Cybersecurity Incidents

Talal Ahmed | Team Lambda

In the past 18 months dozens of major organizations have suffered significant cyberattacks, data breaches, and supply-chain compromises. Notable cases include South Korea's SK Telecom (data leak of ~27 million SIM subscriber records), Australia's Qantas airline (5 million customers' contact records stolen via a third-party platform), Red Hat (570 GB of consulting code exfiltrated from an internal GitLab), and massive ransomware campaigns exploiting software flaws (e.g. Oracle EBS CVE-2025-61882). There were unprecedented "password dumps", a 16-billion credential compilation from various malware, exposed online. Supply-chain attacks spiked too: the NPM Axios package was poisoned by North Korean hackers (backdoored versions published in Mar 2026). Finally, AI firm Anthropic made headlines by accidentally leaking its internal documents, a misconfigured CMS exposed ~3,000 assets (including an unreleased "Mythos" model), and later an npm packaging error published ~500,000 lines of its Claude AI coding agent source code. These incidents reveal recurring themes: unpatched software and weak configurations are common root causes, threat groups like ALPHV/Cl0p and "Scattered Spider" continue global extortion campaigns, and even internal human errors (misconfigurations, packaging mistakes) can be disastrous. In the sections below I analyze ten major recent incidents in depth (with dates, affected organizations, attack vectors, technical details, impact, timelines, and aftermath) and conclude with a comparative summary table, a timeline chart, a flowchart of common attack chains, and lessons learned. Where available, I cite official reports, vendor advisories, or respected analyses. In all cases I note if any detail remains unspecified in public sources.

Incident Analysis

SK Telecom (April 2025)

Incident

On April 18, 2025, South Korea's largest mobile carrier SK Telecom detected unusually high outbound traffic from its network. A joint government investigation found 28 internal servers had been infected with stealthy malware (including 27 instances of the "BPFDoor" rootkit). Attackers exfiltrated 25 categories of USIM (SIM-card) data totaling 9.82 GB (~26.96 million subscriber records); two servers briefly held

customers' personal data (names, DOB, phone, email) in plain text. The breach stemmed from extremely lax security: SK Telecom's intranet had no administrative passwords and outdated operating systems, allowing easy ingress. Root causes cited were poor credential management, failure to encrypt data, and not addressing a prior 2022 intrusion.

Impact

26.96 million IMSI (subscriber identity) records were leaked, risking SIM-cloning and call interception. The government found SK Telecom neglected required protections and reporting duties. SK Telecom later waived early-termination fees for affected customers. The Personal Information Protection Commission fined SK Telecom KRW 134 billion (~$97 million) on Aug 28, 2025, for negligence. The carrier pledged to overhaul security (stronger passwords, encryption, governance) and incurred no system downtime reported.

Timeline & Response

Attackers gained access by April 18, but the breach was only reported to regulators on April 20 (24 hrs overdue). SK Telecom notified by late April; remediation involved isolating malware and scanning ~42,000 servers. By June 2025 the government team concluded their probe.
SK Telecom publicly acknowledged a "grave responsibility" and committed a KRW 700 billion security investment. No further data leaks have been reported, but investigations highlighted the ease of breach.

Aftermath & Lessons

The incident underscores the critical need for basic defenses (account passwords, patches, encryption). SK Telecom is now mandating CEO-level cybersecurity governance, stronger passwords, data encryption, and increased security staffing. This case illustrates that negligence can lead to massive fines and reputational damage. Companies should ensure rapid incident reporting, as delayed notification itself triggered penalties.

Sources

- Reuters
- Official Investigation

16 Billion Password Aggregation (June 2025)

Incident

In June 2025 researchers discovered a **massive credential leak** consisting of ~16 billion usernames and passwords. This was not a breach of a single company, but an aggregation of data from many infostealer malware and past breaches. Security firm Cybernews reported (June 18, 2025) that the dataset spanned ~30 databases, many overlapping, with data "from the last few years including very recent breaches". Some samples dated only a year old, suggesting ongoing malware collection.

