Coalition | SECURITY LABS

coalitioninc.com

Cyber Threat Index 2024

Insights on internet security, cyber risk trends, and critical vulnerabilities

Table of Contents

Executive Summary

Businesses should have an effective way to prioritize vulnerabilities and address other digital risks. Cybercrime is easy and lucrative, which means threat actors will continue to target businesses by leveraging exploitable vulnerabilities and poor cyber hygiene.

The number of vulnerabilities continues to grow exponentially, with thousands announced each month. Unfortunately, businesses tend to optimize for growth, not cyber risk management, and many security and IT teams are stretched thin. So, how can defenders stay ahead of threats and help protect their digital infrastructure from exploitation?

Legacy vulnerability scoring systems commonly lack a practical understanding of real-world exploitation. While the news cycle can provide critical information for defenders, it often lags behind threat actors and can have misleading results.

At Coalition, we believe the solution lies in gathering and analyzing data to produce meaningful recommendations and results. The breadth and depth of data we collect on cyber exposures allow us to make sense of cyber risk and, in turn, share actionable insights with policyholders and the security community to help them better prepare for and respond to cyber threats.

The Cyber Threat Index 2024 is a compilation of data and insights from Coalition Security Labs. Highlights from this year's report show:

- Nearly 35,000 Common Vulnerabilities and Exposures are expected in 2024 — a 25% increase in the rate of discovery compared to the first 10 months of 2023
- Honeypot activity spiked 1,000% more than two weeks before the MOVEit security advisory
- Scans from unique IP addresses looking for risky technologies (like Remote Desktop Protocol) increased by 59%
- More than 10,000 businesses are running the end-of-life database Microsoft SQL Server 2000

This data is derived from Coalition's threat-collection technology. Our findings can help organizations formulate a data-driven strategy for managing cyber risk. Readers will gain insights into how to prioritize vulnerabilities, understand which technologies threat actors are targeting, and compare cyber hygiene across industries.

About Coalition Security Labs

Coalition Security Labs is a team of researchers and analysts committed to addressing tomorrow's most urgent cybersecurity challenges. Protecting the unprotected starts with security. We're invested in helping IT and security professionals prevent and respond to fast-moving cyber risks. Security Labs brings together the innovation and expertise that serves as the foundation to Coalition's data-centric approach to cyber risk management.

SECTION 1

The Growing Surge of Vulnerabilities

Vulnerabilities are one of the top three vectors ransomware actors use to compromise victims, making it essential to understand their impact. Vulnerabilities are primarily tracked as common vulnerabilities and exposures (CVEs), although some may have an incorrect or nonexistent CVE identifier. The volume of vulnerabilities discovered has steadily increased since the 1990s, with the number of CVEs surging 500% since 2016 (Figure 1.1). Notably, this spike followed six years in which the annual volume of vulnerabilities was relatively stable.

Potential drivers of this time-specific surge in vulnerabilities include:

- The rise of bug bounty programs in which vendors directly pay hackers for finding vulnerabilities in their system
- The commercialization of cybercrime and the emergence of marketplaces where threat actors can trade exploit kits, stolen credentials, and access to victim networks
- The rise in the number of entities with the authority to issue a CVE number, which increases the number of public CVEs in a calendar year

The sharp spike in CVEs has led to an increased focus on identifying vulnerable software from both threat actors seeking a means of ingress and defenders trying to protect against exploitation.

Increased Monthly Volume of CVEs, 2010-2023 *(Figure 1.1)*

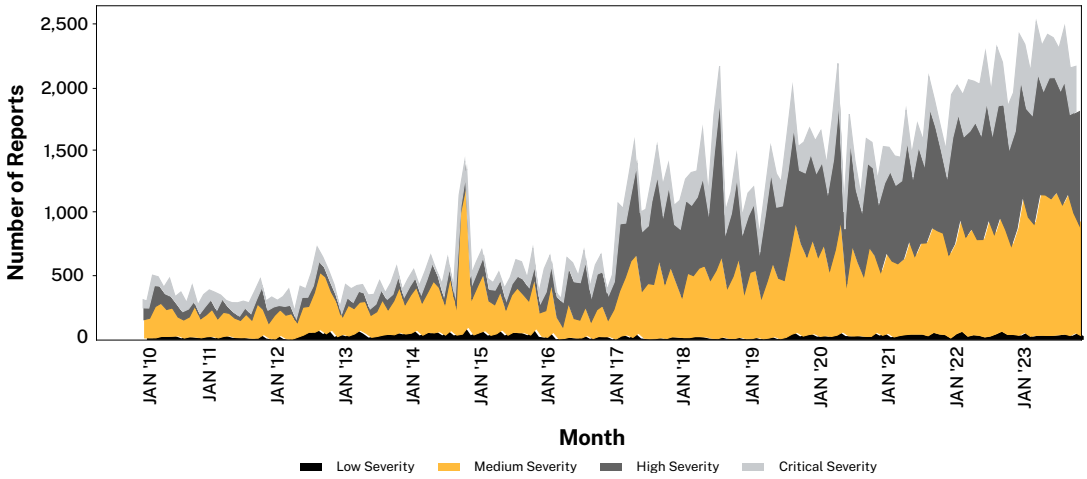

[1]*Amy Stokes-Waters links the increase in CVE IDs to "the rise of bug bounty platforms and programmes offered by either third party organizations like BugCrowd, HackerOne, Intigriti, etc and by tech companies themselves." Meanwhile, Sridhar et al. focus on the increase in entities who can process CVE requests. The number of authorized entities (CNAs) "grew from 22 in 2016 to over 100 in 2019," which coincides with the surge in CVE IDs from 2017 onward.*

Our Cyber Threat Index 2023 predicted the volume of CVEs in 2023, falling short of the average for the first 10 months of 2023 of 2,321 monthly vulnerabilities. While our prediction netted lower than the true number of vulnerabilities per month, our estimation still helped enable us to better anticipate seasonal variations in the volume of notifications we sent to policyholders. Our lower estimation also demonstrates the acceleration in the discovery of vulnerabilities and shows the magnitude of the cybersecurity problem in 2023 as cybercriminals amplify the scale of their attacks.

Monthly CVE IDs Forecast for 2024 *(Figure 1.2)*

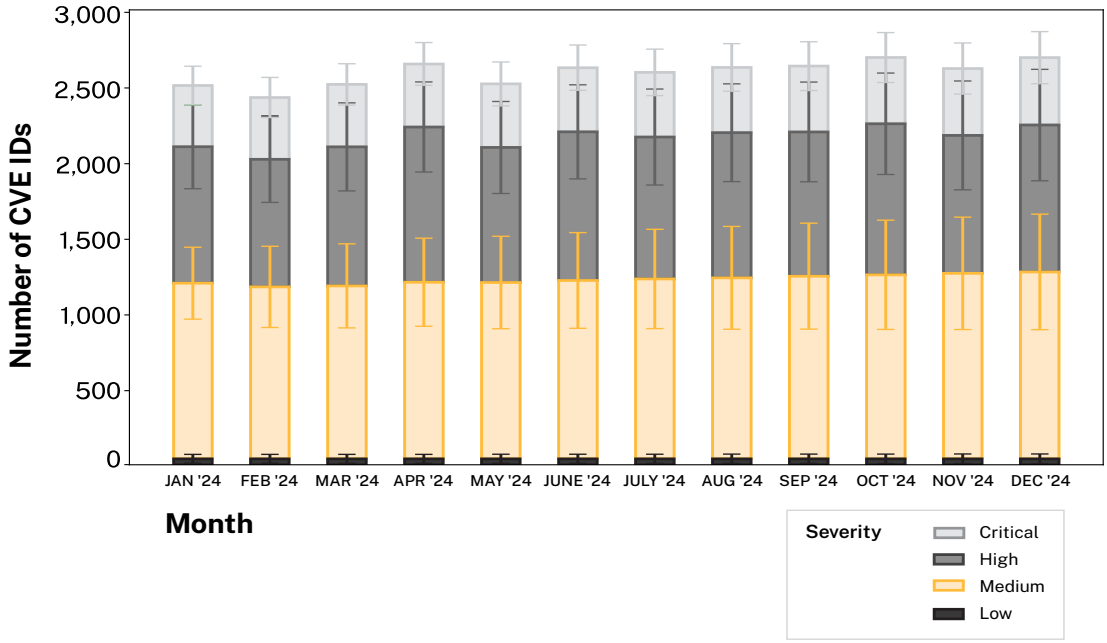

We expect the number of CVEs to increase even more in 2024 (Figure 1.2). We trained an autoregressive integrated moving average (ARIMA) model, commonly used for forecasting time series data, to estimate the number of vulnerabilities for 2024. Our analysis predicts that 34,888 vulnerabilities will be published in 2024, or roughly 2,900 monthly vulnerabilities, a 25% increase from the first 10 months of 2023.

Using an ARIMA model ties into vulnerability forecasting, an emerging science that aims to help defenders estimate future workload regarding vulnerability management.[2] Accurate vulnerability forecasting can help with long-term decisions regarding hiring and potentially short-term decisions like time off, providing the forecasts are sufficiently granular.

It would be challenging for under-resourced businesses to process and action alerts based on every CVE each month. Responding to vulnerabilities will become even more challenging in 2024, when that number nears 3,000 per month, highlighting the need for an effective method to prioritize which vulnerabilities require the most urgent attention.

[2]*Leverett et al. published a paper titled "Vulnerability Forecasting: Theory and Practice" and the lead author went on to co-organize the FIRST Vulnerability Forecasting Technical Colloquium.*

Celebrity CVEs

While relevant, the news cycle for emerging vulnerabilities can have misleading results, as illustrated by several CVEs that made a splash in 2023, either for their disruptive nature or for other factors. We call these "celebrity CVEs."

Let's look at some examples of celebrity CVEs that made news headlines in 2023 and underscore the importance of a data-driven approach to prioritizing mitigation and response.

Exim

The September 27 announcement of a critical vulnerability impacting Exim made waves. The Zero Day Initiative (ZDI) caused panic among security professionals after publishing a security advisory regarding a remote code execution (RCE) vulnerability in the Exim Mail Transfer Agent. An RCE vulnerability would allow a threat actor to extract emails at scale, including sensitive data, or send spam emails from stolen accounts.

To understand the scope of the problem, Coalition's security team scanned policyholders for vulnerable infrastructures. In 330,000 instances, our scans detected the vulnerable configuration just once for CVE-2023-42114 and four times for CVE-2023-42115. Notably, ZDI reserved two CVEs for the vulnerability, but neither was published in the National Vulnerability Database (NVD). Given the number of global Exim installs, the results of an easily exploitable vulnerability combined with a common configuration could have been catastrophic.

