JavaScript and C++ modules to execute remote commands, man-in-the-browser attacks, keystroke exfiltration, screenshots, and credential theft. As with several other initial access trojans, GootLoader started with relatively limited capabilities before undergoing significant development. It is commonly used by IABs and ransomware-as-a-service (RaaS) affiliates.

IcedID

IcedID is just one malware type that resurged as Emotet struggled to overcome law-enforcement activity. It began as a banking trojan, then developed into a malware dropper employed on compromised systems. The new version of IcedID is free of several seemingly unnecessarily features related to banking fraud. This version of the malware loader first appeared in February 2023, distributed directly through thousands of personalized, invoice-themed phishing emails. These messages used Microsoft OneNote attachments (.one) to execute a malicious HTA file that, in turn, runs a PowerShell command to fetch IcedID from a remote resource.

SocGholish Malware Distribution Framework

A trend first observed in 2022 and carrying on in recent months is the use of the SocGholish (aka FakeUpdates) malware distribution framework. This common initial access method has received substantial ReliaQuest attention. SocGholish employs social engineering and drive-by compromise to drop malware on endpoints. It deceives individuals into downloading a fake web-browser update (as seen in Figure 16), which contains an archive file with an embedded SocGholish JavaScript payload.

Once executed, the JavaScript payload establishes a C2 channel to relay system information it has gathered from the compromised endpoint. If the host is found to have been "domain joined" (a method companies use to manage Active Directory users), additional discovery commands are provided and executed to collect more details. If the endpoint and its host environment pique the interest of the threat actor operating the campaign, Cobalt Strike or similar frameworks are typically deployed for post-exploitation objectives.

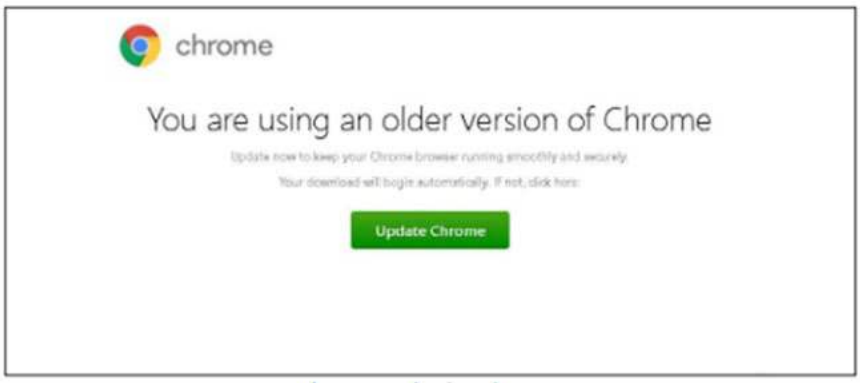

Figure 16: SocGholish fake update message and linked button

Think of SocGholish as, primarily, a preliminary foothold for additional cybercrime groups to follow up after initial access is established. In recent months we have identified and responded to two discrete "hands-on-keyboard" intrusions traced back to a SocGholish compromise. With these two intrusions, we found overlapping artifacts suggesting that the compromises were performed by the same threat actor. During our investigations, network telemetry was found belonging to "Evil Corp" infrastructure, potentially indicating that threat group's involvement. ReliaQuest contained both intrusions by preventing what looked like the threat actor's primary objective: deploying ransomware.

What Steps Can Defenders Take Now?

- IABs thrive on an abundance of susceptible remote services that are externally exposed and have insufficient authentication processes. As this most commonly affects RDP, the most important step to take is to ensure RDP services are not unnecessarily exposed to the internet, and make credentials strong enough to withstand brute forcing and other credential-cracking attacks. Any accounts known to frequently use RDP should also be placed under focused monitoring to identify any possible suspicious activity.

- Use the most secure encryption and authentication methods across remote services, which will depend on network infrastructure and VPN devices in use. Authentication should also include the use of 2FA.

- IABs are known to targeted unpatched VPN devices to solicit initial access. Have administrators prioritize updating VPN devices with the latest patches.

- Consider a Zero Trust security model, which prohibits user access to data by default, and requires users to be consistently authenticated and verified. Although threat actors cannot monitor VPN-encrypted traffic from outside the VPN, if they are able to connect to the VPN, they gain access to any resources connected to that network. It only takes one compromised account or device for an attacker to gain access to VPN-gated data.

Vulnerability Intelligence

Exploiting common security vulnerabilities remains one of the most readily used methods for trying to access a network. This includes IABs, who as identified in our section above, commonly target susceptible vulnerabilities on external facing systems to gain initial access. If you want to vastly improve your cyber resilience to a huge range of malicious actors, improving your vulnerability remediation process is a great place to start.

Given that there were over 24,000 CVEs[2] in 2022, knowing which to prioritize is challenging. Making vulnerability management even more complex is the simple fact that threat actors are not concerned if a vulnerability is brand new or years old. Vulnerability intelligence sits at the intersection of vulnerability management and threat intelligence. Vulnerability management is an ongoing process of identifying, investigating, assessing, reporting, and patching vulnerabilities, and vulnerability intelligence feeds actionable insights into vulnerability management. Intel is only sometimes part of a vulnerability management program, but it provides vital context for to understand how likely it is a given vulnerability will be exploited.

Although some vulnerability management tools include elements of vulnerability intelligence, the security industry is still failing to provide sufficient details about how these flaws are discussed and exploited. Without being able to fully survey your threat landscape, threat actors can continue exploiting vulnerabilities before your security team has had sufficient time to notice and act.

