INTELLIGENCE BRIEFING
Security Command Center
TLP:CLEAR
2026-05-13 18:56 UTC

RSM's Cybersecurity Special Report Finds Middle Market Racing Into AI Faster Than It Can Secure It

GOVERNANCE | HIGH

SCC Item ID	SCC-GOV-2026-0034
Type	Governance
Severity	HIGH
Affected Products	Middle market companies (across various sectors, revenue typically $10M-$1B)
Published	2026-05-13
Discovery Source	Gemini

Executive Summary

RSM US LLP's May 2026 Cybersecurity Special Report identifies a structural governance gap in middle market organizations: AI adoption is outpacing the identity controls, access governance, and policy frameworks needed to manage it safely. Only 35% of executives report using formal AI governance frameworks, leaving the majority exposed to shadow AI risks, ungoverned third-party AI integrations, and AI-augmented credential and phishing attacks. The compounding factor is a perception-versus-reality gap: executives self-report high confidence in their defenses despite these exposures, which delays corrective action.

Technical Analysis

No CVE or CVSS score applies. This is a governance and architectural risk finding, not a discrete vulnerability. The report maps to four MITRE ATT&CK techniques: T1078 (Valid Accounts) - ungoverned AI workloads and integrations frequently operate under overprivileged or shared credentials without MFA or lifecycle controls; T1204 (User Execution) - shadow AI adoption involves employees installing or connecting unsanctioned tools that may exfiltrate data or introduce malicious dependencies; T1195 (Supply Chain Compromise) - third-party AI integrations added without vendor risk assessment expand the software and data supply chain attack surface; T1566 (Phishing) - AI-generated spear phishing lowers the cost and raises the quality of social engineering attacks targeting middle market employees who lack security awareness training calibrated to AI-generated content. No patch exists. Remediation is architectural: AI governance frameworks, identity controls scoped to AI workloads, and shadow IT discovery programs.

Action Checklist

1. Step 1: Inventory. Conduct a shadow AI discovery sweep within 7 days using web proxy logs, DNS query logs, and endpoint DLP telemetry to identify unsanctioned AI tools in use (ChatGPT, Copilot variants, AI writing tools, AI browser extensions). Prioritize tools with outbound data access.
2. Step 2: Detection. Query firewall and proxy logs for outbound traffic to known AI service endpoints (openai.com, anthropic.com, gemini.google.com, character.ai, etc.). Flag uploads or large POST requests. Cross-reference against approved application lists. Alert on new AI domains not in the approved set.
3. Step 3: Eradication. Block unapproved AI tool access at the network layer via web proxy policy. Revoke OAuth tokens granted by employees to third-party AI apps without authorization. Audit service accounts and API keys associated with any AI integrations added in the past 12 months.
4. Step 4: Recovery. Implement a formal AI acceptable use policy with explicit data classification rules (what data may and may not be submitted to AI tools). Establish an AI governance working group with representation from IT, security, legal, and business units. Map approved AI tools and their data access scope.
5. Step 5: Post-Incident. Conduct a gap assessment against NIST AI RMF (AI 100-1) or CISA's AI security guidance. Implement identity controls for AI workloads: dedicated service accounts, least-privilege scoping, MFA where applicable, and quarterly access reviews. Add AI-generated phishing simulation to security awareness training.

IR / Forensic Enrichment

Triage Priority	URGENT
Escalation Criteria	Escalate immediately to legal counsel and executive leadership if the shadow AI discovery sweep (Step 1) reveals that regulated data categories — PII, PHI, PCI cardholder data, or material non-public financial information — were submitted to unapproved third-party AI services, as this likely triggers breach notification obligations under HIPAA, state privacy laws (CCPA, state data protection acts), or SEC disclosure requirements, and the organization's cyber insurance carrier should be notified before containment actions alter the evidentiary record.
Recovery Notes	Post-containment, monitor proxy and DNS logs daily for the first 30 days for any attempts to access newly blocked AI domains via circumvention methods (VPN clients, personal hotspots, DNS-over-HTTPS to bypass corporate resolvers) — treat any circumvention attempts as a separate policy violation incident. Re-run the shadow AI sweep from Step 1 at 30 and 90 days post-recovery to detect newly introduced tools, since the RSM report's finding that only 35% of middle market executives use formal AI governance frameworks means employee-driven AI adoption pressure will continue. Verify that OAuth token revocations from Step 3 held by re-querying the identity provider's enterprise application consent report — revoked tokens can sometimes be re-granted by users with standard permissions if conditional access policies have not been hardened to require admin approval for third-party app consent.

