

2024 Attack Intelligence Report

Caitlin Condon, Director of Vulnerability Intelligence
Stephen Fewer, Principal Vulnerability Researcher
Christiaan Beek, Senior Director of Threat Analytics



CONTENTS

Executive Summary	3
<hr/>	
Big Picture: Threat Climate Change	5
<hr/>	
Vulnerability Exploitation Trends	7
A Note on Exploitation Terminology	7
2023 New Widespread Threats	8
Ground Zero: Pre-Patch Exploitation	11
Countdown to Exfil: File Transfer Hacks	13
Silver Linings	15
Other 2023 Exploited Vulnerabilities	16
State-Sponsored Threat Activity	18
Ransomware	18
2023 Initial Access Vectors	22
Life on the Edge: Network Pivots 2020 - 2024	23
Attacker Utilities	25
Vulnerability Classes	27
Programming Language Distribution: 2023 Vulnerabilities	28
Government Guidance on Eliminating Key Vulnerability Classes	30
<hr/>	
Practical Guidance for Defenders	32
Additional Resources	34
<hr/>	
Appendix	35
Notes on Methodology	35
Threat Categorization	36
Ransomware Citations	36
Calculating Time to Known Exploitation (TTKE)	37
Glossary of Terms	37
Attacker Utilities	38
Vulnerability Classes	39
<hr/>	
References	41

EXECUTIVE SUMMARY

Since 2020, Rapid7 has released an annual Vulnerability Intelligence Report with curated vulnerability data and in-depth analyses of exploit trends. In an effort to broaden the scope of this research and offer a more holistic view of the attack landscape, this year's report – renamed The Attack Intelligence Report – combines vulnerability and exploit research with hands-on data from Rapid7's managed detection and response (MDR) division, as well as our threat analytics and emergent threat response teams.

Our 2024 Attack Intelligence Report presents insights and guidance that security practitioners can use to better understand and anticipate modern cyber threats. This year's report highlights multi-year vulnerability and exploit trends in addition to examining recent high-impact attacks and CVEs. This research is based on 210+ vulnerabilities disclosed since the end of 2019, including 60+ exploited vulnerabilities from 2023 and early 2024. See our appendix for additional context on vulnerability selection.



Key findings include:



In 2023, for the second time in three years, more mass compromise events arose from zero-day vulnerabilities than from n-day vulnerabilities. 53% of new widespread threat vulnerabilities through the beginning of 2024 were exploited before software producers could implement fixes – a return to 2021 levels of widespread zero-day exploitation (52%) after a slight respite (43%) in 2022.



Mass compromise events stemming from exploitation of network edge devices have almost doubled since the start of 2023, with 36% of widely exploited vulnerabilities occurring in network perimeter technologies. More than 60% of the vulnerabilities Rapid7 analyzed in network and security appliances in 2023 were exploited as zero-days.



While skilled adversaries are still fond of memory corruption exploits, most of the widely exploited CVEs from the past few years have arisen from simpler, more easily exploitable root causes, like command injection and improper authentication issues.



41% of incidents Rapid7 MDR observed in 2023 were the result of missing or unenforced multi-factor authentication (MFA) on internet-facing systems, particularly VPNs and virtual desktop infrastructure.




Rapid7 Labs tracked more than 5,600 separate ransomware incidents over the course of 2023 and the first few months of 2024. The number of unique ransomware families reported across 2023 incidents decreased by more than half, from 95 new families in 2022 to 43 in 2023.

Rapid7 Labs tracked more than

5,600

ransomware incidents over the course of 2023 and early 2024



in 2023 the number of unique ransomware families decreased by more than half

Big Picture: Threat Climate Change

Over the last several years, Rapid7 researchers have regularly published in-depth analyses of significant vulnerabilities and major cyber incidents, prioritizing attack vectors that have threatened many organizations globally. In 2020, amid what was then considered to be an “outbreak” of critical vulnerability exploitation, our research team began tracking widely exploited CVEs separately from CVEs used in limited, targeted attacks, which were often conducted by a single threat actor. Rapid7’s inaugural [vulnerability intelligence report](#) included just over a dozen of these “widespread threats,” meaning vulnerabilities with many attackers and a large vulnerable target population. At the time, this elevated risk climate was novel, compelling, even alarming.

But 2021 onward [put prior years to shame](#), drawing a rather stark divide between “then” and “now.” Zero-day exploitation soared, reaching a new zenith before leveling off at what looks to be permanently high altitude; the median number of days between vulnerability disclosure and exploitation, which we began tracking several years ago, has stayed in single digits across the CVEs in our annual datasets; widespread exploitation of major vulnerabilities has shifted from a notable event to a baseline expectation; ransomware attacks regularly take entire public-facing systems offline, sometimes for weeks or months at a time; and state-sponsored adversaries have ramped up activity across sectors, using geopolitical conflict as both a rationale and a disguise for espionage, hacktivism, supply chain sabotage, and more.

A hearty and heartfelt disdain for hype is baked into Rapid7 research philosophy. When we talk about a changing threat climate, we require hard evidence — evidence collected and analyzed beyond the binary categorization of “exploited” and “not.” If we’re to say the world has changed, our research ideology demands we produce data to support that claim.

The world has changed. As in previous years, zero-day attacks and widespread exploitation remained common across the vulnerabilities we analyzed in 2023 and the start of 2024. But we’ve also seen a pronounced shift in the way some widespread compromise events play out.

Instead of following the familiar pattern of “many attackers, many targets,” nearly a quarter (23%) of widespread threat CVEs arose from well-planned, highly orchestrated zero-day attacks in which a single adversary compromised dozens or hundreds of organizations in one fell swoop, often leveraging custom tooling like proprietary exploits and backdoors. These aren’t our grandparents’ cyberthreats — this is a mature, well-organized cybercrime ecosystem at work, with increasingly sophisticated mechanisms to gain access, establish persistence, and evade detection.

Anecdotally, these compounding changes in the threat climate have also incited some regressive practices among software producers. In our experience, it’s been increasingly common for vendors to silently patch security issues, withholding advisories and CVE descriptions until days or weeks later. Even then, more vendors appear to be deliberately obfuscating vulnerability details, declining to publish root cause and attack vector information based on an understandable but misguided belief that obscurity deters adversaries and mitigates reputational risk to software producers.

Finally, we’ve also seen the broader security market start to veer even more heavily toward privatization of vulnerability and exploit information, with technical findings often shared in closed loops rather than openly — for profit or otherwise. The rumors of Twitter’s demise exacerbated this trend starting in late 2022, with many disillusioned security community members opting to migrate to an array of alternative platforms and closed intelligence-sharing circles. As of March 2024, industry concern over the future of the National Vulnerability Database (NVD) has sparked new discussions on whether a move toward private ownership of the database might be an improvement over public stewardship.

These are nuanced challenges without simple solutions, and often without institutional or industry support for the human critical infrastructure that so much of our technological ecosystem depends upon — from exhausted security practitioners to beleaguered open-source maintainers and under-appreciated public-sector analysts. Privatization can be a powerful tool, but it’s not a panacea. When we consider it, we should also consider what we may be giving up.

Vulnerability Exploitation Trends

A Note on Exploitation Terminology

Starting in 2023, all of the vulnerabilities we categorize as exploited in the wild in our data are confirmed to have been exploited successfully by adversaries in real-world production environments. This is a departure from previous years, where occasionally we would use third-party honeypot data as a primary source of in-the-wild exploitation of the CVEs in our vulnerability intelligence datasets.

Certain attack-related terms have become more broadly adopted by the industry over the past 18 to 24 months. We've seen a higher number of claims signaling "mass exploitation" or "widespread exploitation" of new vulnerabilities, including in a number of cases where, upon further investigation, "mass exploitation" actually meant a public proof of concept being lazily thrown at the internet, without any exploitable code paths to execute it. This is a particular risk when assessing threat activity around library vulnerabilities or flaws in third-party components. Shared components are implemented in different ways and places in technology stacks, which more often than not makes them remotely inaccessible and difficult to target with one-size-fits-all exploits.

We have frequently observed that honeypot data does not (or cannot) distinguish scanning or unsuccessful exploit attempts from genuine, successful compromise of target systems; this is an understandable technological limitation, but it can also mean clumsy exploit attempts that are unlikely to succeed in real environments may be incorrectly interpreted as "mass exploitation." While scanning activity or exploit attempts (e.g., with publicly available proof-of-concept code) can be reasonable indicators of attacker *interest*, they are rarely indicative of attacker skill. As we know well, "vulnerable" and "exploitable" are not the same thing, particularly without sufficient ability to develop and execute successful attacks outside limited test cases. The **sheer number of honeypots** littering the internet also contributes to inflated prevalence and faulty attack intel.

While some of the vulnerabilities discussed in this report have also been exploited in honeypot deployments, going forward we are changing our practices to only include honeypot data as authoritative evidence of exploitation if we can confirm code execution, payload delivery, or other high-fidelity indicators of successful compromise. This applies to our own honeypot data as well as third-party feeds.

2023 New Widespread Threats

Rapid7 vulnerability researchers prioritize CVEs that are likely to impact many organizations, instead of those likely to affect only a few. We differentiate mass attacks from smaller-scale exploitation; when a vulnerability is exploited to compromise many organizations across many verticals and geolocations, we deem that vulnerability a **widespread threat**. Organizations should expect to conduct incident response investigations that look for indicators of compromise (IOCs) and post-exploitation activity during widespread threat events in addition to activating emergency patching protocols.

