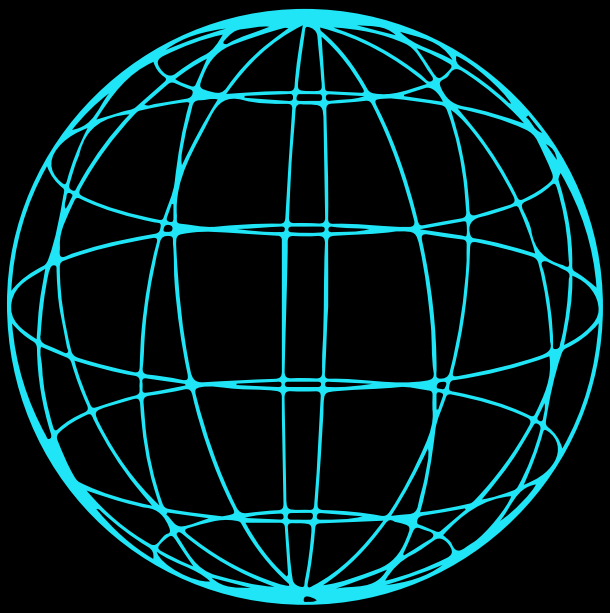

MACHINE-SPEED BREACHES AND THE NEW ERA OF CYBERSECURITY:

AN ANALYSIS OF EMERGING CYBERSECURITY THREATS IN THE AGE OF ARTIFICIAL INTELLIGENCE

MARCH 2026

www.front-code.com

Machine-Speed Breaches in 2026

Identity, AI Agents, Infostealers, and the New Cybersecurity Control Plane

March, 2026

Audience: Security-aware leaders, engineers, and students
Scope: Defensive, non-operational guidance (no offensive instructions)

Front-Code.com

Executive Summary

Cybersecurity in early 2026 is defined by a harsh operational shift: many adversaries are no longer "breaking in" with rare zero-days; they are "signing in" with stolen, replayed, or socially engineered identities. That shift has existed for years, but its speed and scale have changed dramatically as credential theft ecosystems mature and AI compresses attacker workflows.

Incident-response telemetry from Palo Alto Networks Unit 42 illustrates the collapse of time-to-impact. In their 2026 Global Incident Response Report, the fastest 25% of intrusions in calendar year 2025 reached confirmed data exfiltration in about 72 minutes—down from roughly 4.8 hours (285 minutes) in 2024.[1] In other words, a meaningful subset of breaches now complete the "access → lateral movement → data theft" cycle faster than many organizations can convene a war room.

At the same time, organizations are introducing a new class of identities into production: machine identities and AI agents. Service accounts, workload identities, API keys, and autonomous agents frequently receive broad permissions so they can be useful. Threat actors increasingly aim for these identities because they often bypass the friction defenders have added to human logins (MFA, device checks, conditional access).

The hot topic is therefore not a single malware family. It is an architectural imperative: treat identity and authorization as the enterprise control plane, engineer it for machine-speed incidents, and extend the same discipline to AI agents and software supply chains.

This article explains how machine-speed attacks typically unfold; why infostealers and access brokers have become the on-ramp to ransomware and data theft; why LLM-powered systems create distinct security failure modes; and what a practical, prioritized roadmap looks like for the next 90 days.

Key thesis: If you harden identity (human + non-human), constrain authorization, and automate containment, you can reduce both the probability of compromise and the blast radius when compromise inevitably occurs.

Table of Contents

1. The defining change: time-to-impact collapses

Security programs have historically been structured around periodic risk assessments and remediation backlogs. That cadence assumes you can tolerate long dwell times—days or weeks between initial compromise and material attacker actions. Machine-speed intrusions break that assumption.

Unit 42's 2026 incident-response data offers a clear benchmark. The fastest quartile of intrusions reached confirmed data exfiltration in roughly 72 minutes in 2025 (calendar year), down from nearly five hours in 2024.[1] Unit 42 also describes simulated AI-assisted operations that compress the time between compromise and exfiltration to minutes, illustrating how AI can accelerate known techniques even without "novel" exploits.[1]

Time-to-impact collapse changes the defensive optimization target. You still need prevention (patching, segmentation, MFA), but you also need: (a) earlier signals, (b) faster decision loops, and (c) automated containment actions that trigger on high-confidence indicators.

