



2026 State of the Internet Report

The Vast and Volatile Internet



Table of Contents

About This Report	3
Executive Summary	4
Introduction	6
The Shape of the Internet: Ports, Protocols, and Certificates	8
The Long Tail of Vulnerability Response: How Exposed Software Ages	17
Exposure of Industrial Control Systems: Critical Infrastructure on the Internet	34
Threat Infrastructure: An Internet-Wide View of Adversary Tools	47
Exposed AI Infrastructure: The Fast-Growing Attack Surface	62
Conclusion	77
Glossary	80
References	84
AI Disclosure	89

About This Report

A persistent problem in Internet security is that organizations often cannot see their own external network footprint as clearly as an outsider can. Software sits behind cloud providers and CDNs, infrastructure is scattered across vendors, subsidiaries, contractors, and legacy environments. A system can be plainly reachable from the public Internet while its owner remains unclear, sometimes even to the organization responsible for it. Meanwhile, threat actors exploit these same obscuring factors to conceal their own infrastructure, making emerging threats difficult to track and detect.

For defenders, gaps in visibility become gaps in risk assessment and response.

Censys provides an outside-in view of the attack surface. We scan the public Internet continuously, tracking what appears, what disappears, how quickly technologies spread, and which exposures persist for years. This year, we apply that lens across ports and protocols, vulnerability response, industrial control systems (ICS), threat infrastructure, and AI, creating a baseline that shows what is exposed and how it is changing. **Read this report to benchmark your own attack surface** against these broader Internet trends and understand the larger Internet landscape to better anticipate emerging threats.

Executive Summary

1 The Shape of the Internet

Things do not always live on defaults on the Internet. While **infrastructure and management protocols** like DNS, NTP, and SNMP do run almost entirely on their IANA-assigned ports, **access and application** services tend to drift widely: 34% of SSH, 68% of HTTP(S), and 96% of SOCKS proxies run on non-standard ports. This is why simply knowing what's exposed in your network is still so difficult for many security practitioners: detection methods and asset inventories typically only look in default places.

The TLS certificate ecosystem is more distributed than one might expect. Let's Encrypt is the most common CA among certificates observed on Internet-facing hosts, accounting for 24%, with DigiCert, Amazon, and Google Trust Services making up much of the rest. At the same time, **a different certificate model is gaining popularity:** Let's Encrypt made **IP address certificates** generally available for the first time in January 2026, and by June over 60,000 were observed on hosts in Censys.

2 Vulnerability Response

The oldest software seems to be the least likely to ever update. Across four case studies we tracked over two years, the very oldest releases became less prevalent overall, but **middle-aged software** only a few years out of date saw the largest share of updates. The exposure of one enterprise reporting product's 4-year-old release fell 53 percentage points, while its oldest (8-year-old) release budged less than 1 percent.

3 Industrial Control Systems and Critical Infrastructure

As of 2026, North America continues to lead in ICS exposures, representing roughly 38% of the global ICS landscape. We note that Asia's ICS footprint has expanded from 22% of exposures in 2024 to 25% in 2026. In contrast, Europe's proportion of ICS exposure has decreased, from 36.5% of global total in 2024 to 32% in 2026.

Over 80% of ICS hosts also run at least one non-ICS IT service like HTTP, SSH, or FTP. These services increase the attack surface of ICS hosts and introduce more avenues to exploitation. A threat actor doesn't necessarily have to understand the inner workings of Modbus or EIP if a vulnerable web interface or SSH service is exposed on the target host.

Roughly 70% of hosts running ICS devices and services globally are consistently found on consumer and mobile networks. The WHOIS and other host metadata typically point to the telecommunications provider, not the operator. It's trivially easy to discover exposed ICS assets on the Internet, but determining ownership and notifying the owner of exposure is often impossible.

4 Threat Infrastructure and Adversary Tools

China and the U.S. host a plurality of the world's malicious software and services, but the profiles of what they host diverge significantly. China is home to 36.5% of malicious host-level infrastructure like C2s and remote access tooling in contrast to the U.S.'s 16.8%. The U.S., conversely, hosts 45.5% of web-based threats like phishing and other deceptive pages, while China hosts only 20.9% of web-based threats observed by Censys. This suggests operational differences in where actors host their wares.

5 Exposed AI Infrastructure

The Internet-facing AI attack surface is growing quickly and spans far more than models and LLMs. By July 2026, we observed 267,915 unique hosts exposing AI/LLM-related tooling, **up 63% in nine months** from 164,088 in October 2025.

A recognizable AI infrastructure stack is emerging on the public Internet. The exposed systems we detected spanned a wide range of **20+ functional categories**, all forming the layers needed to build and operate AI workflows. This includes model inference, orchestration, LLM gateways, user interfaces, RAG systems, vector databases, observability, and other tools for supporting development infrastructure.

The most widely exposed products are also among the most concerning. n8n, the open-source workflow and AI-agent automation platform, is the most exposed product in our dataset by a wide margin, with **113,926** exposed hosts observed in Censys. Of the hosts running a parseable version, **over three quarters** (76%) fall within the affected range of at least **one recent high-severity vulnerability** (CVSS 9.0+). n8n has weathered a steady stream of critical vulnerability disclosures and supply chain attacks over the past year. We predict that tools like this sitting at the orchestration layer, meaning the platforms that wire agents together and store credentials for connected external services, are becoming a new crucial asset class that will demand stronger hardening and trust practices.

Introduction

At Censys, we continuously scan the public Internet to understand what is exposed and how that exposure changes over time. We scan the entire public IPv4 address space across all 65,535 ports, leverage redirects and the Domain Name System to discover IPv6 systems, and detect more than 200 protocols. We also maintain the largest collection of TLS certificates, drawn from Certificate Transparency (CT) logs and discovered through our own scans. At the time of writing, we have indexed **218 million hosts, 17 billion certificates, and 2.4 billion web properties**.

Maintaining an accurate inventory of Internet-facing assets remains one of the most challenging barriers to securing them. Censys gives us a unique birds-eye view into what is hosted on the Internet and where; and because we collect this data continuously over time, we can do more than just portray a snapshot. We can measure **how the Internet changes, which exposures persist, and how it responds to new technologies and security events as a whole**.

Everything analyzed in this report is visible on the public Internet. We never log into or otherwise improperly access any device or service; we analyze only what is visible from external Internet scanning. Some may find the massive number of connected devices and services visible to scanners surprising, but the Internet is vast — and anything Censys can see, a threat actor conducting reconnaissance can see as well.

Internet exposure of an asset does not, by itself, prove that the asset is insecure or vulnerable. Our measurements cannot tell us whether a host is exploitable, whether compensating controls exist, or whether a system has been compromised.

However, in some specific cases, such as industrial control systems (ICS) and critical infrastructure, Internet exposure itself is a vulnerability, as **these systems should never be exposed to the public Internet**.

Previous editions of this report have focused on a single facet of our scan data, such as the web or ICS. This year, we're taking a broader view across five key areas of Internet exposure with the goal of establishing measurements we can revisit year over year:

-
- 1 The Shape of the Internet:
Ports, Protocols, and Certificates**
 - 2 The Long Tail of Vulnerability Response:
How Exposed Software Ages**
 - 3 Exposure of Industrial Control Systems:
Critical Infrastructure on the Internet**
 - 4 Threat Infrastructure: An Internet-Wide
View of Adversary Tools**
 - 5 Exposed AI Infrastructure:
The Fast-Growing Attack Surface**
-

The first section examines the Internet as a whole, and each section that follows narrows in on a specific category of exposure and what defenders can learn from it.

The Censys Internet Map





Figure 1. Censys's view of hosts exposed on the public Internet (July 2026)

Data via Censys

1

The Shape of the Internet

Ports, Protocols, and Certificates

Censys data allows us to ask questions about the vastness and volatility of the Internet: we can study vulnerability remediation timelines, AI tooling exposures, and where threat infrastructure lives. All of these applied questions are built on the core Censys dataset, comprising hosts, websites, ports, protocols, and certificates. Here we offer a broad but brief characterization of the ports, protocols, and TLS ecosystem that inform the rest of the analysis in this report.

Port Distribution of 25 Most Common Services




Figure 2. These are the most commonly observed services, sorted by highest default port. Those without an IANA-assigned port are located at the top of the graph.

Data via Censys

Ports + Protocols

Protocols typically have one or more standard ports assigned by IANA, the Internet Assigned Number Authority, such as port 80 (TCP) for plain HTTP, port 443 (TCP) for TLS-encrypted HTTP (HTTPS) or port 53 (both UDP and TCP) for DNS. **Internet infrastructure and management protocols like DNS, NTP, and SNMP run almost entirely on their IANA-assigned ports (99.5%, 100%, 99.9% respectively)**: just over 300 out of 5.7 million hosts serving the NTP protocol, used for clock synchronization, do not run the service on its IANA-assigned port 123.

This makes sense: for example, when setting up a DNS client on a device, unless otherwise specified, DNS will assume port 53. The small minority of hosts running DNS (or NTP or SNMP) on a non-standard port are likely test set-ups or may be misconfigured.

For SMTP the situation is slightly different: while some use cases automatically assign the standard port, others offer users a choice, which increases the instances of SMTP using non-standard ports. The process of looking up a mail server for an email address (through the DNS MX record) doesn't involve the choice of a port, so for that use, SMTP always uses TCP port 25. However, for email submission by a mail client to an email provider's mail server (which will then handle the further delivery of the email), a port can be chosen in the mail client and providers often offer additional ports to bypass port blocking. This is why, other than port 25 (44%) and the IANA-assigned submission ports 465 (16%) and 587 (29%), we see 11% of hosts running SMTP on a different port.

Interestingly, POP3 and IMAP, protocols used by mail clients to download emails, similarly allow for manual port selection within a mail client. However, both are largely seen on their IANA-assigned ports (POP3 at 99.1% and IMAP at 99.3%).

The situation is very different for access and application protocols. For example, an SSH client can choose any port

when connecting to a server and in HTTP, a non-standard port can be set in the URL ([http://www\[.\]example\[.\]com:1234/](http://www[.]example[.]com:1234/)). This explains why 34% of SSH instances are on a non-standard port (not port 22), along with 68% of HTTP(S) services on ports other than 80 and 443.

For SOCKS proxies, the situation is even more extreme. SOCKS proxies have effectively abandoned their IANA port: **96% of SOCKS services we observe are not on port 1080** (see queries at end of section). The de facto standard is now 10801 (34%), followed by ports 43080 (5.9%) and 21 (5.7%), the latter being the IANA-assigned port for FTP. A scanner or detection rule checking only port 1080 for SOCKS proxy exposure would thus be operating on 4% of the actual population.

An interesting case is CWMP, the protocol used to auto-configure and monitor customer-premises equipment (CPE) such as routers, modems, and network gateways. A small minority of the hosts running this protocol run it on the IANA-assigned port 7547. Outside this standard, port choices reflect regional and carrier differences: Telecom Italia operates 97% of CWMP on port 7170, Bell Canada runs 99.7% of CWMP on ports 50011–50014, and ASMedi (Morocco) and Orange's West African subsidiaries (including Sonatel) together hold 94% of CWMP on port 50805. In practice, port number can be a sufficient selector for devices that a given carrier manages via CWMP.

96%

of SOCKS services we observe
are not on port 1080

Port Distribution by Protocol and Continent




Figure 3. This chart represents port distribution of the 25 most commonly observed services across different continents. Data via Censys

The horizontal axis maps the port number on a logarithmic scale. Darker circles on each subplot represent IANA-assigned ports, while lighter circles (regardless of size) represent non-standard ports.

TLS + Certificates

Over the past fifteen years, browser vendors, certificate authorities, and the CA/Browser Forum have pushed Internet traffic toward encryption to prevent its interception by hostile actors. TLS (Transport Layer Security) plays a major role in this by encrypting many Internet protocols, HTTP in particular: a URL that starts with <https://> uses HTTP on top of a TLS-encrypted connection.

TLS doesn't just take care of the encryption; it also ensures no adversary in the middle can intercept the traffic by decrypting it and then re-encrypting it. It does so through public-key certificates.

When a client attempts a TLS connection, the server presents a certificate and cryptographically signs a specific string with the corresponding private key. The client verifies this signature, verifies that the certificate is valid for the hostname (or, occasionally, IP address) it is connecting to, and verifies the certificate is cryptographically chained to a certificate in its own "store" of root certificates.

Those root certificates are (sometimes indirectly) linked to entities called Certificate Authorities (CAs), who are the ones issuing the certificates. This used to be fairly costly, which is why only a small minority of websites used HTTPS. However, Let's Encrypt and a few other CAs issue certificates for free and automate the renewal process (certificates are only valid for a limited time period).

