CROWDSTRIKE

LESSONS FROM THE FRONT LINES:

A PLAYBOOK FOR OUTPACING MODERN ADVERSARIES

In this Guide

Security outcomes are rarely determined by a single alert or a single control failure. Adversaries string together what looks harmless in isolation: a trusted identity, a misconfigured permission, a third-party integration, an overlooked out-of-date edge device, an unmonitored control plane. The intrusion does not announce itself. It blends in, crosses domains, and forces hard decisions when time is limited and certainty is scarce. That is the operating reality for modern defense.

Lessons from the Front Lines: A Playbook for Outpacing Modern Adversaries is built for this reality. Anchored to the CrowdStrike 2026 Global Threat Report, it takes the year's most consequential adversary themes and converts them into a practical playbook for prevention and readiness. The guidance reflects what CrowdStrike's frontline responders, threat hunters, and intelligence teams see in active intrusions and what CrowdStrike Professional Services teams apply to help organizations harden environments, validate controls, and reduce the gaps adversaries repeatedly exploit.

This paper is organized around six adversary themes. Each theme includes what is being observed in the field, the preventative steps that most reliably reduce exposure, and how response unfolds when prevention is tested. The intent is not to add another checklist. It is to help teams focus effort where it changes outcomes: shrinking the number of trusted paths attackers can abuse, making cross-domain movement harder to sustain, and improving the speed and clarity of containment and eviction when an incident is already underway.

The Role of Intelligence in Modern Threat Prevention

Identifying and stopping the evasive adversary is challenging, but it is not impossible. The adversary themes in this guide reflect the intelligence insights surfaced through CrowdStrike Counter Adversary Operations, which combines threat intelligence, managed threat hunting, and trillions of telemetry events from the AI-powered CrowdStrike Falcon® platform to detect, disrupt, and stop evasive adversaries.

CrowdStrike delivers these capabilities through CrowdStrike Falcon® Adversary Intelligence and CrowdStrike Falcon® Adversary OverWatch™. Falcon Adversary Intelligence provides real-time, personalized intelligence on adversaries, campaigns, and vulnerabilities, helping organizations understand who is targeting them, how they operate, and where they are most exposed. Falcon Adversary OverWatch delivers expert-led, proactive threat hunting across endpoint, identity, cloud, and third-party data, identifying hidden and cross-domain activity that automated tools alone may not catch. Together, these offerings strengthen security operations by turning adversary insight into decisive action.

The sections that follow build on this perspective, combining adversary intelligence with frontline incident response and consulting experience to translate observed tradecraft into practical prevention and readiness guidance.

Table of Contents

ADVERSARY NAMING CONVENTIONS

BEAR	RUSSIA
BISON	BELARUS
BUFFALO	VIETNAM
CHOLLIMA	DPRK (NORTH KOREA)
CRANE	ROK (REPUBLIC OF KOREA)
HAWK	SYRIA
JACKAL	HACKTIVIST
KITTEN	IRAN
LEOPARD	PAKISTAN
LYNX	GEORGIA
OCELOT	COLOMBIA
PANDA	PEOPLE'S REPUBLIC OF CHINA
SAIGA	KAZAKHSTAN
SPHINX	EGYPT
SPIDER	eCRIME
TIGER	INDIA
WOLF	TÜRKIYE

Adversary Themes

AI Threat Duality: The Adversary's Advantage and Enterprise's New Attack Surface

Threat Overview

In 2025, the threat landscape reached a critical turning point as adversaries transitioned from AI experimentation to direct operational use. This shift created a dual threat: adversaries weaponizing AI to enhance attacks while simultaneously targeting AI systems themselves. By leveraging AI, adversaries evolve their tradecraft to accelerate attack speeds and lower technical barriers to entry, enabling both nation-state and eCrime threat actors to scale operations with unprecedented efficiency, autonomy, and reach.

Frontline Trends and Observations

CrowdStrike responders have observed adversaries using AI to develop and refine social engineering, including deepfakes designed to match convincing or legitimate personas. Activity linked to the North Korean-nexus adversary FAMOUS CHOLLIMA reflects this shift toward more polished deception.

In multiple cases, artifacts consistent with persona deepfake and social engineering AI tooling were surfaced in endpoint evidence sources. This reinforces the importance of durable endpoint visibility when attackers experiment quickly and iterate their tradecraft.

Responders have also seen malicious use of enterprise AI tooling, including internally and externally hosted large language models (LLMs), to fabricate or modify company records. In these incidents, traces of LLM interaction were observable in hosting and provider logs. That visibility is critical because the most consequential activity may occur in the AI platform and its surrounding control plane, not only on the endpoint.