But in reality, attackers could only exploit these vulnerabilities if the mail servers used a specific configuration. Despite headlines that sparked substantial attention — Critical vulnerabilities in Exim threaten over 250,000 email servers worldwide — we opted against manual outreach. Instead, we notified the impacted policyholders via Coalition Control™, our cyber risk management platform.

The Exim vulnerability helps illustrate how we think about systemic risk. When a vulnerability in a product with a significant presence is disclosed, it can look like a catastrophic cyber event. In reality, most vulnerabilities only impact a subset of assets. Our Active Cyber Risk Model accounts for these nuances while leaving room for the small possibility of a wide-scale exploit that affects all configurations and versions of a product.

MOVEit

The June 2023 vulnerability in Progress Software's MOVEit received widespread attention, not least because of the high-profile victims. The Cl0p ransomware gang maintained a public list of victims, including the Shell oil company, multiple airlines, consultancies, and universities, slowly releasing victim names over weeks.

The MOVEit vulnerability highlights how waiting until news breaks about active vulnerabilities means missing the opportunity to prevent infections or contain them. The relevant advisory was released on Sunday, May 31, 2023, by Progress Software. The next day, the Cybersecurity and Infrastructure Security Agency (CISA) released an advisory urging organizations to mitigate the vulnerability. In contrast, major news reports broke nearly two weeks after the initial disclosure.

Coalition provided manual outreach to policyholders on June 1, 2023, the first working day after the initial advisory. But this raises the question of whether it would be possible to react sooner. Rapid7 detected the same web shell installed on multiple customers' systems dating back to May 27, 2023, though this was only discovered after the Progress Software advisory prompted Rapid7 to scan.

Coalition honeypots detected initial scans related to the MOVEit vulnerability as early as November 2022. We identified evidence of early reconnaissance by searching for "human.aspx" in the payload of inbound packets to our honeypot infrastructure. This provides evidence that a threat actor is trying to connect to MOVEit technology, which could allow them to map out potential targets when scanning hosts at scale. A similar string, "human2.aspx", was used by Rapid7 to detect MOVEit exploitation.

MOVEit Honeypot Data *(Figure 1.3)*

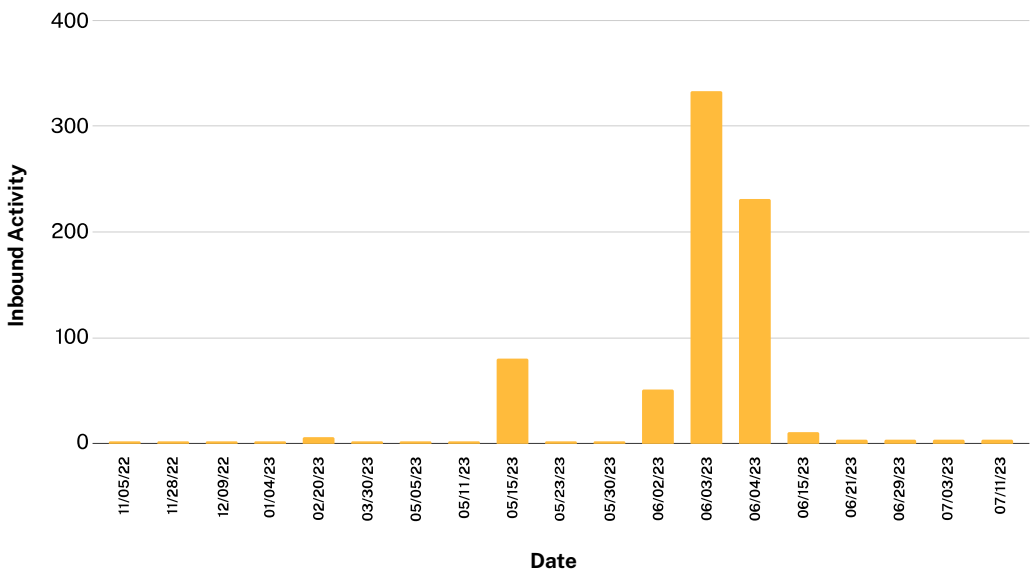

> Roughly **16 days** before Progress published its security advisory, Coalition honeypots detected a **1,000% spike** in scans for MOVEit technology.

Looking at our honeypot data over time shows low activity between November 2022 and the initial spike in mid-May 2023 (Figure 1.3). Roughly 16 days before Progress published its security advisory, our honeypots detected a 1,000% spike in scans for MOVEit technology. A second spike began on June 1, 2023, when security companies began scanning the internet to understand the scope of the problem.

The speed with which the Cl0p ransomware group exploited the MOVEit vulnerability highlights how, at times, even the most relevant information comes too late. For these reasons, using honeypot data as a source that feeds into automated vulnerability prioritization is an exciting prospect.

Our long-term goal is to associate inbound traffic payloads with specific technologies, which would make a spike in activity related to a given technology a feature in our prioritization model. Machine learning (ML) would then assign a weight to this feature. Ultimately, this would allow us to use honeypot traffic to identify celebrity CVEs before they hit the news — providing companies with the opportunity to take action before threat actors.

Citrix Bleed

On October 10, 2023, Citrix released a security bulletin for a sensitive information disclosure vulnerability in their NetScaler application. This software provides virtual private network (VPN) functionality, which enables a secure connection between the end-user and the server. Citrix software is popular — with over 400,000 customers worldwide, including many Fortune 500 businesses.

At the time of disclosure, Citrix was not aware of any exploits in the wild. While Citrix internal teams had identified the vulnerability, they had not observed any threat actors publicly exploiting it to gain access.

Given the nature of the vulnerability and the widespread adoption of the software, developing an exploit was attractive to threat actors. Coalition notified affected policyholders immediately after the initial disclosure, one week before Citrix announced it had identified exploitation in the wild. Over the next several weeks, the situation continued to evolve, and patching became critical (Figure 1.4).

Citrix Bleed Timeline *(Figure 1.4)*

OCT.10	Citrix publishes a bulletin announcing the vulnerability and an available patch.
OCT.10	Coalition sends out notifications for policyholders to patch.
OCT.17	Citrix first detects threat actors exploiting the vulnerability in the wild.
NOV.21	CISA publishes advisory and security researcher Kevin Beaumont publishes evidence that the Lockbit ransomware gang is actively exploiting the vulnerability prioritization.

While it made sense for Citrix to avoid alarmism and wait and see whether the vulnerability would be exploited at scale, there were clear early warning signs, shown by honeypot activity as early as July (Figure 1.5). All of this underlines the importance of proactive vulnerability prioritization.

Citrix Bleed Exploitation Activity *(Figure 1.5)*

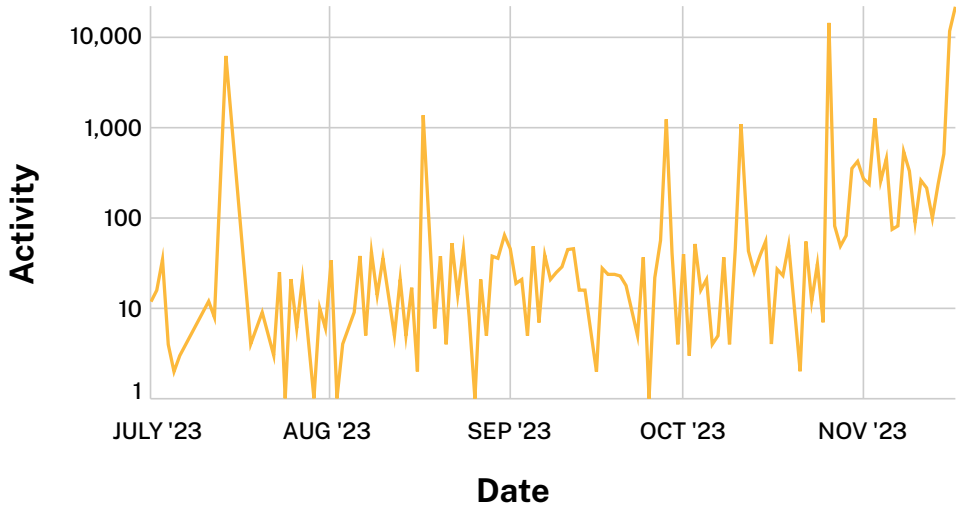

Zero-day vulnerabilities have received significant attention over the last year, but the Citrix Bleed vulnerability reminds us that many threat actors still build exploits for vulnerabilities where the vendor has already issued a patch.

The timing of the Citrix Bleed vulnerability further underscores the need for a dynamic vulnerability scoring system. The initial disclosure minimized the potential severity of the vulnerability. By the time news broke that exploits were available in the wild, ransomware gangs had already moved into action.

A policyholder following Coalition's notifications could have implemented a patch six days before Citrix announced exploits were possible and 42 days before CISA warned ransomware gangs were capitalizing on the vulnerability. Despite the vulnerability's initial lack of attention, we were confident about the need to patch because we have repeatedly observed past attempts to exploit Citrix products in our scan data.

Anti-celebrity CVE

Examining these celebrity CVEs has led us to question whether other vulnerabilities deserved more attention in 2023, an “anti-celebrity CVE,” if you will. Choosing one vulnerability collapses a lot of nuance. Ultimately, instead of a single anti-celebrity CVE for 2023, we determined the most fitting anti-celebrity security concern of 2023 was self-hosted IT infrastructure.

This choice focuses on the reality that many organizations lack the resources and threat intelligence to prioritize patching their infrastructure in a timely manner, leading to a disproportionate risk of suffering an adverse cyber incident or a cyber insurance claim.

Take the October Confluence vulnerability as an example. After scanning policyholders, we found that the majority of their systems were not vulnerable or had already been patched (Figure 1.6). However, some organizations with either an on-premises server or a self-managed data center still required a nudge toward patching. An even smaller minority ignored our notification, making this group the long tail of vulnerability remediation.

Those with unpatched on-premises Confluence assets represent a fraction of our impacted policyholders. However, our experience with other vulnerabilities has shown us that policyholders who do not consistently upgrade their technology or apply patches are more likely to be compromised. In fact, Coalition’s claims data showed that policyholders with even one unpatched critical vulnerability were 33% more likely to experience a claim.

We see this story repeatedly: Most of our energy goes toward supporting policyholders in patching self-hosted infrastructure. Most of these policyholders respond to an alert indicating they are safe from the vulnerability, but some still fail to patch for various reasons.