Certificate Authority Market Share

It may seem surprising that, despite offering free certificates and renewal automation, Let's Encrypt accounts for less than a quarter of market share. Of the certificates seen on hosts, only 23.9% were issued by Let's Encrypt, followed by DigiCert (14.7%), Amazon (13.8%), Google Trust Services (10.1%), and Sectigo (7.5%).

This is interesting alongside observations of how big Amazon's footprint is: these certificates are primarily on Amazon's hosts (Amazon Web Services in particular). Similarly, the majority of Google-issued certificates are found on Google hosts. DigiCert and Sectigo certs are issued to hosts on other providers or CDNs, such as Akamai (whose customers [use DigiCert certificates](#), in addition to Let's Encrypt).

Certificate Authority Market Share (Web PKI)




Figure 4. Let's Encrypt leads CA market share, providing nearly a quarter of the certificates observed on hosts.

Data via Censys

Certificate Signing Algorithms

One does not need a certificate authority to generate a certificate; this can be done using a simple command. A browser won't trust these so-called self-signed certificates, but they can still be used as a placeholder (they are often installed by default on a web server), or in some cases a client can be explicitly told to trust them. Nearly a quarter (24.1%) of the certificates on hosts we see are self-signed, while only 52.4% have a certificate from a publicly trusted CA.

The remaining 22.9% are device or appliance certificates from customer-premises equipment (CPE), network appliances, and NAS devices. These are typically hosted on services that such devices regularly connect to: the certificates are signed by the manufacturers themselves and trusted by the devices. If someone were to somehow connect to the hosts from a web browser, they would see an error.

SHA256-RSA remains the dominant signing algorithm for browser-trusted certificates, found on 61.7% of certificates observed in this analysis. It's followed by ECDSA-SHA384 at 23.3%, ECDSA-SHA256 at 13.4%, and SHA384-RSA at a distant 1.5%.

Certificate Signing Algorithms: Browser-Trusted Leaf Certificates




Figure 5. SHA256-RSA is the most common signing algorithm for browser-trusted certificates, used on nearly two thirds of certificates observed on hosts.

Data via Censys

Certificate Validation Level Trends

At the time of analysis, only 0.02% of the publicly trusted CA-issued certificates were so-called Extended Validation (EV) certificates. These certificates are issued after a longer (and more expensive) vetting process. Initially, web browsers added a green URL bar to indicate a site using an EV certificate. However, after [research](#) indicated users didn't notice it, the UI artifact was abandoned in 2018 and 2019. Their very low prevalence in 2026 suggests that organizations have deemed EV certificates all but obsolete.

Certificate Validation Level: Publicly Trusted CA Certs

1 square $\approx$ 16M certs




Figure 6. Low user adoption has rendered Extended Validation certificates all but obsolete.

Data via Censys

IP Based Certificates

As of September 2, 2026, Censys had more than 17 billion certificates stored in its database. Most of these certificates are valid for one or more hostnames, but certificates can also be issued for an IP address. This is rarer since addresses get reassigned far more often than names do, and many CAs have historically declined to issue for them at all.

In most cases, there is no need for an IP address itself to have a TLS certificate. People and devices almost always connect to a hostname, and if that connection is intercepted or misrouted (through a [BGP hijack](#), for example) the hostname's certificate already prevents adversary-in-the-middle attacks. However, there are some cases where a connection to an IP address needs to be secured. Let's Encrypt provides [several examples](#) of use cases for IP address certificates, including:

- ✦ Encrypted access to websites that don't have a domain name.
- ✦ Securing access to home IoT devices that don't have a domain name.
- ✦ Securing connections between backend cloud servers.

Let's Encrypt [issued](#) its first certificate for an IP address in its [staging environment](#) in July 2025. Such certificates were made [generally available](#) on January 15, 2026. Our data shows that the number of certificates issued for IP addresses began growing ten days earlier, on January 5, 2026 (see queries at end of section). That day, the number of IP address certificates recorded by Censys rose from 3,878 to 9,793, while the number of them observed being used on a host increased from 1,095 to 3,380.

As the chart below shows, the number of such certificates showed a gradual increase, exceeding more than 100,000 valid certificates on most days in June 2026 (over 60,000 of which were observed in use on a host). More than 99% of these were issued by Let's Encrypt. The four-day cycle starting early May is caused by a one-off peak of certificates for new IP addresses followed by a four-day renewal cycle.

Daily Issuance of Certificates for IPv4 Addresses, First Half of 2026




Figure 7. Certificate issuance for IPv4 addresses grew gradually over the first few months of 2026. By June 2026, daily issuance exceeded 100K certificates on most days; 99% of these were issued by Let's Encrypt.

Data via Censys

Takeaways

Security tooling that assumes devices and services run on their default ports misses a non-trivial subset of Internet-exposed assets.

34% of SSH services run on a port other than default 22; 68% of HTTP services run on ports other than 80 and 443; 96% of SOCKS services run on a port other than 1080. However, infrastructure and management protocols, like DNS, NTP, and SNMP, are the exception—99.5%, 100%, and 99.9% (respectively) of these services are found on their IANA-assigned ports.

Explore the Data in Censys Platform

[QUERY 1] `host.services: (protocol = "SOCKS" and not port = 1080)` [🔗](#)

[QUERY 2] `cert.ever_seen_in_scan = "true" and cert.parsed.issuer_dn: "Let's Encrypt" and cert.names =~ "^[0-9]{1,3}\\.[0-9]{1,3}\\.[0-9]{1,3}\\.[0-9]{1,3}$"` [🔗](#)

To filter to just the certificates generated on a specific day:

`and cert.parsed.validity_period.not_before = "yyyy-mm-dd"`

2

The Long Tail of Vulnerability Response

How Exposed Software Ages

Vulnerability research typically focuses on the first few weeks after an incident appears in the news cycle. Censys [Rapid Response](#), our research function tracking the Internet-facing impact of emerging vulnerabilities, works on that timescale too: it captures what exposed populations do within days of an incident, when hosts change fastest and the security community is paying attention. What happens afterward is much less studied. Does software eventually upgrade over months or years? How quickly? And how much do those upgrades seem to coincide with public security incidents associated with the product, like disclosures of new vulnerabilities or active exploitation?



Four Corners of Vulnerability Response

To answer these questions, we analyzed four products which we had about two and a half years of detailed exposure and version data for: **Progress Telerik Report Server**, **Jenkins**, **Cleo MFT**, and **Microsoft SharePoint Server**.

These four products will not represent the entire software ecosystem as a whole. They are case studies that are meaningful to compare precisely because they are so different from each other in terms of purpose, scope of exposure, and the types of network environments they're found in. That makes them useful for investigating the range of behaviors we see in long-term vulnerability responses across different software populations.

Product	Purpose	Exposure Scale	Typical Environment
Telerik Report Server	Enterprise reporting / BI	Very low (hundreds)	.NET and internal BI environments
Jenkins	Open-source CI/CD	Very high (~70K)	DevOps and platform engineering
Cleo MFT ¹ (Harmony, VLTrader, LexiCom)	Managed file transfer	Low (600–1,300)	Supply-chain-heavy industries
SharePoint Server ²	Collaboration and document management	High (~10K)	Enterprise, government, education, healthcare

For each product, we tracked Internet-facing population and version distribution from January 2024 through May 2026, giving us enough history to distinguish short-lived reactions from longer-term maintenance patterns, and compared those changes with major vulnerability events.

Declining exposure on the Internet is not necessarily evidence of improved security posture. In this analysis, we treat **exposure decline** and **version adoption** as complementary signals: a host disappearing may reflect decommissioning or a change in Internet visibility, while movement to a newer version is stronger evidence of a security response. We look for one or both of these behaviors to constitute a significant "response":

- + **Exposure response:** Hosts disappear from the Internet.
- + **Upgrade response:** Hosts remain exposed but move to a newer release.

The goal is not to establish causation, but to add more detail to the software lifecycle changes we observe in these products and examine how real-world events might coincide with (and help explain) those shifts.

1 We use "Cleo MFT," or Cleo Managed File Transfer, as the product name to collectively represent Cleo Harmony, VLTrader, and LexiCom. All three are built on the same underlying VersaLex platform but target different deployment scales and include different capabilities. Because the vulnerabilities we studied affected all three and shared the same fixed-version numbering, we treat them as one product family.

2 This exposure is restricted to self-hosted Internet-facing SharePoint Server deployments.

All Analyzed Products Are Declining in Exposure

In keeping with the overall theme of this report, the Internet is volatile, and Censys's overall view of it has fluctuated over the past two years. In order to isolate just the exposure trends in these specific products, we normalize their host counts over changes in the underlying dataset.

The Internet-facing exposure of all four products decreased over the observation period, though some much more dramatically than others. Two shrank by **more than half**: Telerik Report Server declined by **67.1%** and Microsoft SharePoint declined by **52.6%** by May 2026.

67.1%

decline in Internet-facing Telerik Report Server exposures over two years

Internet Exposure Has Declined Across All Four Products




Figure 8. Exposure depicted as an indexed change: 100 represents starting exposure level; 50 represents a 50% decline in normalized host counts, for instance.

Data via Censys

This is a marked decline, and not something we always see in exposure trends. More often, we see exposed populations persisting on the Internet far longer than they likely should (some examples of this phenomenon can be found in the ICS section of this report, for instance).

Product	January 2024 IP Count	May 2026 IP Count	Normalized % Change
Telerik Report Server	391	113	▼ -67.1%
Microsoft SharePoint	23,783	9,920	▼ -52.6%
Jenkins	110,861	74,215	▼ -23.9%
Cleo MFT	1,307	645	▼ -43.9%

Despite some volatile fluctuations in exposure, especially in Telerik Report Server, the broader decline across all four products is gradual and sustained. That makes it unlikely that any one security incident explains the overall trend. Instead, the data points to a mix of software lifecycle changes, infrastructure turnover, and hosts becoming less publicly observable in general. Those changes in Telerik Report Server give us a useful example of what one such infrastructure turnover event can look like up close.




Figure 9. Telerik Report Server: AWS infrastructure churn

Data via Censys

Most of the exposed Telerik Report Server we observed comes from a single UK-based AWS population of EC2 instances in **eu-west-2**. This group of hosts was very ephemeral: over the observation window, we observed **975 unique IPs** associated with it, and only a fraction were visible at any one time.

Telerik Report
Server Historical
Exposure

#	CVE	CVSS
1	CVE-2024-1800	8.8 (H)
2	CVE-2024-4357	6.5 (M)
	CVE-2024-4837	5.3 (M)
3	CVE-2024-4358*	9.8 (C)
4	CVE-2024-6327	9.8 (C)
5	CVE-2024-8015	7.2 (H)
	CVE-2024-7292,-7293	8.8 (H)
	CVE-2024-7294	6.5 (M)
6	CVE-2024-7295	6.2 (M)
7	CVE-2025-0556	6.5 (M)




Figure 10. The initial post-disclosure decline coincides with the UK-based AWS population upgrading to a newer release, but the continued drop through late 2025 reflects hosts disappearing from the scan surface entirely

Data via Censys

These hosts initially looked like they upgraded after [CVE-2024-4358](#), a critical authentication-bypass vulnerability in Telerik Report Server disclosed in May 2024: by July, **87 of 159 visible hosts, about 55%, had moved to the newly released 10.1 branch**. That grew to 87% by September. But the story did not end with a clean migration to newer releases. After rebounding to roughly **180 hosts in November**, the population shrank rapidly rather than moving onto the newer 10.2 and 10.3 fix branches. By May 2026, only **five hosts remained**, all still on 10.1. From the outside, this looks less like continued patching and more like an **upgrade-and-then-disappear event**, possibly tied to an infrastructure refresh. This is a useful reminder of why we treat Internet exposure counts carefully. **The Internet is volatile**: a falling host count can reflect infrastructure being retired, rebuilt, moved, or simply becoming less visible to an outside scanner.

Older Versions Are Disappearing, but Slowly: The Oldest Are the Longest to Linger

As mentioned earlier, exposure numbers are only one way to measure how these populations respond. Going further by tracking the specific software versions that exposed hosts are running over time gives us a clearer view into whether those populations are actually moving forward.

Let's first look at how the version distributions for each product change on their own, irrespective of any real-world events.

How Software Aged Over Time³




Figure 11. Products with vendor-managed updates (SharePoint Server, Cleo MFT) consolidated around fewer versions over time, while self-hosted tools (Jenkins, Telerik) sprawled across dozens of releases. In every case, a stubborn population of the oldest versions remained exposed.

Data via Censys

³ Because versioning schemes differ across products, we group SharePoint Server by edition, Cleo MFT and Telerik Report Server by major.minor branch, and Jenkins by LTS (Long Term Support) baselines. Jenkins Weekly, the fast-moving release track, appears separately as the persistent orange band at the top of its chart.