The malicious use of AI tools, whether by external adversaries or nefarious insiders, represents a substantial and growing threat that organizations will increasingly confront in 2026 and beyond.

Prevention Playbook

With adversaries increasingly using AI in victim environments to augment their attacks, organizations need to exercise control over what AI tools are present and what they can do.

DRIVE A TOP-DOWN AI STRATEGY

Monitoring AI requires visibility across multiple planes within an organization's environment, including browsers, endpoints, and cloud environments. Unifying visibility and setting clear objectives for AI tooling across these planes requires an organization-wide AI governance strategy. Documenting approved and unapproved tools is an essential step in establishing AI security governance. Organizations should identify approved AI tooling and models, while blocking unapproved AI tools to limit shadow AI such as OpenClaw. Additional measures include discovering and inventorying AI usage across the environment and documenting use cases.

> "Documenting approved and unapproved tools is an essential step in establishing AI security governance.

MONITOR HOW AI TOOLS ARE USED

Adversaries are exploiting victims' AI tooling as a means for collection. Unauthorized data access from AI systems, whether intentional or not, can be reduced by defining ownership of data within the environment, classifying it, and implementing clear access policies. Organizations should implement strong identity verification procedures for access to AI tools and connected resources. They should also apply established access control principles to AI systems, including least privilege and Zero Trust for users accessing information through AI tooling.

Sensitive data connected to an unmonitored AI model increases the risk of misuse and may trigger regulatory reporting requirements if abused. Organizations should consider integrating AI detection and response (AIDR) capabilities into their security stack to monitor internal AI tooling use. Prompts and outputs should be continuously monitored and logged to support root cause analysis during an incident. Immutable logs should also be maintained for model training data and configuration.

TEST IT BEFORE TURNING IT ON

Organizations should protect AI systems from exploitation by establishing criteria to assess the security of third-party vendors. When selecting a new AI product, document and secure the tool's dependencies and require secure configurations. Organizations should establish change control processes for AI tooling configuration, including for enabling AI features in resources such as SaaS platforms. AI models should undergo red teaming and adversarial testing to evaluate risk, and employees should complete AI-focused security training before being granted access to AI tools.

To defend against adversaries such as FAMOUS CHOLLIMA seeking employment at a victim organization, HR personnel and hiring managers should complete awareness training on potential red flags. Examples include applicants with AI-generated social media profiles and photos, and interviewees using LLMs to answer questions during interviews. Organizations should also be prepared to respond to incidents originating from or involving AI tools, with incident response roles and responsibilities defined and business continuity plans in place to minimize operational disruption.

How CrowdStrike Helps in a Crisis

As adversaries expand their use of LLMs and agentic AI tooling such as OpenClaw, the likelihood of misconfigured or unsecured deployments increases. This can create opportunities for direct or indirect prompt injection, as well as broader abuse of AI-connected workflows.

In an active incident, CrowdStrike responders focus on two priorities in parallel: quickly identifying and securing AI tooling that could be exploited, and reviewing AI tool and LLM interaction logs to determine what may have been manipulated by an external adversary or insider threat. To do this, responders combine endpoint forensics with CrowdStrike Falcon® Forensics for rapid triage (or traditional forensic imaging for deeper analysis), and leverage CrowdStrike Falcon® for IT to identify AI tooling in use and CrowdStrike Falcon® AI Detection and Response (AIDR) to help identify and prevent prompt injection attacks against AI tooling.

Ransomware Adversaries Expand Cross-Domain Tradecraft

Threat Overview

Ransomware adversaries are increasingly avoiding well-protected systems and attacking the gaps instead. In 2025, the most effective ransomware groups gained access through social engineering, cloud accounts, and unmanaged infrastructure, then encrypted data in places security teams often can't see or control. By moving across cloud, identity, and virtualization environments, attackers were able to steal data and disrupt operations without triggering traditional endpoint defenses. This approach helped ransomware remain one of the most damaging threats of the year and will continue to drive attacks in 2026.

Frontline Trends and Observations

CrowdStrike responders continue to see ransomware intrusions that pivot quickly across domains to reach high-leverage infrastructure. One common pattern begins in the cloud or on-premises environment, then moves into virtualized infrastructure, most often VMware ESXi. In these cases, adversaries use privileged credentials such as vpxuser and root to maintain persistence through compromised identities, then deploy ransomware built to run on hypervisors. With a single hypervisor often hosting hundreds of virtual machines, ESXi remains a lucrative target for ransomware operators.

