IAM & SECURITY BRIEFING

IAM & Security Weekly Briefing

WEEK OF 2026-05-24 to 2026-05-30

Reporting window: Sunday–Saturday of the prior calendar week (does not include the current in-progress week).

---

1. Executive Summary (TL;DR)

- **The FBI put a name on the year's dominant identity attack.** An IC3 public service announcement warned about **Kali365**, a Telegram-distributed phishing-as-a-service kit (~$250/month per tenant, first seen April 2026) that abuses Microsoft's legitimate **OAuth 2.0 device-code flow** to hijack Entra/Microsoft 365 accounts. The victim completes a real Microsoft login *and* MFA; the attacker walks away with **OAuth access and refresh tokens** that grant persistent, password-less access to Outlook, Teams, and OneDrive. Token theft — not password theft — is now the primary identity kill chain. (BleepingComputer, Cybersecurity Dive)
- **Help-desk vishing is the other half of the playbook.** Analysis of CrowdStrike's *2026 Financial Services Threat Landscape Report* named **"Mutant Spider"** the sector's most active actor, with its signature move being **Teams-based vishing**: impersonate internal IT, talk an employee into resetting credentials and MFA, then register an attacker-controlled device. Reset-the-MFA and steal-the-token are now the two sides of identity compromise. (VentureBeat)
- **A credential-less VPN auth bypass hit CISA KEV. CVE-2026-0257**, an authentication-override-cookie bypass in Palo Alto Networks GlobalProtect (PAN-OS), lets an attacker forge valid cookies and impersonate *any* user without credentials. Rapid7 saw active exploitation from May 17; **CISA added it to the KEV catalog on May 29** with a **June 1 federal remediation deadline**. (BleepingComputer)
- **A single social-engineered employee account exposed ~6M people.** Carnival Corporation began notifying **5.99M individuals** (names, DOBs, government ID numbers) after attackers used social engineering to compromise an **employee account** and exfiltrate files. **ShinyHunters** claimed the breach. Human-targeted identity compromise — not malware — drove a top-tier data theft again. (SecurityWeek)
- **AI agents are being formally reframed as identities.** CSA's *State of AI Cybersecurity 2026* found **92% of security professionals concerned** about AI agents' impact, stressed agents "must be governed as identities, with least-privilege access," and flagged that only **37%** have formal AI policies. Vendors moved in lockstep — **Orchid Security** shipped agent identity-governance tooling, **Integrated Quantum** unveiled the **MASQ** agent-governance framework, and **Didit** raised $6M for AI-native identity infrastructure. (CSA)
- **Phishing-resistant MFA crossed from "best practice" to "mandate." Salesforce** confirmed all privileged users must use **FIDO2/passkeys** (push and TOTP no longer qualify) starting **June 22 (sandbox) / July 1 (production) 2026**, and **Yubico's YubiKey 5 FIPS Series achieved FIPS 140-3**

validation — the only authenticator DoD-authorized to hold both PIV credentials and FIDO2 passkeys on one device. (Salesforce Break, Yubico)

- **The new Verizon DBIR reframed where credentials sit in the kill chain.** In-window analysis of the **2026 DBIR** (31,000+ incidents, 22,000+ confirmed breaches) showed **stolen credentials fell to 13%** as an initial-access vector while **vulnerability exploitation rose to 31%**, humans factored into **62%** of breaches, and **67% of users accessed AI services with non-corporate accounts** on work devices (shadow AI). (Help Net Security)
- **Strategic signal:** every major thread this week — Kali365, Mutant Spider, the Palo Alto bypass, Carnival — routes *through identity* rather than malware, and the defensive response (passkey mandates, ITDR, device-code governance, agent identity controls) is converging on **continuous, token-aware, identity-first defense**.

