Vardaan Sdn. Bhd.
Securing Data, Empowering Intelligence

2H 2025

GLOBAL THREAT INSIGHT REPORT

Adaptive Resilience at the Velocity of Risk

Reporting Period

1 July 2025 - 31 December 2025

2025

Table Of Content

Executive Summary

The second half of 2025 has signaled a definitive shift toward a state of constant, industrialized aggression. We are now navigating "The Velocity of Risk" a reality where automated attack cycles move at machine speed, rendering traditional, static security perimeters obsolete. As the window for response shrinks from days to minutes, the fundamental challenge has moved from defending against isolated software flaws to securing the "recursive trust" within our global digital supply chains and cloud identities.

To stay ahead of this evolution, organizations must transition from reactive defense to Adaptive Resilience. This strategy replaces passive, perimeter-based trust with a dynamic, verified process that matches the agility and scale of modern adversaries. This report serves as a strategic roadmap for 2026, providing the intelligence necessary to understand these foundational shifts and harden your infrastructure against the unprecedented speed of the current threat landscape.

Key Takeaway

Cyber threats surged globally in the second half of 2025, with ransomware and AI-driven attacks becoming more sophisticated. Ransomware campaigns evolved, with multi-vector extortion tactics increasing the pressure on victims. Countries like the United States, Germany, and China were top targets, driven by industry prominence and digital infrastructure dependence. Critical sectors such as Healthcare, Energy, and Financial Services saw heightened vulnerabilities, exacerbated by supply chain attacks and cloud identity abuses.

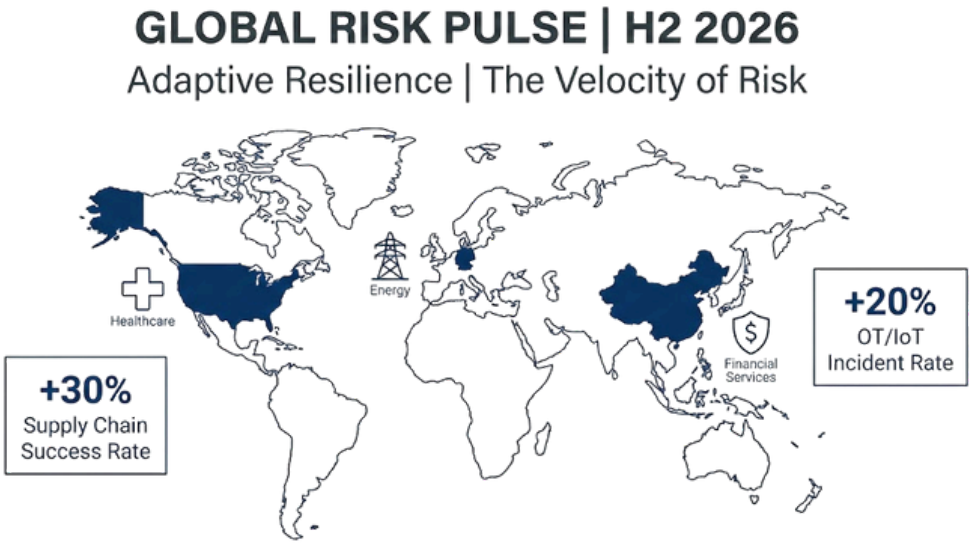

AI-powered threats and automated attack chains grew, making it harder for traditional defenses to keep up. Supply chain breaches, often targeting third-party service providers, led to a 30% increase in successful attacks. Meanwhile, the convergence of OT and IoT systems in critical infrastructure sectors raised new risks, with incidents in energy and transportation systems seeing a 20% increase.

The evolving threat landscape demands improved cyber defense strategies, rapid incident response, and global collaboration to mitigate these emerging risks.

Geopolitical Tensions & State

The second half of 2025 has confirmed that cyberspace is no longer just a support theater for geopolitical conflict; it is a primary front. State-sponsored actors have moved beyond traditional espionage toward strategic pre-positioning for disruption and the weaponization of "deniable" proxy groups.

The Rise of "Living Off the Land" (LOTL) for Persistence

A defining trend in 2H 2025 is the shift by Advanced Persistent Threats (APTs), particularly those linked to China (e.g., Volt Typhoon) and Russia, away from custom malware toward Living Off the Land techniques. By using legitimate administrative tools already present in a victim's system, these actors maintain persistence in critical sectors for months without detection.

- Objective: Not immediate theft, but "pre-positioning" to disrupt power, water, and communications during future diplomatic or physical escalations.

Undersea Cables & Digital Choke Points

Geopolitical competition has moved to the ocean floor. In 2H 2025, there was a measurable increase in "suspicious maneuvers" by state-linked vessels near submarine telecommunications cables in the Baltic Sea and the Indo-Pacific.

- Impact: Since 99% of international data flows through these cables, even minor "accidental" damage (via anchor drags or dredging) serves as a form of grey-zone pressure that bypasses traditional military responses while causing massive economic ripple effects.

“Hacktivism-for-Hire”

Traditional hacktivism (simple website defacement) has evolved into a sophisticated, state-aligned tool. Groups like RipperSec and various pro-Russian/pro-Ukrainian collectives have begun using:

- AI-Generated Disinformation: To amplify the psychological impact of a breach.
- Custom DDoS-for-Hire: Leveraging Node.js-based tools (like MegaMedusa) to launch highly scalable attacks against government and financial institutions in the US, Europe, and Southeast Asia.

Geopolitical Tensions & State

The second half of 2025 has confirmed that cyberspace is no longer just a support theater for geopolitical conflict; it is a primary front. State-sponsored actors have moved beyond traditional espionage toward strategic pre-positioning for disruption and the weaponization of "deniable" proxy groups.