Details & Impact

The leaked “dataset” contains usernames, email addresses and plaintext or hashed passwords, but also some credit cards and personal data in smaller quantities. No new vulnerability was exploited; rather, attackers likely used infostealer Trojans on user machines and then compiled all known stolen data. The scale makes it by far the largest credential compilation to date. Although no single firm was directly breached in this event, any organization with weak password policies or credential reuse could see credential stuffing attacks. The leak underscores how widely user credentials circulate and the enduring value of strong, unique passwords.

Response & Lessons

Since this was researcher-driven analysis rather than a “take down” of a live breach, no patch or notification applied. Cybernews and other outlets warned users to change reused passwords and enable multi-factor authentication. The lesson is that all organizations should assume stolen credentials are in the wild, enforce strong authentication, and monitor for account abuse.

Sources

- Cybernews report

Qantas (June–Oct 2025)

Incident

Qantas Airways disclosed on Sept 11, 2025 that on June 30, 2025 it detected "unusual activity" on a third-party platform used by its contact centers. The attack had accessed a Salesforce-based customer database. Qantas immediately contained the incident; its own airline systems were not breached. The compromised data included names, email addresses, phone numbers, birth dates, and frequent-flyer numbers for *~5 million customers*. Crucially, Qantas confirmed no passwords, PINs, financial data or passport details were stolen.

Attack Vector

Public reports suggest the attackers exploited Salesforce or related credentials, though Qantas has not publicly detailed the initial breach method. The group *Scattered Lapsus$ Hunters* took credit, posting an extortion note on Oct 2025 (demanding ransom for data on ~40 companies). Intelligence firm Intel 471 said Qantas was among ~40 victims of an attack on Salesforce instances, which also hit Gap, Toyota, Disney, IKEA, etc.. Qantas believes an unauthorized third party accessed the third-party CRM during June 2025.

Impact

Up to 5 million Qantas customers had personal data exposed (names, email, FF #s, etc.). While Qantas says frequent-flyer accounts were *not* compromised (passwords and balances intact), the stolen data is enough for personalized phishing. The stolen data was later leaked on the darknet on Oct 11, 2025 after a ransom deadline passed. Qantas obtained a NSW Supreme Court injunction to legally bar use or publication of the data.

Timeline & Response

Detection on June 30, 2025 triggered immediate containment. Qantas set up customer support lines and credit monitoring for affected customers. On Oct 11, 2025 (post ransom deadline), the hacker group publicly leaked 5 M Qantas records. Qantas quickly announced it was investigating the leak and working with cybersecurity experts. Qantas's chief executive later testified the airline had been notified by AWS in early July of a data exposure in an S3 bucket tied to Salesforce, but said no consumer credit or identity documents were involved. (Public detail: exact timeline of AWS notification is partly redacted in filings, but July 2025 was cited in media.)

Aftermath & Lessons

The incident cost Qantas (and insurers) an estimated ~$500 million in insurance payouts according to later reports, and the airline's market value dipped after the disclosure. Qantas now regularly reminds customers never to divulge account info to callers (to prevent phishing fallout). The breach exemplifies supply-chain risk: even an aircraft company is vulnerable via its service providers (CRM platforms). Two key lessons are to secure third-party integrations (strong OIDC, MFA on all enterprise apps) and to maintain an "assume breach" posture. Immediately revoking exposed credentials and aggressively patching/configuring cloud assets are vital.

Sources

- Qantas official advisories
- The Guardian article

Allianz Life (July 2025)

Incident

On July 16, 2025, Allianz Life Insurance Co. of North America reported a breach of a third-party cloud CRM system used by its business in the US. Attackers (identified as the Scattered Spider extortion gang) had stolen customer data from Salesforce-based systems. A Maine AG notice later confirmed 1,497,036 Allianz Life customers, brokers, and select employees were affected.

Attack Vector

The compromise was attributed to unauthorized access of Allianz's Salesforce CRM – likely via stolen credentials or an API exploit. Scattered Spider (a group tied to ShinyHunters) had been mass-targeting Salesforce instances worldwide in mid-2025. The gang later claimed responsibility on their leak site. Allianz says none of its internal systems were breached; the hackers only accessed the vendor's CRM environment.