1.1 A reference model for the machine-speed attacker workflow

A common path in 2025–2026 incidents looks like this (details vary by sector and adversary):

- Credential acquisition (infostealer logs, phishing/vishing, credential stuffing, third-party compromise).
- Initial access using valid accounts (VPN, SSO, cloud console, SaaS admin) or exploitation of an exposed edge device.
- Privilege escalation and persistence (token theft, OAuth app abuse, creation of new accounts/keys, abuse of delegated admin).
- Rapid discovery and data staging (cloud storage enumeration, file shares, database snapshots).
- Exfiltration (often from cloud/SaaS sources where egress looks "normal").
- Extortion (ransomware deployment and/or "data leak" pressure).

1.2 What changes for defenders

Machine-speed incidents force three practical changes:

- Instrument identity like you instrument networks: baseline behavior, detect anomalies, and treat high-risk sign-ins as urgent events.
- Make containment policy-driven: auto-revoke sessions, disable accounts, quarantine devices, and rotate secrets when risk crosses a threshold.
- Engineer resilience: offline backups, segmentation, and recovery drills must be treated as production capabilities, not paperwork.

2. Identity is the control plane

In cloud-first enterprises, identity mediates nearly every important action: reading customer data, deploying code, creating infrastructure, authorizing payments, and granting access to suppliers. That is

why attackers increasingly target identity systems, identity proofs, and authorization pathways—especially those that bypass the friction of the normal user login.

Multiple datasets converge on the same story. Microsoft's Digital Defense Report 2025 commentary notes that more than 97% of identity attacks are password attacks and highlights growth in identity-based attacks in the first half of 2025.[2] Verizon's 2025 DBIR emphasizes credential abuse and states that stolen credentials are involved in a large majority of breaches within its "Basic Web Application Attacks" pattern.[3] Unit 42 similarly reports that identity-related weaknesses appear in a very large share of incident-response investigations.[1]

The strategic implication is simple: identity security is not a subset of IT administration. It is core cyber risk management.

2.1 Identity engineering: beyond IAM checkboxes

Identity-centric defense requires an engineering stance across four domains:

- Authentication strength and recovery: phishing-resistant MFA where it matters, and hardened recovery flows that cannot be socially engineered.
- Authorization hygiene: least privilege, separation of duties, and time-bounded access (especially for administrators).
- Non-human identity governance: service accounts, API keys, certificates, and agent identities with minimal privileges and short lifetimes.
- Telemetry and response: identity logs, token events, privileged actions, and automated revocation/rotation playbooks.

2.2 Why "valid accounts" is such a powerful tactic

From an attacker's perspective, using a valid identity collapses multiple defensive layers. It can evade simplistic controls (because the user is "authorized"), reduce noisy exploit indicators, and blend into normal application flows—particularly in SaaS and cloud environments where activity is API-driven.

3. Infostealers and the access broker economy

Infostealer malware has become the industrial-scale input mechanism for identity compromise. These malware families are optimized to extract browser-stored passwords, session cookies/tokens, autofill data, and sometimes cryptocurrency wallets. They are distributed via phishing, malvertising, cracked software, fake installers, and social engineering campaigns.

Mandiant's M-Trends 2025 report explicitly highlights the growing use of infostealer malware to enable intrusions using stolen credentials.[4] The FBI and CISA issued an advisory on LummaC2 (an infostealer) describing how it exfiltrates sensitive information and threatens organizations across critical infrastructure sectors.[5] High-profile disruptions and domain takedowns show how central infostealers have become to the cybercrime ecosystem.[6]

Infostealers do not only create “compromised passwords.” They create complete session material, device fingerprints, and saved secrets—allowing replay attacks even when password hygiene is strong.