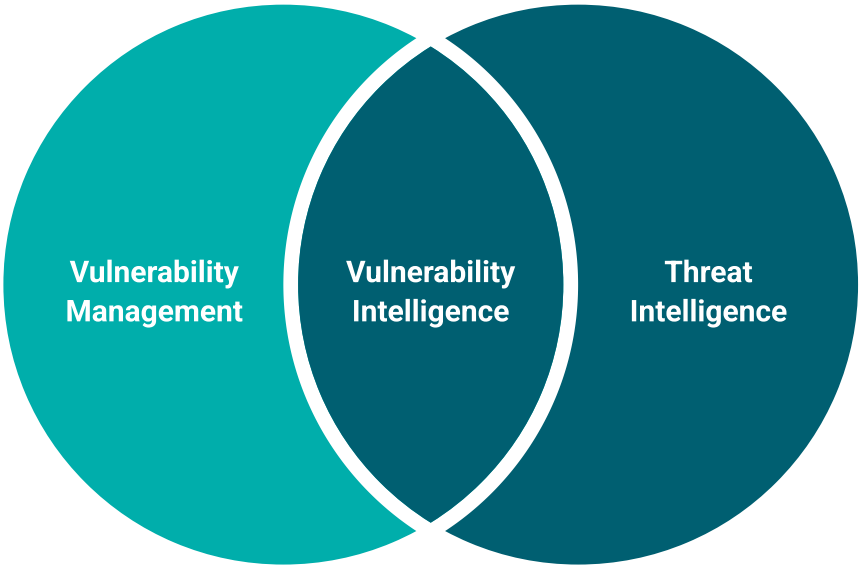

Figure 17: Relationships of vulnerability intelligence

Many companies' efforts to triage and remediate vulnerabilities are, at best, insufficient. Most vulnerability-management efforts rely on the Common Vulnerability Scoring System (CVSS) to prioritize patching. Like any numerical representation of risk, it is helpful at a glance, but does not give a complete picture of the risk of exploitation and the potential impact each vulnerability could have on your organization.

Delays in detection and reporting often mean that CVSS scores are not as timely as we would like. CVSS is maintained by the National Vulnerability Database, which can be slow to analyze and triage vulnerabilities. What's more, the scores lack any indication of the likelihood of exploitation. To effectively prioritize, teams must often scour many sources for clues. And CVSS scores are not dynamic; although the threat landscape changes regularly, the scores often do not reflect these changes. Vulnerability intelligence is the missing piece of the picture, enabling security teams to go beyond CVSS scores.

[2] Common Vulnerabilities and Exposures: Publicly disclosed cybersecurity vulnerabilities and exposures, listed in a dictionary with an identification number, a description, and at least one public reference

Where the Risks Lie

Security teams require deeper insights than mere CVSS scores or media hype, which fail to convey the exploitability of a particular vulnerability and, in turn, fail to recognize a true risk. Understanding the true risk of exploitation can only be achieved through continuous monitoring and a consistent, methodical process that takes into consideration several risk factors. Is there any evidence of exploitation in the wild? Is there a reliable and ready-to-use proof of concept (PoC) that would enable a relatively unskilled attacker to exploit? What kinds of actors are interested in the vulnerability? Has there been a link to nation-state activity, or chatter about the bug on cybercriminals' platforms? Of course, a determination of risk is very much specific to your company and business composition, but these are the factors that should be primary considerations when establishing vulnerability risk.

We answered the same questions to identify the vulnerabilities that represented the greatest risk to organizations during the reporting period. We cross-referenced those flaws with the ten technologies most commonly used by ReliaQuest customers; see Figure 18.

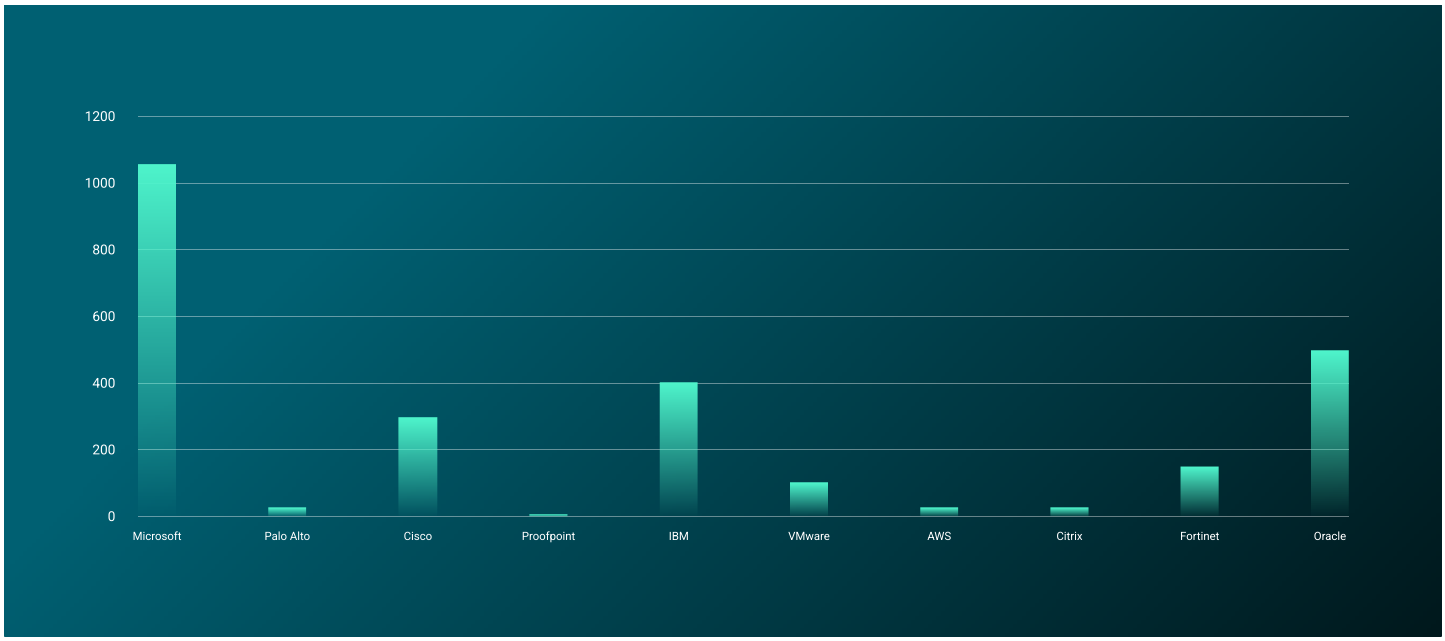

Figure 18: Number of CVEs found in top technologies used by ReliaQuest

As you would expect, the most commonly used technologies were overwhelmingly those of Microsoft, followed by Oracle, IBM, and Cisco. Next, we identified those vulnerabilities announced in the reporting period that represent the greatest risk.