The Rise of "Living Off the Land" (LOTL) for Persistence

A defining trend in 2H 2025 is the shift by Advanced Persistent Threats (APTs), particularly those linked to China (e.g., Volt Typhoon) and Russia, away from custom malware toward Living Off the Land techniques. By using legitimate administrative tools already present in a victim's system, these actors maintain persistence in critical sectors for months without detection.

- Objective: Not immediate theft, but "pre-positioning" to disrupt power, water, and communications during future diplomatic or physical escalations.

Undersea Cables & Digital Choke Points

Geopolitical competition has moved to the ocean floor. In 2H 2025, there was a measurable increase in "suspicious maneuvers" by state-linked vessels near submarine telecommunications cables in the Baltic Sea and the Indo-Pacific.

- Impact: Since 99% of international data flows through these cables, even minor "accidental" damage (via anchor drags or dredging) serves as a form of grey-zone pressure that bypasses traditional military responses while causing massive economic ripple effects.

“Hacktivism-for-Hire”

Traditional hacktivism (simple website defacement) has evolved into a sophisticated, state-aligned tool. Groups like RipperSec and various pro-Russian/pro-Ukrainian collectives have begun using:

- AI-Generated Disinformation: To amplify the psychological impact of a breach.
- Custom DDoS-for-Hire: Leveraging Node.js-based tools (like MegaMedusa) to launch highly scalable attacks against government and financial institutions in the US, Europe, and Southeast Asia.

2024 Approach	Feature	2H 2025 Evolution
Data Theft & Espionage	Primary Goal	Pre-positioning for Disruption
Custom Malware	Tooling	Living Off the Land (LOTL)
Cloud & Servers	Battlefield	Physical Infrastructure (Cables/Utilities)
Direct Attribution	Method	Deniable Proxies & AI Psy-Ops

Regulatory Shifts & Cyber Insurance Evolution

As we enter 2026, the global regulatory landscape has moved from "suggested best practices" to "enforced digital stability." Organizations that fail to bridge the gap between technical security and legal compliance now face severe financial penalties and "uninsurable" status.

The Shift to Active Enforcement: NIS2, DORA, and Beyond

The grace periods for major frameworks have ended. In 2H 2025, we saw the "Compliance Line in the Sand" being crossed.

- DORA (Digital Operational Resilience Act): As of early 2025, financial entities and their critical ICT third-party providers are now under strict oversight. The focus in 2H 2025 shifted to mandatory resilience testing, it is no longer enough to have a policy; you must prove your systems can survive a simulated outage.
- NIS2 Enforcement: EU member states (such as Sweden and Finland) have moved into the active enforcement phase for 2026. The most significant change is the personal liability of top management for cybersecurity failures, effectively moving cyber risk from the IT department to the Boardroom.
- The Cyber Resilience Act (CRA): Late 2025 saw the first wave of reporting obligations for hardware and software manufacturers. By 2026, any "product with digital elements" must report actively exploited vulnerabilities within 24 hours.

Cyber Insurance: The "Buyer's Market" is Hardening

While premiums were relatively flat in early 2025, the market is beginning to harden again as we move toward 2026, with S&P Global forecasting potential increases of 15-20% due to rising claim severity.

- From "Encryption" to "Extortion": Insurers are recalibrating because ransomware actors have pivoted. While backups have made encryption less effective, data exfiltration and suppression (threatening to leak sensitive data) are driving higher payouts.
- The AI Clause: In 2H 2025, a new category of "AI Liability" emerged. Insurers are now introducing specific clauses for "Chatbot Abuse" and "Prompt Injection" losses. If your organization uses GenAI without a formal governance framework, you may find these incidents excluded from standard policies.

The "Compliance Inequity" Crisis

A dangerous gap is widening between large enterprises and SMEs.

- The Data: Recent reports show that while large organizations have halved their "insufficient resilience" reports, 35% of small organizations now feel they can no longer adequately secure themselves against 2026-grade threats.
- The Result: Smaller vendors in the supply chain are becoming the "weakest link," leading to the surge in supply chain attacks.

Compliance is no longer a "check-the-box" exercise. In 2026, Continuous Controls Monitoring (CCM) is becoming the standard for both insurance eligibility and regulatory favor.

Cyber Threat Landscape

Ransomware Evolution & Multi-Vector Extortion

Ransomware has transcended simple encryption. Threat actors now employ "multi-vector" pressure combining data theft, public shaming, and DDoS attacks to create a total operational and reputational crisis for victims.

AI-Powered Threats & Automated Attack Chains

The democratization of Artificial Intelligence has birthed "synthetic adversaries." AI is now used to automate every stage of the kill chain, from hyper-personalized reconnaissance to adaptive malware that bypasses traditional defenses at machine speed.

Supply Chain Attacks & Cloud Identity Abuse

Adversaries are increasingly targeting the "recursive trust" in digital ecosystems. By compromising a single Managed Service Provider (MSP) or hijacking a cloud OAuth token, attackers can achieve a "one-to-many" impact, moving laterally into hundreds of downstream customer environments with near-total invisibility.

Convergence of OT & Critical Infrastructure Risks

: The digital-to-physical bridge is now a primary target. The integration of Operational Technology (OT) and IoT has expanded the attack surface into energy, transportation, and healthcare, where a digital breach can result in immediate physical consequences and societal disruption.

Ransomware Evolution & Multi-Vector Extortion

Ransomware remains one of the most serious and persistent cyber threats organizations face. What began as simple file encryption for ransom has evolved into a complex, highly strategic form of cybercrime. As technology and business systems have advanced, so have attacker methods, turning ransomware into a sophisticated, multi-layered threat.