3.1 The access broker business model

Infostealer output is monetized through access brokers: actors who buy, package, and sell initial access (VPN credentials, SaaS tenant access, cloud accounts, RDP endpoints) to ransomware crews, fraud groups, and espionage actors. CrowdStrike’s 2025 threat reporting notes increased access broker activity and significant growth in vishing as a credential theft method.[7] CrowdStrike’s 2025 Global Threat Report discusses broader adversary trends—especially identity abuse and social engineering—reinforcing the same access-broker dynamics.[8]

3.2 Defensive moves that specifically counter infostealers

- Adopt phishing-resistant MFA for privileged and remote access paths to reduce the reuse value of stolen passwords.[11]
- Reduce session lifetime and enforce device posture checks for sensitive applications; stolen cookies should not be universally reusable.
- Detect token replay and unusual device fingerprints; alert on “new device + sensitive action” combinations.
- Treat unmanaged endpoints as hostile: require managed devices for admin consoles and high-risk SaaS operations.

4. Ransomware and extortion still pay (and why identity is upstream)

Ransomware remains the most visible monetization layer of cybercrime because it converts access into cash quickly. But the pipeline is increasingly identity-driven: initial access via valid accounts, followed by privilege escalation, data theft, and then extortion—often with or without encryption.

Sophos’ State of Ransomware 2025 reporting notes that organizations are increasingly stopping attacks before encryption, but ransomware remains damaging and costly. Sophos also reports that a substantial portion of victims still pay ransoms, and that backup-based recovery (while crucial) is not a perfect substitute for prevention and containment.[9]

On the payment side, Chainalysis reported that ransomware payments in 2024 were approximately $813.55 million, representing a decrease versus 2023, but still indicating a large and persistent extortion economy.[10] The lesson is not that ransomware is “going away,” but that its economics adapt: groups shift to data theft, “double extortion,” and faster operations when the environment demands it.

CISA’s ransomware guidance emphasizes fundamentals: offline backups, regular restore testing, segmentation, patching known exploited vulnerabilities, and phishing-resistant MFA.[11] Those recommendations are effective precisely because they reduce time-to-impact and blast radius.

4.1 Why ransomware operations are accelerating

Four forces are pushing ransomware operations toward speed and professionalism:

- Access as a commodity: infostealers and brokers reduce the effort needed to get in.
- Automation and templating: scripts and frameworks let operators scale lateral movement and disable controls.
- Hybrid cloud targets: exfiltration from cloud storage can look like normal business activity.
- AI-assisted social engineering: faster personalization for phishing/vishing and more "professional" extortion messaging.[1]

4.2 Minimum viable ransomware resilience

If you must choose a small set of actions that measurably improve resilience, prioritize:

- Offline backups and tested restores (including identity systems and critical SaaS exports).
- Phishing-resistant MFA for admins, remote access, and backup administration paths.[11]
- Segmentation between user networks, servers, backups, and management planes.
- Fast isolation capability: ability to disable accounts and isolate endpoints automatically when high-confidence indicators appear.

5. MFA bypass and session theft: why "MFA" is not a single control

As MFA adoption rose, attackers shifted to bypass strategies that do not look like "password cracking." Common patterns include real-time phishing proxies (which relay MFA), social engineering that pressures users to approve prompts, SIM swapping for SMS-based MFA, and theft of session tokens from browsers.

The key operational insight is that "MFA" is not a single control. Some factors are phishable or socially engineerable (codes and prompts), while phishing-resistant MFA (e.g., FIDO2/WebAuthn) is substantially harder to replay remotely. CISA and the FBI repeatedly recommend phishing-resistant MFA in ransomware advisories.[11]

Therefore, MFA modernization should be risk-based: move the highest-blast-radius identities to phishing-resistant MFA first, then expand coverage.

5.1 An MFA modernization ladder

Maturity	Description
Level 0	Passwords only. High risk.
Level 1	Passwords + SMS/OTP. Better, but still phishable and vulnerable to SIM swap.
Level 2	Push MFA with number matching and risk-based prompts. Useful, but still socially engineerable.
Level 3	Phishing-resistant MFA (FIDO2/WebAuthn, passkeys, smart cards) for privileged and remote access.

Level 4	Continuous access evaluation: device posture + user risk + behavior + token binding.

Treat the ladder as an implementation plan, not a taxonomy: schedule upgrades for the accounts that can do the most damage first.

6. AI agents and machine identities: the next permission crisis

Generative AI is changing enterprise architecture by embedding LLMs into workflows that have real-world effects: customer support, internal automation, data analysis, and incident response. When these systems are granted agency (tool access), they become operational identities—often with broad privileges so they can be effective.