Rapid7 researchers tracked more than 30 **new vulnerabilities** that were **widely exploited** in 2023 and the start of 2024. More than half (53%) of the following CVEs arose from zero-day exploits. Rapid7 MDR has observed exploitation of many of the below vulnerabilities in customer environments.

Broadly exploited vulnerabilities that drove compromises across many verticals and target organizations in 2023 include:

CVE-2023-0669 Fortra GoAnywhere MFT Remote Code Execution	CVE-2023-3519 Citrix NetScaler ADC/Gateway Remote Code Execution	CVE-2023-2868 Barracuda Email Security Gateway Remote Command Injection
CVE-2023-42793 JetBrains TeamCity CI/CD Server Authentication Bypass	CVE-2023-24489 Citrix ShareFile Improper Access Control	CVE-2023-29059 3CX Supply Chain Compromise
CVE-2023-34362 Progress Software MOVEit Transfer SQL Injection	CVE-2023-20269 Cisco ASA and FTD Unauthorized Access	CVE-2023-46604 Apache ActiveMQ Remote Code Execution
CVE-2023-40044 Progress Software WS_FTP Server Deserialization of Untrusted Data	CVE-2023-20198 Cisco IOS XE Web UI Privilege Escalation	CVE-2023-26360 Adobe ColdFusion Improper Access Control
CVE-2022-47986 IBM Aspera Faspex Unauthenticated Remote Code Execution	CVE-2023-20273 Cisco IOS XE Web UI Command Injection	CVE-2023-22515 Atlassian Confluence Server and Data Center Broken Access Control
CVE-2023-4966 Citrix NetScaler ADC/Gateway Buffer Overflow	CVE-2023-46805 Ivanti Connect Secure and Policy Secure Authentication Bypass	CVE-2023-22518 Atlassian Confluence Improper Authorization

<p><u>CVE-2023-28771</u></p> <p>Zyxel Multiple Firewalls OS Command Injection</p>	<p><u>CVE-2023-32315</u></p> <p>Ignite Realtime Openfire Path Traversal</p>	<p><u>CVE-2022-47966</u></p> <p>Zoho ManageEngine Unauthenticated Remote Code Execution</p>
<p><u>CVE-2023-27532</u></p> <p>Veeam Backup & Replication Remote Code Execution</p>	<p><u>CVE-2023-38831</u></p> <p>RARLAB WinRAR Code Execution</p>	<p><u>CVE-2022-36537</u></p> <p>ZK Framework Information Disclosure (ConnectWise R1Soft Server Backup Manager Remote Code Execution)</p>
<p><u>CVE-2023-27350</u></p> <p>PaperCut NG Improper Access Control Vulnerability</p>	<p><u>CVE-2023-24880</u></p> <p>Microsoft SmartScreen Security Feature Bypass</p>	<p><u>CVE-2022-44877</u></p> <p>CentOS Web Panel Unauthenticated Remote Code Execution</p>
<p><u>CVE-2023-3722</u></p> <p>Avaya Aura Device Services OS Command Injection</p>	<p><u>CVE-2023-22952</u></p> <p>SugarCRM Remote Code Execution</p>	<p><u>CVE-2022-46169</u></p> <p>Cacti Command Injection</p>

Two Citrix NetScaler ADC/Gateway vulnerabilities (**CVE-2023-3519** and **CVE-2023-4966**, disclosed in July and October respectively) drove incidents that spanned the back half of 2023 and extended into early 2024. **CVE-2023-42793**, a critical authentication bypass in JetBrains TeamCity CI/CD software, was disclosed in September 2023 and exploited by **Russian** and **North Korean** state-sponsored threat actors, prompting **bulletins** from global intelligence agencies months after a patch was released. It was the second major supply chain attack vector that broke news in 2023, after a desktop application from telecommunications company 3CX was **found to have been backdoored** (CVE-2023-29059) as part of a **suspected** North Korean threat campaign.

Adobe ColdFusion **CVE-2023-26360** was initially disclosed in March 2023 as having been exploited in a limited, targeted fashion, but the vulnerability served as an **initial access vector** in multiple campaigns throughout the year, including a successful **attack** on U.S. government servers. A **zero-day** remote code execution vulnerability in SugarCRM (**CVE-2023-22952**) was used to deploy webshells and cryptominers in addition to providing initial access **to AWS cloud environments**). Attackers **exploited CVE-2023-46604**, a **zero-day** remote code execution flaw in Apache ActiveMQ, to drop at least two different kinds of ransomware, as well as **webshells**, cryptominers, and **rootkits**.

While widely exploited vulnerabilities have become commonplace over the last few years, the past 15 months have seen a significant shift in attacker behavior during widespread compromise events. Before 2023, the most common attack pattern we observed for opportunistically exploited vulnerabilities was “many attackers, many targets” — namely, an initial wave of low-skilled exploit attempts, frequently looking to deliver cryptominers or webshells, followed by more adept ransomware group and/or APT exploitation.

Starting in 2023, however, we've seen an increase in mass compromise events where initial exploitation was orchestrated and executed by a **single motivated threat actor** using complex zero-day exploit chains and/or custom implants.

A number of 2023's large-scale attacks followed this pattern:

- The ClOp ransomware group used new zero-day exploits to target two popular file transfer solutions, MOVEit Transfer ([CVE-2023-34362](#)) and **GoAnywhere MFT** ([CVE-2023-0669](#)), in highly orchestrated smash-and-grab campaigns that resulted in data exfiltration and extortion for hundreds of organizations around the world, prompting breach notifications to **tens of millions** of consumers. Both attacks were exceedingly well-planned and executed; the MOVEit Transfer attack, which began over a holiday weekend in the U.S., may have been the culmination of **nearly two years** of threat actor reconnaissance and testing.
- In a truly wild series of updates over the course of several weeks starting in May 2023, Barracuda Networks **disclosed an incident** where a single adversary used a zero-day command injection exploit ([CVE-2023-2868](#)) to compromise a large swath of Email Security Gateway (ESG) appliances with a custom **backdoor** so persistent that the vendor finally told customers to decommission physical devices entirely. In late December 2023, the company **disclosed** a second zero-day vulnerability ([CVE-2023-7102](#)) that had also been exploited by attackers.
- In October 2023, Cisco Talos **shared information** on a pair of Cisco IOS XE zero-day vulnerabilities ([CVE-2023-20198](#), [CVE-2023-20273](#)) that had been exploited by an as-yet-unattributed threat actor to deploy a custom implant christened "BadCandy." The implant had **reportedly been deployed** on tens of thousands of devices before the adversary modified it to **evade industry detection** – the implant is now on at least its third iteration.
- Investigation into a "**suspected APT**" attack on Ivanti Connect Secure and Policy Secure gateways in January 2024 revealed a zero-day exploit chain ([CVE-2023-46805](#), [CVE-2024-21887](#)) that adversaries had used to compromise vulnerable devices, after which they deployed webshells and **backdoored legitimate files**. Thousands of gateways **remained vulnerable** to follow-on CVE disclosures as of mid-February. U.S. government agencies **published a joint advisory** on February 29 emphasizing that threat actors were able to deceive Ivanti's Integrity Checker Tool (ICT), resulting in a failure to detect compromise; the advisory notes that "The authoring organizations strongly urge all organizations to consider the significant risk of adversary access to, and persistence on, Ivanti Connect Secure and Ivanti Policy Secure gateways when determining whether to continue operating these devices in an enterprise environment."


We've seen plenty of low-skill, opportunistic attacks during 2023 and early 2024 – so-called script kiddies haven't magically disappeared. But overall, the skill and sophistication observed in incidents like those above have trended much higher than in years past. While all the vulnerabilities above pertained to either network edge devices or file transfer technologies, concern over shifting attacker behavioral patterns has also made its way into discussions about supply chain security and insider threats – and for good reason, as key 2023 supply chain **attack vectors** and **major incidents** showed.

Ground Zero: Pre-Patch Exploitation

Between the end of 2020 and the end of 2021, large-scale incidents that resulted in the compromise of many organizations more than doubled; those statistics have never returned to pre-2021 levels. But even more concerning is that widespread zero-day threat events quintupled (and then some) in 2021 and have become a mainstay since then.


In 2023, for the second time in three years, more mass compromise events arose from zero-day vulnerabilities than from n-day vulnerabilities.

Widespread Threat CVEs 2020-2024



Since 2021, Rapid7 researchers have tracked the time between when vulnerabilities become known to the public and when they are (reliably) reported as exploited in the wild. This window, which we call “Time to Known Exploitation,” or TTKE, has narrowed considerably in the past three years, largely as a result of prevalent zero-day attacks. Zero-day flaws accounted for 43% of the known-exploited CVEs we’ve reported on since January 2021, with 55% of those vulnerabilities having been exploited within a week of public disclosure and 60% within two weeks. For comparison, zero-day flaws comprised under a quarter of our **2020 Vulnerability Intelligence Report** data, with 30% of vulnerabilities exploited within a week and 32% within two weeks.


Time to Known Exploitation 2020-2024



Average time to known exploitation tends to be a less useful metric when such a large proportion of TTKE values are zero. Nevertheless, the **average** time to known exploitation is just over **22 days** for CVEs in our data whose TTKE values are known. One day is the **median** time to known exploitation for our cumulative annual datasets.