Forensic Artifacts	Web proxy access logs (Squid, Zscaler, Bluecoat, or equivalent) filtered for outbound POST requests to AI service API endpoints (api.openai.com/v1/chat/completions, api.anthropic.com/v1/messages, generativelanguage.googleapis.com) with request body size field — large outbound POST bodies are the primary artifact of data submission to AI tools rather than passive browsing | Identity provider OAuth consent and token grant logs (Azure AD Audit Log event type 'Consent to application' or Google Workspace Admin Reports API oauth_token events) showing which employees granted which data scopes (mail, drive, calendar, contacts) to which third-party AI applications and when — this is the definitive record of delegated data access exposure | DNS query logs with client IP attribution for resolutions of AI service FQDNs (openai.com, anthropic.com, huggingface.co, replicate.com, character.ai, perplexity.ai) — enables per-device and per-user attribution of AI service access even when HTTPS prevents content inspection | API key and secrets exposure artifacts: .env files, application configuration files, CI/CD pipeline environment variables, and code repository commit history (use truffleHog or git-secrets to scan for high-entropy strings matching OpenAI sk-, Anthropic, and Google AI API key patterns in repositories accessible to employees) | DLP system alert history and quarantine logs for outbound email or file transfers to external recipients that were flagged during the same period as AI tool usage — AI-augmented phishing and data exfiltration via AI tools often co-occur with other data handling policy violations and the correlation establishes blast radius

Per-Action IR Details

Step 1: Inventory — conduct an immediate shadow AI discovery sweep using web proxy logs, DNS query logs, and endpoint DLP telemetry to identify unsanctioned AI tools in use (ChatGPT, Copilot variants, AI writing tools, AI browser extensions). Prioritize tools with outbound data access.

NIST Phase: Detection Analysis

Reference: NIST 800-61r3 §3.2 — Detection and Analysis: identifying the scope of adverse activity through log correlation and asset enumeration before containment decisions are made

Controls: NIST SI-4 (System Monitoring) — monitor for unauthorized outbound data flows to AI service endpoints, NIST AU-2 (Event Logging) — ensure proxy, DNS, and DLP events are being captured with sufficient fidelity to detect AI tool usage, NIST AU-6 (Audit Record Review, Analysis, and Reporting) — actively review logs for AI-domain traffic patterns, not just retain them, CIS 2.1 (Establish and Maintain a Software Inventory) — identify unauthorized AI browser extensions and desktop apps installed on endpoints, CIS 1.1 (Establish and Maintain Detailed Enterprise Asset Inventory) — establish which devices and users are generating AI-bound traffic before scoping containment

Compensating: For teams without enterprise DLP or SIEM: (1) Export DNS query logs from your recursive resolver (Windows DNS debug logging: enable via dnscmd /config /logFilePath and /loglevel 0x8100; parse with Get-Content dns.log | Select-String 'openai|anthropic|gemini|character.ai|claude'); (2) Pull Squid or pfSense proxy logs and grep for AI domains: grep -Ei 'openai\.com|anthropic\.com|gemini\.google\.com|character\.ai|cohere\.com|huggingface\.co' /var/log/squid/access.log; (3) Use osquery on Windows endpoints: SELECT name, path, description FROM programs WHERE name LIKE '%AI%' OR name LIKE '%Copilot%'; also query SELECT * FROM chrome_extensions WHERE name LIKE '%AI%' OR name LIKE '%GPT%'; (4) On Windows, use PowerShell to find recently installed browser extensions: Get-ChildItem 'C:\Users*\AppData\Local\Google\Chrome\User Data\Default\Extensions' -Recurse -Filter manifest.json | Select-String 'ai|gpt|copilot' -List.

Evidence: BEFORE sweeping, preserve: (1) Raw proxy/firewall logs showing outbound connections to AI service domains with source IP, username, timestamp, bytes transferred, and HTTP method — large POST requests to /v1/chat/completions (OpenAI API) or /api/generate endpoints indicate data submission, not just browsing; (2) DNS query logs with client IP resolution showing which internal hosts resolved AI service FQDNs; (3) DLP alert history for outbound transfers flagged against data classification policies (look for PII, financial, or IP-tagged content in outbound flows to *.openai.com, *.anthropic.com); (4) OAuth token grant records from identity provider (Azure AD audit log Event 'Consent to application' or Google Workspace Admin audit for 'Authorize API client') showing employees authorizing third-party AI apps to access corporate data; (5) Browser history or installed extension manifests on endpoints where

DNS/proxy hits are concentrated.