Unit 42's 2026 report discusses how AI can improve attacker outcomes (e.g., hyper-personalized social engineering and faster scripting) and highlights that many breaches still stem from preventable weaknesses like identity and configuration issues.[1] The market is responding: Reuters reported in January 2026 that CrowdStrike announced an acquisition of identity security startup SGNL to strengthen continuous identity controls, explicitly framing the move in the context of AI-driven threats and autonomous AI access patterns.[12]

The risk is not "the AI becomes evil." The risk is that an attacker can steer an agent, steal its credentials, or exploit excessive privileges—turning a helpful automation into a high-speed data access layer.

6.1 Design rules for safe AI agents

- Give every agent a dedicated identity. No shared API keys. Separate dev/test/prod identities.
- Scope permissions narrowly to required API endpoints and data domains. Avoid wildcard admin roles.
- Allow-list tools/functions the agent can call; deny-by-default and constrain parameters.
- Gate irreversible actions (data export, payments, privilege changes) with explicit approvals and logging.
- Instrument agent behavior: prompts, retrieved sources, tool calls, and outputs stored in tamper-resistant logs.

6.2 The non-human identity inventory you probably don't have

Most organizations can list their employees faster than they can list their machine identities. Start by enumerating:

- Service accounts in cloud and SaaS platforms (including CI/CD and automation bots).
- API keys and tokens in code repositories, secret stores, and integration platforms.
- Certificates and workload identities used by microservices and Kubernetes workloads.
- OAuth applications and delegated permissions in productivity suites.

7. LLM application security: OWASP's Top 10 and the UK NCSC warning

LLM security has matured into a distinct domain. OWASP's Top 10 for Large Language Model Applications is a practical starting taxonomy that includes prompt injection, insecure output handling, training data poisoning, model denial of service, and supply chain vulnerabilities.[13]

In December 2025, the UK National Cyber Security Centre (NCSC) published a warning titled "Prompt injection is not SQL injection (it may be worse)." The NCSC argues that LLMs lack a robust internal separation between instructions and data, and that this property can undermine traditional mitigation expectations.[14]

This implies a design constraint: you must assume residual prompt injection risk. Security is achieved by limiting blast radius and validating outputs, not by assuming the model will always follow the "right" instruction.

7.1 Controls that meaningfully reduce LLM/agent risk

- Treat model output as untrusted input. Validate, sanitize, and constrain before any execution or data write.
- Implement retrieval access control: the model should only retrieve data the calling identity is permitted to access.
- Use human-in-the-loop for high-risk actions and create clear escalation paths.
- Design for observability: log prompts, retrieval sources, tool calls, and model outputs with correlation IDs.
- Red-team the system with prompt injection and tool misuse tests before production rollout.

7.2 A simple mental model: "confusable deputy"

A useful framing from the NCSC discussion is to treat the model-driven component as a potentially confusable deputy: it can be tricked into acting on attacker-controlled content. Your job is to ensure that, even if it is confused, it cannot do irreversible harm without additional safeguards.[14]

8. Cloud and SaaS sprawl: permissions, posture, and misconfiguration

Cloud migration redistributed perimeter and privilege. The boundary now includes identity providers, SaaS tenants, CI/CD systems, and API surfaces. Misconfiguration remains high-impact because it is common and it scales: a single permissive policy can expose large datasets.

Mandiant's M-Trends 2025 report highlights cloud migration risks and notes that exploits were the most common initial infection vector observed in 2024 investigations.[4] ENISA's Threat Landscape 2024 identifies ransomware and threats against data and availability as prime threats in the EU, reinforcing that cloud incidents are part of systemic risk rather than isolated "cloud mistakes."[15]

Cloud defense therefore hinges on two engineering disciplines: authorization hygiene (least privilege) and continuous posture management (configuration as code, drift detection, and audit-ready logging).

8.1 The three layers of cloud identity defense

Layer	Goal and examples
Prevent	Least privilege, short-lived credentials, conditional access, hardened admin workflows.
Detect	Audit logs, anomaly detection for token use, alerts on privilege escalation and mass data access.
Respond	Automated session revocation, key rotation playbooks, tenant-wide containment procedures.