The following chart examines exploited and widely exploited vulnerabilities that Rapid7 has included in annual research datasets over the past four years, along with the percentage of these vulnerabilities that were exploited as zero-day flaws. Since our vulnerability classification and selection methodologies have necessarily become stricter and more prescribed over time, the 2023 data below is a relatively conservative analysis of today’s exploitation trends.

Attack Scale and Speed Trends




Our takeaway from this data is that organizations are under considerable pressure not only to patch new vulnerabilities quickly and continually reduce internet-facing attack surface area, but also to implement compensating controls and detection strategies that minimize an adversary's ability to achieve objectives after that adversary has already gained access to a target network. Technologies like endpoint detection and response (EDR) are key components of a defense-in-depth strategy, but we believe that business leaders should be aware that combating and preventing modern cyberthreats continues to require human expertise in addition to technology. More than ever, burnout and brain drain on security teams compound risk from well-resourced, motivated adversary operations.

Countdown to Exfil: File Transfer Hacks

CI0p attacks on GoAnywhere MFT [CVE-2023-0669](#) and MOVEit Transfer [CVE-2023-34362](#) have dominated news headlines for the past year, in no small part because of the steady stream of breach notifications that went out to tens of millions of global consumers in 2023. But GoAnywhere MFT and MOVEit Transfer weren't the only file transfer technologies exploited by financially motivated adversaries over the past year and a half. In fact, these incidents were enough of a pattern in 2023 that Rapid7 researchers created a new **"smash-and-grab"** attacker utility category for half a dozen file transfer CVEs (see *Attacker Utilities* later in this report).

CVE-2022-47986, a YAML deserialization **issue** in IBM's Aspera Faspex data transfer solution, was exploited by both **ransomware** and Iranian **state-sponsored** threat actors (Rapid7 researchers **analyzed the vulnerability** and had "more trouble preventing [the application] from crashing than actually exploiting it"). **CVE-2023-40044**, another deserialization issue in Progress Software's WS_FTP secure file transfer product, **came under attack** via **multiple exploit chains**, including for attempted **ransomware** deployment. Interestingly, deserialization is overrepresented as a root cause across smash-and-grab exploits — a deserialization issue was also at the root of GoAnywhere MFT **CVE-2023-0669**, and while MOVEit Transfer **CVE-2023-34362** itself is a SQL injection flaw, a .NET deserialization issue was a key part of the **full remote code execution attack chain**.

Citrix ShareFile **CVE-2023-24489**, which very narrowly made our widespread threat list this year, saw a brief spike in **honeypot exploitation activity** in August 2023 that tapered off faster than expected, possibly because the vendor **allegedly disabled access** on the vulnerable component until the patch had been applied. Still, by early 2024 the vulnerability had been cited in enough ransomware incident reports (public example **here**) to demonstrate that unpatched controllers remain tempting targets. The final bug we categorized as a smash-and-grab opportunity appears to be unexploited at time of writing — **CVE-2023-43177**, an unauthenticated remote code execution issue in CrushFTP for which **a nifty exploit chain** is publicly available.



These technologies are a boon for financially motivated adversaries in ransomware and extortion campaigns because of the sensitive data they can store; moreover, many of the attacks Rapid7 has observed on file transfer products have been executed quickly, with attackers getting in, stealing data, and getting out within minutes or hours rather than days or weeks. While there can be room for interpretation on data breach notification regulations generally, an attack ending in data exfiltration from a file transfer application has a lot less wiggle room as far as regulatory reporting requirements go.

Those reporting requirements were also part of the reason that security firms, media, and regulators were able to quantify victims from these types of attacks so granularly, especially when it came to collateral damage. Counting the number of consumer breach notifications that go out as a result of a major cybersecurity incident may not be the best way to quantify attack impact, but it's highly effective for tracking so-called "blast radius."

Silver Linings

File transfer vendors seem to have gotten understandably spooked by the MOVEit Transfer and GoAnywhere MFT hacks perpetrated by CI0p (and the public relations nightmare that followed for many victims). But The Year of the File Transfer Hack™ wasn't without its silver linings, sparse as they may have been. In Rapid7's own **experience** disclosing **vulnerabilities** to file transfer **vendors** in 2023, **they were exceptionally responsive**, displaying a high sense of urgency throughout the disclosure process and in most cases shipping fixes for new vulnerabilities within weeks instead of months. In several cases, these companies patched reported vulnerabilities in less than half the time most firms take.

More broadly, there's been evidence that some file transfer technology vendors are using recent CI0p attacks as an impetus for maturing vulnerability disclosure and product security practices — for example, by speeding up remediation SLAs, establishing formal disclosure mechanisms for external security researchers, and implementing more frequent and transparent patch release cycles. These are positive indicators that may help accelerate and streamline vendor vulnerability responses in the future, alongside strong proactive measures that seek to identify new vulnerabilities before adversaries do.

Other 2023 Exploited Vulnerabilities

The following vulnerabilities are known to have been exploited in the wild in either 2023 or 2024 but as of February 2024, did not have enough technical evidence of large-scale attacks to be included in our widespread threat list.

<u>CVE-2023-46747</u> F5 BIG-IP Configuration Utility Authentication Bypass	<u>CVE-2023-36845</u> Juniper Junos OS EX and SRX PHP External Variable Modification	<u>CVE-2023-38035</u> Ivanti Sentry Admin Portal Authentication Bypass
<u>CVE-2023-29298</u> Adobe ColdFusion Access Control Bypass	<u>CVE-2023-7102</u> Barracuda Email Security Gateway Arbitrary Code Execution	<u>CVE-2023-49103</u> ownCloud Graph API Critical Information Disclosure
<u>CVE-2023-38203</u> Adobe ColdFusion Deserialization of Untrusted Data	<u>CVE-2022-21587</u> Oracle E-Business Suite Remote Code Execution	<u>CVE-2023-28432</u> MinIO Information Disclosure
<u>CVE-2023-33246</u> Apache RocketMQ Remote Command Execution	<u>CVE-2023-21839</u> Oracle WebLogic Server Remote Code Execution	<u>CVE-2023-37580</u> Synacor Zimbra Collaboration Suite Cross-Site Scripting
<u>CVE-2023-41265</u> Qlik Sense Enterprise HTTP Tunneling Vulnerability	<u>CVE-2023-20867</u> Broadcom VMware Tools Authentication Bypass	<u>CVE-2023-29357</u> Microsoft SharePoint Server Elevation of Privilege
<u>CVE-2023-47246</u> SysAid Path Traversal	<u>CVE-2023-20887</u> Broadcom VMware Aria Operations for Networks Command Injection	<u>CVE-2023-23397</u> Microsoft Outlook Elevation of Privilege
<u>CVE-2023-1671</u> Sophos Web Appliance Command Injection	<u>CVE-2023-34048</u> Broadcom VMware vCenter Server Out-of-Bounds Write	<u>CVE-2023-36884</u> Microsoft Windows Search Remote Code Execution
<u>CVE-2023-41179</u> Trend Micro Apex One Arbitrary Code Execution	<u>CVE-2023-35078</u> Ivanti Endpoint Manager Mobile Authentication Bypass	<u>CVE-2023-28252</u> Microsoft Windows Common Log File System Driver Elevation of Privilege
<u>CVE-2023-27997</u> Fortinet FortiOS Heap-Based Buffer Overflow	<u>CVE-2023-35081</u> Ivanti Endpoint Manager Mobile Path Traversal	
<u>CVE-2022-41328</u> Fortinet FortiOS Path Traversal	<u>CVE-2023-35082</u> Ivanti Endpoint Manager Mobile and MobileIron Core Authentication Bypass	

Network edge devices, application development and delivery technologies, and IT security management systems make up most of the exploited vulnerabilities above, but a few outliers merit a mention. Microsoft SharePoint [CVE-2023-29357](#) started out on our “impending threat” list but was exploited in the wild before this report was finalized. A pair of information disclosure vulnerabilities (MinIO [CVE-2023-28432](#), ownCloud [CVE-2023-49103](#)) offered attackers access to cloud secrets, including credentials. An HTTP tunneling vulnerability ([CVE-2023-41265](#)) in Qlik Sense, a data analytics platform, was **exploited** by the Cactus ransomware group in November 2023 for initial access to corporate environments.

Vulnerabilities that are exploited in targeted zero-day attacks often have some of the more scintillating backstories, and this crop is no exception. Of the 28 CVEs in the above list, 15 were exploited as zero-day vulnerabilities (54%). Some of the notable inclusions:

- [CVE-2023-34048](#), a memory corruption issue in VMware vCenter Server that **according to Mandiant** was exploited by UNC3886, a Chinese espionage group, for more than a year before it was discovered;
- [CVE-2023-28252](#), a privilege escalation vulnerability in Microsoft’s CLFS drivers that was discovered during a **Nokoyawa ransomware campaign**;
- [CVE-2023-36884](#), a Windows vulnerability that Microsoft **indicated** was used in both targeted espionage and opportunistic phishing campaigns with Ukraine-related lures;
- [CVE-2023-23397](#), a critical elevation of privilege (NTLM hash leak) bug in Microsoft Outlook that Microsoft said was exploited by a **Russia-based APT** for nearly a year in attacks on **government entities**, critical infrastructure providers, and military suppliers;
- [CVE-2023-47246](#), a path traversal vulnerability in SysAid servers whose zero-day exploitation **Microsoft attributed** to LaceTempest — the threat actor known for distributing ClOp ransomware.