---

2. Top IAM & Security News

1. FBI WARNS OF "KALI365" PHAAS HIJACKING MICROSOFT 365 VIA OAUTH DEVICE-CODE FLOW

The FBI's IC3 issued a public service announcement on the **Kali365** phishing-as-a-service platform — distributed via Telegram (~$250/month per tenant), first seen April 2026 — that abuses Microsoft's **legitimate OAuth 2.0 device-code authorization flow**. Victims are lured (fake Adobe Acrobat Sign, DocuSign, SharePoint emails) into entering an attacker-supplied device code on an authentic Microsoft verification page; MFA fires on the victim's own device while the attacker captures the resulting **access and refresh tokens**. The kit ships AI-generated lures, automated campaign templates, and real-time victim dashboards, and has hit manufacturing, education, insurance, financial, healthcare, and government targets across North America, Europe, the Middle East, and Africa. - **Why it matters:** This defeats MFA *at the identity-provider layer* — a single consent yields long-lived, password-less access across an entire SSO-connected M365 estate. Because the flow rides legitimate Microsoft infrastructure, it looks normal to victims and to many detection tools. The mandatory control is **Conditional Access to block/restrict the device-code flow** (excluding break-glass accounts), plus OAuth grant inventories and token-level monitoring. - **Source:** BleepingComputer, Cybersecurity Dive, Infosecurity Magazine

2. PALO ALTO GLOBALPROTECT AUTH-BYPASS (CVE-2026-0257) EXPLOITED, ADDED TO CISA KEV

CVE-2026-0257 is an authentication bypass in PAN-OS GlobalProtect portals/gateways. When a device reuses the same certificate for HTTPS and authentication-override cookies, an attacker can extract the public key and **forge valid override cookies for any user account** — establishing an unauthorized VPN session with no credentials. Rapid7 observed successful exploitation across multiple customers beginning May 17, 2026, with follow-on waves on May 18 and May 21. CISA added the flaw to its **Known Exploited Vulnerabilities catalog on May 29**, with a **June 1 federal remediation deadline**, and the severity was raised from Medium to High once exploitation began. - **Why it matters:** An unauthenticated identity-bypass on an internet-facing VPN gateway is a textbook initial-access vector that lets an attacker impersonate any user before any IAM control engages. Patch immediately and rotate the affected device certificates so previously-forgeable cookies can't be replayed. - **Source:** BleepingComputer

3. CARNIVAL BREACH EXPOSES ~6M PEOPLE AFTER SOCIAL-ENGINEERED EMPLOYEE ACCOUNT

Carnival Corporation began notifying roughly **5.99M individuals** that their data was stolen in an incident identified April 14, 2026. Attackers gained access to an **employee's account via social engineering**, pivoted into company systems, and exfiltrated files containing names, addresses, dates of birth, email addresses, phone numbers, and **government-issued ID numbers**. The extortion group **ShinyHunters** claimed responsibility, alleging theft of 8.7M records; Carnival is offering 24 months of credit monitoring. - **Why it matters:** A single compromised human identity was the entire breach's initial-access vector — no malware required. It reinforces the week's theme: help-desk verification hardening, phishing-resistant MFA, and least-privilege on employee accounts are the controls that actually move the needle against large-scale data theft. - **Source:** SecurityWeek

4. CROWDSTRIKE: "MUTANT SPIDER" DOMINATES FINANCIAL SERVICES WITH RESET-MFA VISHING

Analysis of CrowdStrike's **2026 Financial Services Threat Landscape Report** identified **Mutant Spider** as the sector's most active threat actor, with its primary technique being **voice phishing over Microsoft Teams** — operators impersonate internal IT support, convince employees to reset credentials and MFA, then register their own devices on the corporate network. The coverage pairs this with the FBI's Kali365 warning, framing the two dominant patterns as "remove MFA via social engineering" and "capture the post-auth token." - **Why it matters:** Adversaries targeting high-value sectors have stopped stealing passwords. The defensive answer is stronger **help-desk identity verification**, **device-registration controls**, and **continuous session/behavioral monitoring** that doesn't end at login-time MFA. - **Source:** VentureBeat

5. MFA PROMPT BOMBING: WHY A SECOND FACTOR ALONE ISN'T SAVING YOU

A widely-shared analysis dissected **MFA prompt bombing** (push fatigue): attackers who already hold valid credentials spam push prompts and pair them with vishing calls posing as IT support to coerce approval. Push notifications carry no context about request origin, device, or legitimacy, making errant approval likely — the same technique behind the 2022 Cisco breach. Recommended defenses: replace push with **phishing-resistant factors** (FIDO2 keys, hardware tokens, number matching), screen Active Directory against breached-password datasets, and apply Conditional Access on geography, device posture, and login timing. - **Why it matters:** Push fatigue plus help-desk vishing remains a leading account-takeover route (the Scattered Spider / ShinyHunters playbook). It's the practitioner-level case for the broader 2026 shift away from push/OTP toward phishing-resistant authentication. - **Source:** The Hacker News