8.2 SaaS incidents: why logs and governance matter

SaaS systems are often the fastest path to sensitive data because they already consolidate documents, tickets, finance workflows, and identity. If your SaaS audit logs are not collected, correlated, and monitored, then exfiltration may look like normal file access until it is too late.

9. Secure-by-design and secure-by-default: measurable risk reduction

"Secure-by-design" only matters if it measurably reduces exploitability. CISA's Secure by Design Pledge asks technology providers to make concrete commitments, including building roadmaps to transition toward memory-safe languages and reducing entire vulnerability classes.[16] CISA also publishes a list of product security bad practices, such as hardcoded credentials and use of insecure cryptographic functions.[17]

For operators, the same principles apply internally: your organization is also a software producer through custom applications, scripts, infrastructure-as-code, and automation. Secure-by-default baselines, automated patch verification, and guardrails in CI/CD can reduce the number of preventable identity and configuration failures that dominate real incidents.

9.1 Controls that reduce vulnerability classes

- Eliminate default credentials and enforce strong authentication for administrative interfaces.
- Use parameterized queries and prepared statements to eliminate injection classes in internal services.
- Adopt memory-safe languages for new high-risk components where feasible; isolate legacy components.
- Harden update mechanisms: signed updates, protected build pipelines, and verified deployment provenance.

10. Software supply chain security: SBOMs, SSDF, and SLSA provenance

Modern applications are assembled from thousands of components, build steps, and third-party services. This creates systemic exposure: a compromise in a widely used dependency, maintainer account, or build pipeline can cascade into many downstream victims.

Three frameworks help structure a credible supply-chain program:

• SBOMs (Software Bill of Materials) create visibility into what is inside software. In August 2025, CISA published "2025 Minimum Elements for a Software Bill of Materials," updating guidance to help organizations manage software risk.[18]

• NIST's Secure Software Development Framework (SSDF) defines a set of fundamental practices for secure development and for mitigating software vulnerability risk.[19]

• SLSA (Supply-chain Levels for Software Artifacts) defines progressively stronger controls to prevent tampering and improve artifact integrity and provenance.[20]

Supply-chain security is not about collecting PDFs. It is about creating verifiable integrity signals you can trust when you deploy software.

10.1 What "good" looks like

- Require SBOMs from critical vendors and produce them for your own software.
- Implement signed builds and verified provenance; protect CI/CD with strong identity controls and approvals.
- Continuously map new vulnerabilities to your SBOM inventory to identify real exposure quickly.
- Adopt SSDF-aligned gates: threat modeling for critical systems, security testing in pipelines, and secure code review for sensitive changes.

11. Regulation in the EU: NIS2, DORA, and accountability

Regulation increasingly shapes cyber priorities in Europe. The NIS2 Directive expands risk-management and incident reporting expectations across essential and important entities.[21] In January 2026, the European Commission proposed targeted amendments intended to simplify compliance and increase legal clarity, indicating that the regulatory implementation is still evolving.[21]

For security leaders, the most relevant consequence is governance: security is treated as business risk with management accountability. Identity governance, supplier management, and incident response readiness become not only security best practices but compliance necessities.

11.1 Why identity-centric controls map well to compliance

- Role-based access control and least privilege are auditable and reduce systemic exposure.
- Strong authentication and secure onboarding/offboarding reduce credential and insider risk.
- Centralized logging (identity + cloud control plane) improves detection and incident reporting quality.
- Third-party access governance reduces supplier-driven breach paths.

12. Post-quantum cryptography: crypto-agility as a resilience capability

Quantum computing is not the main driver of most 2026 incidents, but it is a long-horizon risk for data that must remain confidential for many years. In August 2024, NIST released the first three finalized post-quantum cryptography standards (FIPS 203, 204, and 205).[22]

The practical goal is crypto-agility: the ability to inventory cryptographic usage, swap algorithms and key sizes, and coordinate upgrades across systems and suppliers without destabilizing operations.

12.1 Minimal PQC migration approach

- Create a cryptographic inventory (protocols, certificates, libraries, HSM/KMS usage, embedded devices).
- Classify data by confidentiality lifetime: what must remain secret for 5–20 years?
- Prioritize externally exposed TLS and critical internal PKI for migration planning.
- Demand PQC roadmaps from key vendors and cloud providers; align contracts where possible.
- Pilot hybrid approaches where supported and test compatibility and performance.