Coalition Policyholders with Vulnerable Confluence Assets *(Figure 1.6)*

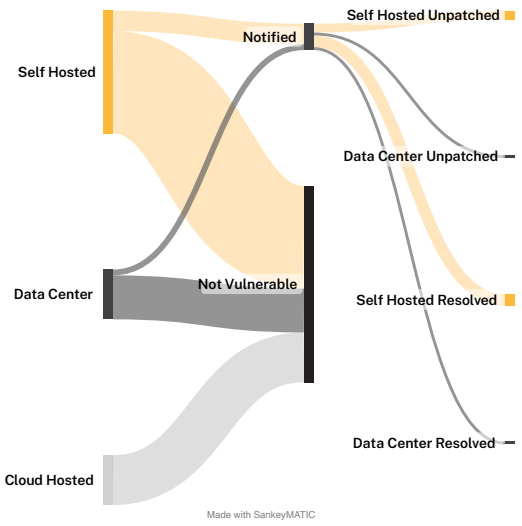

SECTION 2

Vulnerability Scoring

Defenders need a timely, objective method for scoring vulnerabilities. In many cases, exploits are already available to threat actors before a CVE is published, which means threat actors often have a head start on defenders (Figure 2.1).

For a significant minority, exploits became publicly available before the CVE was published. An even higher fraction had exploits privately available before publication.

CVEs with Publicly Available Exploits *(Figure 2.1)*

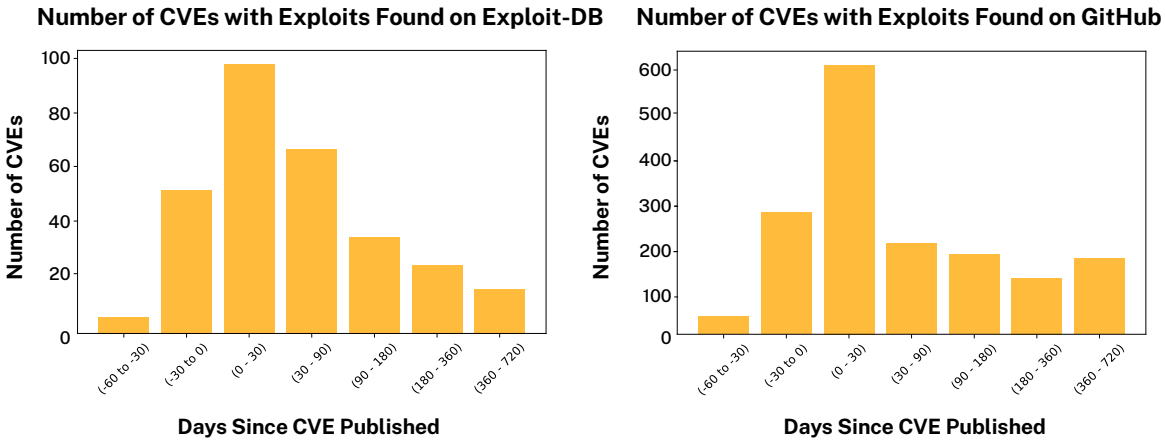

The delay in CVE scoring often means that defenders face two uphill battles regarding vulnerability management. First, they need a prioritization method to determine which of the thousands of CVEs published each month they should patch first. Second, they must patch these CVEs before a threat actor leverages them to target their organization.

The Common Vulnerability Scoring System and Others

The most seasoned candidate for vulnerability scoring is the Common Vulnerability Scoring System (CVSS). CVSS scores vulnerabilities based on their impact and exploitability by considering the potential for confidentiality, integrity, availability breaches, and ease of exploitation. The scores range from 0 to 10, with higher scores indicating more severe vulnerabilities.

CVSS was never intended for risk-based prioritization, even though some compliance bodies recommend using it for that purpose. The core problem is that CVSS scores vulnerabilities without considering the broader context around how the software is deployed in the real world.

Threat actors most often focus on software exposed to the public internet or widely used products that are easy to exploit. Instead of considering the commonality of a vulnerable software configuration, CVSS creates a disconnect between severity and the actual likelihood of exploitation in the wild.

CISA provides an alternative to CVSS with the Known Exploited Vulnerabilities (KEV), a list of vulnerabilities known to have been exploited in the wild. KEV helps focus attention on relevant risks because a small minority of vulnerabilities are ever exploited in the wild. However, KEV should not be used in isolation; its coverage is not global because of its focus on protecting U.S. institutions and there are some delays in publication.

The Coalition Exploit Scoring System

To address these limitations in existing approaches, Coalition developed our own approach to scoring vulnerabilities. The Coalition Exploit Scoring System (Coalition ESS) provides a systematic approach to extracting insights from disparate datasets: security advisories published by security vendors, CVSS scores, CISA's KEV catalog, and other sources.

Our philosophy is data-driven pragmatism. None of these data sources are perfect in isolation; however, all of these sources contain some signal. The problem of vulnerability prioritization is how to weigh and combine these different data sources, a problem well-suited to machine learning.

Coalition ESS is an early source of truth for security risk managers. The Coalition ESS model scans the descriptions used in newly published CVEs and compares them to previously published vulnerabilities to predict the likelihood of exploitability, generating two scores:

1. **Exploit Availability Probability (EAP):** the likelihood that an exploit will be made publicly available, which means the code for the exploit is readily available on the internet for threat actors to leverage in their attacks
2. **Exploit Usage Probability (EUP):** the likelihood that threat actors will actually use an exploit to execute a large-scale cyber attack

Coalition ESS scores are dynamic and updated as more information becomes available, with accompanying histories of scores and changes over time. This is a departure from traditional approaches like CVSS, where scores often remain static after issuing.

CVSS is Reborn, Not Dead

Known issues with CVSS scoring have led various authors to argue that "CVSS is dead." While CVSS is no longer the sole approach to scoring vulnerabilities, it still contributes significant value by becoming an input to other prioritization methods.

The contribution of CVSS to Coalition ESS can be quantified using a statistical approach called SHapley Additive exPlanations (SHAP). This analysis allows us to evaluate how a specific feature of the model contributes to predictions. The SHAP value will be lower if that feature is less important.

Applying this technique to the newest Coalition ESS model reveals that CVSS-derived features contribute 30% of the total SHAP score. This clearly shows that CVSS is not dead but instead can be a core part of modern vulnerability prioritization. Pushing the "CVSS is dead" narrative risks undermining a core feature in other predictive models.

> Coalition ESS provides a systematic approach to **extracting insights from disparate datasets**: security advisories published by security vendors, CVSS scores, CISA's KEV catalog, and other sources.

Analyzing the predictive power of individual features is a window into the future of Coalition ESS and vulnerability prioritization more broadly. Coalition ESS seeks to combine existing data sources to create actionable and timely notifications for our customers.

Future iterations of Coalition ESS will use AI to integrate novel data sources, such as real-time honeypot payload data. This innovation is the future of how we defend our policyholders and provide businesses with the resources and insights to protect themselves from cyber threats.

Prioritizing Patches: The Way to Go

CISA has begun to critique the "patch faster" model, noting that it does not account for adversaries' capabilities. While we agree with the underlying premise — that many technology vendors fail in their responsibility to secure their technology — aggressively critiquing patch management risks undermining faith in one of the most effective security processes for small and medium businesses.

Vulnerability management is plagued by information overload. It is unreasonable to expect under-resourced defenders to apply every patch available from every vendor. Most can reasonably only apply a handful of the most critical patches every month. The core problem lies in identifying which vulnerabilities meet the appropriate level of criticality. With Coalition ESS, businesses can face this problem head-on with a free tool to help reduce the pain of identifying which vulnerabilities are worth immediate action.

For the foreseeable future, data-driven vulnerability prioritization will remain a defender's best bet, which means, arguably, the appropriate model is not "patch faster" but "patch smarter using better insights" about the risk environment around them.

SECTION 3

Honeypot Data

While vulnerabilities are a major risk for businesses, there are multiple other ways threat actors can compromise a network. When software is misconfigured or exposed to the public internet, it signals weak security controls or unprotected infrastructure is likely in place.

Threat actors crawl every global IP address, looking for vulnerable software to target for easy-to-execute cyber attacks. Coalition maintains an extensive network of honeypots to gain insights into threat actor behavior.

Our honeypots span multiple locations, listening to internet activity around the globe. We have hundreds of honeypots operating at any given time and regularly deploy new honeypots to understand emerging threats.

The honeypots are enticing to threat actors: they are configured with multiple vulnerabilities and run outdated software and appliances. Operating these sensors provides important clues about what threat actors are scanning on the internet and the weaknesses they are finding.

The web traffic received by our honeypots is divided into three categories: benign, malicious, and unknown. Benign traffic is usually research-oriented and generated by internet scans performed by security companies, search engine crawlers, or universities. Malicious traffic is generally from threat actors trying to exploit specific vulnerabilities or traffic from known malicious actors like botnets. Unknown traffic is neither identifiably benign nor malicious. Often, this traffic is a novelty and categorized later.

Benign Traffic: Search Engines

Search engines crawl the web for updates, and threat actors often impersonate these benign activities, which makes it sometimes challenging to separate benign and malicious web traffic.

To increase trust in their scanning activity, search engines publish information on how they crawl the web. Google has 82 unique “user agents,” communicated via the User-Agent HTTP header describing the scanning entity’s application, operating system, vendor, version, etc. (Table 3.1).

Unique User Agents Associated with Specific Search Engines *(Table 3.1)*

	SEARCH ENGINE	COUNTRY	UNIQUE USER AGENTS
1	Baidu	China	85
2	Google	United States	82
3	Yahoo	United States	16
4	Bing	United States	9
5	DuckDuckGo	United States	4
6	Sogou	China	4
7	Yandex	Russia	3
8	ExaLead	France	1

This transparency prevents operators from unintentionally blocking search engines, which could result in a website not being displayed in search results. A threat actor can replicate one of the 85 user agents used by Baidu, a Chinese search engine. Threat actors impersonating search engines is one potential explanation for why we observe so much traffic from Bing's Bingbot (Figure 3.1).

Another example of benign traffic comes from firms who scan the internet to help organizations manage their attack surface, such as CENSYS (see Table 1.2) or ourselves in writing this report.

