Kiteworks

TOP 5

Top 5 Data-Layer Takeaways in the Verizon 2026 DBIR

The 2026 Verizon Data Breach Investigations Report reframes the breach story as a data-exchange story. Two-thirds of stolen data is "emails, plans, and reports." Third-party involvement reached 48% of breaches. Shadow AI is now a top three insider behavior with a 4x year-over-year jump. Vulnerability exploitation has overtaken credential abuse as the #1 initial access vector. The DBIR makes the strongest external case yet for unifying email, file sharing, MFT, SFTP, data forms, APIs, and AI data access under a single governed control plane.

1. Shadow AI Tripled in 12 Months and Is Now a Top-3 Insider Behavior Cycle

Regular AI use on corporate devices jumped from 15% of employees last year to 45% this year -- a tripling in 12 months. More consequentially, Shadow AI is now the third most common non-malicious insider action in DLP datasets, a fourfold year-over-year increase. This is no longer a fringe behavior to monitor; it is a mainstream data-egress channel that has scaled faster than any governance program. The DBIR has empirically established the exact failure pattern that governed AI data access is designed to address -- uncontrolled employee-to-LLM data flow at enterprise scale.

2. Source Code and Research IP Are Flowing Into Ungoverned LLMs

Across 858,440 DLP events targeting GenAI tools, source code was the most common data type submitted to external AI by a large margin, followed by images and structured data. In 3.2% of DLP policy violations, research and technical documentation was uploaded to unauthorized AI systems. Separately, 67% of users access AI services from non-corporate accounts on their corporate devices, and more than 15% of users at the average company have unauthorized AI browser extensions installed -- passive collectors that read every page a user visits. The data-loss story isn't customer PII flowing to chatbots; it is the engineering, research, and proprietary work product that defines competitive advantage, leaving through both prompt-paste and silent browser-extension channels.

3. Third-Party Breaches Surged 60% — and the Salesloft Drift Cascade Is the Template

Third-party-involved breaches rose 60% year-over-year, reaching 48% of all breaches. The defining campaign of 2025 was the Salesloft Drift/ Salesforce OAuth token compromise, which cascaded into customer data theft at Google, Zscaler, Cisco, and others. Only 23% of third-party organizations fully remediated missing or improperly secured MFA on cloud accounts, and 37% had an admin account with MFA disabled on an IaaS offering. Every connected AI tool, MCP integration, and AI-enabled SaaS plugin is a new third-party data path -- the Salesloft Drift pattern (tokens stolen from one vendor, used to exfiltrate from another) is the template for what AI-integration breaches will look like at scale.

4. Internal Data — Emails, Plans, Reports — Is the #1 Stolen Asset

Internal data was the most stolen data type in 67% of breaches. Verizon's framing is direct: this is "emails, plans, and reports -- the kind of material you'd expect to be lying around once an attacker strolls in." Credentials were taken in 28% of breaches; personal data in 23%. The same content that constitutes "internal data" in the DBIR taxonomy -- per the Verizon 2026 Data Breach Investigations Report -- is exactly what employees are voluntarily pasting into ungoverned LLMs and what AI agents will be asked to access at scale. The same data attackers are stealing post-breach; employees are exporting pre-breach.

5. The Data Doesn't Support a Ban — It Supports Governed AI Access

Less than 2.5% of AI-assisted attacks involved rare or novel techniques; the median AI-assisted technique already had 55 known existing malware examples. AI is making attackers faster at known things, not unlocking new attack classes. Meanwhile, 45% of employees are regular AI users on corporate devices and 60% of malicious-insider breaches are now driven by convenience -- the same motivation that defeated "don't email work to personal accounts" a decade ago. Prohibition fails. What the DBIR data supports is sanctioned AI access through a governed control plane -- policy enforcement, redaction, and tamper-evident audit across email, file share, SFTP, MFT, data forms, APIs, and AI.