[CVE-2023-37580](#) was a reflected cross-site scripting (XSS) bug disclosed in Zimbra Collaboration, a popular attack target, in June 2023. According to Google’s Threat Analysis Group (TAG), the vulnerability was used in **at least four APT campaigns** targeting government organizations in Greece, Moldova, Tunisia, Vietnam, and Pakistan. Google’s analysis includes a note on how “regular exploitation of XSS vulnerabilities in mail servers also shows a need for further code auditing of these applications, especially for XSS vulnerabilities.”

Finally, a spate of attacks on Ivanti Endpoint Manager Mobile (formerly MobileIron Core), a mobile device management solution, garnered media and industry attention starting in July 2023 for three separate vulnerabilities disclosed over a

10-day period. The first (**CVE-2023-35078**) was a critical authentication bypass that had been used in a **zero-day attack** on a dozen Norwegian **government ministries**. Four days after CVE-2023-35078 was **disclosed** publicly and a patch released, Ivanti published a **second advisory** on **CVE-2023-35081**, an arbitrary file write issue that had been exploited in the wild and could be chained with CVE-2023-35078 to bypass administrator authentication. In the course of analyzing **CVE-2023-35078**, Rapid7 researchers **discovered** that it was still possible for a remote unauthenticated attacker to access API endpoints on an EPMM management server; the issue was assigned **CVE-2023-35082**.

State-Sponsored Threat Activity

Rapid7 researchers gather, analyze, and vet data from a wide range of sources for inclusion in a central threat library that supports Rapid7 products and services. Our data sources include dark web forums, private messaging platforms (e.g., Telegram), and public reports in addition to intelligence from private industry sources and Rapid7's own managed services teams.

This process includes rigorous analysis detailing the suspected source of said campaigns, with associated confidence of attribution. Between January 2023 and March 2024, our analysis identified 188 separate campaigns in which we had a moderate to high confidence that the source originated from APT groups. "Moderate confidence" means that we have a significant amount of evidence that the activity we are observing is similar to what we have observed from a specific group or actor in the past; however, there is always a chance someone is mimicking behavior or conducting a false flag operation.

Based on our analysis, we concluded that across the top 15 groups, 30% of attacks attributed to APT groups originated from Chinese-backed campaigns. Campaigns from Russian-affiliated groups came in a close second, comprising 26% of campaigns conducted by state-sponsored adversaries, while North Korea (14%) and Iran (9%) ranked third and fourth.

Ransomware

Ransomware payments were said to have **topped \$1 billion** in 2023, with groups like Cl0p continuing to employ double extortion attacks that yield huge payouts and have allegedly racked up **hundreds of millions** in profit. The Rhysida ransomware group executed an attack on the British Library that **took down systems** for months, an incident **attributed to LockBit** disrupted British Royal Mail operations, and an estimated 100 million+ American consumers became collateral damage from ransomware attacks **targeting healthcare organizations** in 2023.

Ransomware incident volume hasn't decreased so far in 2024. Already this year, Japanese car maker Nissan disclosed an **Akira ransomware incident** affecting roughly a hundred thousand people, U.S. government agencies **published a joint warning** on defending against Phobos ransomware affiliate operations, and an attack on two healthcare payment providers in France **exposed the data** of 33 million citizens. In February, U.S.-based healthcare payment provider Change Healthcare fell victim to a **catastrophic attack** attributed to the BlackCat (ALPHV) ransomware group; as of late March 2024, the incident was continuing to **disrupt** critical operations, contributing to hefty prescription **backlogs**, payment processing **delays**, and **financial crises** for smaller healthcare providers.

Government action aimed at disrupting ransomware operations has also scaled up. In 2023 alone, the U.S. Department of Justice announced a multi-national **takedown** of Qakbot malware infrastructure, the **seizure** of Hive ransomware servers and sites after a months-long FBI infiltration of the group's network, an international **operation** to disrupt BlackCat/ALPHV ransomware, and the **indictment** of multiple foreign nationals in connection with Conti ransomware and Trickbot malware operations — to name a few.

In October 2023, the U.S. Cybersecurity and Infrastructure Security Agency (CISA) **announced** that they had begun incorporating ransomware information into their Known Exploited Vulnerabilities (KEV) list; while the feature is relatively opaque and doesn't list specific data sources (like the rest of KEV), CISA's strictly conservative approach to adding new vulnerabilities to the KEV helps instill confidence that "Known" ransomware designations are well-vetted.

As of mid-March 2024, the KEV list had 219 CVEs that the agency says are known to be used in ransomware attacks. The following chart shows KEV's ransomware vulnerabilities broken down by CVE identifier year:

CISA KEV Ransomware CVE Distribution



Rapid7 Labs conducted an analysis of 2023 ransomware attacks using data sourced from both external reports and intelligence from Rapid7 MDR. Our combined intelligence sources tracked nearly 5,600 reported ransomware cases between January 2023 and February 2024. That number doesn't account for the many attacks that inevitably go unreported, and therefore is probably much lower on paper than it is in reality.

As Rapid7 Labs **wrote** in January 2024, the number of unique ransomware families these groups used in 2023 decreased by more than half, from 95 new families in 2022 to just 43 in 2023. This implies that the current ransomware families and business models are profitable, and there may not be an immediate need to develop brand new capabilities. The below scatter plot shows the number of ransomware incidents attributed to the top 20 ransomware groups for 2023 and early 2024, based on leak site communications, public disclosures, and Rapid7 incident response data.

Distribution of Posts Over Time by Top 20 Groups in 2023 and 2024



Unsurprisingly, LockBit 3.0, a prolific ransomware-as-a-service (RaaS) operation, continues to be the top group by sheer volume of activity. ALPHV, also known as BlackCat, also maintained a high incidence of activity throughout the year and into 2024, despite a U.S. government-led effort to disrupt the group's operations. The U.S. and UK **announced** in February 2024 that they had disrupted a LockBit ransomware variant. The Justice Department's **December 2023 press release** on ALPHV disruption notes that the FBI made a government-developed decryption tool available to hundreds of organizations, "saving multiple victims from ransom

demands totaling approximately \$68 million.” On March 27, 2024, as part of its “Rewards for Justice” program, the U.S. State Department **announced** a reward of **up to \$10 million** for information leading to the identification or location of BlackCat ransomware gang members.

As mentioned previously, we’ve seen a shift within the past year toward smash-and-grab attacks targeting technologies like file transfer solutions where adversaries sought to quickly gain access to and exfiltrate data from corporate networks. More “traditional” ransomware incidents, where victim data was encrypted, also remained common; the vast majority of ransomware-related incidents Rapid7 MDR responded to in the last year have included encryption, with only a small number of incidents ending in extortion alone.

Exploiting a public-facing application and having a valid account are the top initial access vectors we observed in ransomware and extortion attacks in 2023. Vulnerabilities that are known to have been abused by ransomware groups in 2023 and early 2024 include (but are not limited to) the CVEs below.

CVE-2023-20269 Cisco ASA and FTD Unauthorized Access Vulnerability	CVE-2022-47966 Zoho ManageEngine Unauthenticated Remote Code Execution	CVE-2023-35078 Ivanti Endpoint Manager Mobile Authentication Bypass
CVE-2023-42793 JetBrains TeamCity CI/CD Server Authentication Bypass	CVE-2023-3519 Citrix NetScaler ADC and NetScaler Gateway Unauthenticated Remote Code Execution	CVE-2023-35082 Ivanti Endpoint Manager Mobile Unauthenticated API Access
CVE-2023-0669 Fortra GoAnywhere MFT Remote Code Execution	CVE-2023-4966 Citrix NetScaler ADC/Gateway Buffer Overflow	CVE-2023-38831 RARLAB WinRAR Code Execution
CVE-2023-34362 Progress Software MOVEit Transfer SQL Injection	CVE-2023-24489 Citrix ShareFile Improper Access Control	CVE-2022-21587 Oracle E-Business Suite Remote Code Execution
CVE-2023-40044 Progress Software WS_FTP Server Deserialization of Untrusted Data	CVE-2023-22515 Atlassian Confluence Server and Data Center Broken Access Control	CVE-2023-24880 Microsoft SmartScreen Security Feature Bypass
CVE-2022-47986 IBM Aspera Faspex Unauthenticated Remote Code Execution	CVE-2023-22518 Atlassian Confluence Improper Authorization	CVE-2023-28252 Microsoft Windows Common Log File System Driver Elevation of Privilege
CVE-2023-27532 Veeam Backup & Replication Remote Code Execution	CVE-2023-47246 SysAid Path Traversal	CVE-2022-36537 ZK Framework Information Disclosure (ConnectWise R1Soft Server Backup Manager Remote Code Execution)
CVE-2023-32315 Ignite Realtime Openfire Path Traversal	CVE-2023-41265 Qlik Sense Enterprise HTTP Tunneling Vulnerability	
CVE-2023-27350 PaperCut NG Improper Access Control Vulnerability	CVE-2023-46604 Apache ActiveMQ Remote Code Execution	

2023 Initial Access Vectors

While we expect ransomware groups to continue exploiting new and known vulnerabilities in public-facing applications, there are still fundamental steps many organizations aren't taking to protect themselves from ransomware or other attacks. **41% of incidents** Rapid7 MDR observed in 2023 were the result of missing or unenforced multi-factor authentication (MFA) on internet-facing systems, particularly VPNs and virtual desktop infrastructure. This statistic has remained stable over the past year and a half, clocking in at **39% in 1H 2023** before rising slightly in the back half of the year.