Honeypot Activity Per Search Engine User Agent *(Figure 3.1)*

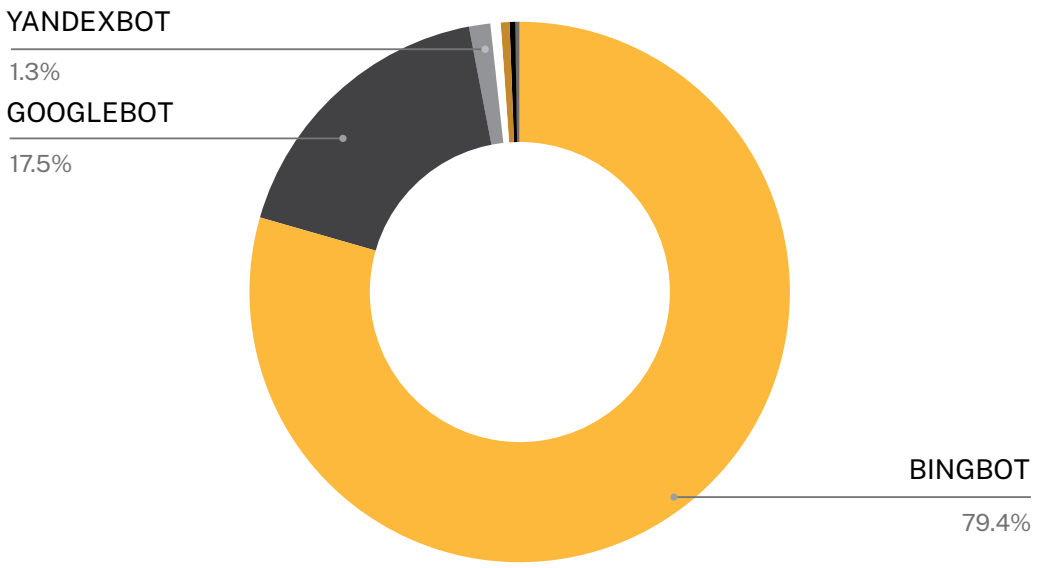

Malicious Traffic: Remote Desktop Protocol

Remote Desktop Protocol (RDP) traffic continued to grow in 2023 despite a lull over the summer, showing that threat actors continue to expect to find internet-exposed assets with open RDP (Figure 3.2). **We detected a 59% increase in unique IP addresses that were scanning for RDP from January 2023 to October 2023.**

When we combine our scan data with our insurance claims data, the risks associated with internet-exposed RDP become apparent. **Coalition data shows that businesses with RDP exposed to the internet are the most likely to experience a ransomware event.**

RDP Scanning Activity by Month *(Figure 3.2)*

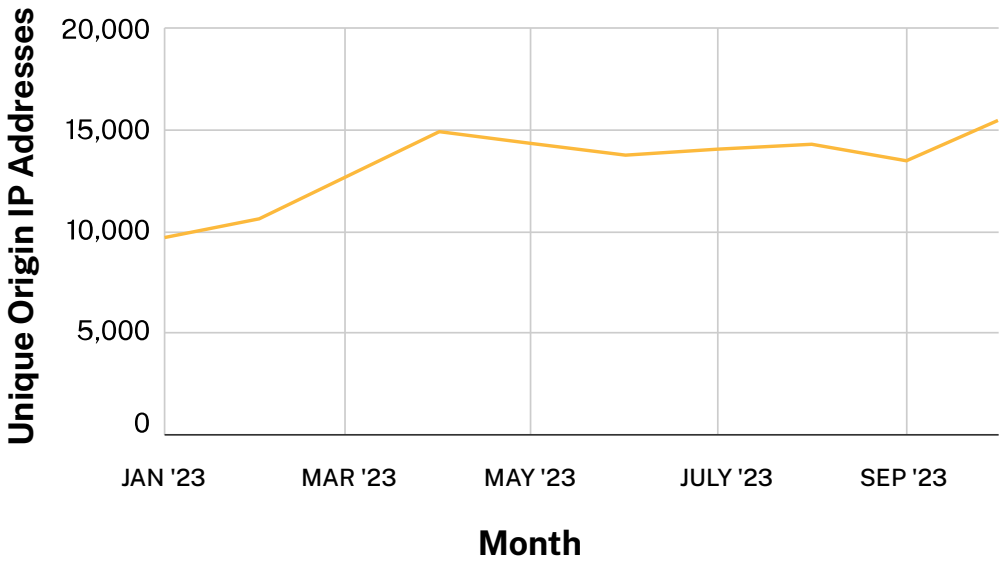

Better Understanding Honeypot Traffic

Coalition provides all honeypot traffic with one or more tags to organize and classify the traffic. The types of tags range from the specific name of the technology or CVE the threat actor is trying to exploit to generic “scanner” traffic, like HTTP or SSH. Below are the top 10 tags for 2023, which show the top types of protocols threat actors seek to exploit (Table 3.2).

In 2023, the first five honeypot tags that threat actors leverage to seek out vulnerable companies remained unchanged from their rankings in 2022, published in our last report. However, RDP comprises an even higher percentage of the traffic in 2023.

Top Scanners *(Table 3.2)*

TAG	PREVALENCE	DESCRIPTION
RDP_SCANNER	75.19%	Scanning for Remote Desktop Protocol
SSH_SCANNER	21.42%	Valid SSH connections
ICMP_ECHO_REQUEST	1.15%	Ping event
HTTP_SCANNER	0.81%	Valid HTTP connections
SSL_SCANNER	0.53%	Valid SSL connections
CENSYS	0.26%	A specific internet-wide scanning company
PROXY_SCANNER	0.14%	Scanning for open proxies
HTTP_REFLECTION	0.08%	Potential DDoS attack
VOIP_SCANNER	0.07%	Scanner for VOIP protocol
SIP_SCANNER	0.07%	Scanner for initiation of VOIP protocol
SMB_SCANNER	0.05%	Scanner for SMB Protocol often affiliated with the exploitation of Microsoft Windows
Other	0.20%	All remaining traffic

The Future of Honeypots

While honeypots provide a wealth of information on threat actor behavior, they also create a great deal of noise. One of the core problems with identifying the needle in the haystack of honeypot data is the volume of benign traffic, which makes determining malicious traffic challenging.

This challenge was perfectly illustrated by the MOVEit vulnerability. In retrospect, the honeypot data we collected allowed us to plausibly identify when the threat actor began searching for vulnerable systems, but the volume was so low relative to other internet traffic that we did not initially detect the anomaly, let alone flag it as malicious.

Honeypot data will continue to inform our threat researchers on the broader security landscape and give us a deeper understanding of the technologies threat actors target. We are slowly rolling out generative AI-enabled tagging rules (Figure 3.3), enabling us to rapidly review and categorize honeypot traffic. Enhanced traffic tagging will allow us to make better sense of anomalous honeypot traffic in real time.

Coalition's Honeypot Tagging Rules Over Time *(Figure 3.3)*

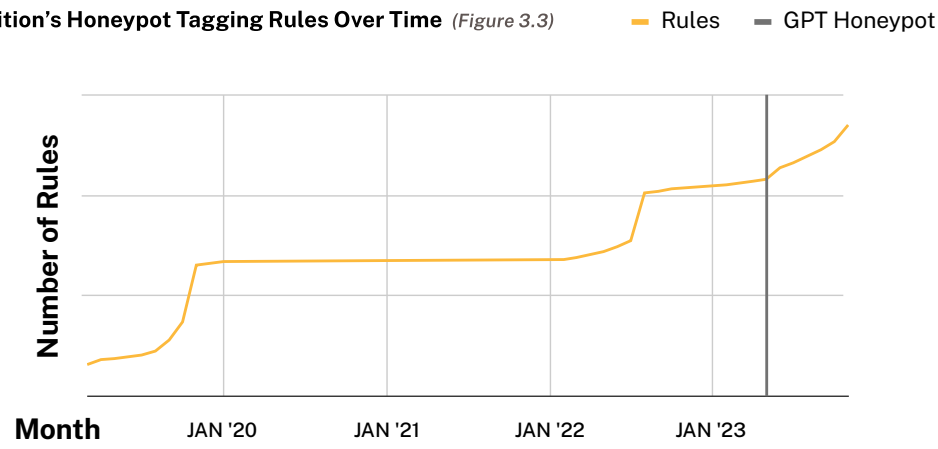

SECTION 4

Internet-exposed Services

Scanning the internet is challenging because of scale, heterogeneity, and efforts to thwart mass scanning. The internet comprises more than five billion IPv4 and IPv6 addresses across tens of thousands of products, each configured differently.

Using Coalition's proprietary Active Data Graph, we routinely examine the riskiest exposures on the public internet. When looking at cyber risk in aggregate, a big piece of the puzzle is understanding what defenders, often unintentionally, expose to the internet.

At the micro level, policyholders are most impacted by commoditized cybercrime in which the threat actors scan the internet for vulnerable or misconfigured services. Identifying the services threat actors are targeting allows manual outreach to impacted businesses before they are compromised.

At the macro level, examining these exposures allows us to assess the potential for a catastrophic cyber event. This might look like self-replicating malware that infects thousands of organizations, as in NotPetya in 2017, or it could look like a cloud outage at a popular service provider. Coalition keeps track of these potential aggregation technologies and vendors (ATVs). We routinely scan for cloud services, including payments, content delivery network (CDN), email, hosting, and so on, as well as non-cloud technologies that attackers might exploit.

Open Ports

Coalition's Active Data Graph is enhanced with information from 246 ports each month and another 461 ports every other month. The core value of port scanning lies in searching for easy routes to compromise and helping our policyholders and businesses reduce that attack surface. Our honeypot data shows that attackers routinely scan random IP addresses — there is no other reason to try to connect to our honeypots — to see if specific ports are open or services are running. Our scans attempt to identify targeted open ports before threat actors do.

The top open ports identified by our scans were related to web traffic, networking equipment management, internet cameras, and time synchronization (Figure 4.1).