Risk Calculation

To enhance the risk assessment of CVEs, we have developed an automated scoring tool that effectively integrates intrinsic CVE data and the insights our Vulnerability Intelligence team gathers and analyzes continuously. Our scoring methodology considers four critical factors for each vulnerability: the availability of reliable exploits, its technical impact, its automation potential, and its possible business impact. By emphasizing these elements, our approach addressed the limitations of traditional scoring systems, which may not accurately portray a vulnerability's fluid and dynamic risk.

We applied our methodology to all CVEs reported in 2022 for the top ten vendors identified, pinpointing the top five critical vulnerabilities for each. The vulnerabilities identified on the next page in blue show the CVEs referenced by the United States Cybersecurity and Infrastructure Security Agency (CISA) as having been exploited in the wild.

Microsoft	CVE-2022-26809
	CVE-2022-41040
	CVE-2022-41082
	CVE-2022-34721
	CVE-2022-41080
Palo Alto	CVE-2022-0024
	CVE-2022-0031
	CVE-2022-0028
	CVE-2022-0025
	CVE-2022-0026
Cisco	CVE-2022-20699
	CVE-2022-20708
	CVE-2022-20700
	CVE-2022-20923
	CVE-2022-20705
Proofpoint	CVE-2022-46332
	CVE-2022-46333
	CVE-2022-46334
	CVE-2022-25294
	CVE-2021-31608
IBM	CVE-2022-22425
	CVE-2022-40752
	CVE-2021-38869
	CVE-2021-38945
	CVE-2021-38969

VMWare	CVE-2022-22965
	CVE-2022-22954
	CVE-2022-22963
	CVE-2022-22947
	CVE-2022-22955
AWS	CVE-2022-29972
	CVE-2022-25809
	CVE-2022-41828
	CVE-2022-25165
	CVE-2022-23511
Citrix	CVE-2022-27518
	CVE-2022-27510
	CVE-2022-27513
	CVE-2022-27516
	CVE-2022-27511
Fortinet	CVE-2022-40684
	CVE-2022-42475
	CVE-2021-44168
	CVE-2022-33872
	CVE-2022-33874
Oracle	CVE-2022-21587
	CVE-2022-22963
	CVE-2022-22947
	CVE-2021-35587
	CVE-2022-31813

Table 2: Highest risk CVEs found in products most used by ReliaQuest customers, by technology vendor

Vulnerabilities Deep Dive

Threat actors were busy over the reporting period, exploiting new and sometimes old vulnerabilities in assets that were left unpatched and exposed to the public internet. ReliaQuest has responded to multiple incidents where initial access was gained by exploiting vulnerabilities. Read on for a few examples.

Spring4Shell RCE Vulnerability

CVE-2022-22965 is a critical remote code execution (RCE) flaw discovered in the widely used Spring Core framework for Java in March 2022. The vulnerability was disclosed as a bypass of the patch for CVE-2010-1622, enabling attackers to target the Spring Web MVC (Model-View-Controller), a component of the Spring Framework used for developing web applications. Exploits available in the wild typically involved forcing the application to write a malicious .jsp file to the web server, which could be executed to gain remote command execution.

Although patched versions were released quickly and the conditions for exploiting the vulnerability were limited, the widespread use of the Spring Framework made this a notable concern for the entire infosec community. During the crisis, the Vulnerability Intelligence team provided valuable insights and monitoring, which helped detect and monitor several evidence of exploitation attempts in the wild. Spring4Shell was identified as having the highest risk score when calculating our scoring methodology.

Log4Shell Vulnerability

In 2022 ReliaQuest responded to the infamous Log4Shell vulnerability (CVE-2021-44228), found in the Apache Log4J 2 Java library. The flaw left the vast majority of public-facing web applications easily exploitable. In just three days, starting with a successful exploit of Log4Shell, the attackers gained a significant foothold in our customer's environment. This was an extreme example of many incidents ReliaQuest investigated involving Log4Shell, but should serve as a reminder that attackers are opportunistic. There has been steady chatter on dark-web forums throughout 2023 of Log4Shell being repeatedly exploited.

Oracle EBS Vulnerability

The Oracle EBS vulnerability (CVE-2022-21587) prompted ReliaQuest to remain on high alert across our customer base. The flaw allowed an attacker to upload arbitrary files on devices, which resulted in significant interest from threat actors. On Russian-language cybercriminal forums, we observed multiple posts from users discussing the vulnerability and sharing PoCs (including one shown in Figure 19).

Figure 19: PoC for CVE-2022-21587 shared on cybercriminal forum

ReliaQuest quickly created emergency detection rules to inform customers when the flaw was being exploited before Oracle released a patch; in this way we could quickly enable customers' security teams to make informed decisions on threats pertinent to their environments.

Fortinet Authentication Bypass Vulnerability

CVE-2022-40684 was one of the more notable vulnerabilities disclosed by Fortinet in 2022. Urgent mitigation guidelines were published, as adversaries were exploiting the flaw to bypass authentication and log on to the vulnerable systems as an administrator. What followed was a unique series of events in which exploited assets were updated to the recommended firmware and left with an account named fortigate-tech-support. Due to the sensitive nature of this activity, ReliaQuest strongly recommended that customers reload the updated firmware and restore the configuration backup. At the time of this writing there is no intelligence behind the exploitation, account creation, or update to the vulnerable assets.

What Steps Can Defenders Take Now?

Several best practices can help you guard against vulnerability exploitation attempts.

- To accurately identify your current risk, you first need to understand what you own and who is responsible for its upkeep. A lack of asset visibility remains one of the biggest problems in business in 2023. Maintain an up-to-date asset inventory or configuration management database and regularly compare it to the scope of their vulnerability management program. Any differences between the asset inventory and scanning scope should be addressed quickly to reduce or remove visibility gaps.

- The Center for Internet Security recommend running vulnerability scans at least once every two weeks. While this frequency is likely greater than typically run at some organizations, it will reduce the likelihood of newer, high risk, vulnerabilities being missed, as would be the case on monthly or quarterly scans.