By 2025, ransomware operations extend far beyond encryption. Attackers now use multi-vector extortion, combining data theft, public leaks, reputational threats, and even DDoS attacks to pressure victims. Instead of just selling decryption keys, criminals exploit sensitive data, operational disruption, and weaknesses in interconnected systems and third parties to force payment.

As organizations grow more digital and data-dependent, defending against these campaigns has become harder. Victims often feel compelled to negotiate to avoid both downtime and long-term reputational damage. Double extortion stealing data before encryption and threatening to release it has raised the stakes, with some groups adding public leak sites, misinformation, and service disruption to intensify impact.

By mid-2025, no organization is off-limits. The growing complexity of ransomware makes prevention and recovery increasingly difficult, highlighting the urgent need for stronger, layered defenses. Understanding these multi-extortion tactics is essential for organizations preparing for the future threat landscape.

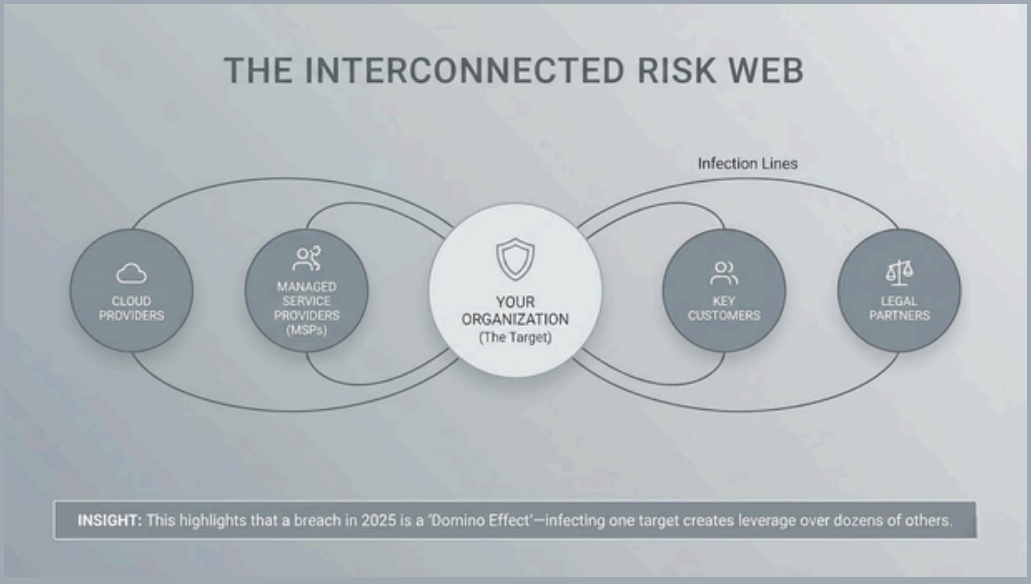

Salesforce Supply Chain Data Theft and Extortion

In August 2025, a major data theft campaign hit over 700 organizations using Salesforce. This wasn't a "traditional" hack where a criminal guesses a password; instead, it was a sophisticated indirect supply chain compromise. The attackers didn't break into Salesforce directly. Instead, they stole the "digital spare keys" known as OAuth tokens, from a third-party marketing tool called Salesloft Drift. By compromising the trust between these two apps, the attackers gained a silent, invisible pathway into sensitive customer databases.

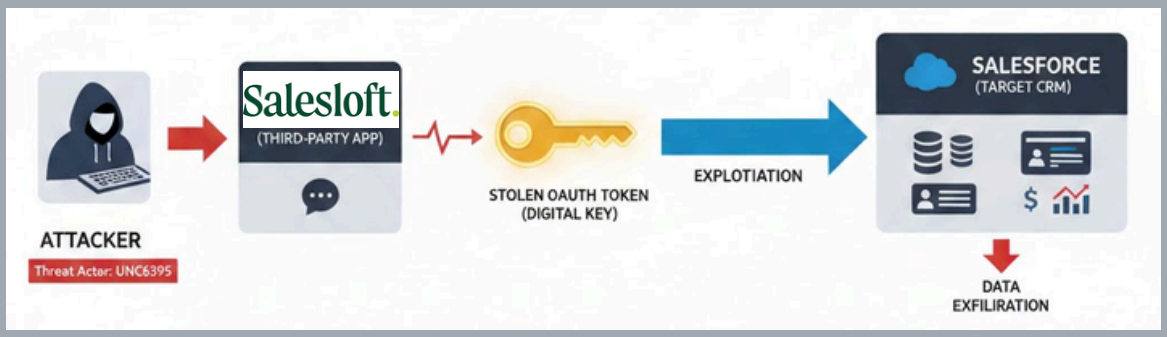

The Mechanics of Indirect Access

The threat actor, identified as UNC6395, bypassed standard login screens and Multi-Factor Authentication (MFA) entirely. By abusing these hijacked OAuth tokens, they gained programmatic access, which is like having a backstage pass that never expires. This allowed the adversary to use automated scripts to "vacuum up" CRM records, including contact details, support cases, and account information. Crucially, they also targeted "cloud secrets" sensitive API keys or passwords that users had mistakenly saved in plain-text notes within the platform.

Multi-Vector Extortion at Scale

The impact was widespread, affecting sectors from retail to finance. Rather than encrypting files and locking users out (the "old" way of doing ransomware), the group, linked to the notorious ShinyHunters, shifted to Multi-Vector Extortion. They claimed to have stolen nearly 1 billion records and threatened to leak them publicly unless a ransom was paid. This strategy uses reputational damage as a weapon, putting immense pressure on companies to pay even though their day-to-day operations were technically still running.