Top 20 Ports Found with Services Open *(Figure 4.1)*

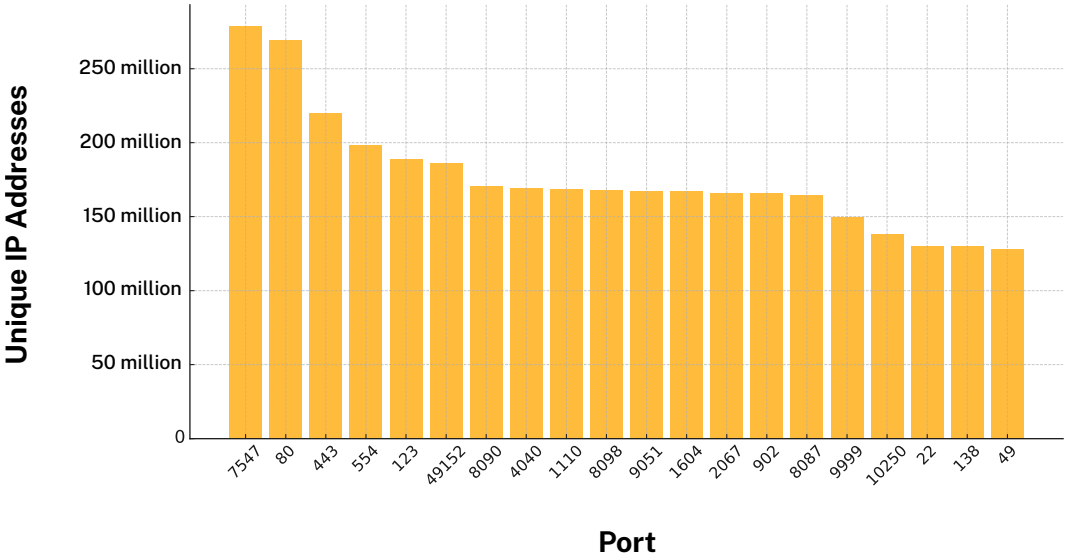

Web Services

Web services are ubiquitous and can be prime targets for threat actors. Defenders need to understand the technologies these services run on because compromising a web server can be a stepping stone to accessing resources deeper into the victim's network, such as internal databases storing sensitive data. Alternatively, threat actors may use the compromise to distribute malware to users who access content from the web server.

The top four web servers have not changed since 2023 (Table 4.1). Express and Kestrel overtook LightTPD and OpenResty. NGINX and Apache were the leading web servers due to their reliability, ability to scale, ability to handle high volumes of traffic, and because both are open-source.

Top Web Servers *(Table 4.1)*

NAME	UNIQUE IP ADDRESSES
Nginx	22,429,212
Apache	15,018,497
IIS	4,532,279
Microsoft HTTPAPI	1,944,066
Express	1,497,952
lighttpd	1,397,360
Kestrel	917,849
OpenResty	820,562
LiteSpeed	538,020
Next.js	404,329

Remote Management Technologies

Exploiting misconfigured remote management technology remains a rampant risk. Misconfigured services, like RDP, allow a threat actor to gain complete control over a device and download or deploy malware.

Microsoft Terminal Services runs most remote management services (Figure 4.2). The percentage of remote services using Microsoft Terminal Services rose from 89.98% to 92.3% from early 2020 to October 2023. This means there are over 3.5 million exposed Microsoft Terminal Services assets on the Internet.

Remote Management Technologies *(Figure 4.2)*

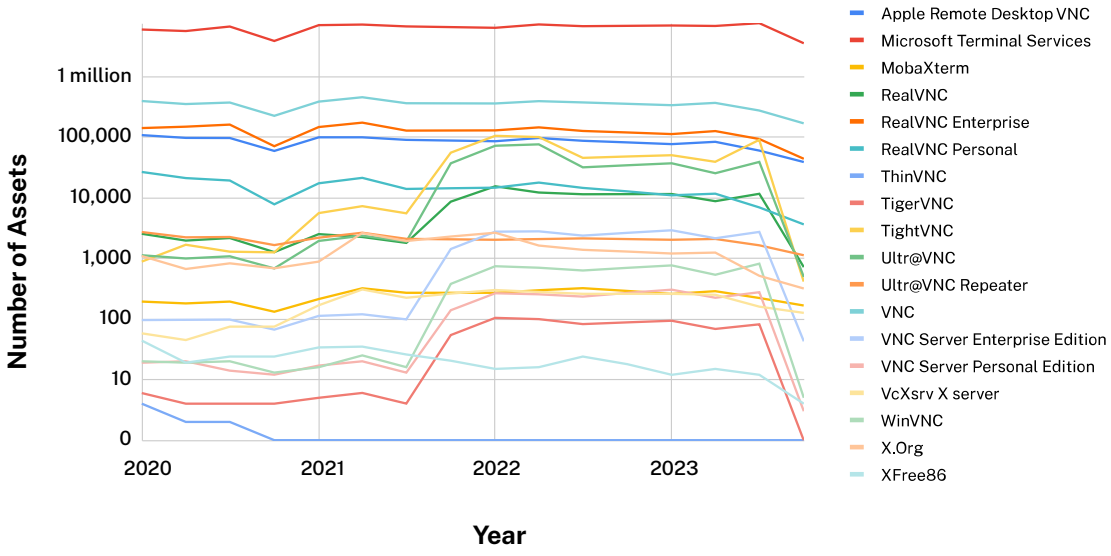

Exposed Databases

Internet-exposed databases are another potential target for threat actors as they often contain sensitive customer or business data. If our scans pick up an asset, then threat actors can also identify these assets at scale, and many of these databases are protected by credentials that can be easily brute-forced. When threat actors successfully exploit weak credentials or unpatched vulnerabilities, they can directly query the database, which could involve submitting malicious requests to extract sensitive data.

Unpatched or EOL versions of databases are a prime target for threat actors. **Our scans found over 100,000 EOL Microsoft SQL servers, including over 10,000 running Microsoft SQL Server 2000, released in 2001.** We find it concerning that so many EOL and unpatched Structured Query Language (SQL) servers remain exposed to the internet because Coalition claims data showed that businesses using EOL software were three times more likely to experience a claim.

Microsoft SQL Servers on the Internet *(Figure 4.3)*

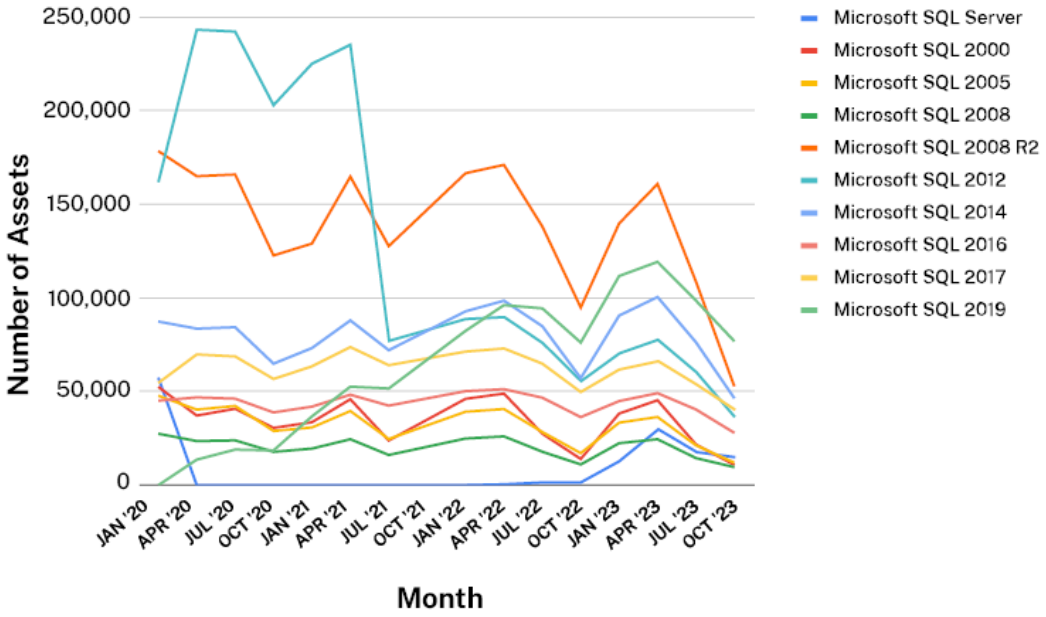

Seeing millions of exposed remote management technologies (Figure 4.2) or tens of thousands of vulnerable SQL servers (Figure 4.3), could create concern over a possible "catastrophic" cyber event. Indeed most business leaders reportedly believe in the inevitability of such an event. However, there has been little agreement on what constitutes a catastrophic cyber event.

Coalition remains confident that cyber is one of the most knowable perils, with more data available than any other type of risk. We believe leveraging this data is the best way to quantify this risk. We will continue to use our Active Data Graph in conjunction with our other threat intelligence services to uncover what technologies threat actors are actively seeking to exploit. This helps us understand what technologies are exposed to the internet and, therefore, what matters for quantifying individual and systemic risk.

SECTION 5

Industry Deep Dives

An easy way for Coalition to see the value in scanning is to gather data on real-world companies and examine their risk in aggregate. We randomly sampled our policyholders in each of the following industries:

- Consumer Services
- Financial Services
- Healthcare
- Professional Services
- Real Estate
- Technology

Using our Active Data Graph, we can see that variations across each industry translate into different types of risk.

Technology firms have the lowest share of cloud assets, typically provided by Amazon Web Services (AWS) and other niche providers. In contrast, Microsoft Azure is the most common cloud provider for all other industries. One possible reason for this difference is the ubiquity of Windows servers and devices in traditional infrastructures, typically managed by Active Directory. Adopting Azure helps these IT teams more easily manage their entire environments.

The Professional Services and Real Estate industries have the lowest number of digital assets, while the Technology sector has the highest. Real Estate, Consumer Services, and Financial Services have the highest number of data leaks per company, while Healthcare and Professional Services have the lowest. For all firms, the most common types of data leaked are personally identifiable information (PII), email addresses, and passwords.

Healthcare has the highest frequency of security checks, which Coalition performs. These are tags we set in place to identify potentially risky technology. The frequency of security check findings indicates how often our scans detected risky technology during the policy period. The findings range from having publicly accessible logins for web panels and content management systems (e.g., WordPress or Atlassian) to unpatched CVEs for a critical vulnerability, such as Citrix Bleed. Security checks are updated regularly to provide reliable insights into our policyholders' risk posture.