Initial Access Vectors



Source: Rapid7 Incident Response data (January 2024)

Rapid7 MDR analysts have also noted that threat actors overall have accelerated their operational timelines, particularly during ransomware incidents and extortion attacks. While adversary dwell time varied substantially across incidents Rapid7 MDR has investigated the past year, our analysts have observed that it's becoming more common for attackers to move from initial access to data exfiltration in minutes or hours rather than days or weeks.

The end of this report contains practical security guidance to mitigate against the most common types of attacks Rapid7 incident responders see.

Life on the Edge: Network Pivots 2020 - 2024


When Rapid7 published its inaugural vulnerability intelligence report in 2020, our research team created a specific attacker utility category for vulnerabilities in technologies that functioned as **network pivots**, providing opportunities for external attackers to gain internal network access. VPNs, firewalls, security gateways, and most network appliances fall into this category, along with several other internet-facing technologies. These flaws are often used in conjunction with local code execution vulnerabilities or network protocol bugs to escalate privileges or move laterally across corporate networks.

Network edge technologies are essential for the operation of many modern networks, providing connectivity and security features for both corporate and home environments. But these devices also represent a significant weak spot in our cybersecurity defenses, as years of high-profile exploitation have demonstrated so clearly. Attackers of all stripes and motivations have incentive to target these devices, and ransomware groups and state-sponsored adversaries have both shown strong interest in n-day and zero-day exploit opportunities in these systems.

2023 saw a surge in attacks on network appliances by various threat actors, including state-sponsored groups and ransomware groups such as **CI0p, Inc, Bl00dy, Akira, Play, LockBit**, and more. Payloads shared publicly in 2023 spanned botnet **malware** and **cryptominers** alongside new **webshells** and **proprietary backdoors**.

Mass compromise events stemming from exploitation of network edge devices have almost doubled since the start of 2023, with 36% of all widely exploited 2023 vulnerabilities occurring in network perimeter technologies. Network edge CVEs are also disproportionately represented across zero-day attacks in our data – for the past three years, more than 60% of the vulnerabilities Rapid7 analyzed in network and security appliances each year were exploited as zero-days.


Network Pivot Exploitation Trends (2020 - 2024)



Network pivot vulnerabilities have comprised 24% of exploited vulnerabilities in our cumulative dataset for the past four years, and a quarter of all widespread threats Rapid7 researchers have tracked since 2020. For comparison, as best we can tell, approximately 19% of the CISA KEV consists of vulnerabilities in network edge devices or security gateways, about half of which were disclosed (and exploited) from 2020 onward – making network pivot prevalence in our data potentially a bit higher than in CISA’s data.

Below is the attacker utility distribution for the past four years of **widespread threats**:

Attacker Utilities 2020-2024 (Widespread Threats)




Detecting intrusions on network appliances is notoriously fraught. The capabilities for logging and threat detection vary widely across devices, depending on the manufacturer and model. Some devices may not log key events or be able to detect malicious activities effectively. The variety of firmware versions and operating systems – many of them proprietary to the manufacturer – also add complexity, as does the fact that firmware or the entirety of a product may be encrypted or obfuscated. Each device can require a unique approach to monitoring and protection, presenting challenges for security teams looking to develop a comprehensive strategy that covers the full spectrum of appliances in use.

Attacker Utilities

If you're joining us for the first time, we map two additional types of metadata in our vulnerability dataset in addition to threat status and time to known exploitation. The first piece of metadata our researchers define when analyzing **emergent threats** is **vulnerability class**, which is useful for making initial assessments about relative exploitability and available tooling. The second type of metadata we add is **attacker utility**, which describes what an attacker can hope to gain as a result of successful exploitation.


Network edge devices like VPNs and security gateways have been high-value, high-frequency targets in both zero-day and n-day attacks over the past few years. This is abundantly obvious from our 2023 dataset, in which network pivot vulnerabilities tower over other attacker utility categories in both sheer number and incident impact.

2023 Attacker Utilities



If we expand this to cover all annual intelligence report vulnerabilities 2020 onward, the distribution shifts somewhat but still maintains strong network pivot and IT security management compromise figures in addition to remote code execution:

Attacker Utilities and Threat Status 2020-2024



Generic remote code execution vulnerabilities (i.e., CVEs that don't map to a more specific bucket of utility, like supply chain compromise or smash-and-grab extortion) are predictably the largest category of attacker utility from the past four years, followed by network edge device compromise (network pivot). IT security management compromises, which map to CVEs that enable remote code execution or takeover of network or security management solutions, have been well-represented among zero-day exploits the past few years (e.g., Trend Micro Apex One [CVE-2023-41179](#), Ivanti Sentry [CVE-2023-38035](#), Zoho ManageEngine ADSelfService Plus [CVE-2022-28810](#)) and have also made up a fair share of widespread threats (e.g., Cacti [CVE-2022-46169](#), Zoho ManageEngine [CVE-2022-47966](#)).

While most of the vulnerabilities in this report affect on-premise technologies, we're starting to see cloud environment compromise show up more frequently in our threat analytics. Many of the CVEs we include in these reports also affect products that have cloud-based versions available (e.g., Atlassian Confluence [CVE-2023-22515](#) and [CVE-2023-22518](#)), though typical vendor advisory wording tends to imply that cloud-hosted versions of software are "not affected," presumably because updates have been applied universally and there's no


action the customer needs to take. As researcher Will Dormann **pointed out in 2023**, “not affected” is potentially a misnomer — there’s nothing magical about cloud-hosted software that prevents adversaries from compromising organizations with zero-day exploits before the vendor knows they exist, and to imply that orgs were never at risk from attacks that pre-dated patch availability is misleading at best.

Vulnerability Classes

As in previous years, our dataset for this report skews heavily toward server-side vulnerabilities in popular enterprise technologies, many of which have previously been targeted by a range of threat actors. The way we compile our data also means that certain vulnerability classes are less represented than others; for example, 75% of the vulnerabilities Google included in their 2023 **“in-the-wild Oday”** list arose from memory corruption issues, versus a mere 5% of the vulnerabilities in our data. If we were to expand our dataset to include mobile operating systems and client-side issues like browser flaws, our root cause distribution would undoubtedly change quite a bit.

Adversaries overall have demonstrated a clear penchant for well-understood, easily accessible vulnerability classes that support stable, reliable, and often trivial exploit development.

Vulnerability Classes and Threat Status 2020-2024




Fully 75% of the CVEs we've included in our annual vulnerability intelligence datasets the past four years have arisen from two superclasses of root cause: **improper access control** issues, like authentication bypasses, improper cryptographic implementations, and remotely accessible APIs; and **injection** flaws, which in our classification schema include root causes like server-side request forgery (SSRF) and improper input validation in addition to SQL injection and command injection. **Deserialization** vulnerabilities, despite having only a single CWE (CWE-502) to their name, are disproportionately represented among high-profile and widely exploited CVEs.

About 80% of the **memory corruption** CVEs Rapid7 researchers have analyzed since 2020 have been network edge device vulnerabilities and operating system-level vulnerabilities (i.e., Windows vulns) – most of them (58%) zero-days. This isn't terribly surprising when we consider that APTs tend to favor both memory corruption exploits and these types of target systems (i.e., OSes, network appliances). But plenty of APT and state-sponsored attacks the past few years have leveraged simpler, more easily exploitable vulnerabilities, too, like authentication bypasses or SQL injection issues in enterprise web applications.

Programming Language Distribution: 2023 Vulnerabilities

Rapid7 researchers looked at the programming languages used across 60 vulnerabilities from 2023 (see *2023 New Widespread Threats and Other Exploited Vulnerabilities* for 2023 CVEs). Among those 60 CVEs, only eight languages are represented (we could not determine language used for three CVEs). Since C and C++ are both memory-unsafe languages and are analyzed in similar ways during root cause analysis, they are combined as C/C++. Java and C/C++ account for two-thirds of the notable vulnerabilities Rapid7 researchers analyzed over the past 15 months, with C# and PHP distant runners-up. Ruby, Go, and Lua are outliers, with a single CVE apiece.

Languages of Primary 2023 CVEs




of Vulnerabilities (n=60)

Unsurprisingly, Java is the most common language we see, representing 43% of our 2023 CVE data – almost double that of the second most common language group, C/C++ (23%). Given Java’s continued popularity as an enterprise application development language, we’d expect Java and C/C++ to remain the most prominent languages among critical enterprise software vulnerabilities for the foreseeable future.

One programming language that’s notably absent from our dataset is Rust – a memory-safe, general purpose programming language that offers strong protections against memory corruption-based vulnerability classes, such as buffer overflows or use-after-free vulnerabilities. The adoption of Rust is still in its infancy, with limited (but increasing) amounts of production code being shipped in operating systems, web browsers, and so on. We haven’t seen any Rust-based enterprise software applications in production in the course of our research, but it’s a good bet that Rust will form an important part of secure software development in future years.

Vulnerability Classes and Language of Primary 2023 CVEs



If we look at the vulnerability classes affecting each language, we see **memory corruption** as a root cause only in a small percentage of C/C++-based vulnerabilities. While some of this is down to the way we compile our data, it also shows that attackers are favoring easier and more reliable vulnerability classes when exploiting enterprise software. At the top of those easy and reliable classes is **injection**, which is the most common superclass we see across all languages represented in our data. In fact, command injection vulnerabilities have been so common the past few years that Metasploit released a **new “fetch” payload** in 2023 to simplify command injection exploitation.