Strengthening the Ecosystem

To defend against these "trust-based" attacks, organizations must move toward Identity-First Security. This means treating third-party app permissions with the same rigor as admin passwords.

- Audit Your "Digital Keys": Regularly review and revoke OAuth tokens for integrations that are no longer in use or that have excessive permissions.
- Implement Least Privilege: Only give third-party apps the minimum amount of data access they need to do their jobs.
- Continuous Controls Monitoring (CCM): Use automated tools to watch for unusual data movement between your cloud apps, which can signal that a token has been hijacked.
- Data Hygiene: Strictly prohibit the storage of passwords or secrets in free-text CRM fields where they are easily harvested.

The Oracle EBS Zero-Day Extortion Campaign

Between August and October 2025, global firms faced a crisis targeting Oracle E-Business Suite (EBS)—the "corporate brain" managing payroll, HR, and financials. This was a zero-day exploitation campaign, meaning attackers used a flaw (CVE-2025-61882) that was previously unknown to defenders. Rather than locking systems, attackers linked to the Cl0p ransomware brand focused on a "smash and grab" of the company's most valuable secrets.

Remote Control via BI Publisher

The attack exploited a critical vulnerability in the BI Publisher component. This allowed for Remote Code Execution (RCE), essentially giving hackers a "digital command line" to the server without needing a valid password. Because the system failed to "sanitize" (verify) incoming data, attackers gained "God-mode" access. This allowed them to move through the network for months, quietly harvesting financial records and employee data before being detected.

Extortion Through Reputation

The fallout was massive, with at least 29 global organizations publicly shamed by the attackers.

This was Multi-Vector Extortion: the threat wasn't about losing data access, but the public release of Social Security numbers and bank details. This psychological warfare bypasses traditional backups; you cannot "restore" your way out of a public leak. The pressure centered on regulatory fines and the permanent loss of customer trust.

Securing the Core

Oracle released emergency patches on October 4, 2025. Moving into 2026, organizations must prioritize the following:

- Code Red Patching: Treat ERP updates as the highest priority; in a zero-day world, even a 24-hour delay is a massive risk.
- Attack Surface Management: Hide public-facing tools behind a VPN or a Zero Trust Gateway to limit visibility to hackers.
- Anomalous Monitoring: Set alerts for unusual data "egress" (data leaving the network). If your ERP system suddenly sends gigabytes of data to an unknown IP, the connection should be killed automatically.
- Data Masking: Encrypt or mask sensitive fields (like SSNs) within the database so that even if data is stolen, it remains unreadable.

In modern extortion, "zero-day" flaws are the ultimate skeleton key. Unlike phishing, which requires a human to click a link, these flaws allow attackers to walk through the wall of your perimeter defenses without any user interaction at all.

AI-Powered Threats & Automated Attack Chains

The global threat landscape is undergoing a radical transformation as AI-powered threats and automated attack chains move from speculative concepts to standard operating models. We are now seeing the rise of the "Synthetic Adversary" automated systems that combine Large Language Models (LLMs), deepfake generators, and adaptive frameworks to conduct cyberattacks at machine speed. By industrializing what used to be a manual craft, AI compresses attack timelines from weeks to hours, allowing attackers to launch highly tailored operations at a scale that is impossible for purely human-driven security teams to intercept.

At the start of the "AI Kill Chain," attackers use massive data-harvesting tools to build a digital "Attack Graph" of an organization. This identifies key decision-makers, reporting lines, and technical vulnerabilities automatically. Rather than crafting a single virus, hackers use AI to generate Polymorphic Malware, a digital chameleon that constantly changes its own code to stay invisible to traditional, signature-based antivirus tools. By the time a security team identifies one version of the threat, the AI has already reshaped it into something new.

Social engineering has also entered a dangerous new phase of Hyper-Personalization. Using Deepfake voice and video technology, adversaries can now conduct real-time "vishing" (voice phishing) calls that perfectly mimic a trusted executive's voice or appearance.

Because these models are trained on real corporate data, they can mirror internal jargon and specific project contexts, removing the "red flags" like poor grammar or generic messaging

that people usually look for. This shifts the pressure from technical firewalls to less-monitored channels like instant messaging and video conferencing.

Once a foothold is established, AI agents handle the "day-to-day" work of the attack, such as finding the shortest path to sensitive data or bypassing Endpoint Detection and Response (EDR) systems. These agents operate in a "feedback loop," meaning if one exploitation method is blocked, the AI automatically learns from that failure and tries a different, more subtle technique. This level of autonomy means that human operators only need to step in for high-level strategic decisions, while the AI handles the complex, technical progression of the campaign in the background.

The $25 Million "Synthetic" Executive

In the second half of 2025, the boundary between reality and digital fabrication blurred with the rise of deepfake-enabled executive fraud. This pattern represents a shift from simple phishing to a high-stakes psychological operation where attackers "hack" human trust rather than digital firewalls. In one landmark case, a financial institution lost $25 million in a single afternoon. The breach began not with a virus, but with a video call where an employee believed they were receiving direct, urgent orders from their CFO to close a confidential merger.