Figure 11 shows, within each product's population, the share of hosts running relatively older (teal) versus newer (orange) versions over time. From this data, we can clearly see one phenomenon that all products share: **older versions are gradually becoming less prevalent in the exposed population, while newer versions are making up a greater proportion over time.**

It's also interesting to note that the products' proportions of up-to-date population differ widely at the starting line. Telerik, for example, begins the observation window heavily weighted toward older releases, with only about 30% of hosts on the newer half of its version spectrum (8.x+). Cleo is almost the opposite, with 96% already on 5.0+ in January 2024.

To see just how old some of those older version releases really are, we can break the data out into specific version buckets based on relative age. The chart below tracks how each bucket, from oldest to newest, changes as a share of the product's total exposure over time.

Version Cohorts: Change Over 2.5 Years (pp)




Figure 12. Older releases of software became less prevalent while newer ones gained share.

Data via Censys

This reveals that, while the vast majority of older version cohorts decreased in share for each product, they didn't do so uniformly. The **middle-aged versions** of software, typically a few years out of date, were the ones that tended to disappear by the largest percentages, while the very oldest versions are remarkably stagnant. Across all four products, **the oldest cohort changed by roughly 1 percentage point or less over the entire two-and-a-half-year period**. Telerik 7.x's share of its overall exposure, for example, which dates back to 2021, fell by **52.9 percentage points**, while the oldest 4.x cohort barely moved. WSS 2.0 (Windows SharePoint Services) is a legacy version of SharePoint dating back to 2003 that we still see in small numbers on the Internet. It budged the least, decreasing a completely negligible percentage of 0.04 over roughly 2 and a half years.

We can interpret this to mean that, in this case study, **populations are phasing out older versions of software, but the very oldest software is disproportionately stagnant**. These oldest cohorts typically make up only a small share of total exposure, so **they persist, but in relatively small numbers**. One possible explanation is that these are older or effectively abandoned instances that are simply unlikely to ever be upgraded.

Routine Maintenance vs. Episodic Upgrades

Now comes the question: **do these changes in version distribution seem to coincide with real-world events**, for example vulnerability disclosures or exploitation activity?

To answer this, we can look at **version upgrade rates over time**. We measure upgrade rate as the percentage of hosts observed across two consecutive monthly snapshots that moved to a newer software version in the later snapshot: so, an upgrade rate of **10%** in a given month means that **10% of returning hosts moved to a newer software version than they were running the month before**. Note that we're measuring **version movement** here, not checking for the presence of specific patches, so an "upgrade" does not necessarily mean a particular vulnerability was remediated. We count an upgrade as a movement to a newer minor release or build. This is a useful metric for measuring specifically the pace and *timing* of upgrade events (rather than necessarily the volume). This metric is computed only for IPs that appear in consecutive snapshots. Across the four products, this recurring subset represents a median **75–97% of each snapshot's population**, so the upgrade rate is representative of the observed fleet, but does not capture the more ephemeral hosts that might disappear after one month.

In Figure 13, looking at the upgrade rate curves by themselves, we can observe some clear differences between the upgrade behavior of each product. A flatter trend is an indicator of **more consistent version upgrades**, and perhaps a sign of automated upgrades, while a spikier trend shows **more episodic upgrades**.

Version Upgrade Rate: % of Returning Hosts that Upgraded Month-to-Month




Figure 13. Long-term upgrade behavior varied dramatically by product.

Data via Censys

Now let's correlate these trends with the timing of two types of events. For the purposes of our analysis we focus on disclosures of one or more vulnerabilities rated **medium severity or higher**, and any subsequent additions of those to [CISA's Known Exploited Vulnerabilities catalog \(KEV\)](#) as an indicator that they're undergoing active exploitation. It's worth acknowledging that each product has a very different vulnerability history and patch cadence. **SharePoint Server had 12 vulnerabilities** meeting our criteria during the observation period, compared with just **two for Cleo MFT**. These events give us useful context for understanding patterns in version adoption and upgrades.

Let's compare and contrast the different update cultures we see.

Jenkins

Jenkins: Upgrade Rate




Figure 14. Despite 9+ CVE disclosures over two years, exposed Jenkins instances upgraded at a remarkably consistent rhythm. Data via Censys

Jenkins shows an extremely consistent version upgrade trend over the observation window. Between 5% and 9% of returning hosts upgrade each month, with little detectable sensitivity to events like vulnerability disclosures. The cadence appears to be driven more by routine maintenance than vulnerability response, which might be in part because of how frequently Jenkins releases updates: [weekly releases](#) arrive roughly every seven days, while a new LTS baseline is selected about every 12 weeks, with point releases landing roughly every four weeks.

Jenkins' peak upgrade rate, **9.1% on November 5, 2024**, occurred roughly three months after [CVE-2024-23897](#), a critical arbitrary file-read vulnerability, was added to CISA KEV on August 19. Three months is a relatively long gap, however, and across the full observation window, including eight other vulnerability disclosures, we don't observe any seemingly related dramatic upgrade events. Overall, the Jenkins population we studied appeared to upgrade routinely, regardless of security events.

We see much more pronounced upgrade events in the other products, particularly Cleo MFT.

Cleo MFT

Cleo MFT shows almost the opposite update culture of Jenkins. During the first nine months when no major events occurred, upgrade rates stayed very low, hovering at a median of just **1.9% of persisting hosts upgrading month to month**. That changed with [CVE-2024-50623](#), a critical remote code execution (RCE) vulnerability rated 9.8 CVSS that was disclosed in late October 2024. Shortly after, upgrade activity jumped to **38.4% on November 5, 2024**.

Cleo MFT: Upgrade Rate




Figure 15. Cleo MFT instances were the most reactive out of those we studied. 63% of returning hosts upgraded after reports of two actively exploited vulnerabilities. Data via Censys

What complicated this particular Rapid Response was that **CVE-2024-50623 was fixed in 5.8.0.21, but about a month later, that release was found to remain vulnerable to a separate, related issue, CVE-2024-55956**. Cleo released 5.8.0.24 to address the second vulnerability. Around the same time, **exploitation activity was reported**: CVE-2024-50623 was added to CISA KEV on December 13, followed by CVE-2024-55956 on December 17. The CIOp ransomware gang later claimed responsibility for exploitation of this vulnerability targeting multiple organizations. After 5.8.0.24 became available, the upgrade rate climbed even higher to **63% by January**. That level of responsiveness is encouraging, and may partly reflect the nature of Cleo MFT as a more specialized enterprise product, likely maintained by dedicated IT or infrastructure teams.

The two spikes seemed to be correlated with different events: November was a patch-driven response to a vendor advisory, and then January was an exploitation-driven response, with a ransomware group, two KEV entries, and a failed first patch all pushing at once. The exploitation-driven response was roughly 1.6 times the size of the patch-driven one.

Interestingly, upgrade activity then quickly fell back to the same near-1% cadence seen before the vulnerabilities and plateaued there. One thing to point out is that most Cleo hosts were already on relatively newer releases at the start

of the observation window, with 96% on 5.0 or higher. This might help explain this low baseline alone: there was still substantial spread across versions within the 5.x release line and room to upgrade to newer versions, but likely much less of an incentive. This is likely created by a software culture that **doesn't incentivize routine upgrade activity**, followed by **sharp bursts around major security events** that increase the urgency of patching. Anecdotally, this is likely what much of software maintenance looks like across other products: if it isn't clearly broken, don't fix it.

SharePoint Server

SharePoint shows a much higher baseline upgrade rate than the other products, with a baseline median of 26.1% of returning hosts upgrading routinely each month. An upgrade event here includes instances that move to a newer build, not just a newer edition. Similar to Jenkins, this may partly reflect SharePoint's release cadence. Microsoft ships monthly cumulative updates for SharePoint Server. That is much more frequent than Cleo or Telerik, though still less frequent than Jenkins. Vulnerability disclosures generally do not appear to drive upgrade events far beyond this already high baseline. For example, when CVE-2024-38094, an RCE vulnerability rated CVSS 7.2, was added to CISA KEV, the upgrade rate moved only modestly from **22.5% to 25.4%**.

Two Software Products, Two Opposite Maintenance Cultures

5–9%

of returning Jenkins hosts moved to a newer version every month, with or without a vulnerability disclosure

63%

of returning Cleo MFT hosts moved in a single month after two of its vulnerabilities were actively exploited

The most compelling evidence of a potential reactive upgrade burst appears around [CVE-2025-53770](#), the unauthenticated RCE zero-day associated with the ToolShell chain. About 16 days after it was added to KEV on July 20, 2025, the upgrade rate rose to roughly **36%**. SharePoint later reached its peak of about **40% in April 2026**, though that spike does not line up cleanly with any event we are tracking and could reflect other maintenance cycles. Overall, SharePoint appears to upgrade routinely, with some acceleration when major security events occur.

Microsoft SharePoint: Upgrade Rate




Figure 16. Self-hosted SharePoint Server upgrades climbed steadily from 20% to 40% over the time period.

Data via Censys

Telerik Report Server

Telerik tells a similar story to Cleo MFT, but with less dramatic swings. We see **four** distinct upgrade spikes, ranging from 9.4% to 22.4% of returning hosts upgrading, that appear to line up with **four waves** of vulnerability disclosures in the first eight months of the observation window. The largest upgrade event, **22.4% on June 25, 2024**, also falls close to notable security activity. This occurred just 12 days after **CVE-2024-4358 was added to CISA KEV**. After those events subside, upgrade activity falls back to a much lower baseline. These patterns suggest a **moderate sensitivity to security events**, followed by a clear decay back toward baseline once those events pass. In other words, similarly to Cleo, **Telerik tends to upgrade around when new exploitation is reported**, but not much otherwise. This may also be a reflection of Telerik's slower release cycle compared to the other products under analysis here, on the order of every three to four months.

Telerik Report Server: Upgrade Rate




Figure 17. A mid-2024 flurry of vulnerabilities, one added to KEV, drove Telerik Report Server's peak upgrade rate of 22.4%, but the response was brief.

Data via Censys

Product	Median upgrade rate	Peak upgrade amount	Release cadence	Update culture
Jenkins	7.3%	9.1%	Weekly + LTS ~4 wks	Upgrade routinely, regardless of events
SharePoint Server	26.1%	39.7%	Monthly (cumulative)	Upgrade routinely, and at an accelerated rate when events occur
Telerik Report Server	2.8%	22.4%	Quarterly (~3 months)	Upgrade when events occur, not much otherwise
Cleo MFT	1.9%	63.0%	~2–3 months ⁴	Upgrade when events occur, not much otherwise

There's an interesting relationship here between each product's **routine baseline rate of upgrades** during calmer periods and **how dramatic its peak upgrade event was**. The two almost look like inverses of each other. Cleo MFT, which had the lowest routine upgrade activity, had the largest upgrade event, while Jenkins, with the steadiest routine movement, underwent the smallest.

This intuitively makes sense: devices without a regular maintenance rhythm will have to address more upgrades all at once when a major event happens, while those with a routine upgrade cadence have already absorbed much of that change over time.

It's also encouraging that **all four applications show some evidence of active upgrades**, albeit in small proportions of their overall population. Three show clear increases in upgrade activity around high-priority vulnerabilities, while Jenkins, which shows little sensitivity to such disclosures, is nonetheless upgrading steadily as part of its normal maintenance cycle.

4 Cleo's support portal is gated, so version timelines were reconstructed from a combination of cached vendor release notes PDFs, public security advisories, NVD CPE entries

Takeaways

Software lifecycles are a very nuanced phenomenon to measure. The dent that the aftermath of a security incident leaves in the overall security posture of a software population can look very different depending on how a product is normally maintained: some populations turn over steadily, while others change very little until a major event creates a visible burst of upgrades. These products are not representative of the entire software ecosystem, but together they illustrate how differently that long-term response can unfold.

01



Over the long term, exposed software populations are gradually shifting toward newer versions.

Across all four case studies, older releases became less prevalent while newer versions gained share, but this movement was mostly among versions landing in the middle of the age spectrum. The very oldest versions barely changed as a proportion of exposure, persisting in small but remarkably stable levels.

02



Baseline upgrade frequency appears to reflect differences in how products are deployed.

The most widely deployed products with more frequent release cadences showed relatively steady version turnover from month to month. Smaller, more specialized products with slower release cycles had much lower levels of routine upgrade activity. This seems to reflect different cultures surrounding each product that correlate with their vendor's maintenance rhythm.

03



Products seem to either upgrade steadily or in bursts.

Products with lower levels of baseline upgrade activity showed the largest spikes in version upgrades around major vulnerabilities or new exploitation activity, while products that upgraded continuously showed much smaller changes, likely having absorbed more version change gradually over time.