Impact

Personal data of ~1.5 million people was stolen: names, addresses, dates of birth, and Social Security numbers. Importantly, Allianz reported no financial account numbers or medical data was involved. The exposure is the second major US healthcare insurance breach (following UnitedHealth's 2024 hack of Optum). Allianz is providing two years of free ID theft monitoring to affected individuals.

Timeline & Response

Allianz first detected the incident in mid-July 2025 and contained it, noting it "was limited to the third-party CRM". It notified regulators and the Maine AG (as required). By early October 2025 Allianz notified affected persons and regulators as the exact scope became clear. The company has said it is working with forensic investigators; no ransom payment was reported. Remediation included revoking old CRM tokens, enforcing stronger auth, and closely auditing all third-party integrations.

Aftermath & Lessons

This breach highlights the systemic risk of cloud/SaaS compromises. Organizations must monitor and secure all API integrations (MFA, IP restrictions) on platforms like Salesforce. It also underscores the value of rapid breach notification and pre-arranged support plans (Allianz's immediate offer of credit monitoring). Finally, the incident shows that diversified attackers (ransomware/ad extortion groups) will increasingly target high-value data in the cloud.

Sources

- SecurityWeek reporting
- Reuters

Mailchimp (July 2025)

Incident

On July 26, 2025 the Everest ransomware gang posted claims of having infiltrated Mailchimp (an email marketing platform, now owned by Intuit). They posted an entry on their leak site claiming to have stolen 943,536 lines (~767 MB) of internal company and client data. The trove reportedly included user emails, domain names, tech-stack details, and other CRM data for some customers.

Attack Vector

Everest gave few technical details. Security analysts speculate it may have been a compromised Mailchimp employee or contractor account, or a breach of a connected CRM database. However, Mailchimp (Intuit) stated it found “no evidence” of a data breach in its systems. Cybersecurity observers noted the leaked sample was tiny compared to Mailchimp’s total data volume, suggesting it may have been a single client’s data or even a scammed scraping.

Impact

The published data appears limited; no broad customer data was confirmed lost. No extortion demand was made public (Everest did not post a ransom note in this case). Despite noise on social media, industry experts called it a “breadcrumb” leak with little substantive value. However, the incident alarmed many customers who were unsure what was taken. Mailchimp did reach out to any potentially affected customers with guidance.

Aftermath & Lessons

As of August 2025 Intuit reiterated that Mailchimp systems were secure. There is no public evidence of widespread compromise. Nevertheless, this episode highlights how even smaller leaks can damage trust. Marketing firms handling client data should segregate customer accounts and regularly audit access logs. It also underscores verifying claims: many ransomware groups claim big attacks that sometimes turn out to be exaggerated.

Sources

- Cybernews analysis

Oracle EBS (Aug–Oct 2025)

Incident

In mid-2025, a new zero-day (later tracked as CVE-2025-61882) in Oracle's E-Business Suite (EBS) was weaponized by a threat actor using the CL0P data extortion brand. Beginning in early August 2025, attackers exploited this flaw to gain remote code execution on customer EBS servers. In late September they launched a massive phishing campaign to dozens of organizations, claiming to have stolen sensitive documents from victims' Oracle EBS environments.

Attack Vector

The vulnerability was in the EBS "Template Preview" functionality (OA_HTML/SyncServlet) allowing SSRF and XSL injection. Google's Threat Intelligence group reports this was likely exploited from Aug 9, 2025 onward before Oracle issued emergency fixes. Attackers combined SSRF, CRLF injection and XSLT to execute malicious Java payloads (e.g. a "GOLDVEIN" downloader and "SAGEWAVE" servlet filter) on victims' servers. Over weeks they established persistent access and exfiltrated data from dozens of organizations.

Impact

At least dozens (if not hundreds) of companies were targeted. The CL0P gang publicly extorted victims on its dark web site, warning of data publication if unpaid. To date, however, no large victim list has been posted (CL0P typically waits weeks). Victims are believed to include multiple industries (finance, healthcare, manufacturing). The actual data stolen is unclear, but is presumed sensitive corporate data from EBS (invoices, HR info, etc.).