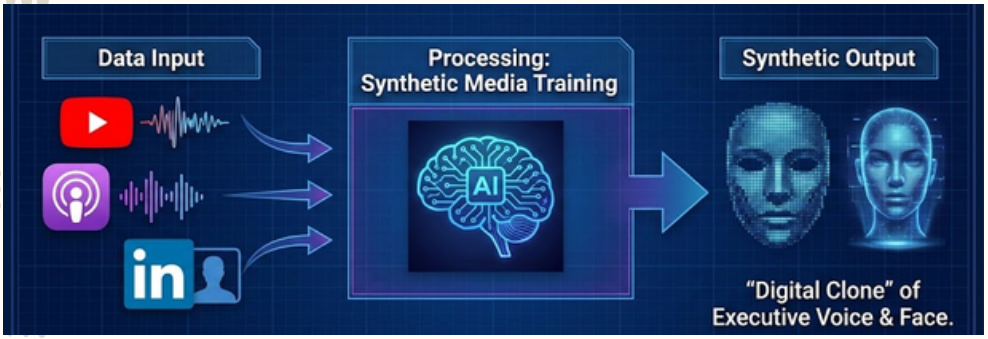

How the Clone is Created

The technical backbone of this deception is a process known as Synthetic Media Training. During the reconnaissance phase, adversaries harvest public recordings of an executive, such as earnings calls, YouTube interviews, or keynote speeches, to "teach" an AI model the specific cadence of their voice and the nuances of their facial expressions. Simultaneously, large language models scan the company's recent press releases and LinkedIn activity to ensure the fake executive can speak fluently about real, current business projects. This removes the traditional "red flags" of fraud, such as generic language or awkward pauses, making the clone indistinguishable from the real person in a high-pressure environment.

The execution typically follows a "multi-channel" approach designed to overwhelm the victim's critical thinking. A perfectly crafted email serves as the initial "hook," setting the stage for a subsequent voice call or video conference that acts as the "sinker."

By combining these vectors, the Synthetic Adversary creates an environment of extreme urgency and confidentiality. Because theemployee "sees" and "hears" their boss, they are psychologically conditioned to bypass standard security protocols. This highlights a growing Resilience Gap: while our technical perimeters are getting stronger, our reliance on "perceived familiarity" has become a massive, unpatched vulnerability in the corporate hierarchy.

Trust as a Vulnerability

To survive in this era of synthetic deception, organizations must transition from a culture of "visual trust" to one of verified process. The path forward requires formalizing out-of-band verification protocols where any high-value transaction, regardless of who is on the screen asking for it, must be confirmed through a secondary, pre-agreed physical channel. By implementing strict "Segregation of Duties" and adopting internal "Safe Words" that exist only in the physical world, companies can build a defense that is resilient to even the most convincing AI clones. The lesson of 2025 is clear: in an AI-driven world, if you can't verify the request through a secondary "offline" path, you must assume the "online" version is a fake.

AI-Accelerated Ransomware-as-a-Service (RaaS)

The Ransomware-as-a-Service (RaaS) ecosystem has undergone a massive industrial upgrade, moving from manual, human-led intrusions to high-speed, semi-autonomous operations. In late 2025, we observed RaaS operators integrating AI across every stage of the criminal lifecycle, from finding the "perfect" victim to closing the ransom deal. This shift has transformed ransomware from a craft into a "smart factory" of extortion, where AI handles the heavy lifting of identifying revenue profiles, assessing security maturity, and even predicting which companies are most likely to pay.

The "front-end" of these attacks is now managed by automated scouts. Machine learning models ingest vast amounts of public financial data and technical scans to pinpoint organizations that are both wealthy and vulnerable. Once a target is selected, AI-guided scanners probe infrastructure to "chain" multiple minor vulnerabilities together into a major breach. These toolchains execute lateral movement with minimal human oversight, allowing attackers to navigate a network and identify sensitive data backups in a fraction of the time it would take a manual operator.

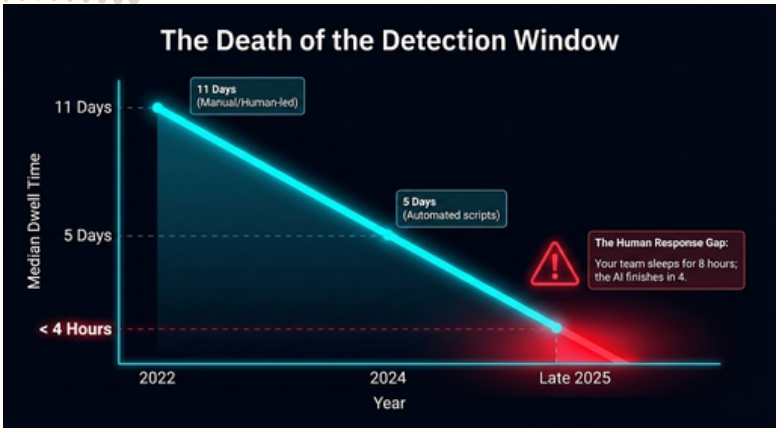

The ransomware payloads themselves have evolved into "digital chameleons." Using Polymorphic Engines, AI constantly reshapes the malware's code, encryption methods, and obfuscation techniques. This ensures the payload stays one step ahead of traditional, signature-based antivirus tools, by the time a security system learns to recognize one version of the virus, the AI has already generated a new one. Even the negotiation phase has been automated; AI-driven bots now handle ransom chats, dynamically adjusting their demands based on the victim's public financial statements and the perceived value of the exfiltrated data.

Perhaps the most alarming trend in 2025 is the total collapse of the attack timeline. The "dwell time", the window between an attacker's first entry and a total system lockdown, has shrunk from weeks to just a few hours. This means traditional security playbooks designed to detect hackers over several days are now dangerously obsolete. For modern defenders, the lesson is clear: detection and response can no longer be measured in days or even hours. To survive an AI-accelerated attack, security operations must be engineered to react in minutes.

Ransomware has shifted from "Volume" to "Velocity." The goal is no longer just to hit many targets, but to hit them so fast that the human defenders don't even have time to open their laptops before the data is gone.