INDUSTRIES

Consumer Services

DISTINCT TECHNOLOGY (PER COMPANY)	CLOUD-HOSTED ASSET RATIO	SECURITY CHECK FINDINGS FREQUENCY	AVERAGE CVE CRITICALITY	DISTINCT DATA LEAKS (PER COMPANY)
31.55 ASSETS	84.35%	14.70%	9.63 OUT OF 10	14.56

CLOUD PROVIDERS

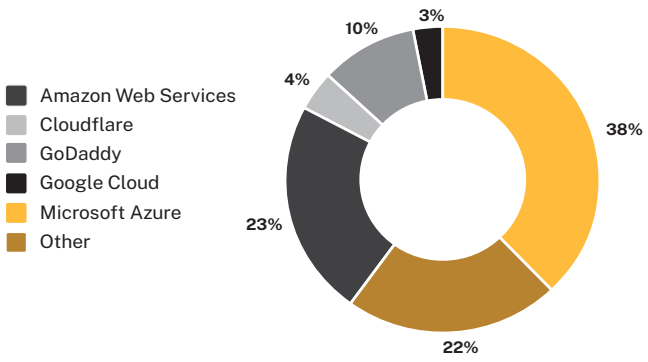

EMAIL PROVIDERS

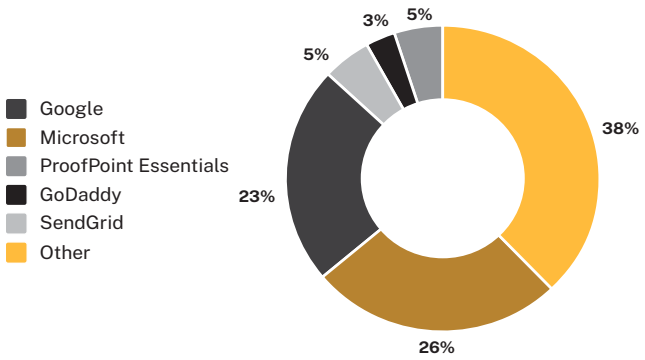

TYPES OF DATA LEAKED

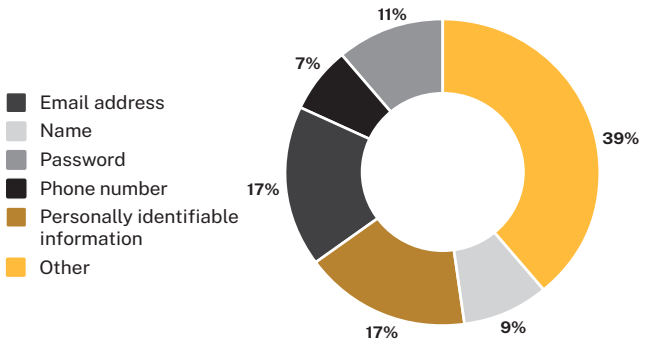

TECHNOLOGIES

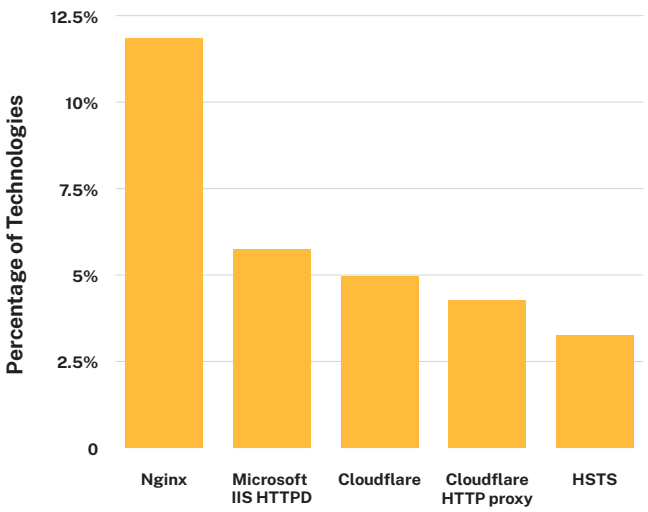

SECURITY CHECK FINDINGS

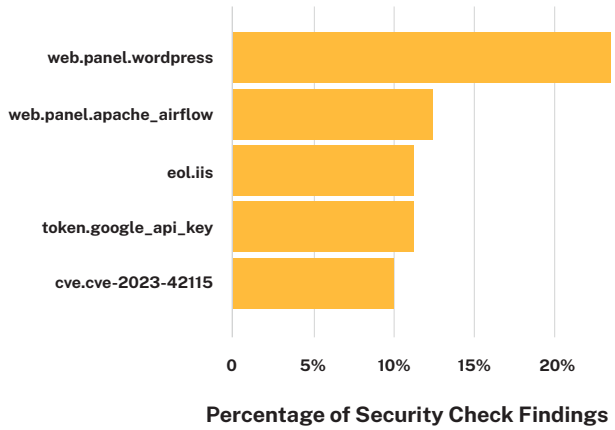

INDUSTRIES

Financial Services

DISTINCT TECHNOLOGY (PER COMPANY)	CLOUD-HOSTED ASSET RATIO	SECURITY CHECK FINDINGS FREQUENCY	AVERAGE CVE CRITICALITY	DISTINCT DATA LEAKS (PER COMPANY)
29.74 ASSETS	72.64%	13.92%	9.57 OUT OF 10	14.92

CLOUD PROVIDERS

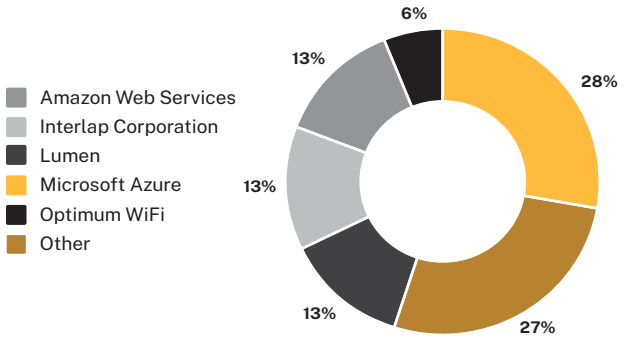

EMAIL PROVIDERS

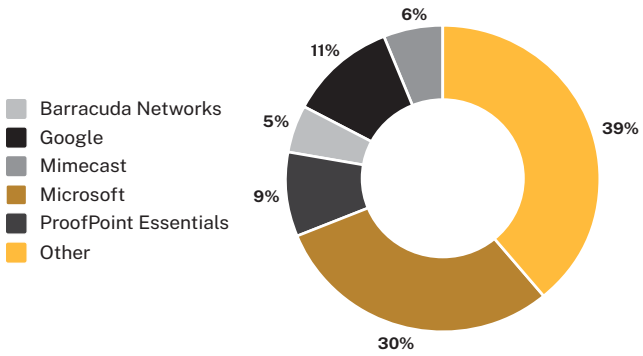

TYPES OF DATA LEAKED

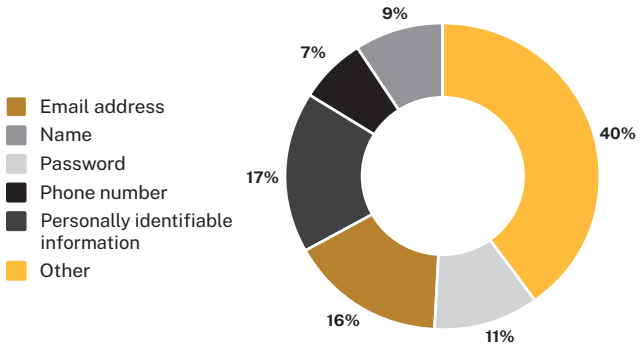

TECHNOLOGIES

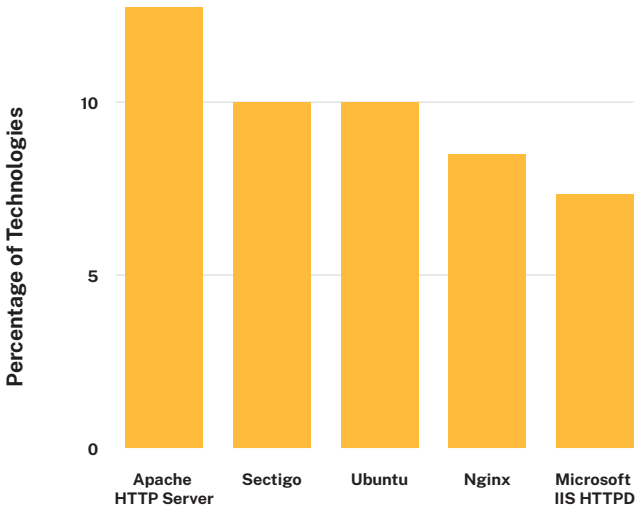

SECURITY CHECK FINDINGS

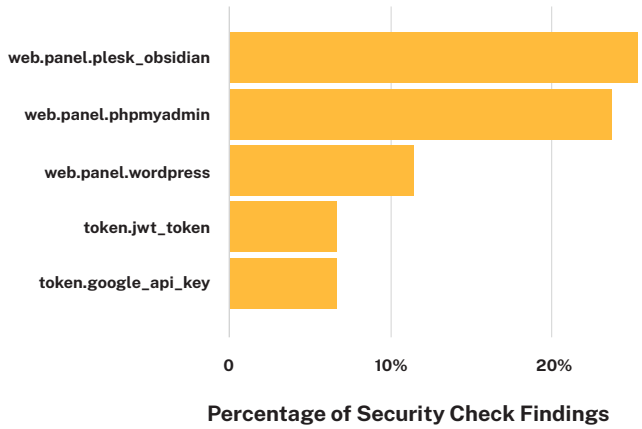

INDUSTRIES

Healthcare

DISTINCT TECHNOLOGY (PER COMPANY)	CLOUD-HOSTED ASSET RATIO	SECURITY CHECK FINDINGS FREQUENCY	AVERAGE CVE CRITICALITY	DISTINCT DATA LEAKS (PER COMPANY)
31.32 ASSETS	79.91%	23.35%	9.15 OUT OF 10	11.34