Step 2: Detection — query firewall and proxy logs for outbound traffic to known AI service endpoints (openai.com, anthropic.com, gemini.google.com, character.ai, etc.). Flag uploads or large POST requests. Cross-reference against approved application lists. Alert on new AI domains not in the approved set.

NIST Phase: Detection Analysis

Reference: NIST 800-61r3 §3.2 — Detection and Analysis: correlating network telemetry against approved baselines to distinguish sanctioned from unsanctioned AI activity and estimate data exposure scope

Controls: NIST SI-4 (System Monitoring) — implement continuous monitoring for outbound traffic to AI service IP ranges and domains not on the approved application list, NIST AU-3 (Content of Audit Records) — ensure firewall and proxy logs capture HTTP method, request body size, destination FQDN, source user identity, and bytes transferred — not just connection metadata, NIST AU-7 (Audit Record Reduction and Report Generation) — generate reports filtering specifically for POST requests exceeding a defined threshold (e.g., >10KB) destined for AI endpoints, CIS 7.1 (Establish and Maintain a Vulnerability Management Process) — treat identification of ungoverned AI data flows as a risk finding requiring tracked remediation, CIS 4.4 (Implement and Manage a Firewall on Servers) — enforce egress filtering rules that can detect or block traffic to unapproved AI service endpoints at the perimeter

Compensating: For teams running pfSense, OPNsense, or iptables: (1) Create an alias list of known AI service CIDRs and domains (OpenAI: 104.18.x.x ranges; Anthropic: verify current IPs via dig anthropic.com) and apply a logging rule before any permit rule to capture hits; (2) Use Zeek (formerly Bro) in tap/span mode for HTTP inspection: zeek -i eth0 /opt/zeek/share/zeek/site/local.zeek then query conn.log and http.log: cat http.log | zeek-cut host method request_body_len | grep -Ei 'openai|anthropic|gemini' | awk '$3 > 10000'; (3) Build a Sigma rule targeting proxy logs — condition: keywords containing AI service domains AND request_method='POST' AND cs-bytes > 10240; deploy via sigmac to convert to your log platform's query syntax; (4) For smaller environments with no proxy: configure Windows Firewall audit logging (auditpol /set /subcategory:'Filtering Platform Connection' /success:enable /failure:enable) and query Security Event Log for Event ID 5156 (network connection permitted) filtered on destination IPs matching AI service ranges.

Evidence: BEFORE alerting or blocking, capture and timestamp: (1) Firewall flow records showing source internal IP, destination AI service IP/FQDN, protocol (TCP/443), bytes sent vs. received — a high sent/low received ratio on POST requests to /v1/chat/completions strongly indicates data exfiltration to AI rather than passive browsing; (2) TLS SNI fields from firewall or proxy logs confirming destination FQDN even over HTTPS (most next-gen firewalls and transparent proxies log SNI without decryption); (3) User-agent strings from proxy logs — API clients (curl, Python requests, LangChain) versus browser user-agents indicate whether an employee or an automated integration is sending data; (4) Timestamp correlation between large outbound AI API calls and internal file access events (Windows Security Event ID 4663 — object access) to identify what data was accessed immediately before submission; (5) OAuth application consent logs from Azure AD (sign-in logs filtered on 'Application' = third-party AI app names) or from Google Workspace Admin SDK Reports API showing data scope granted (drive.readonly, mail.readonly, etc.).

Step 3: Eradication — block unapproved AI tool access at the network layer via web proxy policy. Revoke OAuth tokens granted by employees to third-party AI apps without authorization. Audit service accounts and API keys associated with any AI integrations added in the past 12 months.