Timeline & Response

Google/Mandiant began warning in early October 2025. Oracle released patches on Oct 2 and emergency fixes on Oct 4, and urged all EBS customers to apply updates. Victims scrambled through Sept 29 - Oct 2025 to detect and remove the implants. Some US agencies (e.g. NIH, DHS) later reported attacks on their SharePoint or SAP systems, possibly related.

Aftermath & Lessons

This campaign shows the potency of chaining unknown zero-days with extortion. Organizations should deploy virtual patching (WAF rules) immediately upon public disclosure of attacks, not wait for official patches. After-the-fact, companies were advised to audit Oracle logs for unusual TemplatePreview requests or new "TMP*/DEF*" entries (signs of exploit). Mandiant and Google released IOCs (C2 server domains, YARA rules, etc.) for defenders. The key lessons are to minimize attack surface (disable unused EBS modules), enforce strict network segmentation for ERP servers, and maintain up-to-date patching.

Sources

- Google/Mandiant advisory
- The Hacker News

Red Hat Consulting GitLab (Oct 2025)

Incident

On Oct 2, 2025 Red Hat confirmed that an unauthorized party had gained access to one of its internal GitLab instances used by the Consulting division. The breach was discovered only recently; the intruders had "accessed and copied some data" from this GitLab before being expelled. An extortion group calling itself the Crimson Collective claimed to have stolen ~570 GB of compressed data from ~28,000 repositories, including "Customer Engagement Reports" (CERs) containing network details and credentials for Red Hat clients.

Attack Vector

Red Hat has not disclosed exactly how the hackers got in. The Crimson Collective posted that they found unredacted tokens and database URIs in source code/CER files, which they used to pivot into customer systems. Red Hat's public statement merely said they "detected unauthorized access" and immediately launched an investigation. The breach was limited to a self-managed GitLab, and did not affect Red Hat's public software distribution or supply chain.

Impact

Sensitive consulting materials were stolen. Customers are worried because stolen CERs can contain server configurations and credentials. Known CERs of major companies (Bank of America, T-Mobile, U.S. Navy, etc.) appeared on Telegram lists posted by the hackers. Red Hat says there's no evidence this affects its open-source code integrity. The incident has raised concerns because Red Hat's reputation relies on supply-chain trust.

Timeline & Response

The breach was reported publicly Oct 2 after the Crimson Collective leak. Red Hat said it immediately removed attacker access and isolated the instance. They are notifying consulting customers "directly" if affected. As of mid-Oct 2025, Red Hat said it was investigating but believed no further impact. Security experts advise any customer using Red Hat services to rotate all credentials and scan for misuse. Red Hat also recommended clients monitor their infrastructure for any suspicious login attempts using leaked CER info.

Aftermath & Lessons

This incident highlights that even trusted vendors can be hit via "insider" code leaks. Customers of any consulting firm should treat internal engagement docs as sensitive credentials. Red Hat has urged developers to never store real tokens/passwords in code. It also shows the importance of network segmentation: the GitLab was presumably on Red Hat's internal network, but still the attacker exfiltrated hundreds of gigabytes. Moving forward, Red Hat plans enhanced monitoring of its GitLab environments and stricter data-loss prevention.

Sources

- Red Hat official blog

- BleepingComputer

Anthropic "Mythos" CMS Leak (Mar 2026)

Incident

On March 26, 2026, Fortune News revealed that Anthropic (developer of the Claude AI assistant) had inadvertently left ~3,000 internal files publicly accessible on an unsecured content-management system (CMS). These files included draft blog posts, images, PDFs and other materials not meant for public view. Notably among them were details of an unreleased next-generation AI model dubbed "Mythos," described internally as "the most capable model it has yet trained".