- When high risk vulnerabilities are reported as being exploited in the wild, use this information to push critical patching through to your exposed assets. Prioritize patching vulnerabilities identified by the US CISA as being exploited in the wild.

- Ensure the collection of metrics accurately displays the lifecycle of risk reduction. A well-designed vulnerability management program can help an organization visualize how security risks are being addressed and paint a vivid picture of progress over time. Presenting vulnerability metrics to the board and senior leadership will demonstrate continual improvement and ROI on vulnerability management efforts or highlight the need for additional investment.

- Vulnerability management platforms discover known vulnerabilities and potential exploits, while breach and attack simulation capabilities highlight configuration weaknesses, detection and prevention gaps, and architectural issues. Organizations should ensure that an effective response and recovery plan is properly evaluated through tabletop exercises and is tested periodically and adjusted as the threat landscape, people, systems and business processes change. By combining threat and vulnerability management, organizations can increase their security confidence and decrease their overall risk.

Ransomware Intelligence

So far in our report we've identified the risk posed by IAB activity, who target security flaws—including exploiting vulnerabilities—to facilitate access to susceptible networks. Who do IABs sell these accesses to? Who are the customers of these facilitators? Well one of the biggest customers are ransomware operators, who can quickly use an IAB listing to compromise a network to extort a company for huge sums of money.

Ransomware poses the biggest cyber threat to organizations globally. Ransomware operators aim to compromise their targets' machines, encrypting all files of value to cause significant operational disruptions. As well as encrypting data, since 2020 ransomware groups have been exfiltrating sensitive data and threatening to disclose it if a ransom is not paid: a tactic commonly known as double extortion, that can end with the stolen data posted on dedicated data-leak sites hosted on the dark web. In light of that threat, the risk attached to ransomware attacks can no longer be entirely mitigated by refusing to pay or having secure data backups.

Ransomware Attack Kill Chain

Ransomware attackers are agile, resourceful, and adaptable. Their techniques change, but the process typically follows a consistent path, as shown in Figure 20. It often starts with reconnaissance and/or the discovery of susceptible networks. The attacker can employ numerous techniques, such as using network scanning tools to identify network shares and other network information. As identified in the previous section, they also might use the services of IABs to conduct the necessary work of finding susceptible targets.

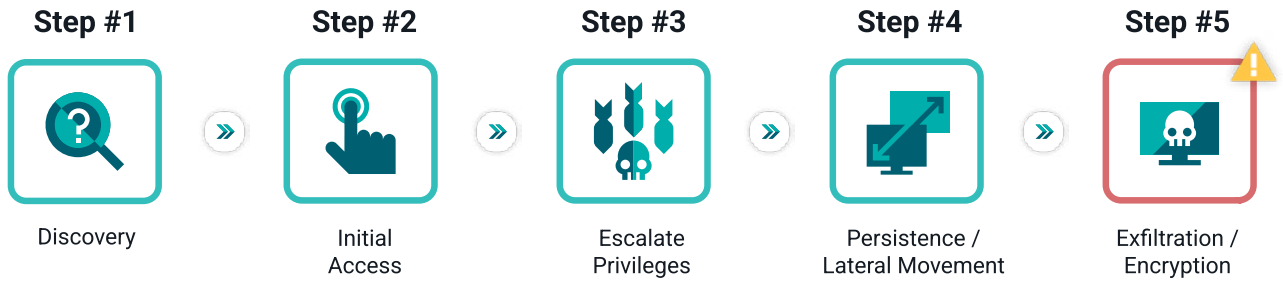

Figure 20: Five typical steps of a ransomware attack

Once a susceptible target has been identified, the attacker can attempt initial access, via phishing emails, exploitation of external-facing vulnerabilities, or, as mentioned, abusing remote services software. Whatever the method, the attacker's goal is to gain a foothold in a target's network, without alerting network defenders.

At this point the ransomware attacker can shift focus to escalating their privileges or creating new domain or administrator accounts via Active Directory to move laterally in the system. The objective would be to gain access to sensitive data or other target systems.

The next stage is to maintain persistence on the network, for example by installing remote-access software or abusing a run-key to execute a C2 beacon. The attacker will also move laterally, if needed to achieve their objective, or to access as many machines as possible before the encryption stage.

Most Targeted Sectors

Figure 21 highlights the total number of victims named on ransomware data-leak sites each month. This data is cataloged and shared with ReliaQuest customers, to identify trends and provide warnings of activity targeting certain sectors and countries.

