Presented by

Trellix ADVANCED RESEARCH CENTER

THE CYBERTHREAT REPORT

November 2026
Executive Summary

The past six months have seen a fundamental structural shift in the global threat landscape. Modern security leaders can no longer evaluate cyber risk through the lens of isolated breach events or localized threat actor campaigns. Instead, the current environment is defined by the convergence of nation-state intelligence tradecraft, market-driven cybercriminal operational efficiency, and machine-speed autonomous execution. As state-sponsored groups prioritize long-term stealth and non-disruptive espionage, commercial cybercriminals are optimizing extortion economics through AI and exploiting systemic architectural gaps in enterprise workflows.

This executive summary offers key takeaways and a high-level summary of the November 2026 CyberThreat Report findings, covering the latest global threat actor trends, geopolitical activity, and emerging AI agent exploits. For industry threat breakdowns, technical exploitation trends, and specific APT attacks, dive into the complete CyberThreat Report.

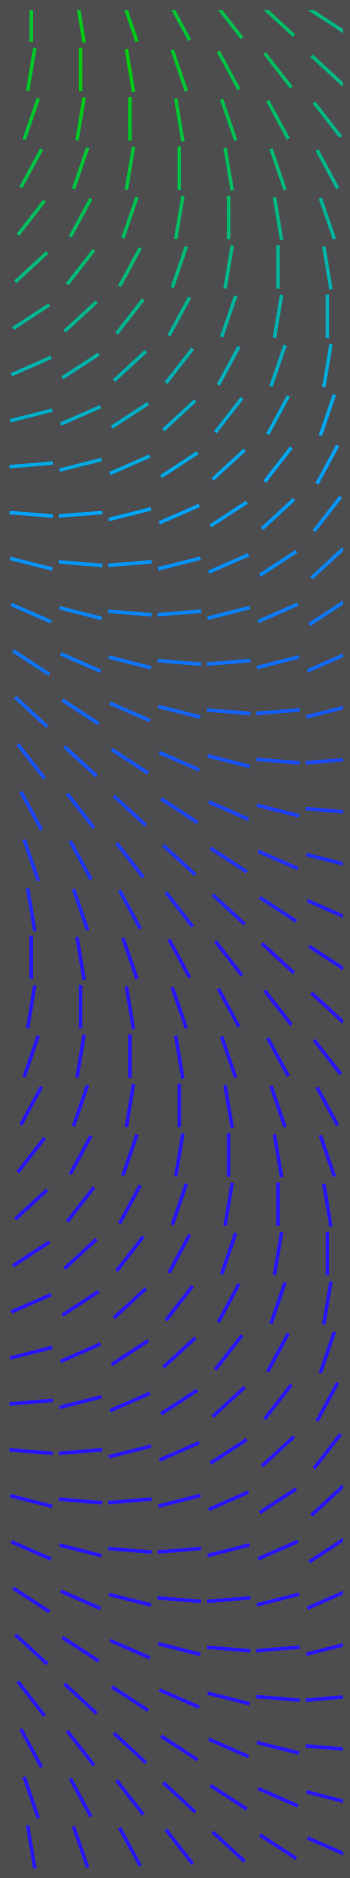

METHODOLOGY

Trellix and experts from our Trellix Advanced Research Center (ARC) compile the statistics, trends, and insights that comprise this report from a wide range of global sources, both open and captive. The aggregated data is fed into Trellix Insights and Trellix ATLAS. Leveraging AI, machine learning, automation, and human expertise, the team cycles through an intensive, integrated, iterative process—normalizing data, analyzing information, and developing insights meaningful to cybersecurity leaders and SecOps teams on the front lines worldwide.

A LOOK INTO THE EVOLVING APT LANDSCAPE

Analysis of advanced persistent threat activity (APT) shows a heavily concentrated focus by nation-state actors on critical infrastructure. The dominance of a two-nation axis, alongside a diverse cohort of threat groups, underscores a systemic risk where espionage operations target vital operational networks and communications backbones.

Key findings:

- **APT targeting is dominated by a two-nation axis,** with the United States and Türkiye jointly representing over three-quarters of all detections.
- **Transportation and shipping and telecom** together absorb over half of all detected campaign volume, reflecting an intense focus on intelligence collection across logistics, global supply chains, and communications backbones.
- **The top 10 actors collectively represent 56% of all detections,** pointing to a core group of highly active, well-resourced operators surrounded by a long tail of 287 lower-volume actors.