A second common pattern starts with identity. Adversaries compromise a single sign-on (SSO) identity, then move between on-premises environments and cloud-based SaaS applications with minimal friction. In one case, eCrime threat actors initially compromised on-premises infrastructure and then hijacked a hybrid identity solution by stealing the directory synchronization agent configuration to create a malicious agent that decoded user credentials into plaintext. In other investigations, eCrime actors used voice phishing (vishing) to persuade victims to reset passwords, enroll unauthorized multifactor authorization (MFA) devices, or approve one-time MFA prompts that enabled access. As organizations centralize access to SaaS applications through SSO, that convenience is increasingly being exploited as a cross-domain access path for ransomware adversaries.

Prevention Playbook

Cross-domain ransomware campaigns thrive where controls and ownership break down across identity, infrastructure, and cloud. The most effective prevention strategies focus on closing those gaps before an adversary can chain them together.

PRIORITIZE IDENTITY HARDENING, ESPECIALLY FOR PRIVILEGED ACCESS

Hardening the identity layer is a key element in stopping, or at least slowing, cross-domain attacks. Organizations should start by identifying and reducing over-privileged accounts and groups, expanding monitoring, and increasing protections for the privileged accounts that remain. Organizations can limit privilege escalation by using just-in-time access, such as Privileged Identity Management (PIM) in Entra ID, using less privileged administrator accounts for day-to-day tasks like desktop support and server administration, and vaulting credentials in a privileged access management (PAM) tool.

To prevent adversaries from leveraging privileged accounts for cross-domain attacks, organizations should implement identity isolation across environments. For example, use dedicated cloud-only Entra ID accounts for Entra ID and Azure administration rather than hybrid identities synchronized from on-premises Active Directory. Similarly, establish separate identity providers, such as a dedicated Entra ID domain, for managing virtualization environments. This isolation helps ensure that an adversary who compromises domain administrator credentials cannot automatically gain control over virtualization infrastructure, creating an effective boundary between these environments. The CrowdStrike Cloud Security Assessment can assist in identifying account hardening opportunities.

TREAT INFRASTRUCTURE ACCESS AS A HIGH-RISK CONTROL PLANE

Infrastructure access is a repeated point of failure in ransomware intrusions. Many organizations do not adequately secure core IT infrastructure components such as backups, storage, and virtualization platforms, often giving them the same level of protection as individual servers. Attacks on these systems can have a far broader impact. Organizations should treat them as high-value control plane systems by applying additional controls and limiting access, including within the IT organization.

Ransomware adversaries frequently target service accounts and cached administrator credentials to achieve privilege escalation and lateral movement. Flat networks and overly broad administrator privileges allow them to quickly pivot into core IT infrastructure, often with devastating results.

Hardening programs should include vCenter, ESXi, on-premises backups, storage appliances, and other common adversary targets. Key measures include isolating management interfaces on a dedicated VLAN, blocking most outbound traffic, limiting access to a small number of IT administrators, requiring MFA, and applying system-specific, vendor-recommended security controls. By layering these controls, organizations can limit the damage from ransomware adversaries even if they obtain privileged access. For organizations that are uncertain how to proceed with hardening different control planes, CrowdStrike's Pulse Services can provide ongoing guidance to help identify and protect the targets adversaries focus on most.

VALIDATE DEFENSES AGAINST REAL CROSS-DOMAIN ATTACK PATHS

Organizations should use ransomware defense assessments, red team and blue team exercises, or adversary emulation engagements to scrutinize programmatic and technical defenses against these intrusions. These engagements help teams identify where assumed controls break down in practice, including gaps in monitoring between domains, inconsistent enforcement of privileged access, and missing detection and response coverage in virtualization and SaaS layers. They also produce a prioritized set of improvements that map directly to the attack paths ransomware operators are using today. Cloud and SaaS security assessments systematically identify misconfigurations, excessive privileges, and risky access patterns across major public clouds and over 60 SaaS platforms.

CONTINUOUSLY IDENTIFY AND REMEDIATE MISCONFIGURATIONS

Organizations should use insights from CrowdStrike Falcon® Identity Threat Protection, CrowdStrike Falcon® Shield, CrowdStrike Falcon® Exposure Management, and other Falcon modules to continuously identify and remediate misconfigurations and attack paths that adversaries can exploit to gain privilege, move laterally, or pivot between on-premises, cloud, and virtualization environments. This includes surfacing poor password hygiene, over-privileged accounts, missing security controls such as MFA, and exploitable vulnerabilities. These tools can help identify exposure points that enable cross-domain footholds, allowing defenders to drive remediation before they are abused. CrowdStrike's Pulse Services offer one option to maintain a regular cadence for identifying and remediating both misconfigurations and the underlying processes that give rise to them. Over time, this continuous cycle reduces easy access paths available to ransomware operators and lowers the likelihood that an endpoint or account compromise expands into a broader, cross-domain incident.