Understanding a product's normal release and maintenance rhythm gives us much more context for interpreting what actually changes in the Internet's long-term security posture long after a zero-day is disclosed.

3

Exposure of Industrial Control Systems

Critical Infrastructure on the Internet


The number of industrial control systems (ICS) exposed to the public Internet has been measurable — and measurably large — for over a decade. Despite sustained research attention and remediation efforts, exposure remains widespread.

ICS protocols, sometimes called OT or automation protocols, enable data exchange between systems in industrial or automated environments. Many were designed without encryption or authentication, modern security necessities that weren't prevalent when the earliest implementations appeared in the 1970s. Despite this, these protocols are not frequently exploited on their own: only a small number of malware variants are known to target ICS protocols, in part because doing so requires specialized knowledge of both the protocol and the physical equipment it controls.

Conversely, human-machine interfaces (HMIs) present a different attack surface and complexity. These graphical interfaces allow operators to remotely monitor and interact with industrial systems.

Unlike ICS protocol attacks, HMI attacks require no system-specific expertise and have resulted in defacement and manipulation of operational values at [water plants](#) and other critical infrastructure sites.

We also find that over **80% of ICS hosts also run at least one non-ICS IT service like HTTP, SSH, or FTP**. These services increase the attack surface of ICS hosts and



Example of an HMI graphical interface

introduce more avenues to exploitation. A threat actor doesn't necessarily have to understand the inner workings of Modbus or EIP if a vulnerable web interface or SSH service is exposed on the target host.

What follows is an examination of the ICS attack surface over the past two and a half years, including how the geography of ICS exposure has shifted, while network footprint has remained remarkably static. We have removed all instances of what appear to be honeypots or hosts running otherwise duplicitous services from this analysis.

The Geography of ICS Protocol Exposures

As of this analysis, **Censys observes an average of 134,000 distinct hosts running ICS services and tooling exposed on the Internet in early 2026, up roughly 4% from about 129,000 in 2024.**

As of 2026, North America continues to lead in exposures, representing roughly 38% of the global ICS landscape. This is consistent with past findings, though we note that Asia's ICS footprint has expanded from 22% of exposures in 2024 to 25% in 2026. In contrast, Europe's proportion of ICS exposure has decreased, from 36.5% of global total in 2024 to 32% in 2026.

Investigating further, we identified that South Korea added roughly 2,360 ICS hosts between 2024 and early 2026 — a 36% increase and the greatest increase observed in Asia. Modbus, a serial communications protocol widely observed in industrial environments, and CIMON PLC, a Korean HMI/SCADA platform, are responsible for nearly all of that growth.

38%

of global ICS exposures are in North America

Global Industrial Control Systems Exposure, 2024-2026




Figure 18. Global population of hosts running at least 1 ICS service Data via Censys

Global Industrial Control Systems Exposure by Continent, 2024-2026




Figure 19. Regional populations of hosts with at least 1 ICS service Data via Censys

South Korea ICS Protocol Trajectory




Figure 20. Modbus and CIMON PLC are key drivers in South Korea's expanded ICS footprint. Data via Censys

Turning to Europe, FOX, the Tridium Niagara building automation protocol, fell 25-32% in France, Italy, and Belgium from 2024 to early 2026. Modbus and EIP, an industrial Ethernet protocol common in factory automation, also dropped across all three countries. Together, France, Italy, and Belgium account for roughly 85% of Europe's ICS exposure decline during the analysis period.

European ICS Decline by Protocol: France, Italy, Belgium (2024-2026 avg)




Figure 21. Modbus, FOX, and EIP are among the protocols that have shown the greatest decline across France, Italy, and Belgium. Note that each subplot is scaled differently.

Data via Censys

ICS Exposure Over Time: France, Italy, Belgium (2024-2026)




Figure 22. While there has been volatility in ICS exposures across Italy, France, and Belgium, since April 2025 exposures in these countries have followed a downward trend. Data via Censys

One possible contributing factor in these declines is the [Network and Information Security Directive 2 \(NIS2\)](#), the European Union's updated cybersecurity directive for critical infrastructure, manufacturing, and a variety of commercial sectors like food, chemicals, and logistics. A replacement of the earlier 2016 NIS1, NIS2 introduced additional security regulations for [ICS environments](#) upon its release in December 2022. EU member states were given autonomy on how and when to sign the requirements into law, though they were given a deadline of October 17, 2024.

Both Italy and Belgium [adopted NIS2](#) by the October 2024 deadline, but as of writing France has not. Organizations operating across the EU may have been preparing for the regulation regardless of whether it's been signed into law in a given country.

Beyond these major changes, much of the global ICS landscape has remained stable over the last two and a half years. The U.S. continues to dominate the exposure landscape, representing 32% of the total landscape — significantly more than any other single country. Though Italy has seen significant decline in hosts exposing ICS services, it still falls within the top 8 countries where we observe exposed ICS.

Industrial Control Systems Exposure by Country (Top 8), 2024-2026




Figure 23. The U.S. hosts significantly more exposed ICS devices than any other country. Turkey and South Korea also trend upward during this timeframe.

Data via Censys

ICS Network Profiles

While the geography of exposed ICS devices has shifted slightly over time, the types of networks where we observe them has remained remarkably stable. Roughly 70% of hosts running ICS devices and services globally are consistently found on consumer and mobile networks, and have been for the last 2.5 years.

This prevalence on mobile networks is often due to the remote locations of these devices in substations and other automated environments, which utilize cellular modems (5G or LTE enabled) for Internet connectivity. **Identifying the owners of these devices presents a significant challenge for researchers**, as host WHOIS and other metadata typically point to the telecommunications provider rather than the owner/operator, **making exposure notifications difficult**.

Explore the ICS Exposure Gap

Learn more about the exposure notification gap for mobile-connected ICS devices and what it means for defenders.

[Read the Blog](#) ➤

Industrial Control Systems Exposure by Network Type, 2024-2026




Figure 24. The types of networks where we observe ICS systems have remained stable over the last two years. Data via Censys

One U.S.-based mobile carrier hosts a tenth of the world's exposed industrial control systems. Verizon's cellular network, Cellco Partnership (CELLCO-PART, AS6167), hosts 10% of all ICS hosts Censys observes globally.

Top 10 Autonomous Systems Hosting ICS




Figure 25. No single carrier dominates the global ICS landscape, but Verizon leads with 10% of observed exposures.

Data via Censys

The pattern of ICS being found on consumer and mobile networks is well-established, but the distribution of networks hosting ICS devices varies significantly by country. In some nations, ICS hosts are concentrated primarily on a single network, while in others, they are more broadly distributed across multiple networks. To illustrate the variation in network footprint, we examine five countries across varying regions below. ICS hosts in Turkey, which has the second highest number of exposures globally, are found largely on Turkcell (81%), the nation's dominant mobile provider. South Korea follows a similar pattern, with 71% of ICS hosts found on Korea Telecom.

In contrast, one network in the U.S. (Verizon's Cellco) hosts a third of all ICS infrastructure in the U.S., with a long tail of other networks where ICS devices are found. Similar profiles appear in China and Canada, with the single largest provider hosting around a third of ICS devices, followed by a long tail of networks. **Homogeneous hosting patterns like these can increase risk:** armed with the knowledge that the bulk of a country's critical infrastructure can be found on a single network, reconnaissance and hunting are made more straightforward for threat actors. Moreover, an attack on such a network could cause widespread outages, making many of a country's industrial systems unreachable.

Network Concentration Variation, 2026




Figure 26. ICS exposure concentrates differently across different countries: Turkey and South Korea's ICS footprints are dominated by single providers, while the U.S., Canada, and China are spread more broadly.

Data via Censys

The IT + ICS Attack Surface

Attacks on ICS protocols themselves are relatively rare, but the attack surface of such systems can dramatically expand when IT services run alongside ICS services on the same host. Many of these devices are networked hardware running embedded software that includes web management panels, remote access services, or both. A threat actor may not understand Modbus or S7, but they can still reach the HTTP interface or the SSH service on the same host, presenting other avenues to potential compromise.

IT / ICS Companions: % of ICS Hosts with a Given Non-ICS Service





Figure 27. HTTP is commonly found alongside a wide range of ICS services, substantially expanding the attack surface of these hosts.

Data via Censys

In Figure 27, we plot the share of each protocol's hosts that also run a given IT service and find that **HTTP is broadly the most common co-occurring service**, appearing on 31-99% of hosts that run at least one ICS protocol. These HTTP services range from router or network device management interfaces and VPNs to industrial protocol management interfaces.

The Tridium Niagara (FOX) building automation framework is almost always (99%) observed with an HTTP service hosting a web management interface for the system. The Niagara platform has over 20 CVEs dating back to 2012, and the web interface represents an entry point to the platform.



Example Tridium Niagara login interface

While the Niagara framework is heavily co-located with HTTP services, Omron's FINS has a very different profile. FINS hosts co-expose FTP (56%), SSH (43%), MySQL (39%), VNC (39%), and Telnet (37%) on the same IP as the Omron PLC runtime. The FINS protocol allows for read/write access to a PLC in manufacturing, food and beverage, and other automated environments. **These additional IT services expand the attack surface and could expose project files, operational data, or offer remote access.**

Human-Machine Interface Exposures

We analyze HMI exposure separately from ICS protocol exposure to more carefully examine the differing attack surface and complexity. HMIs are designed to provide visibility to operators, aiding them in tracking pump levels at water and wastewater treatment facilities, controlling circuit breakers and switching equipment in electrical substations, and adjusting setpoints for pressure in natural gas transmission.

Unfortunately, some HMIs ship with default credentials, making them a simple web search away from compromise. Perhaps the most notable example of this is when Unitronics HMIs made the news in 2023 following a defacement campaign. One of the CISA recommended mitigations was to change the system's password to something other than the default "1111".

When directly exposed to the Internet, this visibility and access that is so convenient for operators also falls into the hands of threat actors as we've seen in numerous incidents over the past several years.



Defacement imagery found on a Unitronics brand HMI in late 2023

As observed in previous research, **the U.S. leads the HMI exposure landscape**, with nearly half of all HMIs observed globally found in the U.S. The remaining top 10 countries for HMI exposure are in North America, Europe, and Australia, a pattern consistent with previous observations.

Top 10 Countries with Exposed HMIs




Figure 28. Collectively, the U.S. and Germany host nearly 70% of Internet-facing HMIs globally. Data via Censys

More than one in four exposed HMIs are found on a single mobile network. While 10% of global ICS protocols are found on Verizon's CELLCO network (CELLCO-PART, AS6167), 27.5% of all HMIs observed globally are also found there. These HMIs have a similar problem to other ICS services hosted on this network: it can be **nearly impossible to identify ownership** based on host metadata alone. HMIs frequently have text that points to a town, city, or other municipal entity, but without that, they're as difficult to attribute to an owner as other ICS infrastructure.

Top 10 Autonomous Systems Hosting HMIs




Figure 29. Verizon and Deutsche Telekom combined host roughly 45% of the world's Internet-exposed HMIs. This is consistent with the U.S. and German lead in HMI exposures. Data via Censys

Takeaways

Over the past two and a half years, ICS exposure has remained broadly stable, apart from minor fluctuations in geography. Exposed ICS systems often sit beyond the purview of traditional asset and exposure management mitigation techniques meant for IT systems; those strategies don't translate cleanly to an industrial environment. This is a major contributing factor in why, despite sustained attention from governments and security practitioners regarding removing ICS systems from the public Internet, we've seen a 4% *increase* in exposed ICS systems over the last several years.

Three findings in this section explain more about why this has happened:

01



ICS devices often sit outside of typical asset management scope.

ICS assets often don't live in an organization's corporate IP space, presenting asset management challenges for operators. Roughly 70% of ICS hosts sit on consumer and mobile networks, not in corporate networks. An operator searching for devices against their registered netblocks will miss their own exposed industrial equipment because they're typically located on cellular modems in a carrier's IP space.

02



Ownership attribution difficulties create an exposure notification gap.

When ICS devices are discovered by researchers or other third parties, there's often no clear indication of who owns the device. With 70% of ICS systems on consumer and mobile networks, the WHOIS and other host metadata point back to the telecommunications company, not the operator. Discovery of these devices is trivial; determining ownership is often impossible.

03



Web interfaces represent a sizable attack surface component of ICS systems.

Of hosts running an ICS protocol, 31-99% also run an HTTP interface. Attacking ICS systems via specialized knowledge of the industrial protocol isn't necessary; threat actors need only find a web application vulnerability to exploit. This also holds true for systems running HMIs, which are often accessible via VNC, RDP, HTTP, or other specialized protocols (e.g., C-more) with default credentials or no authentication required.