TOP 10 AFFECTED COUNTRIES: (APRIL–SEPTEMBER 2026)

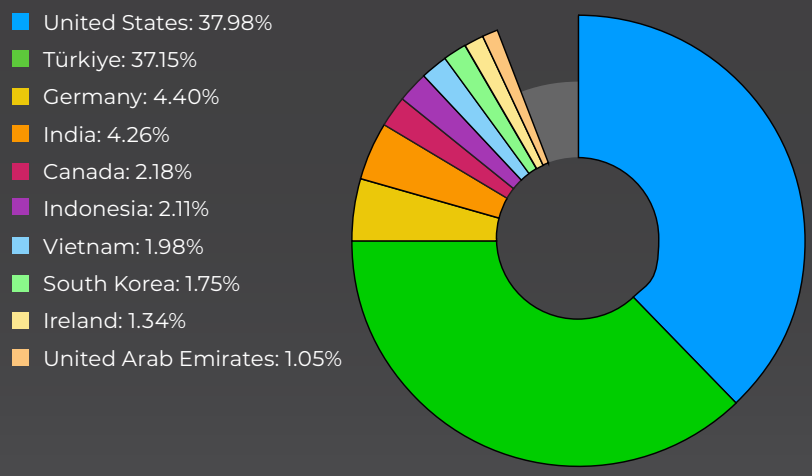

CISO TIPS:

- Establish thorough activity baselines and clear administrative domain segmentation, enforcing strict least-privilege access controls.
- Prioritize ongoing threat hunting to identify abuse of native tools and rigorous auditing of privileged accounts.
- Align threat intelligence, IT, and SOC teams to spot subtle indicators of compromise, including unauthorized lateral movement.

RANSOMWARE: A CROWDED MARKET WITH FRAGMENTED POWER

With 124 active groups, 55 newcomers, and leaders that change from one report to the next, defenses built around specific brands will always lag behind. The steps of an intrusion change far less than the names. In Trellix telemetry over the past six months, the tools most often linked to ransomware were the Windows command shell, PowerShell, PsExec, Net, and Mimikatz, and Data Encrypted for Impact (T1486), which appeared on 96% of ransomware detections.

Key findings:

- **A new number one.** The Gentlemen posted 719 victims (13.1% of all victim posts) and overtook Qilin as the most active ransomware group, less than a year after it first appeared.
- **A record, crowded market.** We observed 5,505 victim posts from 124 active groups, both the highest of any six-month period in our data. Also, 55 groups posted for the first time and accounted for 22% of all victim posts.
- **The US share is falling.** The United States was the target of 41% of geo-identified victim posts, down from 54% in the previous six months, as more groups went after European victims. Trellix telemetry still recorded more unique ransomware campaigns in the US (714) than in any other country.

TOP 5 MOST TARGETED SECTORS (APRIL–SEPTEMBER 2026)

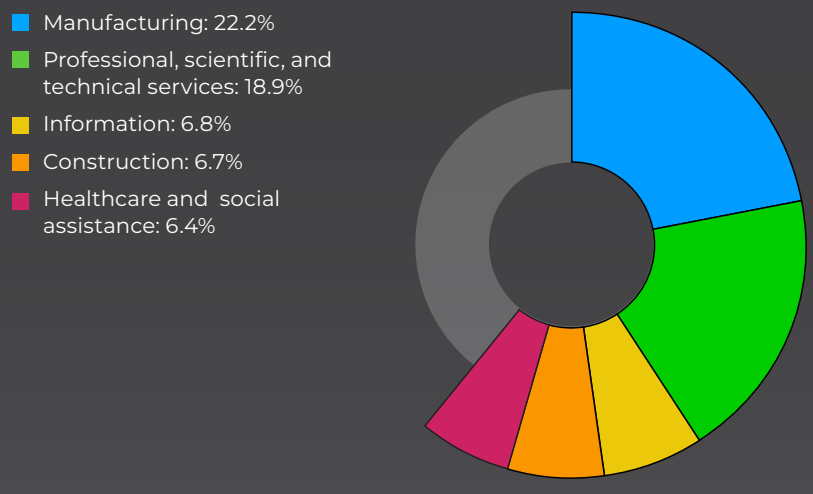

CISO TIPS:

- Make edge-device patching a tracked, time-bound metric reported to leadership.
- Require tamper protection and vulnerable driver blocking on every endpoint.