The second most common vulnerability superclass across almost all languages is **improper access control**. These are logical issues that allow an attacker to access resources they shouldn't be able to. Improper access control is largely language-agnostic – which makes it more difficult to mitigate holistically as a vulnerability class.

Deserialization appears primarily in C# and Java-based vulnerabilities and offers a good example of the ways that programming languages can help prevent entire classes of vulnerabilities. For example, .NET 8 has **disabled the BinaryFormatter** by default – this is often the root cause of deserialization vulnerabilities in C# code bases. Java 17 has also **broken many common deserialization gadgets** by introducing restrictions on reflective access. Enterprise software is historically slow to update major dependencies like Java, which in turn slows the uptake of language-specific protections that prevent deserialization vulnerabilities more broadly in the software ecosystem. Nevertheless, we would hope that deserialization vulnerability prevalence will decrease over time as software producers move to more modern versions of core dependencies with better runtime protections built in.

Government Guidance on Eliminating Key Vulnerability Classes

In the last few months, both **CISA** and the White House Office of the National Cyber Director (ONCD) have **published** missives on the importance of moving away from languages that are susceptible to memory safety issues (e.g., C/C++) in favor of memory-safe languages like Java, C#, Rust, and several others. This is beneficial advice at large, and may be in part directed at critical infrastructure providers whose risk profile has turned more volatile in recent



years; this guidance is also relevant for some network appliances that are written in C/C++. But it's also worth noting that the majority of widely exploited vulnerabilities we've seen since 2021 have been in products already written in memory-safe languages — namely Java and C#.

Non-memory-corruption-based vulnerabilities are more difficult to mitigate at the language level, meaning attackers will always have opportunities to exploit logical issues in software irrespective of the underlying language. In March 2024, CISA and the FBI released **a joint advisory** urging software vendors to perform formal reviews of their products' source code to detect and remediate SQL injection (SQLi) vulnerabilities. This, of course, is also generally good advice, particularly since SQLi vulnerabilities have starred in recent high-profile exploit chains like MOVEit Transfer **CVE-2023-34362** and Fortinet FortiClient **CVE-2023-48788** (the latter of which didn't make the data cut-off for this report). Many SQLi vulnerabilities are also considered “low-hanging fruit” for penetration testers, bug bounty hunters, and product security code reviewers.

With that said, the CISA KEV list appears to have incorporated a mere three SQLi vulnerabilities in 2023, and only one has made it onto the KEV as of March 2024. Rapid7 vulnerability intelligence datasets for the past three years have similarly only included three major SQLi flaws (SonicWall SMA 100 Series **CVE-2021-20016**, Progress Software MOVEit Transfer **CVE-2023-34362**, and Accellion FTA **CVE-2021-27101**), though admittedly all three were broadly exploited as zero-days, and Accellion FTA and MOVEit Transfer were used in ClOp mass extortion campaigns.

All in all, there seems to be some disparity between vulnerability classes that have been the focus of government guidance and vulnerability classes that are disproportionately represented and exploited in enterprise software, like command injection issues. Given CISA's (and other agencies') mission scope, however, it's entirely possible that enterprise software vulnerabilities are not the primary target of government-issued guidelines, even if certain categories of enterprise technology (e.g., network appliances) would also benefit from their application. Enterprise software manufacturers should heed government-issued guidance while also keeping other prevalent vulnerability classes in mind during development, static analysis, and security testing.

Practical Guidance for Defenders

Implement, test, and enforce MFA as a top priority. More than 40% of incidents Rapid7 investigated in 2023 were the result of missing or inconsistent enforcement of MFA, particularly on VPN, VDI, and SaaS products. Rapid7 MDR has also observed an uptick in **MFA push fraud** as a result of notification fatigue. Many MFA vendors offer **number matching** as a way to prevent MFA fatigue. Organizations should implement and enforce MFA everywhere it's supported; if it's not supported on key technologies your organization uses, consider asking your suppliers to make it a part of future roadmaps.

Don't trust, do verify. Organizations should apply principles of least privilege by moving to models that don't assume or imply trust — e.g., by making allowlisting the standard operating procedure rather than blocklisting, implementing granular access control, regularly reviewing (and removing!) users and access levels, and so on.

Be aggressive about continually reducing internet exposure. If it's on the internet, take it off if at all possible. Anything exposed to the public internet — ports, services, interfaces, concentrators, appliances, etc. — should be considered a



target for adversary enumeration and exploitation, as well as a potential target for zero-day attacks. Scan everything, externally and internally; if you can't take it off the internet, then look at vendor guidance on hardening configurations.

A strong proactive vulnerability risk management program is essential – in the cloud, on-prem, and everywhere in between. Robust vulnerability management is one of the foundations for any successful security program. Without the proactive discipline of vulnerability management and strong routine patch management practices, it can be prohibitively difficult for organizations to up-level to effective emergency patching in times of crisis. Having solid inventory and asset management practices is also essential, since it's difficult to act decisively in a crisis without knowing if the affected product is even present in an environment. Identify and catalog your critical and exposed systems, including security boundary devices, internet-facing load balancers, devops tooling and pipeline solutions, and virtualization infrastructure. For more fundamentals, read Rapid7's guidance on security program basics.

Implement “zero-day” patching procedures for critical technologies. Network edge devices are at particular risk of n-day and zero-day exploitation, and vulnerabilities in these devices should be mitigated as soon as vendor-provided patches or workarounds are available. Ensuring that logging is enabled and working will help security teams more effectively hunt for indicators of compromise and other suspicious activity during incidents.

Double (or triple) your backup strategy. Ransomware has reached pandemic levels worldwide, with healthcare and critical infrastructure at particularly high risk. Organizations can bolster ransomware readiness and resilience by maintaining multiple backups, including backups that are offsite and not connected to corporate networks. While this may not stop a ransomware attack, having uncompromised backups gives organizations more options during a major incident and may help deter business leaders from paying a ransom.

One of the most important ways organizations can mitigate catastrophic data theft and extortion risk in the face of a mature, highly motivated threat ecosystem is to **put measures in place to identify and prevent data exfiltration** wherever possible.

This includes (but is not limited to):

- Alerting on or restricting unusually large file uploads; looking for large volumes of traffic to a single IP or domain
- Monitoring for unusual access to cloud storage (ex: Google Drive, SharePoint, ShareFile)
- Monitoring for or adding firewall rules to block known file sharing sites (ex: filetransfer[.]io, anonfiles[.]com, mega[.]nz)

- Monitoring for data transfer utility presence or usage (ex: Filezilla, WinSCP/ Putty, MegaCMD, BITS, etc.)
- Monitoring for data archiving utility presence and usage (ex: 7zip, WinRAR, WinZip, etc.)
- Implementing egress filtering
- Restricting local admin privileges on hosts

Additional Resources

Rapid7 researchers and community members publish vulnerability analysis in Rapid7's open research platform, **AttackerKB**. These analyses often include sample proof-of-concept code and indicators of compromise in addition to exploitation timelines and attack chain analysis. To contribute or subscribe to Rapid7 notifications in AttackerKB, **[create a free account here](#)**.

Rapid7 zero-day vulnerability research is published on a regular basis **[here](#)**.

When a new threat arises, Rapid7 guidance can be found in the **[emergent threats](#)** section of the **[Rapid7 blog](#)**, along with corresponding information for Rapid7 customers. If you are a customer, we'd love to hear your feedback. You can contact your customer success manager (CSM) or technical account manager (TAM), or contact us at **research@rapid7.com**.

Appendix

The data upon which this report is based does not include all CVEs or threats Rapid7 evaluated in any given year, but it does represent a diverse sample of attacker use cases and exploitation case studies, with heavy emphasis on widespread attacks. Our intent is not to imply that any one CVE or vulnerability group is less important than others. Security teams, network administrators, and defenders at large have in-depth understanding of which assets are critical in their environments and how action taken may affect their business priorities. What we offer is an attacker-centric view of the vulnerability landscape that Rapid7 customers and the security community can use to inform the policies and practices that they employ as part of a larger defense-in-depth strategy.

Notes on Methodology

Most CVEs featured in this report were exploited in the wild in 2023 and/or the first two months of 2024. The CVEs we have categorized as exploited in the wild in this report are not the only vulnerabilities to have been exploited during 2023 and 2024. For example, we exclude many small business technologies (e.g., small business routers) that are frequently featured on in-the-wild lists. We also exclude most browser, mobile, and host-based vulnerabilities known to be exploited in the wild (e.g., bugs in Internet Explorer, Chrome, and Firefox, or bugs in iOS and macOS), in addition to exploited vulnerabilities in applications like Adobe Flash Player and Adobe Acrobat. Google Project Zero has a spreadsheet of some other zero-days exploited in the wild in 2023 [here](#), with more of a focus on browser, mobile, and host-based vulnerabilities.

Since the trustworthiness of our data is important, we cite **primary sources** wherever possible for vulnerabilities we've listed as exploited in the wild—that is, we reference firsthand accounts of exploitation from the organizations or individuals who detected, verified, and reported them. Examples of primary sources referenced throughout this paper include U.S. cybersecurity and intelligence agency alerts on known exploitation; security firm analyses of threats and indicators of compromise (IOCs) they've tracked during incident response or other investigations; and vendor advisories that specify exploitation in the wild (this includes CVEs that are disclosed as zero-days).

In the interest of readability, in some cases we also cite articles in security news publications that aggregate disparate reports of exploitation. This is especially useful when certain vulnerabilities are so widely exploited that it is helpful to see them contextualized in a single article. Our goal in citing news sources is to allow readers to understand the volume and impact of exploitation as quickly as possible.