NIST Phase: Eradication

Reference: NIST 800-61r3 §3.4 — Eradication: removing unauthorized access pathways, revoking ungoverned credentials, and verifying that threat vectors specific to shadow AI integrations are closed before recovery begins

Controls: NIST IR-4 (Incident Handling) — execute the eradication phase of the incident handling process to remove unauthorized AI access pathways and revoke ungoverned tokens, NIST AC (Access Control) — specifically enforce least-privilege and need-to-know by revoking OAuth grants that exceed approved data access scope for AI applications, NIST SI-2 (Flaw Remediation) — treat ungoverned AI integrations holding persistent OAuth tokens or API keys as configuration flaws requiring tracked remediation and verification, NIST AU-9 (Protection of Audit Information) — ensure audit logs of revocation actions are write-protected to preserve the eradication timeline for post-incident review, CIS 5.3 (Disable Dormant Accounts) — identify and disable service accounts created specifically for AI integrations that are no longer authorized or actively monitored, CIS 6.2 (Establish an Access Revoking Process) — execute the access revocation process for all OAuth tokens and API keys granted to unapproved third-party AI

applications

Compensating: For teams without a PAM solution or identity governance platform: (1) Azure AD OAuth token revocation via PowerShell: Connect-MgGraph; Get-MgUserAppRoleAssignment -UserId | Where-Object {$_.PrincipalDisplayName -match 'AI|GPT|Copilot'} | ForEach-Object { Remove-MgUserAppRoleAssignment -UserId -AppRoleAssignmentId $_.Id }; also revoke refresh tokens: Revoke-MgUserSignInSession -UserId ; (2) For Google Workspace: use Admin Console > Security > API Controls > App Access Control to revoke third-party app OAuth grants; export the full grant list via gam all users show tokens | grep -Ei 'openai|anthropic|ai' for scoped review; (3) Audit API keys by searching code repositories and configuration files: grep -rE '(sk-[a-zA-Z0-9]{32,}|OPENAI_API_KEY|ANTHROPIC_API_KEY|GEMINI_API_KEY)' /var/www /opt /home --include='*.env' --include='*.config' --include='*.json' 2>/dev/null; (4) Block AI service domains at proxy/firewall by adding a deny rule for the consolidated AI domain list before the default permit rule — test with curl -x proxy:3128 https://api.openai.com to confirm blocking.

Evidence: Before revoking tokens and blocking domains, preserve as forensic record: (1) Full OAuth token grant inventory export from Azure AD or Google Workspace including granted scopes, grant date, last used date, and granted-by user — this establishes the access timeline and data scope that was potentially exposed; (2) API key usage logs from AI service provider dashboards if accessible (OpenAI usage dashboard shows token consumption by API key, timestamps, and model used — export before revoking keys as provider logs may not be recoverable post-revocation); (3) Service account last-activity timestamps and associated process names from Windows Security Event ID 4624 (logon) and 4648 (explicit credential use) filtered on service account names identified during the sweep; (4) Network flow records for the 12-month window showing cumulative data volume transmitted to AI service endpoints per source account — this is the data exposure quantification needed for any regulatory notification assessment.

Step 4: Recovery — implement a formal AI acceptable use policy with explicit data classification rules (what data may and may not be submitted to AI tools). Establish an AI governance working group with representation from IT, security, legal, and business units. Map approved AI tools and their data access scope.

NIST Phase: Recovery

Reference: NIST 800-61r3 §3.5 — Recovery: restoring a controlled operational state by formalizing governance structures and approved pathways so that AI tool usage resumes only under defined, monitored conditions

Controls: NIST IR-8 (Incident Response Plan) — update the IR plan to include AI governance incidents as a defined incident category with specific detection, containment, and escalation criteria, NIST IR-4 (Incident Handling) — re-establish normal operations under new controls, ensuring AI tool access is restored only through approved, monitored channels, NIST SI-12 (Information Management and Retention) — define explicit data retention and handling rules for data submitted to AI tools, particularly for outputs that may contain synthesized sensitive information, CIS 3.3 (Configure Data Access Control Lists) — apply data classification-based access controls to restrict which data categories employees may submit to approved AI tools, CIS 3.2 (Establish and Maintain a Data Inventory) — align the AI acceptable use policy to the existing data inventory so classification labels (PII, PHI, confidential IP) directly map to AI submission rules, CIS 2.2 (Ensure Authorized Software is Currently Supported) — formally designate approved AI tools in the software inventory with version, vendor, data access scope, and review date

Compensating: For teams without a GRC platform or policy management tool: (1) Use the NIST AI RMF Playbook (available at airc.nist.gov) as a free template to structure the AI acceptable use policy — specifically the MAP function for categorizing AI use cases and risk; (2) Create a simple approved-AI registry as a shared spreadsheet with columns: Tool Name, Vendor, Approved Data Classifications, Approved Use Cases, Owner, Review Date, API Access Y/N, OAuth Scopes Granted — review quarterly; (3) Enforce data classification at the endpoint level using Microsoft Purview Information Protection free tier (sensitivity labels on Office documents prevent copy-paste into browser-based AI tools when configured with DLP policies); (4) For the governance working group operating without formal GRC tooling: run monthly 30-minute reviews using a standing agenda: new AI tool requests received, policy exceptions logged, incidents or near-misses since last meeting, approved list changes — document in a shared OneNote or Confluence page with version history enabled.