TOP 5 MOST TARGETED COUNTRIES (APRIL–SEPTEMBER 2026)

- United States: 41.3%
- United Kingdom: 4.5%
- Germany: 4.3%
- Canada: 4.1%
- France: 3.7%

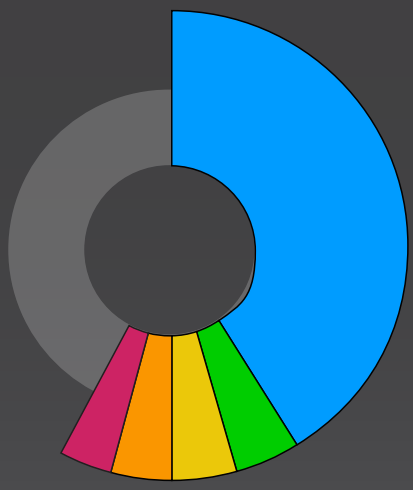

AI IN THE UNDERGROUND: THE EMERGENCE OF AGENTIC MALWARE

The past six months are defined by two developments: AI agents increasingly act with little or no human direction, and agentic malware capable of autonomous decision-making is emerging.

Key findings:

- **AI has evolved** from a drafting aid into an autonomous operator capable of scanning, exploiting, propagating, and executing entire intrusions with minimal or no human direction.
- **Autonomous agents and agentic malware** now independently execute entire intrusions, from targeting and exploit validation to exfiltration and cleanup, with minimal or no human direction.
- **The cybercriminal marketplace has split** into a low-cost, saturated tier of cheap uncensored chatbots and a high-cost tier offering working autonomous agents and automated extortion capabilities.

THE UNSEEN CREDENTIAL SURFACE: HOW AI AGENTS EXPOSE SECRETS

AI coding and operations agents run on a loop of writing, executing, and reading shell commands. While operating systems prioritize keeping secrets encrypted in memory, agents routinely push them into exposed channels. This systemic, accidental exposure forms what we call the unseen credential surface.

CISO TIPS:

- Verify that all internal AI endpoints, API keys, open-source agent frameworks, and third-party model connections are cataloged, restricted, and continuously monitored against unauthorized access or C2 misuse.
- Configure detection models to identify machine-speed autonomous intrusions, model consensus voting, and agentic malware capable of real-time sandbox evasion and target triage.
- Align incident-response playbooks to counter sub-hour, automated intrusions and agentic malware operating without traditional human command-and-control bottlenecks.

Key findings:

- **AI agents** have become one of the enterprise's largest and least protected credential-handling layers—exposed by design.
- **Every exposure is living off the land,** so single-event detection fails.
- **No single event separates the agent from an attacker,** so organizations need to correlate the full chain of credential access, use, and egress back to the agent process that started it.

EXPLOITING THE KNOWN: ENTRENCHED VULNERABILITIES, HEAVILY TARGETED FRAMEWORKS, AND AN EMERGING AI FRONTIER

In April–September 2026, exploit detections point overwhelmingly to known, already-patchable vulnerabilities, led by internet-facing infrastructure such as F5 BIG-IP. Alongside that familiar pattern, more recent application-framework flaws are already generating detection activity on the scale of long-standing bugs. At the leading edge, AI-orchestration and Model Context Protocol (MCP) tooling has become a target class that external research confirms is being exploited, even as our telemetry shows little of it yet.

Key findings:

- **Legacy flaws on internet-facing infrastructure, not zero-days, drive the heaviest exploit-detection volume.** F5 BIG-IP's 2022 authentication-bypass, CVE-2022-1388, is the single most-detected exploit CVE in Trellix NX telemetry (1.2M detections), and F5 BIG-IP is the most-targeted product overall (1.5M detections, 5.0%); the same CVE also ranks among Trellix Helix's top exploit rules.
- **The application-framework flaw, React2Shell, shows heavy detection activity across all telemetry.** The React Server Components pre-authentication RCE CVE-2025-55182 (CVSS 10.0, KEV-listed, known ransomware use) registers 214K detections in NX, 1.02M alerts in Trellix Intrusion Prevention System (IPS), and 6K detections in Trellix Helix—where it is one of the few exploit rules rising quarter over quarter.
- **Within the reporting window, detection volume for most classic N-day exploit rules fell while a small set rose, a shift only period-over-period telemetry reveals.** In Trellix Helix, detections against Apache (CVE-2021-41773), Log4Shell, Palo Alto (CVE-2024-3400), F5 (CVE-2022-1388), and Fortinet (CVE-2022-40684) all declined from Q2 to Q3. At the same time, Atlassian Confluence (CVE-2022-26134, +7.9%) and React2Shell (+7.2%) rose, focusing attention on a short, KEV-listed set rather than the full disclosure flood.