Threat Categorization

In most cases, widespread threats are vulnerabilities under attack by many bad actors. Prior to January 2022, we categorized any CVE that was leveraged by ransomware operators as a widespread threat, since ransomware was generically considered to be an at-scale operation that relies on volume to profit. This policy was changed in 2022 to accommodate evolving statistics on **targeted ransomware**. From 2023 onward, instead of relying on the number of known attackers, we use the volume of **known real-world compromises** to differentiate widespread threats from other exploited vulnerabilities. This change was largely to account for the rising number of mass compromise events orchestrated by a single threat actor (e.g., ClOp).

Threats categorized as “exploited in the wild” are, quite simply, not known to be broadly exploited at time of writing. It is possible that concrete technical evidence of broader exploitation exists but has not been shared publicly. Likewise, while we do not have evidence at time of writing that CVEs in our impending threat category are exploited in the wild, lack of evidence does not mean absence of exploitation.

Ransomware Citations

We use security news articles frequently to document ransomware operators’ use of specific CVEs. Ransomware citations in this report are a binary – either there is credible **technical** evidence of ransomware groups’ usage of a vulnerability or there is not. Lack of confirmation does not mean a CVE has not been used in ransomware operations, only that we have not seen independently verifiable details supporting that conclusion. Credible sources typically include some combination of original analysis, news articles that aggregate primary sources to frame a larger story about risk, and expert commentary on open platforms (e.g., social media, user forums, and other public comments). In general, when a report comes from an individual or a little-known entity rather than a recognized expert, we look for technical information like payloads, observed post-exploitation behavior, threat actor attribution, IOCs, and/or attack chain analysis to support the claim.

Calculating Time to Known Exploitation (TTKE)

Compiling and communicating timelines is one of the most difficult parts of risk assessment. When calculating Time to Known Exploitation (TTKE), wherever possible we will attempt to use the first credible public reference to a vulnerability's existence and the first credible public reference to exploitation in the wild. Often the first and most authoritative source on the existence of a new CVE is a vendor advisory, but in this age of widespread zero-day exploitation and public discourse, community references can pre-date vendor bulletins. The initial two **ProxyNotShell** vulnerabilities ([CVE-2022-41040](#) and [CVE-2022-41082](#)) are an example of this. Rarely if ever do we use sources like the National Vulnerability Database (NVD) for disclosure baseline dates, since those dates tend to be days or weeks behind public (and therefore attacker) knowledge.

Important: The first known report of exploitation is just that—the first known report. It's possible, and in many cases likely, that exploitation began before a public analysis was released. TTKE data should not be taken as evidence that a vulnerability was NOT exploited before the observed date.

Glossary of Terms

Threat Statuses

Widespread threat: A vulnerability or other exploitable condition that is used to compromise many organizations across multiple industry verticals and/or geolocations.

Threat: In cyberthreat intelligence (CTI) parlance, a threat exists when there is an adversary with the intent, capability, and opportunity to act. In the context of vulnerability research, we use "threat" to denote a vulnerability or other attack vector that adversaries have used to exploit real-world production environments, but that has not necessarily been used to compromise large numbers of organizations. "Threat" can also refer to the entity doing the attacking, like an APT or ransomware group.

Impending threat: A vulnerability or other attack vector that is not yet known to be exploited in the wild, but that we believe is likely to be exploited in the future. Most impending threats in Rapid7 research are vulnerabilities in technologies that frequently come under attack (e.g., network edge devices) and/or have known exploits or other tooling available.

Attacker Utilities

Cloud environment compromise: Remote code execution or takeover of cloud accounts, cloud gateways, or API management products. Replaces cloud infrastructure compromise as of 2023.

Information disclosure: The ability to leak sensitive data (e.g., credentials or other secrets, environment variables) and/or enumerate files on a target. Information disclosure vulnerabilities typically do not give an attacker a path to code execution by themselves, but instead function as primitives that enable a secondary part of an exploit chain (e.g., remote code execution). Can aid in turning a post-authentication vulnerability into a pre-authentication vulnerability.

IT security management compromise: Remote code execution or takeover of network, device, and/or endpoint management technologies; remote code execution or takeover of identity and access management solutions, including single sign-on (SSO) and Active Directory (AD) management solutions, or other security products. Replaces network infrastructure compromise as of 2022.

Network pivot: The ability to pivot from an external network to an internal network, most often by exploiting internet-facing systems such as VPNs, firewalls, routers, and other gateway devices. A network pivot gives an attacker visibility into both internal and external traffic and aids in data exfiltration, traffic sniffing, and further attacks within the target network.

Privilege escalation: The ability to run code locally on a system to which the attacker already has some access. Most commonly used to escalate privileges (e.g., by executing code as the user running the vulnerable application). Replaces local code execution as of 2023.

Remote code execution (RCE): Code execution on a remote target. Typically refers to the ability to execute a payload on a target system (e.g., obtain a shell session). Aids in credential stealing, data exfiltration, and so on.

Smash-and-grab extortion: New attacker utility in 2023 data onward; describes a vulnerability in a file sharing platform or secure/managed file transfer application (cloud or on-premises) that allows adversaries to gain access to the target system and quickly exfiltrate potentially large amounts of data. Often targeted in ransomware and extortion attacks.

Social engineering: Encompasses vulnerabilities that typically require a user to click on or preview a malicious attachment for successful exploitation to occur (e.g., email-based or document-based attacks).

Supply chain compromise: New attacker utility in 2023 data onward; refers to a vulnerability that allows for takeover of CI/CD pipelines or other critical software supply chain infrastructure, or a verified incident in which a legitimate piece of software was backdoored or otherwise maliciously modified.

Vulnerability Classes

Deserialization is the process through which an application is able to convert data from a portable format to data types native to its own language. Many modern languages support deserialization, including Java, .NET, Python, and Ruby. The deserialization process can pose a threat to security when the data that is loaded into the native language can be tampered with by a malicious party. Typical attacks involve configuring the data to invoke a method with the arguments necessary to execute an operating system command. This results in command execution in the context of the loading application. Common solutions to this security problem include cryptographically signing the data to ensure its authenticity and utilizing an allowlist of data types that are permitted to be loaded. Associated CWEs: [CWE-502](#)

Improper Access Control refers to a missing or insufficient access control to a particular interface into a system (most often a remotely accessible API). Improper uses of cryptography for the purpose of authentication also fall under this vulnerability class. Common solutions to this problem include proper authentication, authorization, and accounting implementations for all sensitive interfaces, as well as secure management of all related secrets. A non-exhaustive list of associated CWEs: [CWE-285](#), [CWE-200](#), [CWE-287](#), [CWE-732](#)

Memory Corruption is a large category of vulnerabilities that involve the misuse of data through a variety of means to alter memory and produce unexpected behavior. This vulnerability class includes improper boundary enforcement, type confusion, uninitialized data use, and the use of data after it has been freed, to name a few. These vulnerabilities often manifest themselves in languages that are not considered memory-safe. Successful exploitation of memory corruption vulnerabilities can result in arbitrary code execution within the context of the running application, or in an unhandled exception that causes the application to crash and triggers a denial of service (DoS) condition. Common solutions to this problem typically involve additional validation on parameters to key operations, such as those used to load and store data. Successful exploitation of these classes of vulnerabilities has become more complex in recent years due to the variety of mitigation technologies that have been developed for operating systems, compilers, and applications (e.g., kASLR, Control Flow Guard, win32k Type Isolation). A non-exhaustive list of associated CWEs: [CWE-787](#), [CWE-125](#), [CWE-416](#), [CWE-190](#), [CWE-476](#)

Injection is a large category of vulnerabilities involving specially crafted input that is interpreted in a particular way by an associated system. Most commonly seen in web applications, injection attacks are often more specifically labeled by the type of data being interpreted (e.g., SQL, LDAP, OS commands). The root cause of these vulnerabilities is almost always insufficient sanitization on data received from a malicious party. Exploitation of these vulnerabilities tends to be reliable, rarely resulting in service degradation unless intended (such as through SQL or OS commands).

The context under which the logic is executed typically depends on how it is interpreted. In the case of a web application, for example, SQL injection may be executed on a back-end database server, while OS commands are injected on the front-end web server, and JavaScript is executed by the end user's browser. This class of vulnerabilities is therefore unique in that it commonly involves a vulnerability in one system compromising the integrity of others. Common solutions to this problem typically involve implementing strict sanitization on parameters through the use of allowlists. A non-exhaustive list of associated CWEs: **CWE-79**, **CWE-20**, **CWE-89**, **CWE-94**

References

Our sincere thanks to Rapid7 research project manager Cynthia Wyre for compiling this list.