Evidence: Before declaring recovery complete, document and retain: (1) The pre-policy baseline — a timestamped export of all AI tool access observed during the detection phase, serving as the 'before' state against which post-policy monitoring will be compared; (2) Approved AI tool configuration records including OAuth scopes, API key rotation

dates, and data classification restrictions applied — this becomes the audit baseline for future compliance reviews; (3) Evidence of policy distribution and acknowledgment (email delivery receipts or LMS completion records) to satisfy NIST IR-2 (Incident Response Training) requirements and demonstrate due diligence if a regulatory inquiry follows; (4) Network proxy allow-list configuration export showing which AI service endpoints are now explicitly permitted versus blocked, with effective date — preserves the control state at recovery for audit purposes.

Step 5: Post-Incident — conduct a gap assessment against NIST AI RMF (AI 100-1) or CISA's AI security guidance. Implement identity controls for AI workloads: dedicated service accounts, least-privilege scoping, MFA where applicable, and quarterly access reviews. Add AI-generated phishing simulation to security awareness training.

NIST Phase: Post Incident

Reference: NIST 800-61r3 §4 — Post-Incident Activity: conducting lessons-learned analysis, updating detection and governance capabilities, and feeding intelligence back into the preparation phase to reduce recurrence of ungoverned AI adoption patterns

Controls: NIST IR-4 (Incident Handling) — document lessons learned from the shadow AI discovery and update the incident handling procedures to include AI governance incidents as a recurring category, NIST IR-5 (Incident Monitoring) — establish ongoing tracking of AI tool usage metrics as a continuous monitoring capability, not a one-time remediation, NIST IR-2 (Incident Response Training) — incorporate AI-augmented phishing scenarios (deepfake voice, AI-personalized spearphishing) into tabletop exercises and security awareness training cycles, NIST SI-7 (Software, Firmware, and Information Integrity) — implement integrity monitoring for AI integration configurations (API keys, OAuth grants, service account permissions) to detect unauthorized changes, NIST RA (Risk Assessment) — specifically NIST AI RMF AI 100-1 GOVERN and MAP functions: formally assess AI-specific risks including model inversion, prompt injection against internal AI deployments, and data poisoning, CIS 6.3 (Require MFA for Externally-Exposed Applications) — enforce MFA for all employee access to approved AI platforms, particularly those with OAuth access to corporate data, CIS 6.5 (Require MFA for Administrative Access) — require MFA for any administrative access to AI platform APIs, dashboards, or model management interfaces, CIS 5.1 (Establish and Maintain an Inventory of Accounts) — add AI workload service accounts to the enterprise account inventory with explicit owner, scope, and quarterly review date, CIS 7.2 (Establish and Maintain a Remediation Process) — formally track and remediate AI governance gaps identified through the NIST AI RMF gap assessment using the risk-based remediation process

Compensating: For teams without a dedicated security awareness platform or red team capability: (1) AI phishing simulation — use GoPhish (free, open source) to run an AI-themed spearphishing campaign: craft lures impersonating IT announcing a new approved AI tool requiring credential re-entry, or a vendor sending an AI-generated contract for signature via a credential-harvesting link — these mirror real RSM-documented attack patterns against middle market firms; (2) For NIST AI RMF gap assessment without a consultant: download the NIST AI RMF 1.0 Playbook from airc.nist.gov and work through the GOVERN and MAP function questions in a structured workshop with the AI governance working group established in Step 4 — document gaps as risk register entries; (3) For service account MFA where the AI vendor does not natively support it: route API calls through an authenticated forward proxy requiring certificate-based mutual TLS, or use a secrets manager (HashiCorp Vault free tier) to rotate API keys quarterly and log every key retrieval event; (4) Quarterly access reviews without an IGA tool: export service account and OAuth grant lists monthly via PowerShell or gam and diff against the prior month's export — flag any new grants for immediate review: diff <(sort accounts_q1.txt) <(sort accounts_q2.txt).