CLOUD PROVIDERS

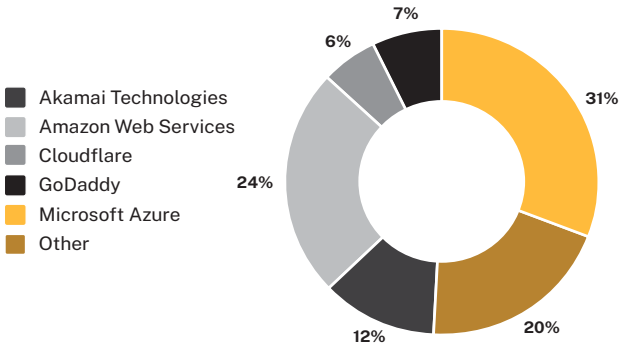

TECHNOLOGIES

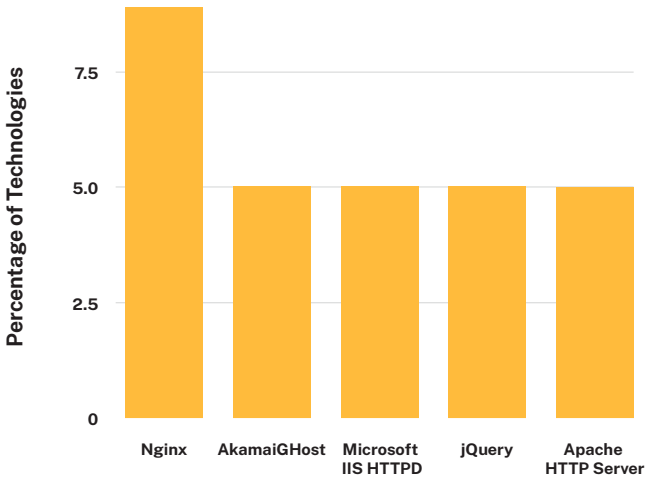

EMAIL PROVIDERS

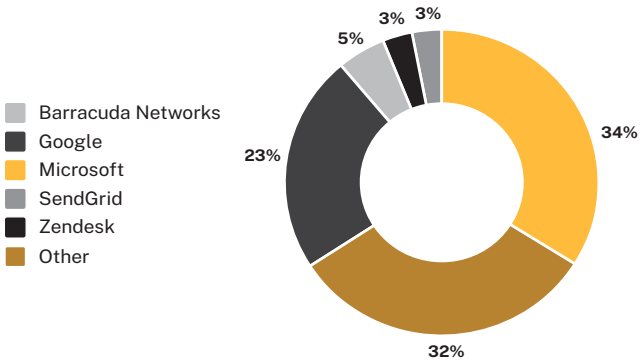

SECURITY CHECK FINDINGS

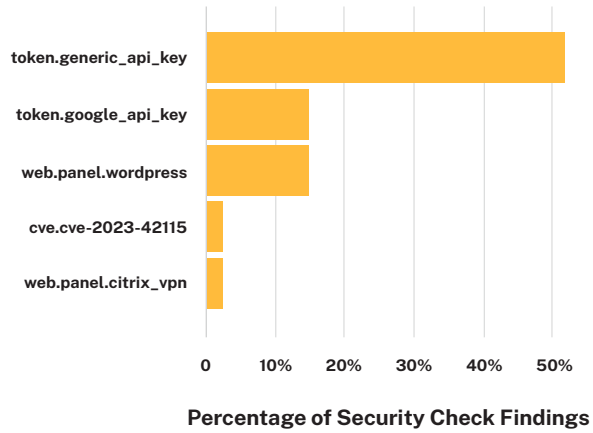

TYPES OF DATA LEAKED

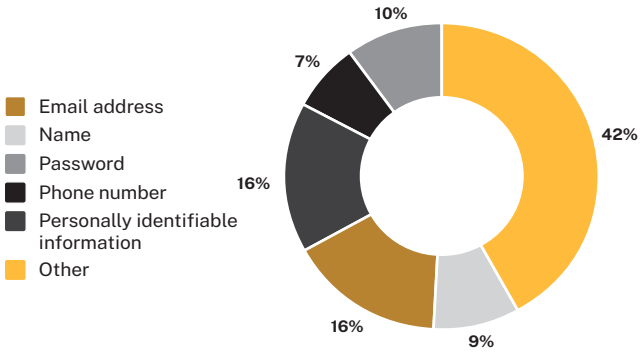

INDUSTRIES

Professional Services

DISTINCT TECHNOLOGY (PER COMPANY)	CLOUD-HOSTED ASSET RATIO	SECURITY CHECK FINDINGS FREQUENCY	AVERAGE CVE CRITICALITY	DISTINCT DATA LEAKS (PER COMPANY)
25.49 ASSETS	80.66%	9.00%	9.75 OUT OF 10	11.41

CLOUD PROVIDERS

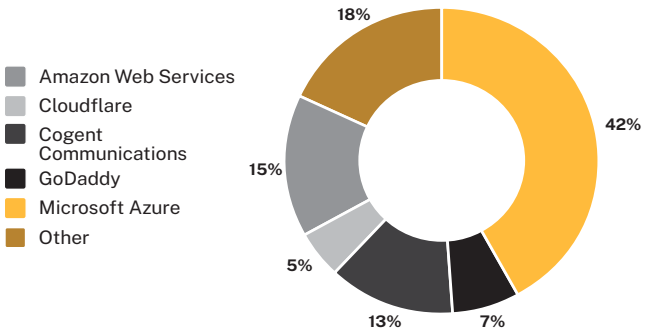

EMAIL PROVIDERS

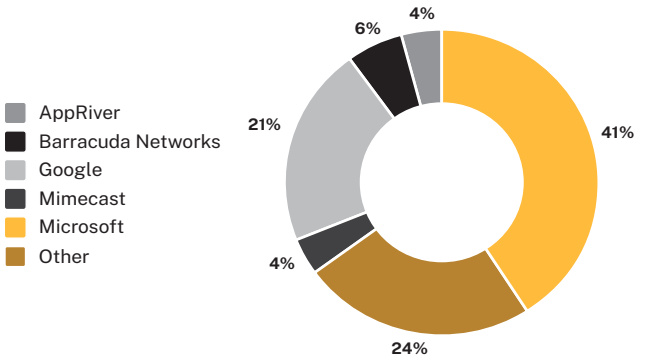

TYPES OF DATA LEAKED

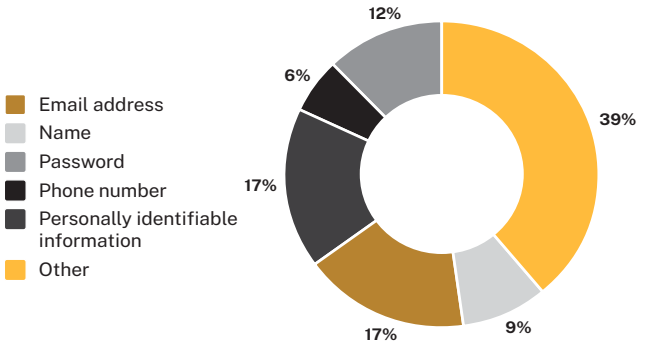

TECHNOLOGIES

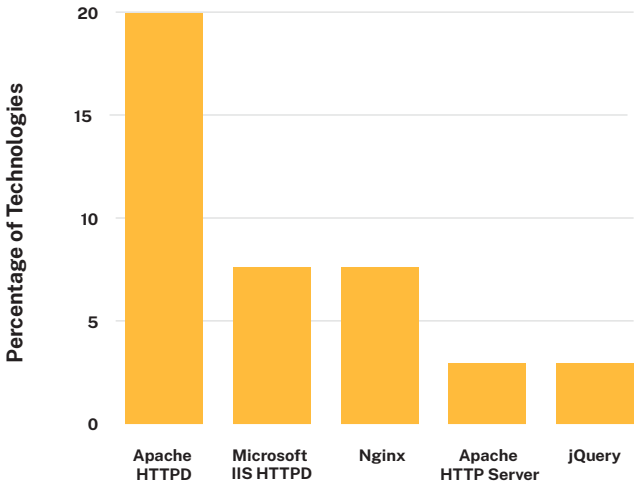

SECURITY CHECK FINDINGS

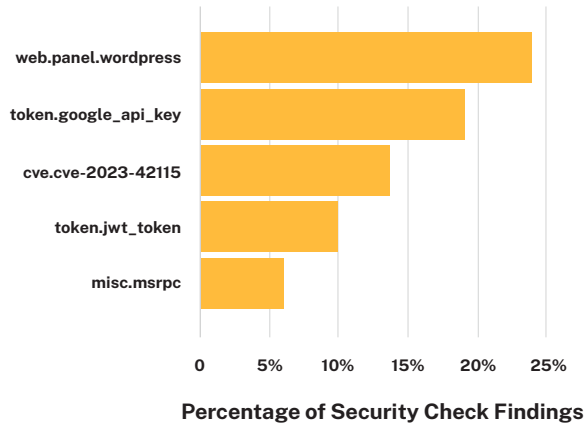

INDUSTRIES

Real Estate

DISTINCT TECHNOLOGY (PER COMPANY)	CLOUD-HOSTED ASSET RATIO	SECURITY CHECK FINDINGS FREQUENCY	AVERAGE CVE CRITICALITY	DISTINCT DATA LEAKS (PER COMPANY)
29.09 ASSETS	81.68%	18.06%	9.75 OUT OF 10	14.76