4

Threat Infrastructure

An Internet-Wide View of Adversary Tools

Threat research often focuses on individual, event-driven artifacts: a new Remote Access Trojan (RAT) is discovered; a novel phishing campaign is observed; an extortion actor targets enterprise organizations. Following such events, reporting often involves a teardown of any obtained malware, analysis of infrastructure, or research about the adversary. These views are critical to our understanding of the mechanics of exploitation.


While individual threat reports examine a single subject in depth, the following analysis examines the aggregate: the full, Censys-visible population of threat actor infrastructure measured over nine months from October 2025 to June 2026.

We examine:

- + Phishing and lure delivery infrastructure
- + Remote access and tunneling tooling
- + Command-and-control (C2) frameworks and post-exploitation tooling

What ties them together is not a shared campaign or adversary usage, but that all of them operate Internet-facing infrastructure that Censys can fingerprint and track. Censys observes threat infrastructure in two ways: host-based threats are detected by scanning an IP address and port, while web-based threats come from fetching the web content served by a hostname.

This analysis is based on 275 threat fingerprints actively maintained by Censys researchers, 162 of which appeared in our data at least once during the study timeframe. We define threats as **Censys-visible infrastructure whose configuration, capabilities, properties, or observable behaviors indicate potential or confirmed use in suspicious or malicious activity**. Malware, of course, falls into this category. However, many of the tools and software we track as threats are **dual-use**, meaning they have benign uses in addition to being adopted by threat actors. This makes the broader context of all services that appear on a host even more important when hunting or investigating.



Host-based Threat Detections, October 2025 - June 2026



Figure 30. All host-based threats sorted by peak host count from October 2025 through June 2026. FRP sits in the top left corner; threats on the top row reach tens of thousands of hosts. Threats in the lower rows often peak in the hundreds or fewer, with many showing sharp bursts rather than sustained presence across the study timeframe.

Data via Censys

Landscape at a Glance

Threat	Hosts	Web Hostnames	Combined	% of total, June 2026
FRP	149,283	60,515	209,798	72.7%
NPS	19,365	6,586	25,951	9.0%
Nezha Monitoring	5,755	4,219	9,974	3.5%
Suspicious Open Directory	3,803	5,054	8,857	3.1%
Fake EZ-Pass	0	7,398	7,398	2.6%
ShellInABox	5,536	327	5,863	2.0%
Interactsh	775	2,418	3,193	1.1%
Gophish	1,052	1,154	2,206	0.8%
Fake Captcha	47	1,936	1,983	0.7%
nps-socks5	1,714	0	1,714	0.6%
Asset Reconnaissance Lighthouse	1,657	10	1,667	0.6%
Acunetix	1,519	59	1,578	0.5%
mitmproxy	1,462	75	1,537	0.5%
Nessus	1,394	88	1,482	0.5%
Vshell	1,179	61	1,240	0.4%

In the table above are the top 15 threats as of June 2026 based on combined host and web footprint. Definitions of each of these can be found in the [Glossary](#) at the end of this report. Some threats are observed only on web infrastructure or host infrastructure, not both.⁵

⁵ In the table above, we see Fake EZ-Pass is not observed on any hosts in this snapshot. Our detection for Fake EZ-Pass is based on web attributes, so we would never expect to see Fake EZ-Pass on a host. Alternatively, nps-socks5 is not observed on any web hostnames, but our detection for nps-socks5 covers both hosts and web.

Dual-Use Tools Top the List of Threats

Considering the table of top threats by total volume above, three dual-use infrastructure tools lead the list:

01		<u>FRP (Reverse proxy / port forwarder)</u>  Exposes internal services through NAT/firewall. It also offers an optional status dashboard and per-client tunnels.
02		<u>NPS (Intranet penetration proxy)</u>  Similar to FRP but a proxy server with central management; clients register to a specific server and are provisioned from that console with different forwarding rules. Offers a more extensive UI to manage the connected hosts.
03		<u>Nezha Monitoring (Server monitoring agent)</u>  Legitimate sysadmin tool, abused for persistent agent deployment.

These three tools, all of Chinese origin, represent 85.2% of the total threat infrastructure tracked by Censys in June 2026. This is driven largely by FRP, which represents 72.7% of tracked threats alone. Given the broad, dual-use nature of FRP and the differing order of magnitude between FRP and the other threats we track, we have excluded FRP from the remaining analysis so as not to misrepresent the landscape in favor of a tool that is not explicitly malicious.

Beyond these three tools, we find a variety of utilities that span the attack lifecycle, including reconnaissance, delivery, tunneling, and C2 capabilities. Acunetix and Nessus are commercial products, and it's challenging to determine whether these are legitimate red team deployments or cracked, unauthorized instances of the tools. Apart from these, the rest of the tooling in the top 15 is open source or freely available.

Where Threat Infrastructure Lives

It's tempting to study the geography of threat infrastructure and make assumptions about the actor's identity, intent, and origin based on where the infrastructure is located. However, threat actors don't pick a country; they pick a provider. Hosting and infrastructure providers are selected based on several key attributes:

1 Cost and functionality

- ✦ Inexpensive or free tiers for disposable infrastructure
- ✦ Ability to provision infrastructure via API for larger scale operations

2 Permissiveness

- ✦ Abuse-tolerant jurisdictions and hosting providers (e.g., bulletproof networks)
- ✦ Providers with minimal customer vetting
- ✦ Cryptocurrency payment options

Geography is important because it shapes detection patterns, not because it heavily shapes deployment or hosting decisions.

Here we combine host- and web-based infrastructure to understand where the totality of Censys-visible threat infrastructure is found. **China and the U.S. dominate** across the full analysis timeframe, though their trajectories are starkly different. **China grew 12% while the U.S. grew 84%**, driven by web activity. Hong Kong ranks third throughout, growing 28% and more closely following China's trajectory than that of the U.S.

The sharp increase observed from December to January in the U.S. is primarily driven by [Fake Captcha ClickFix-style infrastructure](#); these pages nearly tripled from December to January. Beyond January, the sustained increase observed in the U.S. is due to Fake EZ-Pass lure pages, which experienced a 70% increase from January to June 2026.

Threat Infrastructure - Host IPs + Web Hostnames by Country




Figure 31. Though they have experienced different trajectories over the study timeframe, the U.S. and China lead the global threat landscape.

Data via Censys

The Host/Web Inversion

For host-based threat infrastructure, such as C2 servers, proxies, and remote access tooling, China (17,622, 36.5%) leads by more than two to one over the U.S. (8,104, 16.8%). Together, China and Hong Kong represent almost half of host-based threat infrastructure as of June 2026.

Host-Based vs. Web-Based Threat Geography




Figure 32. China and the U.S. are global leaders in threat infrastructure volume, but the composition of what they host differs. China holds 36.5% of all threat IPs but only 20.9% of web hostnames, while the U.S. is the opposite, with 16.8% of IPs and 45.5% web hostnames. Data via Censys

Conversely, examining web-based threats, which comprise primarily phishing and other deceptive pages, **the U.S. leads with 13,674 hostnames (45.5%) to China's 6,280 (20.9%)**. We note the possibility that the web threat disparity between the U.S. and China may be partly language-driven. Our detections for English-language web threats are more robust than for non-English languages, which may impact the observed gap.

However, the host disparity between China and the U.S. remains, suggesting there is some operational difference in where actors host their wares.

Network Concentration

We also observe this pattern when examining networks where threats are found.

How Threat Activity Splits Across Networks





Figure 33. Threat actors use different providers for different purposes: networks represented on the right side of the chart tend to host more web-based threats, while those on the left are home to more host-based threat infrastructure. Data via Censys

Figure 33 illustrates, for each network, how much its threat infrastructure skews toward web or host dominance. A label of 2x web means the network's share of threat-labeled web hostnames is 2 times its share of threat-labeled host IPs. A label of 2x host means the reverse: its host share is 2 times its web share. "Same" at the center means the two shares are equal.

A network's position in Figure 33 points to the type of threat it tends to attract. Web-heavy networks like Cloudflare, Amazon, and Linode tend to host more phishing and other social engineering that's presented directly to potential victims. Host-heavy networks like Chinanet and Alibaba tend to host more operator tooling, which potential targets are never intended to see.

A Tale of Two Web-Based Threats


Many of the tools discussed previously are designed to operate covertly and quietly, allowing threat actors to obtain and maintain persistence on remote systems. We shift now to discuss two threats that work by appearing directly in front of users and encouraging them to interact.



Fake EZ-Pass

Fake EZ-Pass is an SMS-based phishing scheme that impersonates U.S. toll collection services including EZ-Pass, SunPass, and Peach Pass. Recipients receive a text message claiming they have an unpaid toll balance with a link to what purports to be the legitimate payment form. In reality, the page is actor-controlled infrastructure that steals entered payment credentials and personal information.

Beyond toll collection services, similar lures have impersonated the Department of Motor Vehicles (DMV) and the Department of Transportation (DOT) for various U.S. states, as well as shipping services like USPS. These SMS-based phishing campaigns have been attributed to the Smishing Triad, a Chinese language phishing-as-a-service ecosystem active since at least 2023. Examples here are U.S.-focused, but targeting is not limited to the U.S. The activity has international reach, and has also targeted the banking sector, other shipping outlets, and e-commerce platforms, among others.



Fake Captcha

Fake Captcha is the broad label applied to infrastructure associated with ClickFix-style delivery. ClickFix is a social engineering technique that presents users with a fake CAPTCHA challenge, typically on a compromised or malvertising page. The page instructs the user to press a combination of keys that ultimately executes a PowerShell command on their machine. The command downloads and installs an infostealer, often StealIC, Lumma Stealer, or AMOS. Unlike Fake EZ-Pass, this technique is used by a range of threat actors and is not associated with a specific group.

During our analysis timeframe of October 2025 to June 2026, Fake EZ-Pass peaked in November with 10,374 instances observed, then consistently declined until reaching its trough in April 2026 with just over 4,400 instances observed. Since then, it's grown again, with just over 8,000 instances in May and 7,300 in June 2026.

Fake Captcha follows a different trajectory altogether, and on a slightly different scale. While we observed over 10,000 instances of Fake EZ-Pass at its November peak, we only observed 5,667 Fake Captcha web instances (253 hosts) during its peak in January 2026.

Dominant Web-Based Threats, October 2025 - June 2026




Figure 34. Fake EZ-Pass and Fake Captcha have different trajectories over the same timeframe. Spikes in web presence (Fake Captcha) don't always have a related spike in hosts.

Data via Censys

Fake EZ-Pass and Fake Captcha each have a strong presence in the U.S., though the networks and hostnames where these threats are found tell different tales about how the infrastructure is obtained and managed. Cloudflare is prominent across both sets of infrastructure, which is more about Cloudflare's ubiquity in the web space than deliberate actor choice. Apart from Cloudflare, Fake EZ-Pass is found heavily on U.S.-based providers Amazon, Linode, Automatic ([WordPress\[.com\]](https://WordPress.com)), Panamanian VPS provider Private Layer, and China's Tencent, all of which offer inexpensive infrastructure that can rapidly scale and be torn down.

Fake Captcha is also commonly found on Cloudflare, but additionally, it appears across U.S. providers Oracle, Network Solutions Hosting, Unified Layer, France's OVH, and Japan's GMO Internet. This group of networks is a bit more varied than where we observe Fake EZ-Pass, pointing to more opportunistic infrastructure abuse. Examining web hostnames for each further builds this distinction.

Geography and Hosting for Fake EZ-Pass and Fake Captcha, June 2026




Figure 35. The U.S. hosts the majority of Fake EZ-Pass and Fake Captcha, though each threat's network footprint varies. Amazon and Cloudflare combined host 67.8% of Fake EZ-Pass instances, while Fake Captcha is more distributed, with 40.1% on Cloudflare, followed by a long tail of hosting providers.

Data via Censys

Fake EZ-Pass and Fake Captcha: Infrastructure Models



Figure 36. Fake EZ-Pass domain patterns are more consistent with actor-purchased domains designed to impersonate official sites, while Fake Captcha domains are more consistent with injection into legitimate third-party sites the actor does not control.

Data via Censys

Fake Domains in the Subdomain

Fake EZ-Pass domains are **purchased with the intent to deceive**. The hostnames commonly (36-77%) contain a subdomain with an embedded domain (e.g., **michigan[.]gov[.]gvtd[.]xyz**). This can be particularly effective for users opening the sites on mobile browsers following receipt of a fraudulent SMS: the entire hostname is not visible in the URL bar, but just enough of it can be seen so it appears legitimate.

Prior to April 2026, this full, fake embedded domain in the subdomain (e.g., **michigan[.]gov**, from the above example) was the dominant hostname pattern observed. In April, a decline in this style of hostname helped spotlight a small rise in a subtle variation of this style: rather than being contained in the subdomain, the fake site's top level domain (TLD) is included in the registered domain name, as in **e-zpass[.]com-etcqwz[.]vip**.