Evidence: For the post-incident record and future audit readiness, preserve: (1) The completed NIST AI RMF gap assessment output with identified gaps, risk ratings, and assigned remediation owners — this is the primary artifact demonstrating due diligence against an established governance standard; (2) Phishing simulation campaign results including click rates, credential submission rates, and department-level breakdowns — establishes the human risk baseline for AI-themed social engineering and satisfies NIST IR-2 training effectiveness documentation; (3) Before-and-after comparison of the AI tool inventory and access scope (approved tools, OAuth grants, service accounts) from pre-incident sweep versus post-recovery state — demonstrates measurable risk reduction and supports any regulatory or cyber insurance reporting requirements; (4) Identity control implementation evidence: screenshots or exports of MFA enrollment status for AI platform accounts, service account permission scope from identity provider, and scheduled access review calendar invites with owner assignments — establishes the control

baseline for the first quarterly review cycle.

Detection Guidance

Shadow AI detection relies on network and endpoint telemetry rather than signature-based tools. Query web proxy or firewall logs for outbound connections to AI provider domains and APIs. Review DLP alerts for large data transfers to cloud services. Audit OAuth application registrations in identity providers (Azure AD, Okta, Google Workspace) for third-party AI app grants employees may have self-authorized. Monitor for new browser extensions on managed endpoints, particularly those requesting clipboard or document access. For AI-augmented phishing detection, tune email gateway rules to flag messages with high reading-level consistency and low sender reputation, characteristics often present in AI-generated spear phishing. (Analyst note: AI-generated content detection is an emerging discipline; test these rules in your environment.) Behavioral indicators of credential attacks (T1078): off-hours logins, logins from new geolocations, and high-velocity authentication attempts against accounts without MFA.

Framework Mappings

MITRE-ATTACK

- **T1078** — Valid Accounts
- **T1204** — User Execution
- **T1195** — Supply Chain Compromise
- **T1566** — Phishing

NIST-800-53R5

- **AC-2** — Account Management
- **AC-6** — Least Privilege
- **IA-2** — Identification and Authentication (Organizational Users)
- **IA-5** — Authenticator Management
- **AT-2** — Literacy Training and Awareness
- **CA-7** — Continuous Monitoring
- **SI-3** — Malicious Code Protection
- **SI-4** — System Monitoring
- **SA-9** — External System Services
- **SR-2** — Supply Chain Risk Management Plan
- **SR-3** — Supply Chain Controls and Processes
- **SI-7** — Software, Firmware, and Information Integrity
- **SC-7** — Boundary Protection
- **SI-8** — Spam Protection

CIS-V8

- **6.3** — Require MFA for Externally-Exposed Applications

- **14.2** — Train Workforce Members to Recognize Social Engineering Attacks
- **15.1** — Establish and Maintain an Inventory of Service Providers

HIPAA-SECURITY

- **164.312(d)** — Person or Entity Authentication
- **164.308(a)(5)(i)** — Security Awareness and Training
- **164.308(a)(6)(ii)** — Response and Reporting

SOC2-TSC

- **CC6.1** — Logical access security software, infrastructure, and architectures
- **CC7.4** — Responds to identified security incidents
- **CC9.2** — Manages risks associated with vendors and business partners

ISO-27001-2022

- **A.5.34** — Privacy and protection of personal information
- **A.5.21** — Managing information security in the ICT supply chain

NIST-CSF-2

- **GV.SC-01** — Cybersecurity supply chain risk management program

MITRE ATT&CK Mapping

Technique ID	Technique Name	Tactic
T1078	Valid Accounts	Defense-Evasion
T1204	User Execution	Execution
T1195	Supply Chain Compromise	Initial-Access
T1566	Phishing	Initial-Access

Sources

Source	URL	Tier
gemini	https://www.morningstar.com/news/prs/2026-05-13-prnewswire-rsm-s-cy...	T3
Cyber Risk Has Become A Financial Strategy Issue For The Middle ...	https://www.forbes.com/councils/forbesfinancecouncil/2026/02/13/why...	T3
Intruder Releases the Security Middle Child Report, Revealing How ...	https://www.businesswire.com/news/home/20260312735575/en/Intruder-R...	T3

Source	URL	Tier
7 Silent Security Gaps That Threaten Mid-Sized Businesses in 2025	https://www.nextdimensioninc.com/2025-smb-cyber-security-checklist/	T3
Small and Midsize Businesses Face Greater Cybersecurity Risks ...	https://www.score.org/resource/article/small-and-midsize-businesses...	T3

Generated 2026-05-13 18:56 UTC by TJS Security Command Center