Security research is a community pursuit. This report benefited from the work of many individual researchers and research teams, including but not limited to the work of the folks listed below:

[Adobe Security Bulletin \(2023\)](#)

[Adversary Tactics and Intelligence Team, Deepwatch \(2023\)](#)

[Alexander Martin, The Record, Recorded Future News \(2023\)](#)

[Alexander Marvi, Brad Slaybaugh, Ron Craft and Rufus Brown, Mandiant \(2023\)](#)

[Alexander Marvi, Shawn Chew, and Punsaeen Boonyakarn, Mandiant \(2024\)](#)

[Alex Delamotte, SentinelOne \(2023\)](#)

[Alex Delamotte and Christian Vrescak, SentinelOne \(2023\)](#)

[American Pharmacists Association \(2024\)](#)

[Austin Larsen, John Palmisano, John Wolfram, Mathew Potaczek, and Matthew McWhirt, Mandiant \(2023\)](#)

[Barracuda \(2023\)](#)

[Becky Bracken, Dark Reading \(2023\)](#)

[Benoit Sevens, Google \(2023\)](#)

[Bill Toulas, Bleeping Computer \(2024\)](#)

[Bill Toulas, Bleeping Computer \(2024\)](#)

[boB Rudis, GreyNoise \(2023\)](#)

[Boris Larin, Kaspersky Securelist \(2023\)](#)

[Brendan Watters, Rapid7 \(2023\)](#)

[Caitlin Condon, Infosecurity Magazine \(2024\)](#)

[Caitlin Condon, Rapid7 \(2023\)](#)

[Cara Lin, Fortinet \(2023\)](#)

[CERT Polska, NASK \(2023\)](#)

[Christiaan Beek, Rapid7 \(2024\)](#)

[Ciaran Martin, Ciaran's Crispy Cogitations \(2024\)](#)

[CISA Cybersecurity Advisory \(2023\)](#)

[CISA Cybersecurity Advisory \(2024\)](#)

[CISA Cybersecurity Alerts \(2023\)](#)

[CISA Cybersecurity Alerts \(2023\)](#)

[CISA Cybersecurity Alert \(2023\)](#)

[CISA Cybersecurity Alert \(2023\)](#)

[CISA Cybersecurity Alert \(2023\)](#)

[CISA Cybersecurity Alert \(2023\)](#)
[CISA Fact Sheet \(2022\)](#)
[CISA Known Exploited Vulnerabilities Catalog \(2023\)](#)
[Cisco Talos \(2023\)](#)
[Cisco Talos \(2023\)](#)
[Clayton Zechman, Rapid7 \(2023\)](#)
[Clement Lecigne and Maddie Stone, Google Threat Analysis Group \(2023\)](#)
[Daniel Lydon and Conor Quinn, Rapid7 \(2023\)](#)
[Daniella Silva and Aria Bendix, NBC News \(2024\)](#)
[Devna Bose, The Associated Press \(2024\)](#)
[Dietrich Knauth, Reuters \(2024\)](#)
[Drew Burton, Rapid7 \(2023\)](#)
[Drew Burton, Rapid7 \(2023\)](#)
[F5 \(2023\)](#)
[FBI Internet Crime Report \(2023\)](#)
[Francesco Figurelli and Eduardo Ovalle, Kaspersky Securelist \(2023\)](#)
[Fortra \(2023\)](#)
[Fortinet PSIRT \(2023\)](#)
[Fortinet PSIRT \(2023\)](#)
[Glenn Thorpe, Rapid7 \(2023\)](#)
[Glenn Thorpe, Rapid7 \(2023\)](#)
[Global Threat Intelligence, Fox-IT \(2023\)](#)
[Google Threat Analysis Group \(2023\)](#)
[Hans-Martin Münch, Mogwai Labs \(2023\)](#)
[Health Sector Cybersecurity Coordination Center \(HC3\), U.S. Department of Health and Human Services \(2023\)](#)
[Huntress \(2023\)](#)
[Ionut Ilascu, Bleeping Computer \(2023\)](#)
[Ivanti \(2023\)](#)
[Ivanti \(2023\)](#)
[Ivanti \(2023\)](#)
[Jacob Baines, VulnCheck \(2023\)](#)
[James Nugent, Foti Castelan, Doug Bienstock, Justin Moore, and Josh Murchie, Mandiant \(2023\)](#)
[James Sadowski and Casey Charrier, Mandiant \(2023\)](#)
[Jeff Johnson, Fred Plan, Adrian Sanchez, Renato Fontana, Jake Nicastro, Dimiter Andonov, Marius Fodoreanu, and Daniel Scott, Mandiant \(2023\)](#)
[Jonathan Greig, The Record, Recorded Future News \(2024\)](#)
[Joint Cybersecurity Advisory, CISA \(2023\)](#)
[Joint Cybersecurity Advisory, CISA \(2024\)](#)
[Kaspersky \(2024\)](#)
[Kate Morgan, Google Threat Analysis Group \(2023\)](#)
[KFF Health News, U.S. News and World Report \(2024\)](#)
[Lawrence Abrams, Bleeping Computer \(2022\)](#)
[Levi Broderick \(GrabYourPitchForks\) \(2020\)](#)
[Malpedia \(retrieved 2024\)](#)
[Mark Ellzey, Censys \(2023\)](#)
[Margaret Zimmermann, Palo Alto Unit 42 \(2023\)](#)
[Mathew J. Schwartz, Bank Info Security \(2023\)](#)
[Mathew J. Schwartz, Bank Info Security \(2023\)](#)

<u>Matthew Meltzer, Robert Jan Mora, Sean Koessel, Steven Adair, and Thomas Lancaster, Volexity (2024)</u>	<u>Rapid7 AttackerKB (2024)</u>
<u>Microsoft Defender Threat Intelligence (2023)</u>	<u>Rapid7 AttackerKB (2023)</u>
<u>Microsoft Incident Response, Microsoft (2023)</u>	<u>Rapid7 AttackerKB (2023)</u>
<u>Microsoft Threat Intelligence, Microsoft (2023)</u>	<u>Rapid7 AttackerKB (2023)</u>
<u>Microsoft Threat Intelligence, Microsoft (2023)</u>	<u>Rapid7 AttackerKB (2023)</u>
<u>Microsoft Threat Intelligence, Microsoft (2023)</u>	<u>Rapid7 AttackerKB (2023)</u>
<u>Microsoft Threat Intelligence, Microsoft (2023)</u>	<u>Rapid7 AttackerKB (2023)</u>
<u>Microsoft Threat Intelligence, Microsoft (2023)</u>	<u>Rapid7 AttackerKB (2023)</u>
<u>Microsoft Threat Intelligence, Microsoft (2023)</u>	<u>Rapid7 AttackerKB (2023)</u>
<u>Microsoft Threat Intelligence, Microsoft (2023)</u>	<u>Rapid7 AttackerKB (2023)</u>
<u>Microsoft Threat Intelligence, Microsoft (2023)</u>	<u>Rapid7 AttackerKB (2023)</u>
<u>Mike Whitehead (2023)</u>	<u>Rapid7 AttackerKB (2023)</u>
<u>MITRE (2017)</u>	<u>Rapid7 AttackerKB (2023)</u>
<u>MITRE (2022)</u>	<u>Rapid7 AttackerKB (2023)</u>
<u>Norwegian National Security Authority (2023)</u>	<u>Rapid7 AttackerKB (2023)</u>
<u>Office of Public Affairs, U.S. Department of Justice (2023)</u>	<u>Rapid7 AttackerKB (2023)</u>
<u>Office of Public Affairs, U.S. Department of Justice (2023)</u>	<u>Rapid7 AttackerKB (2023)</u>
<u>Office of Public Affairs, U.S. Department of Justice (2023)</u>	<u>Rapid7 AttackerKB (2023)</u>
<u>Office of the Spokesperson, U.S. Department of State (2024)</u>	<u>Rapid7 AttackerKB (2023)</u>
<u>Peter Gimus, Trend Micro (2023)</u>	<u>Rapid7 AttackerKB (2023)</u>
<u>Project Zero, Google (2019)</u>	<u>Rapid7 AttackerKB (2023)</u>
<u>Rapid7 (2020)</u>	<u>Rapid7 AttackerKB (2023)</u>
<u>Rapid7 (2021)</u>	<u>Rapid7 AttackerKB (2023)</u>
<u>Rapid7 (2023)</u>	<u>Rapid7 AttackerKB (2023)</u>

TAKE COMMAND OF THE ATTACK SURFACE FROM ENDPOINT TO CLOUD

About Rapid7

Rapid7 is creating a more secure digital future for all by helping organizations strengthen their security programs in the face of accelerating digital transformation. Our portfolio of best-in-class solutions empowers security professionals to manage risk and eliminate threats across the entire threat landscape from apps to the cloud to traditional infrastructure to the dark web. We foster open source communities and cutting-edge research—using these insights to optimize our products and arm the global security community with the latest in attacker methodology. Trusted by more than 11,000 customers worldwide, our industry-leading solutions and services help businesses stay ahead of attackers, ahead of the competition, and future-ready for what's next.



PRODUCTS

Cloud Security
XDR & SIEM
Threat Intelligence
Vulnerability Risk Management

Application Security
Orchestration & Automation
Managed Services

CONTACT US

[rapid7.com/contact](https://www.rapid7.com/contact)

To learn more or start a free trial, visit: <https://www.rapid7.com/try/insight/>

The information provided in this report is intended for informational purposes only and Rapid7 makes no warranties, express or implied, regarding the suitability of the content for any specific purpose. The content within this report is based on data and findings available up to the date of its publication, which is mentioned within the document.

The information contained herein is provided "as is," and readers are advised to use their own discretion when applying the information to their specific situations. Furthermore, any third-party sources, tools, or software mentioned in this report are included for informational purposes only. Rapid7 does not take responsibility for the accuracy, functionality, or security of these external resources.

Rapid7 is not liable for any damages, losses, or consequences that may arise from the use of the information provided within. This includes but is not limited to direct, indirect, incidental, or consequential damages related to actions taken based on the content of this report.

Any reproduction, distribution, or unauthorized use of this report's contents without explicit permission from the authors and publishers is strictly prohibited.