Dotted Variant	Hyphenated Variant
michigan[.]gov[.]xdep[.]xyz	michigan[.]gov-yxs[.]cfd
e-zpass[.]com[.]gsng[.]xyz	e-zpass[.]com-etcqwz[.]vip
ardot[.]gov[.]gwsb[.]xyz	michigan[.]govbdc[.]vip

Fake Captcha names are largely **stolen**. They're often hijacked from small businesses and other websites around the world, and don't appear to be limited to a specific sector. However, in some cases, we observe crypto-related domain names which may be intentionally purchased and stood up to execute a scam.

TLD Divergence

Another defining characteristic of the Fake EZ-Pass hostnames is their use of novelty TLDs like **.xyz** and **.shop**. Nearly 60% of Fake EZ-Pass instances are found on **.xyz** TLDs.

Two TLDs — **.cc**, the TLD for Australia's Cocos Islands, and **.cn**, China's TLD — stand out on this list of otherwise novelty TLDs. In the past, **.cc** has been linked to malicious domain registrations. Given the Smishing Triad's Chinese affiliation, the **.cn** TLD is not unexpected.

Purchased or Stolen: Fake EZ-Pass and Fake Captcha TLD Patterns




Figure 37. Fake EZ-Pass TLDs reflect a deliberate, actor-curated TLD selection, while Fake Captcha TLDs mirror the broader web and reflect the campaign's reliance on compromised legitimate domains.

Data via Censys

The majority of Fake Captcha instances observed by Censys, however, are found on legacy TLDs like **.com**. Some country-specific TLDs are observed as well, such as **.br**, **.uk**, and **.ru**. This is likely opportunistic and based on whatever sites are taken over, rather than being specifically targeted for a given TLD. This also likely explains the much longer tail of Fake Captcha TLDs (over 300) compared to that of Fake EZ-Pass (29).

(Lack of) Subdomain Depth

While the previous two domain name observations have largely applied to Fake EZ-Pass sites, we turn to Fake Captcha for this phenomenon. Fake Captcha hostnames largely have minimal subdomain depth when compared to Fake EZ-Pass. Two thirds (66%) of Fake Captcha sites are at the root domain or **www**. In contrast, a classic fake EZ-Pass site might look like **michigan[.]gov[.]xdep[.]xyz**, with a subdomain "depth" of two: the root domain is **xdep[.]xyz**, while **michigan** is a subdomain of **gov** in this example.

Takeaways

Geography and network distribution are two common lenses through which researchers and practitioners study threat infrastructure. These are often used as quick heuristics for determining whether something is likely suspicious or benign. A dual-use tool, like NPS, found on a bulletproof hosting provider might be assessed differently than the same tool on a network with a better reputation for addressing abuse. Similarly, a U.S.-based organization suddenly seeing unexpected traffic from a Chinese network might interpret that differently than a traffic spike from elsewhere.

01



China and the U.S. host a plurality of the world's malicious software and services, but the profiles of what they host diverge significantly.

China is home to 36.5% of malicious host-level infrastructure like C2s and remote access tooling in contrast to the U.S.'s 16.8%. The U.S., conversely, hosts 45.5% of web-based threats like phishing and other deceptive pages, while China hosts only 20.9% of web-based threats observed by Censys. This suggests operational differences in where actors deploy their tooling.

02



Subtle shifts in actor behavior can break less robust detection logic.

Fake EZ-Pass's evolution from a dotted variant of the hostname to a hyphenated variant is a trivial change for the actor to make, but it enables circumvention of brittle, overly specific detection rules.

5

Exposed AI Infrastructure

The Fast-Growing Attack Surface

Infrastructure related to artificial intelligence is quickly becoming a major part of the Internet-facing attack surface. For security teams, this creates a new category of exposed assets to inventory, manage, and harden.

Not all of this exposure is a problem: an AI-powered customer support chatbot application, for example, is built to be publicly reachable. The trouble is that **we're increasingly seeing exposures of the plumbing that supports user-facing AI applications**: model endpoints, vector databases, inference APIs, and orchestration frameworks that were likely only ever meant to communicate with the application itself.


Part of this trend comes down to how easy AI has made it to build and self-host software. Developers can move from an idea to a working application in an afternoon using coding assistants, a practice often called "vibe coding." These workflows tend to prioritize getting something working quickly over secure deployment practices, leading to permissive configurations that inadvertently expose the underlying stack.

In other cases, organizations may not even realize how AI tools entered their environment in the first place - until something goes wrong. Third-party services can route traffic through AI APIs without the knowledge of their users. A coding assistant on one employee laptop can create connections to external model endpoints that bypass traditional application architectures entirely. This is increasingly described as **"shadow AI"**, essentially the AI equivalent of shadow IT.

In both of these cases, the importance of having a continuous inventory of these assets remains crucial.

In this section, we use Censys Internet-wide observations to examine how this emerging attack surface is taking shape and ask:

- ✦ What kinds of AI-related infrastructure are appearing on the Internet?
- ✦ How quickly are they growing?
- ✦ Which parts create the greatest security concern?



It's Not All Chatbots

Part of the visibility problem is the sheer variety of technologies we now lump under the word “AI”. Does AI mean only large language models? Machine learning systems? Does it include the applications built around them? What are we actually talking about when we refer to the “AI attack surface”?

Rather than debating this abstractly, we inventoried the AI-related products that we observed to be Internet-facing in our dataset and considered what they actually do and where they sit within an organization's larger infrastructure. By “AI-related product” we mean software **directly involved in building, running, connecting, or operating AI capabilities**. That generally excludes the models themselves, which aren't network-facing, and refers instead to the application, orchestration, and monitoring infrastructure around them.

As of July 2026, we observed **267,915 IPs** exposing one or more of **44 distinct AI-related products**.⁶ These span a remarkably wide range of functions, from AI workflow automation platforms and local model inference services to LLM gateways, proxies, and retrieval-augmented generation (RAG) systems.

Top Exposed AI-related Software




Figure 38. The exposed AI landscape is dominated by one product: n8n, which connects AI models to other apps and services. Data via Censys

⁶ Explore this data in Censys: [host.services.software.type="AI"](#)

Despite that variety, exposure is highly concentrated in just five widely adopted products that appear on 82% of all observed hosts. **n8n**, a workflow automation platform, is by far the most commonly exposed at nearly 114,000 hosts, followed by **OpenClaw** (~54,000), a personal AI-agent assistant; **New API** (~24,000), an LLM gateway; **Ollama** (~23,000), a local model inference server; and **OpenWebUI** (~19,000), a chat frontend.

Top Categories of Exposed AI Software




Figure 39. Workflow automation and AI agents, more than chat interfaces or model servers, make up most exposed AI software. Data via Censys

The long tail of products holds some interesting insights as well. It includes tools that probably don't immediately come to mind when picturing the “AI attack surface”.

Long Tail of Exposed AI Software Categories




Figure 40. Vector databases, which hold the data behind AI applications, lead the next tier of exposed AI software categories. Data via Censys

The AI Attack Surface

This includes:



Vector databases

like Qdrant and Weaviate, which store embeddings used by retrieval-augmented generation systems.



Distributed compute infrastructure

such as Ray Dashboard, used to manage workloads for model training and inference.



Training data-labeling platforms

like Label Studio for data annotation.



LLM observability platforms

like Langfuse designed for monitoring prompts, completions, traces, and costs.



ML experiment-tracking systems

such as MLflow and ClearML, designed for tracking model parameters during development.



Model-serving infrastructure

such as X inference and BentoML that is used for both deploying models and exposing them through APIs for inference.

What this all indicates is that **there's already a characteristic AI infrastructure stack emerging on the Internet**. It's not just exposed LLM interfaces or models, but all of the other layers of supporting infrastructure used to build and operate AI applications and workloads, including separate tools dedicated for orchestration, inference, data management, observability, and more specific use cases.

Each layer interacts with different systems and different kinds of data, which means that the security impact of exposure depends heavily on what role that component plays in the stack.

Where This Infrastructure Lives

While we observe these products being hosted across the globe, exposure concentrates in a few countries. The United States leads with 62,687 hosts, followed by China at 42,511 and Germany at 31,222. Together, these three countries account for just over half of the Internet-facing AI infrastructure we observed.

Top Countries Hosting Exposed AI Software




Figure 41. The United States leads any single country's exposure, but six of the top ten are in Asia. IP geolocation shows where servers are hosted, not necessarily where operators are based. Data via Censys

The U.S.'s high exposure here is consistent with its outsized role in global cloud and hosting infrastructure. The high concentration in China as well is notable but not altogether surprising. The Chinese government has made AI adoption a policy objective through its "AI+ Initiative," which encourages integrating AI into technology used across virtually every sector and industry. That national push is reinforced by local programs offering financial incentives for AI compute resources.

The broader Asia-Pacific presence here is noteworthy as well: South Korea, Singapore, Hong Kong, Japan, and India all rank among the ten countries with the most exposed AI infrastructure. Among these, South Korea's exposure shows an unusually uniform footprint: half of the exposed AI hosts geolocated there (7,133) are systems running just Ollama and n8n together in identical configurations: on the same versions of each tool, the same non-standard ports, and the same preloaded models. Patterns like this suggest potential regional differences in how AI infrastructure is deployed, and point to the role that common templates or installers probably have in spinning up multiple tools at once.

Networks tell a similar story, with cloud providers being the popular choice for hosting this infrastructure.

Top 10 Networks Hosting Exposed AI Software




Figure 42. Amazon AWS leads, but no single cloud provider dominates AI hosting, with a mix of Chinese clouds and European hosting filling much of the top ten.

Data via Censys

Amazon leads at 31,968 hosts (11.9%), then Tencent Cloud at 26,437 (9.9%), Hetzner at 23,844 (8.9%), DigitalOcean at 22,393 (8.4%), and Alibaba Cloud at 18,378 (6.9%). What's interesting here is that exposure seems to be spread more uniformly across the top five hosting providers than we typically see: **no single provider seems to have a clear concentration of exposure**. The remainder spreads across a long tail of smaller VPS operators.

This paints a clearer picture of where on the Internet we see this infrastructure being stood up.

Next we take a look at answering the question: how did this exposure change over time?

A Fast-Growing and Volatile Attack Surface

AI-related infrastructure exposures rose more than 63% over the nine-month period we tracked. 267,915 distinct IPs now expose one of 44 detected products to the public Internet, up from 164,088 in October 2025. When we normalize for fluctuations in the size of the broader Censys scanning dataset over the same period, the increase is trivially higher, at 64.1%.

Note that MCP endpoints are excluded from this longitudinal analysis because we started scanning those in May 2026, making a comparable nine-month baseline unreliable.

AI Software Exposures on the Internet Grew 63% Over Nine Months




Figure 43. This rapid growth and subsequent fluctuations demonstrate how quickly the Internet landscape can shift.

Data via Censys

Exposure grew especially quickly in early 2026, peaking at nearly **300,000 hosts in early April**. When we break this down by specific technologies, we can observe how some products grew steadily in exposure while others skyrocketed in adoption almost overnight. Figure 44 depicts how the top 8 most exposed products we tracked changed over time.

Historical Growth of Top 8 Exposed AI Products




Figure 44. Most of the growth in AI exposures comes from the increasing emergence of n8n, OpenClaw, and New API deployments. Data via Censys


These aggregate counts are useful to understand how the landscape has changed at a birds-eye view. When we contextualize these numbers further, we discover very different deployment and security stories. Two of the largest populations we observed here, OpenClaw and n8n, illustrate different ways of understanding “AI security posture” and what we might expect in the future of AI security.

The volatile orange line that emerged in early January reflects the exposure of OpenClaw control panels, which became one of the more viral AI security stories of 2026.

OpenClaw and the Product Hype Cycle

OpenClaw is an open-source, self-hosted AI agent designed to act autonomously on a user's behalf, which means users can connect it to their personal accounts to complete functions such as responding to emails or posting on social media. Its Control UI gives the operator privileged access to the agent and everything wired into it, meaning it absolutely should not be exposed directly to the public Internet.

The project became instantly popular, and users exposed their control panels with alarming speed. The first release including this control panel component ("clawdbot v2026.1.5") was published on January 5, 2026. Six days later, we detected just **two control panels** exposed online. On January 20 there were 17. By February 3, there were **17,355 hosts**, and on April 7 the count peaked at **85,478 hosts** before falling to 54,375 by July 7.



Example of an exposed OpenClaw control panel

The Meteoric Rise of OpenClaw




Figure 45. OpenClaw exposures rapidly rose and fell within months.