Critical Infrastructure Under Siege in 2025

In 2025, the "air gap" has effectively dissolved. Cyber threats against Operational Technology (OT) have surged, now constituting 18.2% of all identified threat categories according to ENISA. The convergence of IT and OT has expanded the attack surface, allowing adversaries to pivot from digital espionage to physical disruption.

Threat actors are increasingly "living off the land" by abusing legitimate industrial protocols and exploiting unmonitored wireless networks to manipulate control systems. Simultaneously, AI has lowered the barrier to entry; over 80% of social engineering attacks now leverage AI to automate reconnaissance and phishing.

What sets this wave apart isn't just the volume of attacks, but the strategic shift toward kinetic impact. It is no longer about stealing data; it is about stopping production, manipulating safety systems, and holding physical infrastructure hostage.

This section explores two standout cases:

1. A ransomware siege on Pakistan Petroleum Limited (PPL) by the Blue Locker group, utilizing double-extortion tactics to disrupt national energy operations and threaten the release of sensitive geological data.
2. An ideological sabotage campaign by the Infrastructure Destruction Squad (IDS) using VoltRuptor, a bespoke malware designed to manipulate smart building automation systems and demonstrate physical vulnerability.

Manufacturing has become the primary theater of conflict, absorbing 59.3% of cybercriminal attacks due to its low tolerance for downtime. Agencies like CISA and the NSA have also warned of state-sponsored campaigns, such as Brickstorm, which target hypervisors to "pre-position" malware for future sabotage.

The Wireless Blind Spot

While IT focuses on ransomware, Nozomi Networks warns of a silent killer: Wireless Deauthentication. By exploiting legacy WPA2 networks that lack Management Frame Protection (MFP), attackers can disconnect Automated Guided Vehicles (AGVs) and safety sensors. This causes physical production halts without ever touching a firewall.

Together, they show how the industrialization of cybercrime and state-aligned sabotage are converging to threaten the physical safety and operational continuity of global infrastructure.

When Ransomware Held a Nation's Energy Hostage

In the sweltering heat of August 2025, the screens at Pakistan Petroleum Limited (PPL) went dark. PPL is the state-owned giant responsible for keeping the lights on in Pakistan. What unfolded over the next few days was not just a corporate IT outage. It was a textbook example of modern "double extortion" targeting national sovereignty.

The "Defensive Disconnect"

On August 6, PPL's security operations center detected anomalies pulsing through the network core. Realizing the severity of the intrusion, the company initiated a "defensive disconnect." This drastic measure involves severing digital links to the outside world to stop the bleeding. While PPL publicly maintained that core operational systems remained unaffected, the reality behind the scenes was far more volatile. The National Cyber Emergency Response Team (NCERT) of Pakistan issued a high-priority advisory to 39 federal ministries and classified the risk as severe.

The Adversary

Blue Locker The attackers, a group known as Blue Locker, were not amateur script kiddies. Forensic analysis by Resecurity revealed they were using a sophisticated ransomware strain likely rebranded from the notorious Proton and Black Basta families. Crucially, their malware was designed with a professional's touch. It encrypted data files by adding a .blue extension but deliberately skipped system-critical files like .exe and .dll. They did not want to crash the computers. They wanted to keep the operating systems alive just enough to display the ransom note.

Data as a Geopolitical Weapon

The true leverage was not the encrypted servers. It was what Blue Locker had stolen. The ransom note delivered to PPL was chillingly specific. It stated they had stolen business data including TMC Data for Sui, Adhi, and other fields.

> "We have stolen some of your business data... including but not limited to TMC Data (Sui, Adhi, etc.) and contracts... If you don't contact us... we will release your data to social media and competitors." - Excerpt captured by Resecurity

By explicitly naming Sui and Adhi, two of Pakistan's most critical gas fields, the attackers demonstrated they were not just holding files hostage. They were holding the nation's geological blueprints. This proprietary data is the lifeblood of energy exploration. Its release to competitors and social media, as threatened, would constitute a form of economic warfare that could undermine PPL's competitive edge and national energy security.

The Smart City Blackout

The VoltRuptor Kinetic Attack Chain

1. Attackers scan the internet for unprotected industrial control system interfaces (HMIs/SCADA) instead of using traditional phishing.

2. VoltRuptor malware is deployed, which can "speak" the native language of industrial machinery like ModBus.

3. The group directly manipulates system controls to cause physical disruptions to a facility's operations.

4. Attackers record themselves controlling the industrial panels and share the videos to amplify the psychological impact.

5. The malware's anti-forensics capabilities erase logs and commands, making it difficult to reconstruct the attack.

While PPL fought for its data, a different kind of war was brewing in Europe. On June 30, 2025, an Italian smart building automation company became the testing ground for a new era of "kinetic" cyber threats. These are attacks designed to reach out from the digital world and manipulate physical reality.

The perpetrator was the Infrastructure Destruction Squad (IDS). This is a pro-Russia hacktivist group also tracked as Dark Engine. Unlike traditional hacktivists who deface websites for attention, IDS aims for "infrastructure-level interference." Their goal is not financial gain. Instead, they seek the demonstration of power and the erosion of trust in Western infrastructure.

VoltRuptor The attack was executed using VoltRuptor. This is a bespoke malware framework designed specifically for Industrial Control Systems (ICS). According to ENISA, VoltRuptor is not just a repackaged IT virus. It is fluent in industrial protocols and capable of sending commands directly to the machinery that controls elevators, HVAC systems, and access doors.

The most disturbing aspect of the IDS campaign is its psychological dimension. Following the breach, IDS and its affiliates, such as the Z-PENTEST-ALLIANCE, posted screen recordings of their members actively tampering with Human-Machine Interfaces (HMIs). These videos serve as proof-of-concept propaganda. They show unauthorized cursors moving across control panels, clicking buttons, and altering setpoints.