6. "ONLYFANS" 340M-RECORD LEAK SHOWS IDENTITY EXPOSURE IS CUMULATIVE

A threat actor advertised ~**340M "OnlyFans" user records** (usernames, emails, phone numbers, follower counts, linked social profiles) for ~$76,000 in Bitcoin. The seller admitted the data was **not exfiltrated from OnlyFans** but assembled by cross-referencing older breaches (Twitter, Instagram, Spotify) and matching them to profiles; OnlyFans denied any breach and payment-card claims were unverified. - **Why it matters:** Aggregated, cross-referenced breach data can rebuild rich identity profiles that fuel targeted phishing, credential stuffing, and social engineering — *without a fresh platform compromise*. Identity exposure compounds across the breach ecosystem, which is why account-takeover defense and identity-verification workflows can't assume "we were never breached" equals "our users are safe." - **Source:** Hackread

7. NEW INFOSEC PRODUCTS OF THE MONTH SURFACE NHI AND AGENTIC-AI ACCESS GOVERNANCE

Help Net Security's May product roundup highlighted identity-relevant launches: **XM Cyber Continuous Exposure Management for Identities** (credential security and IAM attack-surface reduction tuned for AI-

enabled attackers), **Trust3 AI MCP Security** (a unified trust layer governing non-human/agentic AI identity access to business data and systems), and **LastPass Mobile Smart Scanner** (OCR of credentials into the vault). - **Why it matters:** A useful read on where the IAM/NHI market is investing right now — agentic-AI/MCP access governance and continuous identity exposure management — and a buying-signal datapoint for teams evaluating NHI and ITDR tooling. - **Source:** Help Net Security

---

3. AI, Identity & Emerging Tech

CSA *STATE OF AI CYBERSECURITY 2026*: AGENTS MUST BE GOVERNED AS IDENTITIES

CSA's research found **92% of security professionals concerned** about AI agents' workforce impact, with **61%** naming sensitive-data exposure their top worry — yet only **37%** report having formal AI policies in place (down year over year). The report's framing is the headline for IAM teams: agents frequently receive **broad cross-system permissions** and "must be governed as identities, with least-privilege access and ongoing monitoring." - **Security implications:** Hard survey data to justify treating agents as first-class, least-privilege identities rather than ungoverned service accounts — and to fund the governance gap before agent deployment outpaces controls. - **Source:** Cloud Security Alliance

ORCHID SECURITY TARGETS "AGENT AI AUTHORITY GAP" WITH DELEGATION-CHAIN VISIBILITY

Orchid launched three agent-focused capabilities — **Agentic Enrichment** (maps agents to originating identities, owners, applications, and inherited permissions), **Agentic Observability** (tracks the full delegation chain behind each agent action), and **Agentic Guardrails** (enforces least privilege and identity hygiene). Orchid frames the core risk as the gap between what enterprises *think* is governed and what agents can actually execute, citing that **two-thirds of enterprises already run agents in production** while **67% of non-human accounts are locally managed and invisible to central IAM**. - **Security implications:** Delegation-chain visibility is the under-built control — an agent action is only as safe as the human and machine identities it inherits authority from. Mapping that chain is prerequisite to least-privilege for agents. - **Source:** SiliconANGLE

MASQ FRAMEWORK EXTENDS AGENT GOVERNANCE TO MCP SERVERS AND CONTEXT WINDOWS

Integrated Cyber Solutions (d/b/a Integrated Quantum Technologies) announced **MASQ (Machine Action Security Quotient)**, a governance framework for autonomous agents, and began a patent process. MASQ governs four areas: what actions an agent is authorized to perform, what systems and data it can reach, how it interacts with **APIs, external tools, and MCP servers**, and how sensitive data inside an agent's **context window and reasoning environment** is protected during machine-to-machine interaction. - **Security implications:** Notable for explicitly extending governance to MCP-server interactions and agent context/reasoning data — an emerging machine-to-machine attack surface most IAM programs don't yet inventory. - **Source:** PR Newswire