Attack Vector

This was not a hack but a misconfiguration. Anthropic's CMS apparently defaulted uploaded assets to public access unless manually marked private. Due to human error, unpublished drafts and images (including the Mythos spec sheet) remained reachable via public URLs. A University of Cambridge researcher discovered the issue and alerted Fortune. Anthropic quickly locked down the data after notification.

Impact

Sensitive pre-release information (product roadmap, model capabilities) was exposed. There was no evidence of customer or user data being involved – only Anthropic's own internal documents. The blog post and images in the trove revealed rumored features of Mythos (claimed "step change" in reasoning, coding, cybersecurity). Anthropic had held an invite-only CEO retreat whose agenda and attendees were also in the leak. No financial or personal customer data was compromised.

Timeline & Response

Fortune contacted Anthropic on Mar 26 and the company immediately secured the CMS. Anthropic attributed the exposure to "human error in the CMS configuration" and stressed Claude/AI tools were not at fault. Within hours, Anthropic made the files private. Media covered the story on Mar 26–27, and by end of day Anthropic had begun reviewing logs for who may have accessed the data. The uncovered data remains unpublished, and Anthropic did not confirm any individuals viewed it.

Aftermath & Lessons

Even “leaks” that are not cyberattacks can be damaging. Companies should apply the principle of least privilege – default to private for sensitive internal systems. Use automated crawling or “bug bounties” for internal reconnaissance. Anthropic said it will improve its CMS processes. More broadly, AI firms (which rely heavily on proprietary innovation) must secure their pre-release research rigorously. The incident also illustrates how AI tools (Stanford researchers noted that leaking via CMS is a classic human error) can ironically be mitigated by better engineering reviews.

Sources

- Fortune investigation

Anthropic “Claude Code” Source Leak (Mar 2026)

Incident

On March 31, 2026, AI startup Anthropic accidentally published its internal source code for “Claude Code” (a command-line AI coding assistant) by including a debug source-map file in its npm package release. This exposed the entire client-side codebase (~513,000 lines across 1,906 files) to anyone who installed that package or found the map file. Within hours the code was mirrored on GitHub with tens of thousands of forks.

Attack Vector

No attacker, this was a packaging error. The Node.js runtime tool Bun automatically generated a full .map source map, and Anthropic failed to exclude it from the npm package. That map contained URLs to a zip of all source files hosted on Anthropic’s cloud, which savvy users immediately pulled down. In short, an internal file meant for debugging was pushed to the public npm registry.

Impact

Competitors and developers now have the full architecture of Claude Code: hooks, memory management, internal AI prompt handling, OAuth flows, hidden features, and more. Tens of thousands of developers stared at Anthropic’s code, raising intellectual property and security issues.

There was no loss of user data or keys (per Anthropic), but the leak revealed unreleased features and “roadmap” of the product. Anthropic’s image suffered, as it markets itself as a safety-focused AI lab.

Response

Anthropic quickly identified the cause (human error bundling a .map file). The spokesperson said it was “not a breach” but a release error and that they were “rolling out measures to prevent this”. Anthropic issued DMCA takedown notices to GitHub mirrors, but by mid-April many forks with 50k+ stars remained (the code was essentially irreversibly public). They also warned developers that “leaked” code is proprietary and not safe to run. No regulatory action is known, though leak of private IP might impact IPO valuation.

Lessons

Rigorous internal code reviews are needed for release automation. The incident shows the risk of over-automation (Bun built a map not intended). Companies should routinely audit build artifacts for sensitive content. From a defensive view, users of Claude Code should assume others know its internals and avoid embedding secrets. More broadly, this underscores that insider mistakes can be as harmful as hacks – even engineers must apply strict release hygiene.

Sources

- Axios news
- Zscaler ThreatLabz report

Axios NPM Package Attack (Mar 2026)

Incident

Early on March 31, 2026, attackers published two backdoored versions of Axios (a popular JavaScript HTTP library) to the public npm registry. These malicious versions (1.14.1 and 0.30.4) contained a hidden phantom dependency (@shadanai/openclaw) that installed a cross-platform remote access trojan. The trojan executed platform-specific second-stage payloads on Windows, macOS, and Linux, giving attackers shells on affected developers’ machines.