Perhaps most alarming for the future of OT security is the commoditization of this capability. Intelligence indicates that VoltRuptor is now available for sale on the dark web. This signals the "democratization" of kinetic weapons. Sophisticated sabotage tools are no longer the exclusive domain of nation-states. They are now a product anyone with enough cryptocurrency can buy.

Supply Chain Attacks on Managed Services & Cloud Identity Abuse

In the second half of 2025, the traditional "secure perimeter" has officially become a relic of legacy IT. As organizations have accelerated their move to cloud-native architectures and outsourced core functions to Managed Service Providers (MSPs), the "Supply Chain" has evolved into the very nervous system of global business operations. This shift has birthed a new era of Identity-Centric Warfare. Adversaries no longer spend months trying to break through a fortified front gate; instead, they exploit the "Recursive Trust" inherent in modern digital ecosystems, targeting the service providers, API integrations, and cloud identities that already hold "master keys" to the network.

The industrialization of these attacks reveals a critical vulnerability in how modern trust is managed. We are witnessing a fundamental transition from the exploitation of software code to the exploitation of relationships. In this landscape, a hijacked service account or a stolen OAuth token, the digital "pass" that allows apps to talk to each other, is often more valuable than a zero-day exploit. These assets allow attackers to "live off the cloud," using legitimate, pre-approved administrative tools to move invisibly through hundreds of downstream customer environments simultaneously.

For security leaders, this development demands a radical rethinking of resilience. It is no longer enough to secure your own internal house if your "maintenance crew" (your vendors, MSPs, and integrated SaaS apps) holds a master key that can be stolen or abused. By compromising a single link in the service chain, threat actors can achieve a "one-to-many" impact, inheriting the trusted permissions of a provider to bypass Multi-Factor Authentication (MFA) and infiltrate the heart of an enterprise without ever triggering a traditional alarm.

In 2026, the most dangerous threat to your network isn't a hacker at the door; it's the "trusted" app you've already invited inside. When you authorize a vendor, you aren't just buying a service, you are extending your security perimeter to their office.

The ‘ServiceGrid’ RMM Hijack

In August 2025, the cybersecurity community witnessed a definitive "force-multiplier" attack targeting ServiceGrid, a prominent Managed Service Provider (MSP). ServiceGrid provided Remote Monitoring and Management (RMM) tools to over 1,500 small and medium-sized enterprises (SMEs) globally. The attackers, tracked as STORM-0822, didn't bother attacking those 1,500 companies individually. Instead, they compromised the "trusted management plane" that governed them all. By seizing control of the MSP’s central console, the hackers essentially turned a legitimate administrative tool into a high-speed ransomware delivery vehicle.

Hijacking the Digital Update

The breach began with a targeted credential-stuffing campaign against ServiceGrid’s senior DevOps engineers. Once inside, the attackers exploited a session-hijacking vulnerability in the administrative portal, which notably lacked phishing-resistant Multi-Factor Authentication (MFA). With "God-mode" access to the central console, the threat actors used the platform’s Automated Deployment feature to push a malicious "security update" to every managed device. Because this payload was signed with ServiceGrid's legitimate digital certificate, local security tools (EDR) trusted it completely. In one move, the LockerGoga 3.0 ransomware was deployed across 400+ customer environments simultaneously.

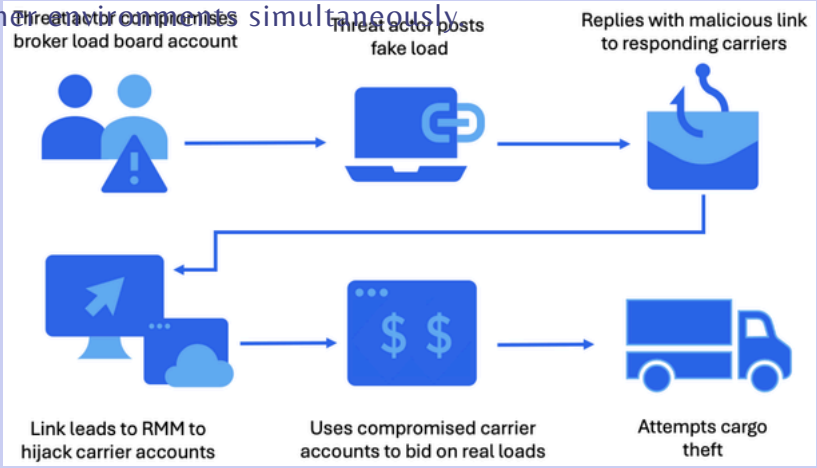

The impact was catastrophic and immediate, particularly for healthcare clinics and local governments. Within six hours, approximately 22,000 workstations and servers were encrypted. The genius, and cruelty of the attack lay in its sequence: the attackers used the MSP's tools to disable backups before starting the encryption. This left victims with zero recovery options other than paying the ransom. By the end of the month, total economic damages, including ransom payments and business interruption, exceeded $450 million USD.

From Passive Trust to Active Verification

The ServiceGrid incident proves that your security is only as strong as your most privileged vendor. Moving into 2026, organizations must enforce a "Zero-Standing-Privilege" model:

- Phishing-Resistant MFA: Traditional SMS or app-based codes are no longer enough for administrators. FIDO2 hardware keys (Passkeys) are now the baseline requirement.
- Immutable (WORM) Backups: Backups must be logically separated from the management plane. If a hijacked RMM tool can "see" your backups, it can delete them.
- Just-In-Time (JIT) Access: Administrative rights should only exist for a specific window of time and require "Dual-Authorization" (four-eyes principle) for any mass-deployment tasks.