DIDIT RAISES $6M FOR AI-NATIVE IDENTITY INFRASTRUCTURE

Didit raised a **$6M seed** (total $7.5M, with Y Combinator, Pioneer Fund, and Orange Collective participating) to build programmable identity infrastructure that verifies people, businesses, and **digital actions taken by AI agents** via developer APIs analyzing 200+ signals across 220+ countries — working toward an identity wallet supporting both human and agent identities. - **Security implications:** Continued investor appetite for the

non-human/agent identity layer, and a marker of the emerging need to authenticate autonomous agents as **transacting entities**, not just humans. - **Source:** SiliconANGLE

Trends crystallizing this week: non-human identities (NHIs) outnumbering and outpacing human ones; autonomous agents acquiring delegated authority faster than governance; identity sprawl into application-local accounts invisible to central IAM; and a shift toward behavioral/continuous identity models (ITDR) to watch what agents and humans actually *do* post-authentication.

---

4. Cyber Threats & Attack Trends

1. OAUTH DEVICE-CODE TOKEN THEFT (KALI365)

- **Attack:** Phishing-as-a-service lures victims to authorize an attacker's device through Microsoft's legitimate device-code flow, capturing OAuth access/refresh tokens.
- **Identity angle:** MFA succeeds on the victim's device but does not protect the attacker's session; the stolen **refresh token** provides persistent, password-less access.
- **Techniques:** Device-code phishing, OAuth token theft, abuse of legitimate IdP infrastructure to evade detection.
- **Real-world example:** FBI IC3 PSA; campaign tracked since April 2026 across multiple sectors and continents.
- **Source:** BleepingComputer

2. RESET-MFA-VIA-VISHING (MUTANT SPIDER)

- **Attack:** Teams-based voice phishing impersonating internal IT to reset employee credentials and MFA, then enroll an attacker-controlled device.
- **Identity angle:** Bypasses authentication entirely by socially engineering the **help-desk / self-service reset** workflow rather than cracking a factor.
- **Techniques:** Vishing, IT-support impersonation, malicious device registration.
- **Real-world example:** Named most-active threat to financial services in CrowdStrike's 2026 Financial Services Threat Landscape Report.
- **Source:** VentureBeat

3. CREDENTIAL-LESS VPN AUTHENTICATION BYPASS (CVE-2026-0257)

- **Attack:** Forging Palo Alto GlobalProtect authentication-override cookies to impersonate arbitrary users and open unauthorized VPN sessions.
- **Identity angle:** Pure authentication bypass — no credentials, no MFA, full user impersonation at the network edge.
- **Techniques:** Public-key extraction from a reused certificate, cookie forgery.
- **Real-world example:** Active exploitation from May 17; CISA KEV listing May 29 with a June 1 federal deadline.
- **Source:** BleepingComputer

4. SOCIAL-ENGINEERED EMPLOYEE ACCOUNT → MASS DATA EXFILTRATION (CARNIVAL / SHINYHUNTERS)

- **Attack:** Social engineering of a single employee account, then lateral movement and bulk file exfiltration.
- **Identity angle:** The compromised human identity *was* the breach — initial access, lateral movement, and data access all flowed from one account.
- **Techniques:** Social engineering, account takeover, extortion.
- **Real-world example:** ~5.99M individuals notified; ShinyHunters claimed responsibility.
- **Source:** SecurityWeek

5. MFA PROMPT BOMBING / PUSH FATIGUE

- **Attack:** Repeated push-MFA prompts plus vishing pressure to coerce approval from a user whose password is already compromised.
- **Identity angle:** Exploits the human approval step of push-based MFA, which carries no request context.
- **Techniques:** Push-fatigue, social engineering.
- **Real-world example:** The technique behind the 2022 Cisco breach; resurfaced as a 2026 advisory.
- **Source:** The Hacker News

---

5. Product Updates & Vendor News

SALESFORCE — PHISHING-RESISTANT MFA MANDATED FOR PRIVILEGED USERS

Salesforce confirmed all privileged users (System Administrators and holders of Modify All Data, View All Data, Customize Application, or Author Apex) must authenticate with **phishing-resistant MFA** starting **June 22, 2026 in sandboxes and July 1, 2026 in production**. Only **FIDO2/WebAuthn built-in authenticators** (Touch ID, Face ID, Windows Hello) and hardware security keys qualify; authenticator apps, push, and TOTP no longer count for these roles. The setting is locked so it can't be disabled, and unregistered admins are blocked at login. - **Why it matters:** A major SaaS platform hard-mandating FIDO2/passkeys for privileged access is a significant forcing function pushing enterprises toward passwordless authentication for their highest-risk identities. - **Source:** Salesforce Break