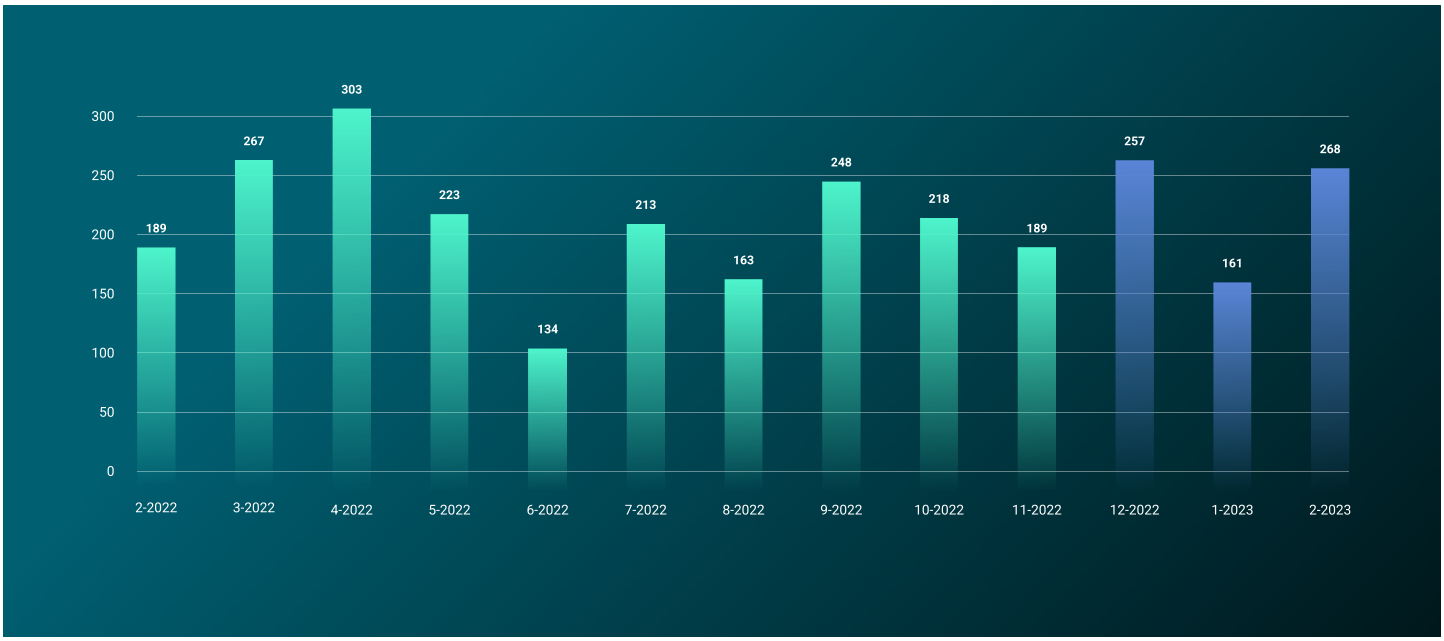

Figure 21: Number of victims named on ransomware data-leak sites, by month

Ransomware activity peaked in April 2022, hit a sharp decline in May, then moved to an average of approximately 214 site posts per month for the second half of 2022. There is no overall observable trend across the reporting period, and although the posting rate fluctuated, ransomware activity remained consistent.

The most commonly targeted sector was, overwhelmingly, industrial goods and services/ manufacturing. One reason to target manufacturing is its inherent susceptibility to outages. As with the construction sector mentioned in an earlier section, manufacturing depends on consistent IT processes. If production is halted, manufacturing ceases to function, and will suffer significantly more than other sectors if there are extended periods of downtime.

Ransomware activity peaked in April 2022, hit a sharp decline in May, then moved to an average of approximately 214 site posts per month for the second half of 2022.