13. A practical 90-day roadmap (with metrics)

Hot topics become useful only when translated into execution. The roadmap below is a practical sequence that improves identity resilience, reduces time-to-impact, and creates durable governance. It is intentionally prioritized: do the things that shrink blast radius and speed up containment before you perfect long-term programs.

13.1 Roadmap table

Time window	Objective	Concrete actions
Days 1–14	Stop obvious identity bleeding	Enforce phishing-resistant MFA for administrators and remote access; disable legacy authentication; harden account recovery; begin privileged account inventory; rotate high-risk secrets.
Days 15–30	Contain blast radius	Implement just-in-time admin access; review and reduce cloud admin roles; lock down service accounts; segment backups; validate offline backups with restore tests.
Days 31–60	Instrument and automate response	Centralize identity + cloud audit logs; implement high-fidelity detections (impossible travel, token replay, mass downloads, privilege escalation); automate session revocation and key

		rotation playbooks.
Days 61–90	Secure software and AI surfaces	Adopt SBOM requirements for critical vendors; add SSDF-aligned controls in CI/CD; define an internal LLM/agent security standard using OWASP Top 10; red-team high-risk agent workflows; run an incident-response drill focused on identity compromise as the initial condition.

13.2 Metrics that prove progress

- Coverage: % of privileged accounts using phishing-resistant MFA (target: 100%).
- Speed: mean time to disable a compromised identity (target: minutes).
- Blast radius: % reduction in high-privilege cloud roles and long-lived credentials.
- Resilience: backup restore success rate and time-to-restore for critical services (tested regularly).
- Supply chain: SBOM coverage for critical systems and vendors (increasing quarter over quarter).
- AI safety: % of LLM/agent workflows with tool allow-lists, output validation, and approval gates for high-risk actions.

Closing perspective

Identity-centric defense is the defining hot topic because it sits at the intersection of cloud adoption, ransomware economics, and AI-enabled automation. The attacker's advantage is speed and reuse: stolen identities, scripted operations, and scalable social engineering. The defender's advantage must become policy, automation, and engineering discipline—especially around identity, authorization, and provenance. If you can make "valid account abuse" hard, visible, and quickly reversible, you will measurably reduce your risk profile even while the threat landscape continues to evolve.

Appendix A: Checklists

A.1 Identity hardening checklist

- Phishing-resistant MFA for admins, remote access, and backup administration.[11]
- Disable legacy authentication and enforce conditional access for sensitive systems.
- Harden account recovery and help desk processes (social engineering-resistant).
- Reduce session lifetime for sensitive apps and revoke on high-risk signals.
- Inventory and govern service accounts, API keys, and OAuth apps; rotate and minimize privileges.
- Centralize identity, SaaS, and cloud audit logs; alert on privilege escalation and mass data access.

A.2 LLM/Agent security checklist

- Threat model untrusted inputs and define what the agent can do (tool scope).
- Allow-list tools; constrain parameters; deny-by-default network egress where feasible.
- Treat outputs as untrusted; validate before execution; use approvals for irreversible actions.
- Log prompts, retrieved sources, tool calls, and outputs with correlation IDs and tamper resistance.
- Test prompt injection and tool misuse in pre-production; assume residual risk.[14]

A.3 Ransomware resilience checklist

- Offline backups and restore tests (include identity systems and SaaS exports).
- Patch and verify internet-facing systems and remote access devices (including firewalls/VPN).
- Segment management planes and backups from user networks.
- EDR coverage and tuned alerts for privilege escalation, mass encryption behaviors, and suspicious admin tooling.
- Incident response runbooks and exercises focused on identity compromise as the initial condition.

Appendix B: Glossary

Term	Definition
AI agent	A system that uses an LLM to plan and execute tasks by calling tools and accessing data.
Access broker	An actor who sells initial access to compromised accounts or networks.
Conditional access	Policy-based access evaluation using user, device, location, and risk signals.
FIDO2/WebAuthn	Standards for phishing-resistant authentication using public-key credentials.
Infostealer	Malware designed to steal credentials, session tokens, and other sensitive data.
Least privilege	Granting only the minimal permissions needed for a role or task.
SBOM	Software Bill of Materials; an inventory of software components and dependencies.
SLSA	Framework for securing build pipelines and establishing artifact provenance.
SSDF	NIST Secure Software Development Framework; a set of secure development practices.
Time-to-impact	Time from initial compromise to attacker objective (e.g., exfiltration).