Attack Vector

The npm account used to publish Axios was hijacked via a stolen GitHub Actions OIDC token (the build workflow did not require reauthentication). In other words, attackers had compromised the CI/CD pipeline's publishing credentials. Once in, they pushed the compromised releases around midnight on Mar 31. Within ~3 hours NPM detected and removed the malicious packages. Security researchers traced the infection chain: the phantom package's post-install script downloaded a RAT from remote servers, enabling reconnaissance and code execution.

Impact

Axios has ~100M weekly downloads, so about 3% of users (millions of developers and applications) pulled the trojanized code during that brief window. This includes cloud environments and developer machines, potentially compromising source code. Several other npm packages (@shadanai/openclaw, @qqbrowser/contextpad) were found distributing the same malware. No widespread reports of successful exploitation have been publicized, but the risk of developer system compromise was high.

Response

The bad versions were unpublished within hours. The Axios maintainers and npm advised all users to update to safe versions and scan for unusual activity. Wiz Security and Sophos analysts posted detailed breakdowns of the payloads. Future mitigation steps include enforcing stricter auth on CI/CD (e.g. time-limited tokens), verifying published packages through multi-sig or secondary reviews, and promoting tools that detect post-install scripts. Developers were urged to monitor dev machines for the GhostSocks/Vidar malware mentioned in analysis.

Lessons

This is a textbook supply-chain attack. It shows that even infrastructure used by millions can be compromised via a single stolen token. Best practices: tie CI tokens to ephemeral credentials, require human approval for public releases, and use reproducible builds that can be verified. Monitoring of package metadata (by npm and security tools) should flag unusual new dependencies. Finally, teams should run endpoint detection tools on developer machines to catch stealthy post-install malware.

Sources

- SecurityWeek report

Lessons Learned

Across these incidents, several common themes emerge:

- **Harden Access and Patch Quickly:** Many breaches exploited missing patches or weak credentials. Organizations must enforce timely patching (e.g. for ERP and CRM systems) and enforce strong password policies with multi-factor authentication. Unpatched software (like SK Telecom's outdated servers or Oracle EBS) often provided an open door.
- **Secure Third-Party Services:** Breaches through vendors/third parties (Qantas's Salesforce platform, Allianz's cloud CRM) demand rigorous vendor management. Ensure least privilege on connected apps, use IP whitelisting, and monitor third-party logs. Regular audits of third-party IAM and API usage are critical.
- **Least Privilege & Encryption:** The SK Telecom and Mailchimp cases show that unencrypted or public data (including source code) can be catastrophic. Encrypt sensitive data at rest (USIM data, customer PII, code repositories) and limit access. Default to private for internal assets (learn from Anthropic's CMS error).
- **Incident Response Preparedness:** Quick detection and containment (as in Qantas/Allianz) limit damage. Maintain IR plans that include legal (breach notifications, injunctions) and customer communication (support lines, monitoring services). Publishing post-incident "lessons learned" (as some did) helps improve community resilience.
- **Supply-Chain Vigilance:** The Axios and Oracle incidents demonstrate that supply-chain is now a critical battlefield. Employ zero-trust architectures, sign software builds, and use SBOMs (Software Bill of Materials) to track dependencies. Monitor popular package repositories and CI/CD pipelines for anomalies.

- **Monitoring & Logging:** Comprehensive logs and anomaly detection (network, process, repository) can reveal stealthy breaches. For example, SK Telecom was ultimately caught by outbound traffic spikes, and Oracle victims needed to check unusual TemplatePreviewPG requests. Ongoing threat intel feeds (e.g. known malicious IPs/domains) should be integrated.
- **Employee Training:** Since phishing was a likely vector for many cases, continuous security awareness is needed. Especially training on social engineering around new threats and how to handle suspicious requests.

In summary, these cases collectively reinforce that *security hygiene* (patching, least-privilege, monitoring) is non-negotiable, and that even minor oversights can scale to global crises. Organizations should treat the post-2023 era of explosive breaches as the new baseline – preparing as if they will inevitably face attacks.