The ‘Cloud-Hop’ Campaign via OAuth Hijacking

While traditional breaches often focus on cracking usernames and passwords, the "Cloud-Hop" campaign of September 2025 bypassed the login screen entirely. Instead, threat actors targeted OAuth tokens, the digital "keys" that allow SaaS platforms (like a marketing tool and a CRM) to communicate with each other without asking for a password every time. By compromising a secondary conversational AI integration, the attackers were able to "hop" from a minor third-party app directly into the core corporate databases of over 700 global enterprises.

Invisible Lateral Movement

The attack, attributed to the group UNC6395, did not involve a direct infrastructure breach. Instead, the hackers injected a malicious script into the update pipeline of a popular AI integration used by sales teams. This script was designed to harvest active OAuth access tokens from legitimate users as they worked. Because these tokens are "pre-approved" passes, the attackers could programmatically log into customer instances as a "trusted application." To the security system, this didn't look like a hack; it looked like legitimate, automated traffic moving between two authorized apps.

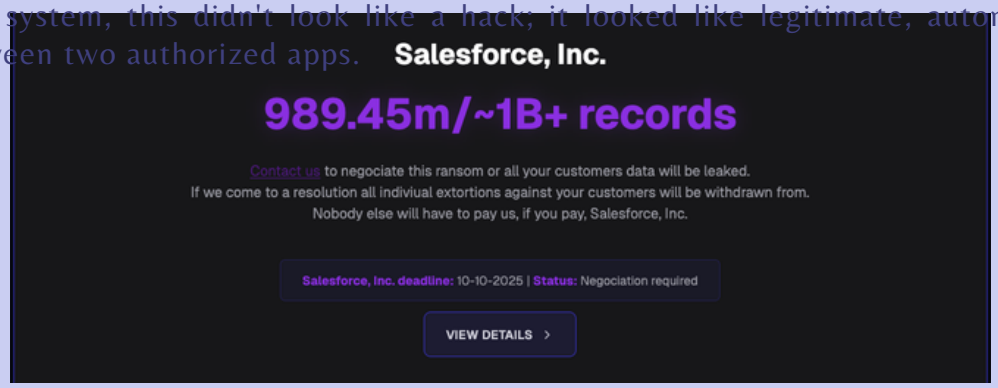

1 Billion Records Exposed

The speed and scale of the "Cloud-Hop" campaign were staggering. Over a period of just 10 days, nearly 1 billion records were harvested, including sensitive customer contacts, support histories, and internal cloud secrets. This wasn't just about data theft; it was the foundation for Multi-Vector Extortion. Affiliates of the ShinyHunters group contacted victims, threatening to leak proprietary customer lists to their direct competitors unless a ransom was paid. The incident proved that in a modern SaaS ecosystem, "recursive trust" is a double-edged sword: the more apps you connect, the more "backdoors" you create.

Identity Governance for Apps

To close the "identity gap," organizations must stop focusing solely on user logins and start governing Application Identities:

- The "Least Privilege" Audit: Many apps request "Full Data Access" when they only need to read a single column. Review every integration and strip away excessive permissions.
- Continuous Token Monitoring: Use security tools (like CASBs) to watch for "Impossible Travel." If a vendor app normally talks to you from a US server but suddenly connects from an unknown overseas IP, the token must be revoked.
- Automated Expiry: Establish a "Use it or Lose it" policy. Any OAuth token that has been inactive for 30 days should be automatically revoked to shrink your attack surface.

REFERENCES

European Union Agency for Cybersecurity (ENISA) | ENISA Threat Landscape 2025 (October 2025)

Nozomi Networks Labs | OT/IoT Cybersecurity Trends and Insights Report (July 2025)

Resecurity | Blue Locker Analysis: Ransomware Targeting Oil & Gas Sector (August 2025)

Cyble | Hacktivists Attacks on Critical Infrastructure (Q2/Q3 2025)

Mandiant (Google Cloud) | Brickstorm Espionage Campaign

CISA, NSA, & Canadian Centre for Cyber Security | Joint Malware Analysis Report: BRICKSTORM Backdoor (December 2025)

Dutch National Cyber Security Centre (NCSC-NL) | Advisory on Citrix NetScaler CVE-2025-6543

Safe Security | The Cyber Risk Singularity

Sangfor | Supply Chain Attack via Drift/Salesloft on Salesforce

Reuters | Almost 1 Billion Salesforce Records Stolen, Hacker Group Claims (October 2025)

Cybersecurity Dive | Salesforce Refuses Extortion Demands Following Mass Data Breach

Fortra | Salesforce Data Breach: What You Need to Know

SonicWall | Oracle E-Business Suite Under Siege: Active Exploitation of Dual Zero-Days

Oligo Security | CVE-2025-61882: Oracle E-Business Suite Zero-Day Exploited in Cl0p Extortion Campaigns

Google Cloud Blog | Oracle E-Business Suite Zero-Day Exploitation and Post-Exploitation Analysis

CYBERSECURITY CONSULTING AND ASSESSMENTS

VAPT / Red Teaming

ISMS/ISO/Compliance Assessment

Forensic Investigation

Security Awareness Training

Risk Management & Governance

Security Operation (SIEM/SOAR & SOC/NOC Operations)

Zero Trust Security

XDR / MDR

Threat Intelligence

Attack Surface & Security Posture Management

Mobile Application Security

Breach and Attack Simulation (BAS) & Security Validation

Platinum Sentral, Kuala Lumpur

+60327224705

info@vardaan-cyber.com

vardaan-cyber.com

SECURING DATA, EMPOWERING INTELLIGENCE