Data via Censys

It quickly became apparent that what started out as a personal project intended for a specific use case was being deployed in record numbers by users who spun it up without fully understanding the security implications. **CVE-2026-25253** in OpenClaw 2026.1.29 is a vulnerability that could leak an OpenClaw authentication token through a crafted link, giving an attacker operator-level control of the gateway and **potentially arbitrary code execution on the host**. It could even affect localhost-bound instances because the victim's browser acted as the bridge. Compromising the gateway could give an attacker access to essentially any of the user's personal accounts that were hooked up to the agent.

China became a major host of the overall OpenClaw control panel footprint very quickly. Its share of observed exposure grew from roughly a quarter to nearly half over the course of February. Interestingly, shortly after this, in March, **Chinese authorities reportedly moved to restrict OpenClaw use** in parts of the state sector, citing security concerns. Exposures in China then rose over 80% over the first half of the month, before cutting down by 52% by July. Geolocation is an imperfect proxy for who operates a host, but these trends do seem consistent with the directive having some effect.

Early Exposures of OpenClaw Control Panels were Concentrated in China




Figure 46. Exposed OpenClaw panels in China made up nearly half the total at the mid-March peak, just after state-sector restrictions were publicly reported.

Data via Censys

OpenClaw eventually became an Internet meme in its own right, but it also reignited a serious and age-old question: **who should own the security model for software**, particularly experimental or open-source tools? Should it be the developer who writes the tool, the end-user who deploys it, or even policymakers? OpenClaw's documentation warned prominently about the risks, and side projects can't always ship mature security controls. The danger is on the other end, when the person wiring an agent into their accounts can't fully assess what they're connecting. For an IT team, the nightmare scenario is an employee granting an agent like OpenClaw access to a corporate mailbox and internal resources with no oversight, access they'd think twice about giving another person.

The immediate impact here probably landed on hobbyists rather than large enterprises, but the broader implication is the speed at which it all happened: **a new AI product can go from obscurity to 85,000 exposed hosts in weeks**, faster than security practices can catch up. The next story may be a product that sits closer to enterprise data.

n8n: The Risks of the Exposed AI Orchestration Layer


n8n is the most widely exposed AI-related tool we observe by a large margin: **113,926 Internet-facing IPs**, more than twice as many as OpenClaw. Note that this measures self-hosted n8n deployments.

One might not think of n8n as an AI tool at first, and a few years ago it wasn't one. It began as a workflow automation platform for connecting services like Slack, GitHub, Google Drive, and databases, but has since increasingly positioned itself as a **popular AI orchestration platform**, with built-in support for agents, vector stores, model providers, and MCP.

That matters from a security perspective because n8n is an execution engine that sits in a **highly privileged spot** between external inputs, workflow logic, stored credentials, and downstream services, so a compromised n8n instance can expose far more than the platform itself.

n8n has also had a very busy year on the vulnerability front. **More than 80 CVEs were published in roughly nine months**, including code execution, credential exposure, file access, SQL injection, and flaws in individual workflow nodes. One of the highest-profile examples was **Ni8mare (CVE-2026-21858)**, an unauthenticated arbitrary file-read vulnerability that researchers showed could be chained into **full instance takeover**.

Risk also extends to n8n's **community-node ecosystem**, a marketplace of npm packages installable directly from the UI. In January 2026, attackers published malicious packages posing as legitimate integrations, with at least eight packages reported to harvest and exfiltrate credentials during workflow execution. Because community nodes run with broad privileges inside the n8n environment, **a malicious node can gain access to sensitive workflow data and credentials used by the services n8n orchestrates**.



Example exposed n8n gateway

Securing n8n is also complicated by the fact that many workflows depend on public webhook triggers, meaning some portion of the service often needs to remain Internet-reachable. That does not mean the n8n service itself should be exposed directly. Traffic should generally be routed through a reverse proxy rather than exposing TCP/5678 to the public Internet. Yet **67% of the exposed n8n instances we observed had port 5678 directly reachable**, and another 970 exposed adjacent ports 5679 or 5680.

89% of exposed n8n hosts reveal a parseable version, giving us a good view into how current these deployments are. The picture is highly fragmented across both 1.x and 2.x releases. n8n also releases software extremely frequently, so age alone understates the gap: the median deployment was **111 days old**, but that translated to being **103 releases behind** the latest version available at the time.

How Many Releases Behind Latest are Exposed n8n Deployments?




Figure 47. The median n8n instance was about 100 releases, or four months, behind the latest.

Data via Censys

That lag matters given the pace of vulnerability disclosure. More than three quarters of versioned n8n hosts, 76,991 (76.1%), fell within the affected range of at least one vulnerability with an NVD-assigned **CVSS score of 9.0 or greater** disclosed between December 2025 and May 2026. Five of those 24 CVEs require additional conditions to be exploitable, such as particular execution modes, enabled node types, or elevated user roles. Even excluding them, however, yields a nearly identical figure of **76,833 hosts (75.9%)**. And those privilege requirements do not necessarily remove the risk: in an exploit chain, an authentication bypass or other initial foothold can make a post-authentication vulnerability reachable.

Figure 48 demonstrates that most are affected by [CVE-2026-44789](#), [CVE-2026-44790](#), and [CVE-2026-44791](#): all three share the same affected range and each individually covers 76,833 hosts, encompassing paths to prototype pollution, arbitrary file access, and remote code execution through nodes that are enabled by default. Ni8mare, mentioned above, presents a more encouraging picture: only 6,963 hosts, or 6.9% of versioned deployments, were running a release within its affected range, though its severity and unauthenticated attack vector make even that fraction alarming.

n8n is a useful example of why the connective tissue of the AI stack deserves particular attention. Orchestration platforms sit between models, credentials, applications, and downstream systems. When one is exposed, outdated, or compromised, the consequences can extend well beyond the platform itself.

n8n Instances Affected by High-Impact CVEs




Figure 48. About 75% of versioned n8n instances are running a version potentially affected by the latest CVE batch.

Data via Censys

Takeaways

The individual products in this dataset will change as new AI tools emerge. New interfaces, orchestration systems, model servers, and deployment frameworks will continue to be adopted as more resources are put into development. We are also watching AI become its own infrastructure stack, an ecosystem including models connected to interfaces, workflow engines, gateways, databases, and other supporting systems. The attack surface grows not simply when another AI application is deployed, but when **each layer supporting that application becomes independently reachable from the Internet**, and therefore something that needs to be secured.

Three findings stand out:

01



The exposure landscape can change extremely quickly.

OpenClaw moved from two observed systems to tens of thousands of hosts within weeks.

02



Available deployment templates have an outsized effect on the exposure surface.

The common exposure of Ollama+n8n together shows how templated deployment mechanisms for AI tools can reproduce the same exposure thousands of times.

03



Risk depends on both a component's function and its current state.

n8n demonstrates what happens when a single product is trusted with the keys to the kingdom across potentially hundreds of services and agents.

Another reminder that Internet exposure of an asset does not, by itself, prove that that asset is insecure. It does, however, quantify and add shape at Internet-scale to an ongoing change in the security landscape: AI is no longer a single new product category on the Internet. It is becoming an infrastructure layer, and its attack surface is taking shape in real time.

6

Conclusion

Conclusion

This is the first version of the State of the Internet Report built to be compared against itself year over year, and it illustrates why **measuring the Internet over longer time frames matters**. The Internet does not change at one speed. As security practitioners, we want to know: how much does aging infrastructure simply stay put, and how quickly can emerging technology reshape what defenders have to worry about? Exposed AI infrastructure rose 63% in nine months, while ICS exposure moved only about 4% over two and a half years. **A single snapshot, whether of a day or a month, would miss much of that story.**

Some of what we measured this year will look very different in 2027, and some of it may barely move. We expect the AI-related product landscape to churn heavily. The oldest deployments of legacy software may sit at the same exposure levels they're at now. The value of measuring both is in seeing those changes, and that lack of change, **against the same dataset over time.**

This report gives us the baseline, and the years that follow will show which parts of the Internet move quickly, which are stagnant, and how the balance between the two changes.

7

Appendices

Glossary

Measurement and Protocols

EOL (End of Life)

The point at which an application or product stops receiving regular fixes and maintenance, including security patches for newly disclosed vulnerabilities

CWMP

CPE WAN Management Protocol, used by ISPs to remotely configure customer-premises equipment (CPE)

DNS (Domain Name System)

Resolves domain names to IP addresses

FTP (File Transfer Protocol)

Unencrypted service for transferring files between remote hosts

HTTP(S) (Hypertext Transfer Protocol)

Exchanges data between web servers and web clients; HTTPS is the TLS-encrypted variant

IMAP (Internet Message Access Protocol)

Allows email clients to interact with mail stored on a server

NTP (Network Time Protocol)

Clock synchronization across networked systems

POP3 (Post Office Protocol)

Used by email clients to retrieve mail from a server

SMTP (Simple Mail Transfer Protocol)

Delivers email from a mail client to a server

SOCKS

Protocol for proxying traffic between a client and a server

TLS and Certificates

Certificate Authority (CA)

A trusted entity that verifies the identity of a certificate requester, signs and issues the certificate, and manages its lifecycle; a publicly trusted CA is one whose roots ship in major browsers

DV Certificate

Domain Validated, issued on proof of control over a domain, often via a TXT record; the most common certificate type in 2026

EV Certificate

Extended Validation, the strictest verification available, requiring confirmation of legal existence, physical address, operational existence, and a callback to a verified phone number; least common as of 2026

OV Certificate

Organization Validated, issued after verifying the legitimacy of the requesting organization

Self-Signed Certificate

A certificate signed by its own key rather than by a CA, providing encryption without the trust associated with a browser-trusted CA

TLS (Transport Layer Security)

Successor to SSL, providing encryption for Internet traffic

ICS Protocols and Services

HMI (Human-Machine Interface)

A graphical interface that lets operators remotely monitor and adjust an industrial process

PLC (Programmable Logic Controller)

The industrial computer that executes control logic on physical equipment

Specific ICS Protocols

ATG (Automated Tank Gauge)

Monitors levels of tank contents, often fuel, over time

BACnet

Building automation and control, covering HVAC, lighting, and building access control

CIMON PLC

Communications protocol for the CIMON programmable logic controller

CODESYS

Hardware-independent automation software used to program and debug PLCs

C-more

Communications protocol for the C-more HMI, which lets operators monitor and interact with industrial control systems

DIGI

Device discovery protocol, often used in industrial settings

DNP3 (Distributed Network Protocol 3)

Communications protocol widely used in North American electric utility systems

E/IP (Ethernet Industrial Protocol)

Exchanges data between device types such as PLCs, HMIs, and controllers, encapsulated inside CIP (Common Industrial Protocol)

FINS (Factory Interface Network Service)

Proprietary protocol for Omron PLCs and HMIs, allowing read/write access to the PLC

FOX

The Tridium Niagara protocol for building automation and control, including HVAC and other facilities management processes

GE SRTP (Service Request Transfer Protocol)

Facilitates communications between GE PLCs and other devices

HART (Highway Addressable Remote Transducer)

Open protocol combining analog and digital communication for industrial systems

IEC 60870-5-104

Part of the IEC 60870 standards series for electrical engineering applications

MMS (Manufacturing Message Specification)

Transfers process information among networked devices in industrial settings

Modbus

Serial communications protocol connecting PLCs, sensors, and other devices in industrial environments

OPC UA (Open Platform Communications Unified Architecture)

Emphasizes interoperability among devices from different manufacturers

PCOM

Proprietary protocol used by Unitronics PLCs

PCWORX

Proprietary protocol for communicating with Phoenix Contact PLCs

ProConOS

Proprietary protocol used by systems running the ProCon operating system

Red Lion Crimson

Software and communications protocol used to configure Red Lion HMIs

S7

Proprietary Siemens protocol used between HMIs and PLCs in industrial environments

WDBRPC (Wind River Debug)

Debug protocol for Wind River's VxWorks real-time operating system (RTOS)

Threats

Acunetix

Dynamic application security testing (DAST) tool for web vulnerability scanning and penetration testing, covering OWASP Top 10, XSS, SQL injection, and out-of-band vulnerabilities

Asset Reconnaissance Lighthouse (ARL)

Automated asset discovery platform providing subdomain enumeration, port scanning, service identification, crawling, screenshot capture, and file-leak detection

Bulletproof Hosting Provider

A host that tolerates malicious or abusive activity on its infrastructure while resisting takedown requests and law enforcement pressure

Fake Captcha

Social engineering infrastructure that presents a fake CAPTCHA or browser error and instructs the user to paste and run a malicious command, typically delivering an infostealer

Fake EZ-Pass

SMS-based phishing pages impersonating U.S. toll collection services to steal payment credentials and personal information