CLOUD PROVIDERS

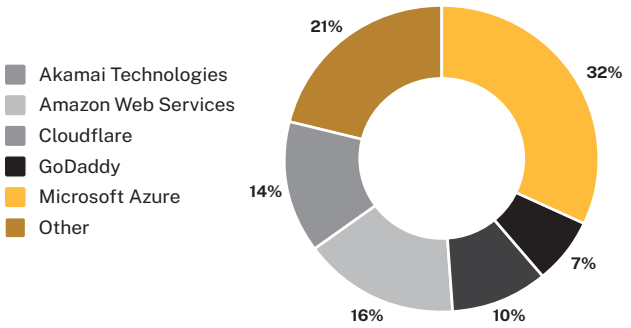

TECHNOLOGIES

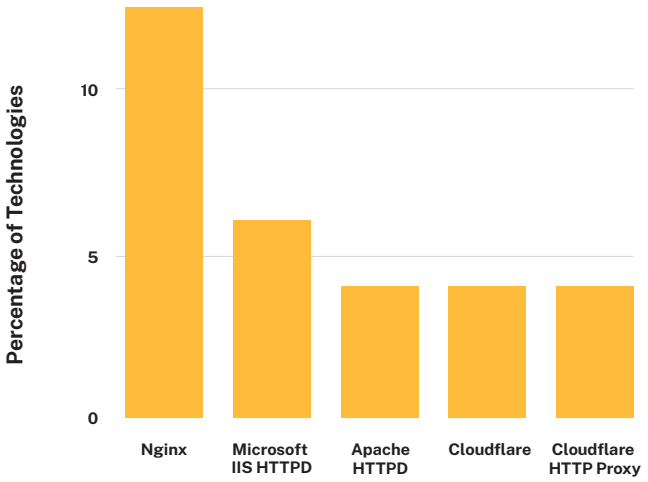

EMAIL PROVIDERS

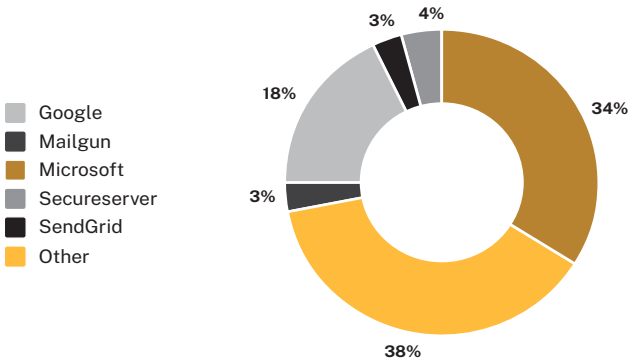

SECURITY CHECK FINDINGS

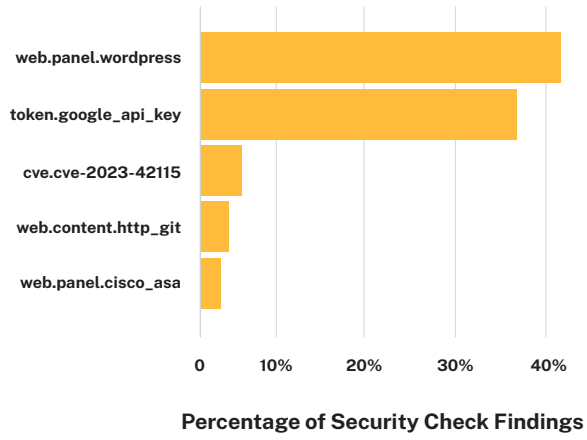

TYPES OF DATA LEAKED

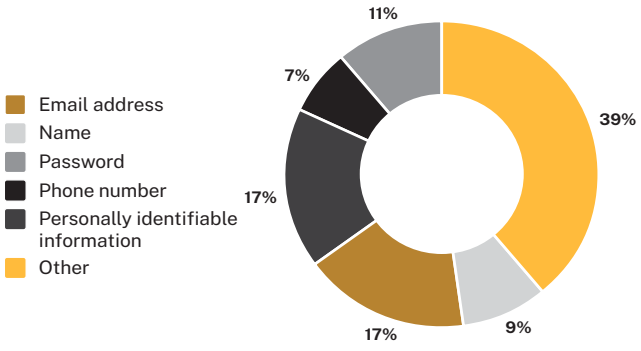

INDUSTRIES

Technology

DISTINCT TECHNOLOGY (PER COMPANY)	CLOUD-HOSTED ASSET RATIO	SECURITY CHECK FINDINGS FREQUENCY	AVERAGE CVE CRITICALITY	DISTINCT DATA LEAKS (PER COMPANY)
44.47 ASSETS	65.69%	10.27%	9.73 OUT OF 10	12.99

CLOUD PROVIDERS

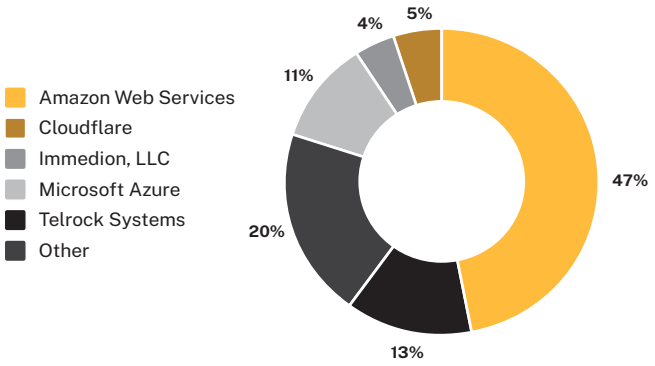

EMAIL PROVIDERS

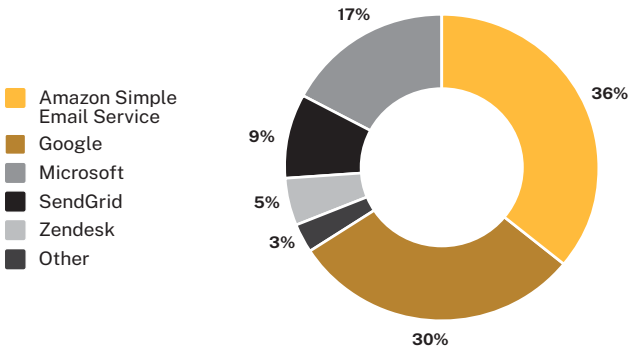

TYPES OF DATA LEAKED

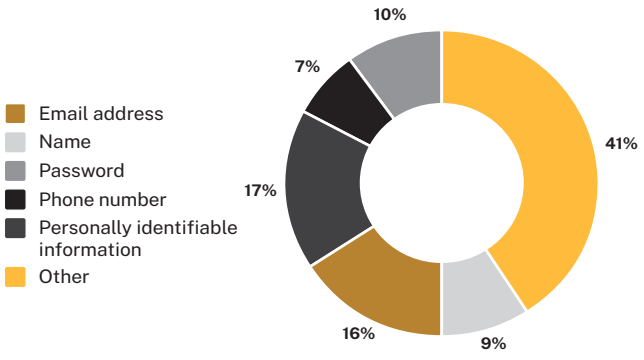

TECHNOLOGIES

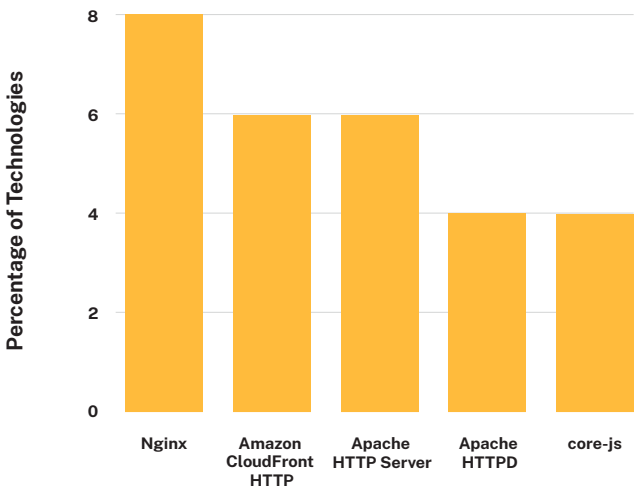

SECURITY CHECK FINDINGS

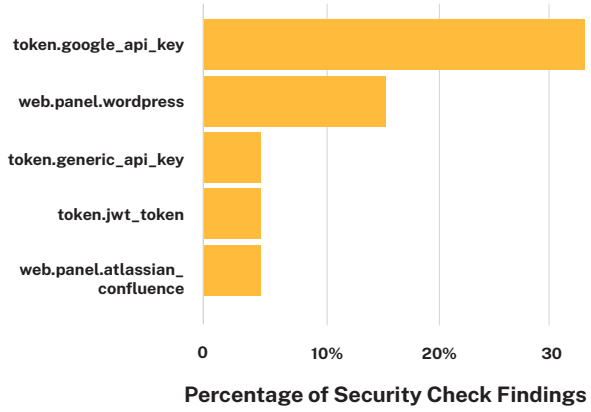

SECTION 5

Finding Better Ways to Manage Digital Risks

The velocity of the digital economy is counterproductive to good cybersecurity hygiene. Every online transaction, every piece of data stored, and every device connected to the internet presents a new risk.

Defenders have a challenging job. They must prioritize thousands of new vulnerabilities discovered monthly and monitor business networks for signs of other cyber attacks. Unfortunately, few teams have the resources to mitigate every emerging risk or respond to all anomalous activity, and most cannot afford 24/7 security. Threat actors know this and use it to their advantage.

One way for security practitioners to work to ensure their networks are protected from fast-moving digital risks is with a managed detection and response (MDR) service. With MDR, businesses don't need to worry about finding the resources or budget to support around-the-clock monitoring. MDR leverages the alerting and detection capabilities of endpoint detection and response (EDR) with human threat hunters who can respond to alerts in real-time.

> Businesses with MDR in place have a **50% faster** mean time to respond (MTTR), dramatically lowering the impact of cyber incidents.[3]

Coalition Managed Detection & Response[4] from Coalition Security Services gives businesses the technology and expertise to help respond and recover faster, minimize impact, and prevent future attacks. Coalition Active Cyber Insurance policyholders may even be eligible for discounts on their cyber insurance premiums[5] if they sign up for our MDR services.

While vulnerabilities are a significant element of cyber risk, they are not the only threat. Threat actors will continue to scan the internet looking for vulnerable technologies, like RDP or EOL software, that provide an easy means of ingress. We believe the best solution is to leverage the right mix of security controls to be prepared to identify suspicious activity and respond before it evolves into a full-scale cyber incident.

Our mission at Coalition is to help protect the unprotected as the world digitizes. We share these insights to help empower cybersecurity defenders to offer businesses stability in the face of growing cyber risks.

[3] *Integrity Research, Managed Detection and Response (MDR) in 20 Cyber Security Statistics*
[4] *Coalition Security Services MDR services are provided by Coalition Incident Response, Inc., an affiliate of Coalition.*
[5] *Terms and conditions apply.*

Methodology

The Cyber Threat Index 2024 is based on Coalition ESS, honeypot, and scan data from January 1 to October 31, 2023.

Vulnerability Scoring

Coalition ESS scores all of the CVEs found in the NVD. The current version uses various features, including CVSS information, vendor and product information, social media chatter, and security advisories.

Scanning the Internet

When scanning, we used Coalition's Active Data Graph, which continuously scans the entire IPv4 space and parts of the IPv6 multiple ports per month. Our data graph first starts a round of TCP-SYN scanning across all IP addresses, followed by service identification and protocol enrichment scanning depending on the port or service being scanned.

Our scanning infrastructure is geo-distributed across multiple countries and providers and uses custom task distribution and scanning modules built in-house. Our Active Data Graph collects information from 246 ports each month and another 461 ports every other month. This includes protocol enrichments for services running on different ports.

Honeypots Listening to the Internet

We have set up an extensive network of honeypots that are geo-distributed across multiple locations and providers. These sensors act as machines that appear unprotected against multiple known vulnerabilities or appear to be running outdated software and appliances. Running these honeypots gives us an idea of what is being scanned on the internet and how attackers are leveraging and exploiting security concerns, including vulnerabilities, to execute their attacks.

Notification Case Study

We included data about our 2023 notification campaign for CVE-2023-22518, which affected Atlassian Confluence assets. This vulnerability potentially impacted hundreds of policyholders. Our working definition of potentially affected was whether a policyholder exposed a Confluence asset to the Internet. We determined that the number of impacted policyholders was significantly smaller than the original pool of hundreds, mainly because most policyholders were running unaffected versions.

coalitioninc.com

55 2ND STREET, SUITE 2500
SAN FRANCISCO, CA 94105