References

[1] Palo Alto Networks Unit 42. (2026). 2026 Unit 42 Global Incident Response Report. https://www.paloaltonetworks.com/resources/research/unit-42-incident-response-report

[2] Microsoft. (2025, Oct 16). Extortion and ransomware drive over half of cyberattacks (MDDR 2025). https://blogs.microsoft.com/on-the-issues/2025/10/16/mddr-2025/

[3] Verizon. (2025). 2025 Data Breach Investigations Report (DBIR) resources. https://www.verizon.com/business/resources/reports/dbir/

[4] Google / Mandiant. (2025). M-Trends 2025 (PDF). https://services.google.com/fh/files/misc/m-trends-2025-en.pdf

[5] CISA & FBI. (2025, May 21). AA25-141B: Threat Actors Deploy LummaC2 Malware to Exfiltrate Sensitive Information. https://www.cisa.gov/news-events/cybersecurity-advisories/aa25-141b

[6] Wired. (2025). Authorities carry out global takedown of Lumma infostealer infrastructure. https://www.wired.com/story/lumma-stealer-takedown-disrupted/

[7] CrowdStrike. (2025, Mar 31). How to Navigate the 2025 Identity Threat Landscape. https://www.crowdstrike.com/en-us/blog/how-to-navigate-2025-identity-threat-landscape/

[8] CrowdStrike. (2025). 2025 Global Threat Report (resources). https://www.crowdstrike.com/en-us/global-threat-report/

[9] Sophos. (2025, Jun 24). The State of Ransomware 2025. https://www.sophos.com/en-us/blog/the-state-of-ransomware-2025

[10] Chainalysis. (2025, Feb 5). 35% Year-over-Year Decrease in Ransomware Payments in 2024. https://www.chainalysis.com/blog/crypto-crime-ransomware-victim-extortion-2025/

[11] CISA. (n.d.). #StopRansomware: Ransomware Guide. https://www.cisa.gov/stopransomware/ransomware-guide

[12] Reuters. (2026, Jan 8). CrowdStrike to buy identity security startup SGNL for $740 million to tackle AI threats. https://www.reuters.com/technology/crowdstrike-buy-identity-security-startup-sgnl-740-million-tackle-ai-threats-2026-01-08/

[13] OWASP. (2025). Top 10 for Large Language Model Applications. https://owasp.org/www-project-top-10-for-large-language-model-applications/

[14] UK National Cyber Security Centre (NCSC). (2025, Dec 8). Prompt injection is not SQL injection (it may be worse). https://www.ncsc.gov.uk/blog-post/prompt-injection-is-not-sql-injection

[15] ENISA. (2024, Sept 19). ENISA Threat Landscape 2024. https://www.enisa.europa.eu/publications/enisa-threat-landscape-2024

[16] CISA. (2024/2025). Secure by Design Pledge. https://www.cisa.gov/securebydesign/pledge

[17] CISA. (2025, Jan 17). Product Security Bad Practices. https://www.cisa.gov/resources-tools/resources/product-security-bad-practices

[18] CISA. (2025, Aug 22). 2025 Minimum Elements for a Software Bill of Materials (SBOM). https://www.cisa.gov/resources-tools/resources/2025-minimum-elements-software-bill-materials-sbom

[19] NIST. (2022). SP 800-218: Secure Software Development Framework (SSDF) v1.1 (PDF). https://nvlpubs.nist.gov/nistpubs/specialpublications/nist.sp.800-218.pdf

[20] OpenSSF. (n.d.). SLSA: Supply-chain Levels for Software Artifacts. https://slsa.dev/

[21] European Commission. (2026, Jan 20). NIS2 Directive (including proposed targeted amendments). https://digital-strategy.ec.europa.eu/en/policies/nis2-directive

[22] NIST. (2024, Aug 13). NIST releases first 3 finalized post-quantum encryption standards. https://www.nist.gov/news-events/news/2024/08/nist-releases-first-3-finalized-post-quantum-encryption-standards