CISO TIPS:

- Treat every endpoint running an AI agent as high-value, credential-handling infrastructure.
- Collapse credential lifetime to short-lived, just-in-time tokens, so that even when an agent pushes a secret down to the command line, disk, or the model context, it expires before anyone can use it.

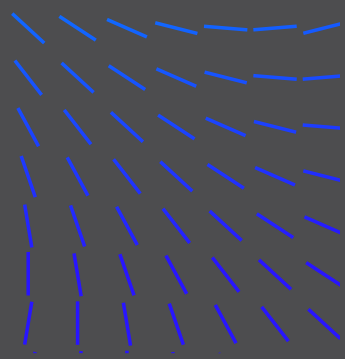

CISO TIPS:

- Prioritize by evidence of exploitation, not by volume of disclosure, and weigh that evidence against your own environment.
- Base remediation urgency on confirmed or predicted exploitation, internet exposure, and asset criticality.
- Measure success by how quickly you mitigate an exploited, internet-exposed flaw on a business-critical system, not by raw patch throughput.

CONCLUSION

In an era defined by autonomous agentic malware, market fragmentation, and persistent edge pressure, maintaining business continuity requires converting strategic threat intelligence into machine-speed containment, continuous architectural hardening, and disciplined exposure management.

Enterprise security posture must evolve beyond signature-centric detection to address an operational environment where human direction is increasingly removed from the attack lifecycle.

By embracing machine-speed defense, organizations can shift from a reactive stance to one of readiness, building cyber resilience for the AI era. For a deeper analysis of the findings, access the complete CyberThreat Report.

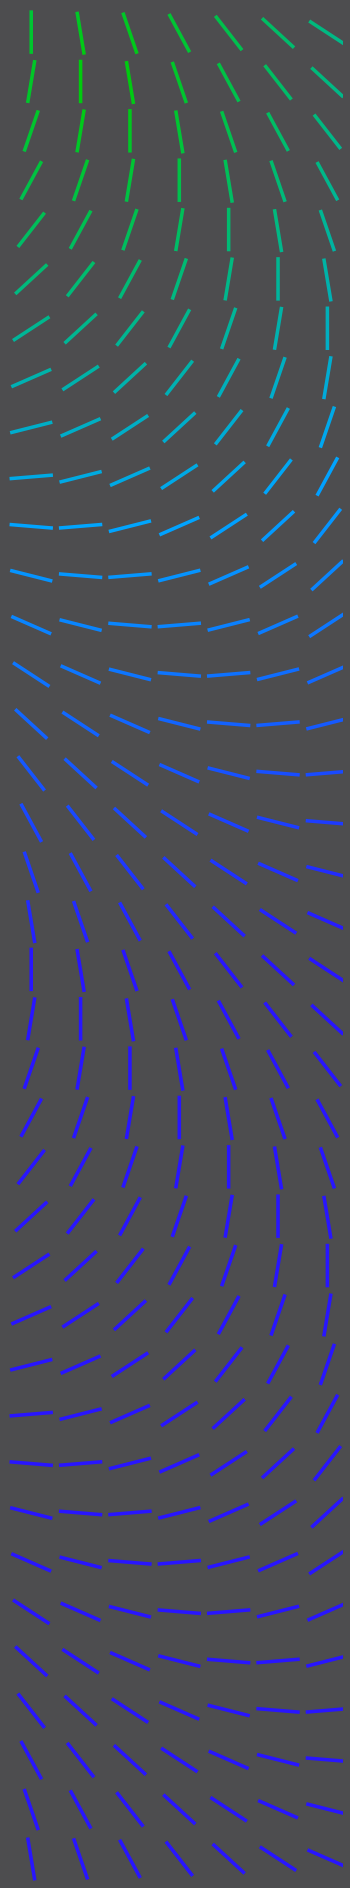

This document and the information continued herein describes computer security research for educational purposes only and the convenience of Trellix customers. Trellix conducts research in accordance with its Vulnerability Reasonable Disclosure Policy I Trellix. Any attempt to recreate part or all of the activities described is solely at the user's risk, and neither Trellix nor its affiliates will bear any responsibility or liability.

Trellix is a trademark or registered trademark of Musarubra US LLC or its affiliates in the US and other countries. Other names and brands may be claimed as the property of others.

Trellix

Copyright © 2026 Musarubra US LLC