YUBICO — YUBIKEY 5 FIPS SERIES ACHIEVES FIPS 140-3 VALIDATION

Yubico's next-gen **YubiKey 5 FIPS Series** received **FIPS 140-3 validation** (Certificate #5291) on 5.7.4 firmware, with expanded algorithm support (RSA-3072, RSA-4096, Ed25519). It remains the only authenticator recognized in U.S. DoD guidance authorized to hold both **DoD PKI/PIV credentials and FIDO2 passkeys** on a single device, alongside OpenPGP and OATH OTP. - **Why it matters:** Lets government and regulated enterprises consolidate smart-card and passkey workflows on one **NIST AAL3** hardware authenticator, lowering friction for high-assurance, phishing-resistant rollouts. - **Source:** Yubico

CROWDSTRIKE — NAMED LEADER AND FAST MOVER IN 2026 GIGAOM RADAR FOR ITDR

CrowdStrike was named **Leader and Fast Mover** in the 2026 GigaOm Radar for **Identity Threat Detection and Response (ITDR)**, with perfect scores in non-human identity security, risk-adaptive access controls, AI-enhanced SecOps, and automated incident response. GigaOm highlighted cross-domain correlation — enriching identity telemetry with endpoint, cloud, and SaaS data via Falcon and Next-Gen SIEM. - **Why it matters:** Reflects the market's shift toward continuous, risk-adaptive authorization spanning human, non-

human, and AI-agent identities — validating identity-first detection as a core enterprise pillar rather than a bolt-on. - **Source:** CrowdStrike

ORCHID SECURITY, XM CYBER, TRUST3 AI, LASTPASS — NHI & AGENTIC-AI LAUNCHES

Beyond Orchid's agent identity-governance suite (Section 3), the May product cycle brought **XM Cyber Continuous Exposure Management for Identities**, **Trust3 AI MCP Security** (governing non-human/agentic access to business data), and **LastPass Mobile Smart Scanner**. - **Why it matters:** The vendor center of gravity is visibly moving toward **NHI and agentic-AI access governance** — the same problem the CSA and Orchid data quantify. - **Source:** Help Net Security — New infosec products of the month: May 2026

---

6. Notable Research & Reports

VERIZON 2026 DATA BREACH INVESTIGATIONS REPORT (DBIR)

In-window analysis of the new DBIR (**31,000+ incidents, 22,000+ confirmed breaches across 145 countries**). Identity-relevant findings: **humans involved in 62% of breaches**; **stolen credentials fell to 13%** as an initial-access vector while **vulnerability exploitation rose to 31%** (now the leader); **67% of users accessed AI services with non-corporate accounts** on work devices (shadow AI); third-party breaches featured in **48%** of breaches (supply-chain breaches up 60%); ransomware in **48%** of breaches. - **Strategic implications:** The relative decline of stolen credentials does *not* mean identity matters less — it means attackers increasingly **bypass** credential checks (vulnerability exploitation, token theft, auth bypass) and that **shadow-AI access** is a fast-growing, largely ungoverned identity surface. Rebalance investment toward edge-device patching, service-account hygiene, and AI access governance. - **Source:** Help Net Security — Lessons from the 2026 DBIR

CSA *STATE OF AI CYBERSECURITY 2026*

Survey-backed report (see Section 3): **92% concerned** about AI agents, **61%** worried about sensitive-data exposure, only **37%** with formal AI policies — a widening preparedness gap as agents acquire cross-system permissions. - **Strategic implications:** The governance gap is the actionable metric — most organizations are deploying agents faster than they can govern them as identities. - **Source:** Cloud Security Alliance

WEEKLY THREAT-INTEL DIGEST — INFOSTEALERS, DBIR, EXPLOITED EDGE FLAWS

Help Net Security's May 24 "week in review" curated the period's identity-adjacent research: the DBIR's initial-access shift, a **PureLogs infostealer** phishing campaign (credential harvesting feeding account takeover), and actively-exploited Microsoft Defender and NGINX vulnerabilities — plus a GitHub compromise via a **poisoned VS Code extension** (a supply-chain/identity-trust angle on developer tooling). - **Strategic implications:** Infostealers remain the upstream feedstock for the credential and session-token theft seen elsewhere this week; developer-tool supply chains are an under-watched identity-trust boundary. - **Source:** Help Net Security — Week in review

---

7. Practical Security Takeaways

1. **Lock down the OAuth device-code flow.** Use Conditional Access to block or tightly scope device-code authentication (exclude only break-glass accounts). This is the single highest-leverage control against the Kali365 pattern.