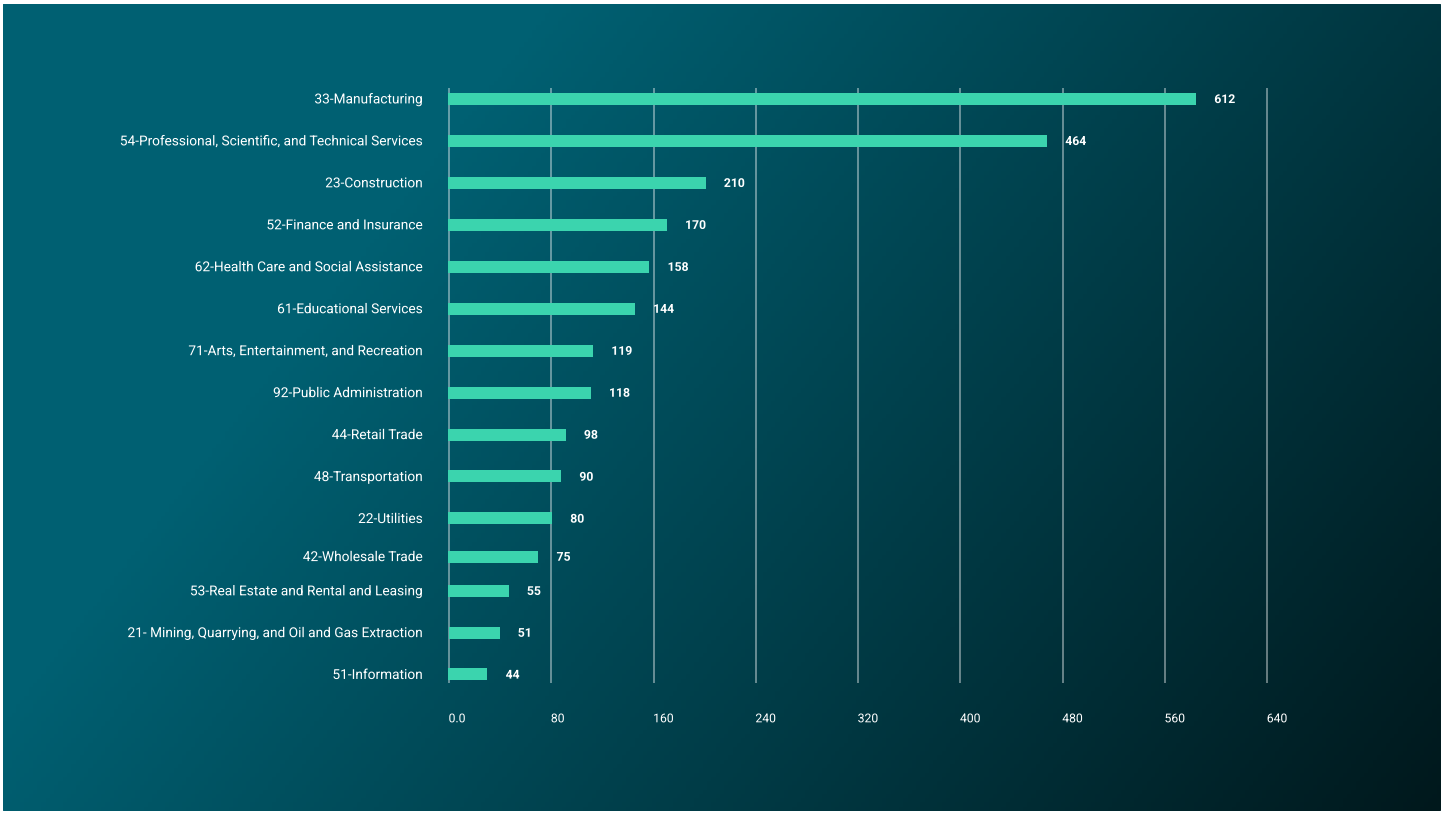

Figure 22: Number of ransomware attacks in the reporting period, by sector

Most Active Ransomware Groups

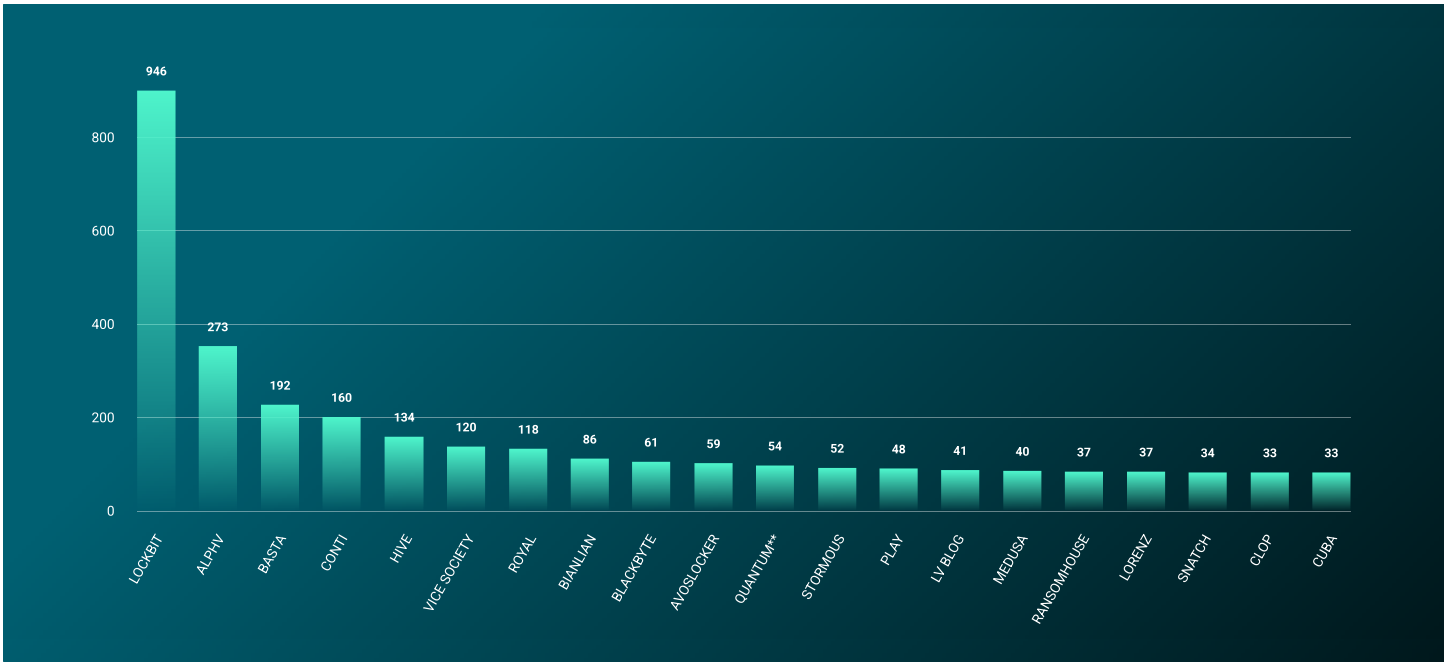

Figure 23: Most active ransomware groups as indicated by number of attacks

The most active group, by far, was LockBit, with 946 posts naming victims on the group's data-leak site. It is common to have such an outlier—similar positions were previously occupied by Conti and "REvil." The demise of the latter two groups was brought about by internal conflict, such as the compromise of Conti's infrastructure and other divisions, or as a result of law enforcers' scrutiny. For LockBit, which has been the most active group since early 2021, the ongoing Russia-Ukraine war may be taking law enforcers' focus away from cybercrime. Many Western law-enforcement agencies are likely more preoccupied with stopping Russian nation-state–sponsored cyber-threat activity.

The war has also had a crippling effect on US-Russian relations. The US, along with several other Western countries, has been providing significant monetary, military, and geopolitical support to Ukraine. Any collaboration between law enforcers in the US and Russia—such as was believed to be instrumental in the January 2022 arrests of several members of REvil—has also likely ceased. With many of LockBit's members likely based in Russia or other CIS countries, opportunities to target the members and infrastructure is probably limited.

As a result, LockBit is likely to have free reign to continue targeting West-based companies; the Russian government would probably view any activity that damages US or European interests as favorable and worthy of ignoring. One of our key predictions at the start of 2023 was that LockBit would continue to lead the way in ransomware activity, and so far that has proven true.

Case Study: LockBit

ReliaQuest investigated an incident involving the LockBit ransomware and its eponymous operating group. Initial access was found to be the result of a SocGholish infection. Following this, Cobalt Strike was loaded on to the host and C2 established. The threat actor began to move laterally in the environment via a combination of Cobalt Strike and RDP. After a few days of movement in the network, they had obtained credentials for a service account with domain administrator permissions.

Then began a two-month-plus period of inactivity, for no clear reason. One theory relates to the timing of the intrusion, which started within a day of the Russian invasion of Ukraine. Geopolitical tension in the region was high following the invasion and included groups and individuals within the cybercrime community. This tension may have played a factor in the long dormancy. However, there was an increase in organizations named on LockBit's data-leak site during those two months, so the war may not have had much impact on operations at the time.

After the dormant period, LockBit returned and resumed its operation by continuing to cement a foothold in the environment. The group moved laterally to additional high-value servers via RDP and compromised additional administrator-level accounts.

Next, LockBit began staging the encryptor file and a copy of psexec on a network. A new Group Policy Object (GPO) was created to launch and execute a BAT file via a scheduled task. The BAT file attempted to halt specific process and services, such as antivirus or endpoint detection and response (EDR), as well as stop the backup service and delete Shadow Volume Copies. It copied the encryptor and psexec from the network share, then used psexec to execute the encryptor. Finally, it was set to clear all tokens from error logs using wevutil.

One unique technique is LockBit's compromise of an account with administrator-level privileges in the organization's EDR console, and use of it to deregister EDR sensors on all hosts in the environment. With defenses fully disabled, a GPO update was pushed, setting off encryption throughout the environment.

What Steps Can Defenders Take Now?

Network Recommendations

- Segment networks: Ensure proper network segmentation of devices so that they can only communicate with other devices needed to support their specific business functions.