FRP

Open-source reverse proxy that exposes local servers to the Internet through NAT or a firewall, supporting HTTP, HTTPS, and TCP

Gophish

Open-source phishing framework used by red teams to simulate phishing attacks and assess security awareness

Interactsh

Out-of-band application security testing (OAST) tool, widely repurposed by attackers since 2021 to validate exploits before use

mitmproxy

Open-source interactive HTTPS proxy that can intercept, inspect, modify, and replay web traffic

Nessus

Commercial vulnerability scanner from Tenable that identifies security weaknesses, misconfigurations, and compliance violations

Nezha Monitoring

Self-hosted server monitoring agent with push notifications, scheduled task execution, and a built-in web terminal, abused for persistent remote access

NPS

Intranet penetration proxy server with centralized management, supporting TCP, UDP, HTTP(S), SOCKS5, and P2P

nps-socks5

Shell-script variant of NPS focused on rapid SOCKS5 proxy deployment

ShellInABox

Web-based shell for executing commands on a remote machine

Suspicious Open Directory

A web-accessible directory exposing security tools, penetration testing utilities, webshells, or other potentially malicious files

VPS (Virtual Private Server)

A virtualized server instance rented from a hosting provider, a common option for self-hosted software

Vshell

Open-source command-and-control (C2) framework written in Go, with implants for Windows, Linux, and macOS

Vulnerability Response

CISA KEV (Known Exploited Vulnerabilities catalog)

A U.S. government list of vulnerabilities with confirmed evidence of active exploitation

CVSS (Common Vulnerability Scoring System)

A framework for rating vulnerability severity on a scale from 0 (minimal impact) to 10 (high impact)

CWE (Common Weakness Enumeration)

A classification system for the software and hardware weaknesses that typically lead to vulnerabilities

NVD (National Vulnerability Database)

A U.S. government repository of publicly disclosed vulnerabilities with standardized metadata such as severity scores and affected components

RR (Rapid Response)⁷

A Censys research program measuring the Internet-facing impact of critical and emerging vulnerabilities

⁷ <https://censys.com/censys-arc/rapid-response-advisories/>

AI Exposure

AI Agent

An AI system that acts on a user's behalf beyond generating content, typically given access to accounts, external tools, or APIs

AI Model

The trained artifact at the center of an AI system, a set of learned parameters that takes an input and produces an output such as generated text, a classification, or an embedding

LLM (Large Language Model)

An AI model trained on very large volumes of text to predict and generate language

LLM Gateway

A proxy that routes requests to one or more model providers, usually holding the API credentials for each

MCP (Model Context Protocol)

An open protocol connecting AI models to external tools and data sources; an MCP endpoint is a network-reachable server exposing those tools to a model

RAG (Retrieval-Augmented Generation)

An architecture that retrieves relevant documents from an external store, such as a vector database, at query time and supplies them to a model as context

Vector Database

Stores embeddings, the numeric representations of text or other data that RAG systems query at retrieval time

References

1

The Shape of the Internet

- 1 **IANA Service Name and Transport Protocol Port Number Registry**
<https://www.iana.org/assignments/service-names-port-numbers>
- 2 **Akamai Community Documentation on SSL/TLS Certificate Chains for Akamai-Managed Certificates**
<https://community.akamai.com/customers/s/article/SSL-TLS-certificate-chains-for-Akamai-managed-certificates>
- 3 **Chromium Project Documentation on Moving the EV Indicator Out of the URL Bar**
<https://chromium.googlesource.com/chromium/src/+HEAD/docs/security/ev-to-page-info.md>
- 4 **Cloudflare Learning Center Explainer on BGP Hijacking**
<https://www.cloudflare.com/learning/security/glossary/bgp-hijacking/>
- 5 **Let's Encrypt Announcement of Its First IP Address Certificate**
<https://letsencrypt.org/2025/07/01/issuing-our-first-ip-address-certificate>
- 6 **Let's Encrypt Documentation on the Staging Environment**
<https://letsencrypt.org/docs/staging-environment/>
- 7 **Let's Encrypt Announcement of Six-Day and IP Address Certificate General Availability**
<https://letsencrypt.org/2026/01/15/6day-and-ip-general-availability>

2

The Long Tail of Vulnerability Response

- 1 **Censys ARC Rapid Response Advisories**
<https://censys.com/censys-arc/rapid-response-advisories/>
- 2 **NVD Entry for CVE-2024-50623 in Cleo Harmony, VLTrader, and LexiCom**
<https://nvd.nist.gov/vuln/detail/cve-2024-50623>
- 3 **NVD Entry for CVE-2024-55956 in Cleo Harmony, VLTrader, and LexiCom**
<https://nvd.nist.gov/vuln/detail/cve-2024-55956>
- 4 **BleepingComputer Reporting on CIOp's Claim of Responsibility for the Cleo Data Theft Attacks**
<https://www.bleepingcomputer.com/news/security/clop-ransomware-claims-responsibility-for-cleo-data-theft-attacks/>
- 5 **NVD Entry for CVE-2024-4358 in Progress Telerik Report Server**
<https://nvd.nist.gov/vuln/detail/cve-2024-4358>
- 6 **NVD Entry for CVE-2024-23897 in Jenkins**
<https://nvd.nist.gov/vuln/detail/cve-2024-23897>
- 7 **NVD Entry for CVE-2024-38094 in Microsoft SharePoint Server**
<https://nvd.nist.gov/vuln/detail/cve-2024-38094>
- 8 **NVD Entry for CVE-2025-53770 in Microsoft SharePoint Server**
<https://nvd.nist.gov/vuln/detail/cve-2025-53770>
- 9 **CISA Known Exploited Vulnerabilities Catalog**
<https://www.cisa.gov/known-exploited-vulnerabilities-catalog>
- 10 **Jenkins Documentation on the Weekly Release Line**
<https://www.jenkins.io/download/weekly/>
- 11 **Jenkins Documentation on the LTS Release Line**
<https://www.jenkins.io/download/lts/>
- 12 **Microsoft Learn Release Notes for SharePoint Updates**
<https://learn.microsoft.com/en-us/officeupdates/sharepoint-updates>
- 13 **Progress Telerik Report Server Release History**
<https://www.telerik.com/support/whats-new/report-server/release-history>

3

Exposure of Industrial Control Systems and Critical Infrastructure

- 1 **Censys 2024 State of the Internet Report**
<https://go.censys.com/rs/120-HWT-117/images/2024SOTIR.pdf>
- 2 **European Commission Overview of the NIS2 Directive**
<https://digital-strategy.ec.europa.eu/en/policies/nis2-directive>
- 3 **SANS Analysis of NIS2 Compliance for OT Environments**
<https://www.sans.org/blog/nis2-compliance-for-ot-strategic-implementation-of-ics-controls>
- 4 **European Cyber Security Organisation NIS2 Transposition Tracker**
<https://ecs-org.eu/nis2-tracker/>
- 5 **NVD Search Results for Tridium Niagara Vulnerabilities**
<https://nvd.nist.gov/vuln/search#/nvd/home?keyword=niagara&resultType=records>
- 6 **CISA Alert on Exploitation of Unitronics PLCs Used in Water and Wastewater Systems**
<https://www.cisa.gov/news-events/alerts/2023/11/28/exploitation-unitronics-plcs-used-water-and-wastewater-systems>
- 7 **CBS News Pittsburgh Reporting on the Municipal Water Authority of Aliquippa Compromise**
<https://www.cbsnews.com/pittsburgh/news/municipal-water-authority-of-aliquippa-hacked-iranian-backed-cyber-group/>
- 8 **CNN Reporting on a Suspected Russian Hacking Group in the Texas Water Cyberattack**
<https://www.cnn.com/2024/04/17/politics/russia-hacking-group-suspected-texas-water-cyberattack>

4

Threat Infrastructure

- 1 **FRP Project Repository**
<https://github.com/fatedier/frp>
- 2 **NPS Project Repository**
<https://github.com/ehang-io/nps>
- 3 **Nezha Monitoring Project Repository**
<https://github.com/naiba/nezha>
- 4 **Censys Research on EtherHiding, Fake CAPTCHAs, and ClickFix Lures**
<https://censys.com/blog/etherhiding-fake-captchas-click-fix-lures-blockchain-backed-payload-delivery/>
- 5 **Censys Research on Fake Toll Payment Infrastructure, "Highway Robbery 2.0"**
<https://censys.com/blog/highway-robbery-2-0/>
- 6 **New Jersey Cybersecurity and Communications Integration Cell Analysis of Toll Violation Text Scams**
<https://www.cyber.nj.gov/threat-landscape/phishing-online-scams/sms-scams/smishing-at-scale-a-deep-dive-into-toll-violation-text-scams>
- 7 **Censys Research Following a USPS Smishing Kit Through DNS Data**
<https://censys.com/blog/following-a-usps-smishing-kit-through-censys-dns-data/>
- 8 **KrebsOnSecurity Reporting on the China-Based SMS Phishing Triad's Pivot to Banks**
<https://krebsonsecurity.com/2025/04/china-based-sms-phishing-triad-pivots-to-banks/>
- 9 **Resecurity Research on Smishing Triad Targeting Banking Customers in Pakistan**
<https://www.resecurity.com/blog/article/smishing-triad-is-targeting-pakistan-to-defraud-banking-customers-at-scale>
- 10 **Unit 42 Analysis of a Global Smishing Campaign**
<https://unit42.paloaltonetworks.com/global-smishing-campaign/>
- 11 **Censys Research on Trust Infrastructure as a Malware Delivery Interface**
<https://censys.com/blog/living-off-the-web-how-trust-infrastructure-became-a-malware-delivery-interface/>
- 12 **Wikipedia Entry on Malvertising**
<https://en.wikipedia.org/wiki/Malvertising>
- 13 **MDN Web Docs Glossary Entry for Top Level Domain**
<https://developer.mozilla.org/en-US/docs/Glossary/TLD>
- 14 **Wikipedia Entry on the .cc Top Level Domain**
<https://en.wikipedia.org/wiki/.cc>
- 15 **MIT Technology Review Reporting on a Pacific Island as a Global Capital of Cybercrime**
<https://www.technologyreview.com/2023/11/02/1082798/tiny-pacific-island-global-capital-cybercrime/>

5

Exposed AI Infrastructure

- 1 **State Council of the People's Republic of China, Guidance on Implementing the "AI+" Initiative**
https://www.gov.cn/zhengce/content/202508/content_7037861.htm
- 2 **OpenClaw Project Blog, "Introducing OpenClaw"**
<https://openclaw.ai/blog/introducing-openclaw>
- 3 **GitHub Security Advisory GHSA-g8p2-7wf7-98mq for CVE-2026-25253 in OpenClaw**
<https://github.com/openclaw/openclaw/security/advisories/GHSA-g8p2-7wf7-98mq>
- 4 **Bloomberg Reporting on China Moving to Limit OpenClaw Use at Banks and Government Agencies**
<https://www.bloomberg.com/news/articles/2026-03-11/china-moves-to-limit-use-of-openclaw-ai-at-banks-government-agencies>
- 5 **NVD Entry for CVE-2026-21858 ("Ni8mare") in n8n**
<https://nvd.nist.gov/vuln/detail/cve-2026-21858>
- 6 **Endor Labs Research on the Supply Chain Attack Against the n8n Community Node Ecosystem**
<https://www.endorlabs.com/learn/n8mare-on-auth-street-supply-chain-attack-targets-n8n-ecosystem>
- 7 **NVD Entry for CVE-2026-44789 in n8n**
<https://nvd.nist.gov/vuln/detail/cve-2026-44789>
- 8 **GitHub Security Advisory GHSA-57g9-58c2-xjg3 for CVE-2026-44790 in n8n**
<https://github.com/n8n-io/n8n/security/advisories/GHSA-57g9-58c2-xjg3>
- 9 **NVD Entry for CVE-2026-44791 in n8n**
<https://nvd.nist.gov/vuln/detail/cve-2026-44791>

AI Disclosure

The findings in this report are based on Censys Internet scan data. We used AI coding assistants to draft dataset queries and write analysis notebooks, and AI editing tools during prose drafting. No finding or interpretation was generated by an AI tool. All conclusions and editorial decisions are our own, and the authors have reviewed and verified the full text.

Questions and Feedback

Media inquiries and requests to cite or feature this data can be sent to press@censys.com.



About Censys

Censys is the authority for Internet intelligence and insights. Delivering the most complete, accurate, and up-to-date global map of Internet infrastructure, Censys provides industry leading solutions for attack surface management, threat hunting, and proactive incident response. Global governments, Fortune 500 companies, and security providers around the world trust Censys to uncover risks faster, respond more effectively, and prevent breaches before they happen.

Contact Us

VISIT
censys.com ➤

CONTACT
hello@censys.com ➤