2. **Hunt for stolen tokens, not just bad logins.** Inventory OAuth grants, monitor for anomalous refresh-token use, and **revoke/rotate sessions** on suspicion — MFA at login does nothing for an already-issued token.
3. **Harden the help desk against vishing.** Require strong, out-of-band identity verification before any credential or MFA reset, and add friction/alerting to **new device registration**. This is the control that stops Mutant Spider and Scattered Spider-style intrusions.
4. **Move privileged users to phishing-resistant MFA now.** Salesforce's mandate is a preview of the direction — migrate admins and high-risk roles to **FIDO2/passkeys or hardware keys** ahead of being forced to.
5. **Patch internet-facing identity/edge devices on a KEV clock.** Apply the Palo Alto GlobalProtect (CVE-2026-0257) fix and **rotate affected certificates**; treat VPN/SSO gateway auth bypasses as emergency change.
6. **Govern AI agents as identities.** Assign every agent an owner, scope least-privilege access, map its delegation chain, and monitor behavior — don't let agents live as ungoverned service accounts.
7. **Audit and rotate non-human identities.** Find application-local accounts invisible to central IAM, assign human ownership per credential, and automate rotation — NHIs are now the fastest-growing and least-governed identity class.
8. **Bring shadow AI under access governance.** With ~2/3 of users reaching AI services via personal accounts on work devices, add AI apps to your SSO/CASB inventory and Conditional Access scope.
9. **Screen credentials against breach data and assume cumulative exposure.** Check AD against breached-password datasets and treat aggregated-leak risk (e.g., the OnlyFans dataset) as live ammunition for credential stuffing and targeted phishing.
10. **Deploy or mature ITDR.** Correlate identity telemetry with endpoint, cloud, and SaaS signals to catch post-authentication attack chains that login-time controls miss.

---

8. Trends to Watch

- **Token theft becomes the default identity attack.** As phishing-resistant MFA spreads, adversaries are shifting decisively from credential theft to **stealing or forging the post-authentication session** (device-code abuse, cookie forgery). Defense moves to token binding, short-lived tokens, and continuous re-validation.
- **AI agents become first-class identities — or first-class liabilities.** The CSA data, Orchid/MASQ launches, and Didit funding all point the same way: agent identity governance (discovery, ownership, least privilege, delegation-chain visibility, MCP-interaction controls) is rapidly becoming a named IAM discipline.
- **From authentication to continuous identity verification.** Push-fatigue, reset-MFA vishing, and token theft all defeat point-in-time authentication, accelerating the shift to **ITDR and risk-adaptive, behavior-aware authorization** that watches what identities do after login.
- **The help desk is the new perimeter.** Social-engineered resets (Carnival, Mutant Spider) make identity-proofing at support and self-service workflows a top-tier control surface — expect more investment in verification tooling and process hardening here.
- **Identity is the breach.** With humans in 62% of breaches and nearly every major incident this week routing through a compromised or bypassed identity, "identity-first security" stops being a slogan and becomes the organizing principle for the security program.