- Monitor external-facing assets: For accidental exposure and out-of-date services, remove any accidental exposure and patch any out-of-date services, prioritizing services that have known vulnerabilities. Threat actors frequently scan the internet for public-facing assets that have an exploitable vulnerability and gain initial access that way.

Internal System Recommendations

- Use application control: Where appropriate and, if possible, only permit the execution of signed scripts. Consider redirecting the default application for JavaScript, Visual Basic, and other executable script formats to open in notepad.exe instead of wscript.exe by default. The use of weaponized script files is used heavily by initial access malware.

- Comprehensive coverage: Ensure coverage is enabled for antivirus/EDR tools within your environment to provide as much visibility as possible to exploit or threat activity. Valuable detection use cases require endpoint logging or visibility.

- Use automatic updates: Apply a software update feature to your computer, mobile device, and other connected devices wherever possible and pragmatic.

Account Recommendations

- Inventory accounts: Service and other privileged accounts in the environment should be accounted for. Ensure that they follow the principle of least privilege (PoLP) and are configured with long, complex passwords. Service accounts are highly targeted in ransomware intrusions, given that they are often configured improperly with domain administrator rights.

- Use standard user accounts: Internal systems should only use standard user accounts, instead of administrative accounts that grant overarching administrative system privileges and do not ensure PoLP.

Cobalt Strike Ransomware Intelligence

So, an IAB has identified your network prone for targeting. Exploited an unpatched system that shouldn't be externally facing, before selling the access onto one of the dozens of ransomware groups looking for their next victim. Our new unauthorized friends are busy hopping around the network looking for methods of maintaining their persistence, while also identifying which systems might be candidates for encryption, or hold data that can be stolen. What is the method used by most groups to coordinate their activity? Enter Cobalt Strike.

The C2 framework that is overwhelmingly used by ransomware groups is that of the legitimate penetration testing tool Cobalt Strike. Its popularity is likely owing to a combination of effectiveness and user-friendliness. Attackers often rely heavily on C2 communications to start and progress attacks, including human-operated ransomware attacks. C2 infrastructure enables attackers to control infected devices, perform malicious activity, and quickly adapt to a targeted organization's environment in pursuit of valuable data and assets.

Breaking this link to C2 infrastructure disrupts attacks—either by stopping the communication completely or delaying its progression, allowing more time for a SOC to investigate and mitigate the intrusion. By default, Cobalt Strike enables payload staging via a valid checksum8 request; a checksum refers to a process of checking a file's integrity. The Cobalt Strike team server will then return a shellcode payload, from which security researchers can extract the payload's configuration. The configuration contains details of how the implant operates, including the C2 address, the C2 port, the spawn to process, and the license ID.

The default configurations of team servers have been well documented by members of the intel community. By searching for unique values in the HTTP response headers, JARM signatures, and default certificates, through network scan data services like Shodan, ReliaQuest can profile potential Cobalt Strike team servers.

With a compiled list of potential team servers, scans can be made in an attempt to retrieve a payload. If a payload is returned, an IP address/domain can be confirmed as a team server, which can then be added to our threat feeds for alerting. These domains/IP addresses can be a high-fidelity indicator of a malicious actor in the environment. Given the ease of collection, these indicators are a great supplement to other behavior-based detections.

What Our Data Tells Us

Table 3 shows the top countries in which Cobalt Strike team servers were hosted; Table 4 shows the top autonomous system numbers (ASNs) used to host these team servers. China hosted the vast majority of the identified Cobalt Strike team servers, followed by the US and Hong Kong. This is unsurprising, given that the abuse of Cobalt Strike is not limited to cybercriminals; it is also used heavily by Chinese nation-state–aligned groups.

Country	Servers Hosted
China	4,830
US	3176
Hong Kong	781
Russia	325
Singapore	176
Lithuania	175
Romania	150
United Kingdom	128
Netherlands	122
Germany	114

Table 3: Number of Cobalt Strike team servers, by host country

China is also, unsurprisingly, widely represented in the ASNs hosting Cobalt Strike team servers, harboring the large majority of the top ten. An autonomous system is a large network or group of networks that have a single routing policy. Each autonomous system is assigned a unique ASN, which is a number that identifies the autonomous system. These are typically owned and operated by a single service provider.

ASN	Count
45090 - TENCENT-NET-AP Shenzhen Tencent Computer Systems Company Limited, CN	2,695
37963 - ALIBABA-CN-NET Hangzhou Alibaba Advertising Co.,Ltd., CN	1,136
14061 - DIGITALOCEAN-ASN, US	674
20473 - AS-CHOOPA, US	390
16509 - AMAZON-02, US	373
8075 - MICROSOFT-CORP-MSN-AS-BLOCK, US	203
14618 - AMAZON-AES, US	190
55990 - HWCSNET Huawei Cloud Service data center, CN	181
132203 - TENCENT-NET-AP-CN Tencent Building, Kejizhongyi Avenue, CN	177
134548 - DXTL-HK DXTL Tseung Kwan O Service, HK	164

Table 4: Number of Cobalt Strike team servers, by ASN

The most common C2 ports can be seen in Table 5; the default ports for HTTP and HTTPS (80 and 443) are the most commonly used for communication.

C2 port	Count
443	4,892
80	3,829
8080	675
8443	541
8090	271
8888	247
8081	183
9999	177
4444	172
8088	167

Table 5: Most commonly used C2 ports during reporting period, by number of instances

Of the Beacon payloads recorded during the time period, most were configured with an IP address for C2. The C2 address is often the same as the team server address. For beacons that used a domain for C2, a majority used content delivery networks, such as those of Tencent, CloudFront, and Azure. The use of these services helps beaconing blend in with legitimate traffic. Table 6 highlights the top registrars used for C2 domains.

NameCheap was the most common registrar of Cobalt Strike team servers, followed by Ename Technology and MarkMonitor. These registrars—which can be seen represented by the stars right of their name—are primarily content delivery networks (CDNs) used for domain fronting. Domain fronting is used to conceal user traffic and is commonly used by threat actors for C2 purposes.