9. Tool / Resource of the Week

Microsoft Entra Conditional Access — device-code flow restriction - **What it does:** A Conditional Access grant/control that lets administrators **block or scope the OAuth 2.0 device-code authentication flow**, the exact mechanism abused by the Kali365 PhaaS kit. Apply it broadly and exclude only emergency-access accounts. - **Why it's useful:** It's a free, native control that directly neutralizes the week's dominant identity attack — turning a "users keep getting phished" problem into a policy decision. Pair it with authentication-transfer policy restrictions and OAuth grant monitoring. - **Link:** Microsoft Learn — Conditional Access: Authentication flows

---

10. Sources

- FBI warns of Kali365 phishing service targeting Microsoft 365 accounts (BleepingComputer) — https://www.bleepingcomputer.com/news/security/fbi-warns-of-kali365-phishing-service-targeting-microsoft-365-accounts/
- FBI warns about PhaaS platform used to access Microsoft 365 environments (Cybersecurity Dive) — https://www.cybersecuritydive.com/news/fbi-warns-phishing-platform-microsoft-365/821105/
- FBI Warns 'Kali365' Phishing Kit Hijacks Microsoft 365 OAuth Tokens (Infosecurity Magazine) — https://www.infosecurity-magazine.com/news/fbi-kali365-phishing-kit-m365/
- Palo Alto GlobalProtect VPN auth bypass flaw now exploited in attacks (BleepingComputer) — https://www.bleepingcomputer.com/news/security/palo-alto-globalprotect-vpn-auth-bypass-flaw-now-exploited-in-attacks/
- Carnival data breach exposed 6 million people (SecurityWeek) — https://www.securityweek.com/carnival-data-breach-exposed-6-million-people/
- The Attack Dominating Financial Services: Reset MFA and Steal the Token (VentureBeat) — https://venturebeat.com/security/attack-dominating-financial-services-resets-mfa-steals-token
- MFA Prompt Bombing: Why Your Second Factor Isn't Saving You (The Hacker News) — https://thehackernews.com/2026/05/mfa-prompt-bombing-why-your-second.html
- OnlyFans mega leak reveals 340M user records, hackers claim (Hackread) — https://hackread.com/hacker-selling-onlyfans-user-records-old-breaches/
- State of AI Cybersecurity 2026: 92% of Security Professionals Concerned About AI Agents (Cloud Security Alliance) — https://cloudsecurityalliance.org/blog/2026/05/27/state-of-ai-cybersecurity-2026-92-of-security-professionals-concerned-about-the-impact-of-ai-agents
- Orchid Security Targets AI Agent Sprawl With New Identity Governance Tools (SiliconANGLE) — https://siliconangle.com/2026/05/28/orchid-security-targets-ai-agent-sprawl-new-identity-governance-tools/
- Integrated Quantum Technologies Unveils MASQ Governance Framework for Autonomous AI Agents (PR Newswire) — https://www.prnewswire.com/news-releases/a-microcap-just-staked-a-claim-in-the-ai-agent-security-land-grab-302785726.html
- Didit Raises $6M to Build AI-Native Identity Infrastructure (SiliconANGLE) — https://siliconangle.com/2026/05/26/didit-raises-6m-funding-build-ai-native-identity-infrastructure/
- Yubico Announces Upgraded YubiKey 5 FIPS Series, Now FIPS 140-3 Validated (Yubico) — https://www.yubico.com/press-releases/yubico-announces-upgraded-yubikey-5-fips-series-now-fips-

140-3-validated/

- CrowdStrike Named Leader and Fast Mover in 2026 GigaOm Radar for ITDR (CrowdStrike) — https://ir.crowdstrike.com/news-releases/news-release-details/crowdstrike-named-leader-and-fast-mover-2026-gigaom-radar-report
- Advancing phishing-resistant MFA (Samsung Business Insights) — https://insights.samsung.com/2026/05/27/advancing-phishing-resistant-mfa/
- What Is Phishing-Resistant MFA? — Salesforce privileged-user mandate (Salesforce Break) — https://salesforcebreak.com/2026/05/27/phishing-resistant-mfa/
- Lessons for organizations from the Verizon 2026 DBIR (Help Net Security) — https://www.helpnetsecurity.com/2026/05/25/lessons-from-verizon-dbir-2026-findings/
- New infosec products of the month: May 2026 (Help Net Security) — https://www.helpnetsecurity.com/2026/05/29/new-infosec-products-of-the-month-may-2026/
- Week in review: GitHub breached via poisoned VS Code extension, critical NGINX flaw exploited (Help Net Security) — https://www.helpnetsecurity.com/2026/05/24/week-in-review-github-breached-via-poisoned-vs-code-extension-critical-nginx-flaw-exploited/
- Microsoft Learn — Conditional Access: Authentication flows (device-code restriction) — https://learn.microsoft.com/en-us/entra/identity/conditional-access/concept-authentication-flows