Registrar	Count
NAMECHEAP, INC.	353
ENAME TECHNOLOGY CO.,LTD. *	330
MARKMONITOR, INC. *	295
GODADDY.COM, LLC	247
OWNREGISTRAR, INC.	171
NICENIC INTERNATIONAL GROUP CO., LIMITED	122
GANDI SAS	103
NAMESILO, LLC	99
HOSTING CONCEPTS B.V. D/B/A REGISTRAR.EU	95
AMAZON REGISTRAR, INC.	86

Table 6: Most commonly used registrars during reporting period, by number of uses

Our data identified the most commonly used "spawn to" processes: temporary processes spawned by the Cobalt Strike implant, which are used to inject code that carries out post-exploitation commands. Each beacon configuration will list a spawn to process for x86 and x64 architecture. However, the process selected is typically the same for both. The default spawn to process is rundll32.exe. This is a great detection opportunity, as many of the top spawn to processes are rarely executed without command line arguments.

Spawn to x64	Count
%windir%\sysnative\rundll32.exe	8,087
%windir%\sysnative\dllhost.exe	1,342
%windir%\sysnative\gpupdate.exe	232
%windir%\sysnative\svchost.exe	195
%windir%\sysnative\WUAUCLT.exe	185
%windir%\sysnative\runonce.exe	184
%windir%\sysnative\regsvr32.exe	149
%windir%\sysnative\WerFault.exe	105
%windir%\sysnative\WerFault -a	73
%windir%\sysnative\choice.exe	45

Table 7: Most commonly used spawn to processes (x64) during reporting period, by number of uses

Spawn to x86	Count
%windir%\syswow64\rundll32.exe	8,087
%windir%\syswow64\dllhost.exe	1,342
%windir%\syswow64\gpupdate.exe	233
%windir%\syswow64\svchost.exe	196
%windir%\syswow64\runonce.exe	184
%windir%\syswow64\WUAUCLT.exe	173
%windir%\syswow64\regsvr32.exe	150
%windir%\syswow64\WerFault.exe	106
%windir%\syswow64\WerFault -a	75
%windir%\syswow64\choice.exe	45

Table 8: Most commonly used spawn to processes (x86) during reporting period, by number of uses

The watermark is the unique license ID for each Cobalt Strike build. Trial versions, cracked versions, and stolen legitimate versions of Cobalt Strike have been leaked and distributed in the wild, making it difficult to attribute based on the watermark. However, it can be helpful when clustering infrastructure with additional configuration settings. Table 9 shows the top watermark IDs seen in the data.

Watermark	Count
1234567890	2,343
0	1,692
305419896	1,216
426352781	1,016
100000	698
1580103824	662
1359593325	542
391144938	431
206546002	387

Table 9: Most commonly used Cobalt Strike watermark ID during reporting period, by number of uses

What Steps Can Defenders Take Now?

- Many of the C2 domains are newly registered domains and are categorized as such by many forward proxies. If your forward proxy solution supports this, consider setting policies to block domains categorized as recently registered.
- Ensure that network telemetry is centrally logged, so that monitoring can be put in place to detect anomalous connections.
- Host-based telemetry offered by EDR technologies play a crucial role in detecting behaviors of Cobalt Strike and other post-exploitation frameworks. Ensure there is significant EDR coverage across the host within your environment, to provide as much visibility as possible

Conclusion

This report aimed to identify trends related to the current cyber threat landscape, looking at data and observations of activity tracked by ReliaQuest in 2022. From our observations we can draw several conclusions based on our data.

- Each sector faces unique challenges, many of which are highly dependent on a company's business model or operating requirements. As identified by our breakdown of GMDRP alerts, some risks will be more pertinent for certain sectors; however steps can be taken to minimize any possible impact. Visibility and context is key - understand what specific threats your business face and apply compensating controls where appropriate.
- The most commonly observed attacker technique was aimed at exploiting external-facing remote services. These were attempted to either initially access and/or persist within a network. This highlights the ongoing problem of sufficiently hardening remote services, which includes the use of Citrix, VPN, and notably, RDP. The exploitation of remote services will continue to represent arguably the most common abuse point for entering your network, by both cybercriminals and nation-state aligned threat actors, across all sectors. Ensure that remote services are not unnecessarily external facing, patched with strong authentication measures in place, using the principle of least privilege to ensure only individuals who need to do their job can access them.

- Taking a patch-all approach to vulnerability management is an ineffective method of tackling vulnerability risk. Adding vulnerability intelligence can guide your security team in tackling the CVEs that represent the greatest chance of causing an impact to your business. Getting a robust, consistent, and repeatable vulnerability remediation program in place can go a long way in raising your overall cyber resilience.

- Initial access malware continues to be delivered via phishing emails, with threat actors adapting their techniques to minimize the effectiveness of organizational controls. This is likely to go hand in hand with the ongoing risk associated with ransomware activity, who often use IAB and associated techniques as the point of entry onto a target network. Increasing resilience to IA malware is best accomplished through a combination of email security controls, group policy to minimize the chance of a malicious file being delivered/opened, and user awareness programs.

- Ransomware remains the biggest risk facing business in 2023. Ransomware actors are agile, resourceful, and capable of reacting to defenders' actions in changing their tactics. It is likely that the ransomware ecosystem will become more saturated in 2023, with the introduction of several new groups. Keeping abreast of the latest developments in TTPs of ransomware activity, in addition to tracking groups known to be targeting your sector, is the best way to stay ahead of the curve from this pernicious activity.

- Use the trends identified in this report to inform your own threat model and act accordingly. It's always better to 'stay left of boom' and act in a proactive manner. Prevention is always a better approach than remediation.

How ReliaQuest can help:

Put our threat intelligence technology to work for you. With the ReliaQuest GreyMatter security operations platform, you can get unparalleled visibility into your entire ecosystem—and beyond.

The GreyMatter Intel capability is fully configurable; use our pre-defined set of threat feeds or even add your own. We'll take that data and return actionable insights on threats and IoCs. And with Digital Risk Protection (DRP), you can be sure your data is safe outside your environment too.

To learn more, visit www.reliaquest.com. Set up a custom demo to walk through your environment and learn how ReliaQuest can help.

If you would like any further information on any of the threats detailed in this report, please contact ReliaQuest's Threat Research team.