ESTABLISH CROSS-DOMAIN VISIBILITY

To detect cross-domain attacks, organizations should ensure their security teams have visibility across endpoints, core infrastructure, cloud, and SaaS. This includes ingesting key log sources from infrastructure systems such as VMware, backups, and storage; cloud environments (both control-plane logs and resource-level telemetry); and identity systems.

CrowdStrike Falcon® Next-Gen SIEM provides purpose-built parsers and detection rules for these diverse data sources, enabling correlation across domain boundaries to identify attack patterns. Enabling additional Falcon modules such as CrowdStrike Falcon® Cloud Security, Falcon Identity Threat Protection, and Falcon Shield can further strengthen behavioral detection across domains by providing the context needed to surface suspicious activity that siloed monitoring often misses.

“ To detect cross-domain attacks, organizations should ensure their security teams have visibility across endpoints, core infrastructure, cloud, and SaaS.

How CrowdStrike Helps in a Crisis

CrowdStrike responds to ransomware and extortion intrusions by running containment, data collection, and eviction in parallel. Using Falcon modules such as Falcon Shield, Falcon Next-Gen SIEM, and Falcon Exposure Management, responders can rapidly collect evidence across endpoint fleets, cloud control planes, and SaaS environments to determine what the adversary has done, what they are doing now, and what must be shut down to prevent impact.

In virtualization-heavy environments, this approach is often the difference between a localized intrusion and widespread business disruption. In one case, Falcon Next-Gen SIEM connectors detected an adversary staging ESXi ransomware inside a victim's VMware infrastructure. With early visibility into that staging activity, responders coordinated rapid containment and eviction actions that prevented encryption of multiple hypervisors and hundreds of virtual machines.

When the intrusion is centered on SaaS environments, the investigative focus shifts to application-layer activity and identity. Falcon Shield supports both historical collection and real-time monitoring, helping responders reconstruct access and actions taken inside SaaS platforms. Paired with Falcon Identity Threat Protection, teams can apply additional controls such as step-up MFA challenges and authentication restrictions to tighten containment and limit follow-on access during eviction.

Some events require an immediate consultative response if the victim organization is facing imminent encryption. These cases often involve an indefensible posture, with limited visibility across endpoints and cloud environments, misconfigured security policies, and unmanaged assets, even as the adversary has already staged ransomware. In these scenarios, CrowdStrike advises the fastest path to prevent catastrophic encryption. This may include disconnecting from the internet, stabilizing the environment, and establishing a more defensible posture using the Falcon platform and strengthened security controls before restoring connectivity.

To keep pace with fast-moving ransomware operators, CrowdStrike Incident Response bundled with CrowdStrike Active Defense Services (ADS) provides near real-time remediation of persistence mechanisms, configuration changes, and other malicious artifacts. This helps sustain containment and supports a streamlined, successful eviction process.

China-Nexus Threat Actors Target Network Perimeter Access for Initial Intrusions

Threat Overview

China-nexus adversaries increasingly gain initial access by exploiting internet-facing edge devices — such as VPNs, firewalls, and gateways — faster than organizations can patch them. These systems are often lightly monitored and typically lack endpoint detection and response (EDR) coverage and consistent logging, allowing attackers to remain hidden after compromise. By rapidly weaponizing newly disclosed vulnerabilities, often within two to six days, these adversaries establish long-term access that supports sustained, large-scale espionage against high-value targets worldwide.

Frontline Trends and Observations

In many investigations, CrowdStrike responders observed China-nexus adversaries exploit network edge devices such as load balancers, VPNs, and firewalls, then pivot, often within hours, into the victim's virtualization environment. Once there, the adversary established ghost VMs as an initial foothold and operating infrastructure, maintaining access for months, and in some cases, years.

Consistent with China-nexus operations, the tradecraft was fast, efficient, and methodical. Responders observed multiple persistence mechanisms, including webshells, remote access trojans, and network proxies, alongside stolen identities and session cookies used to broaden access. Ghost VMs provided the adversary with additional operational security and helped sustain long-term control.

This perimeter-driven access enabled China-nexus adversaries to steal sensitive data, including client and customer information, source code, and intellectual property. In many cases, the activity spanned on-premises, cloud, and SaaS environments, using a range of compromised identities, increasing both the scope and complexity of response requirements.

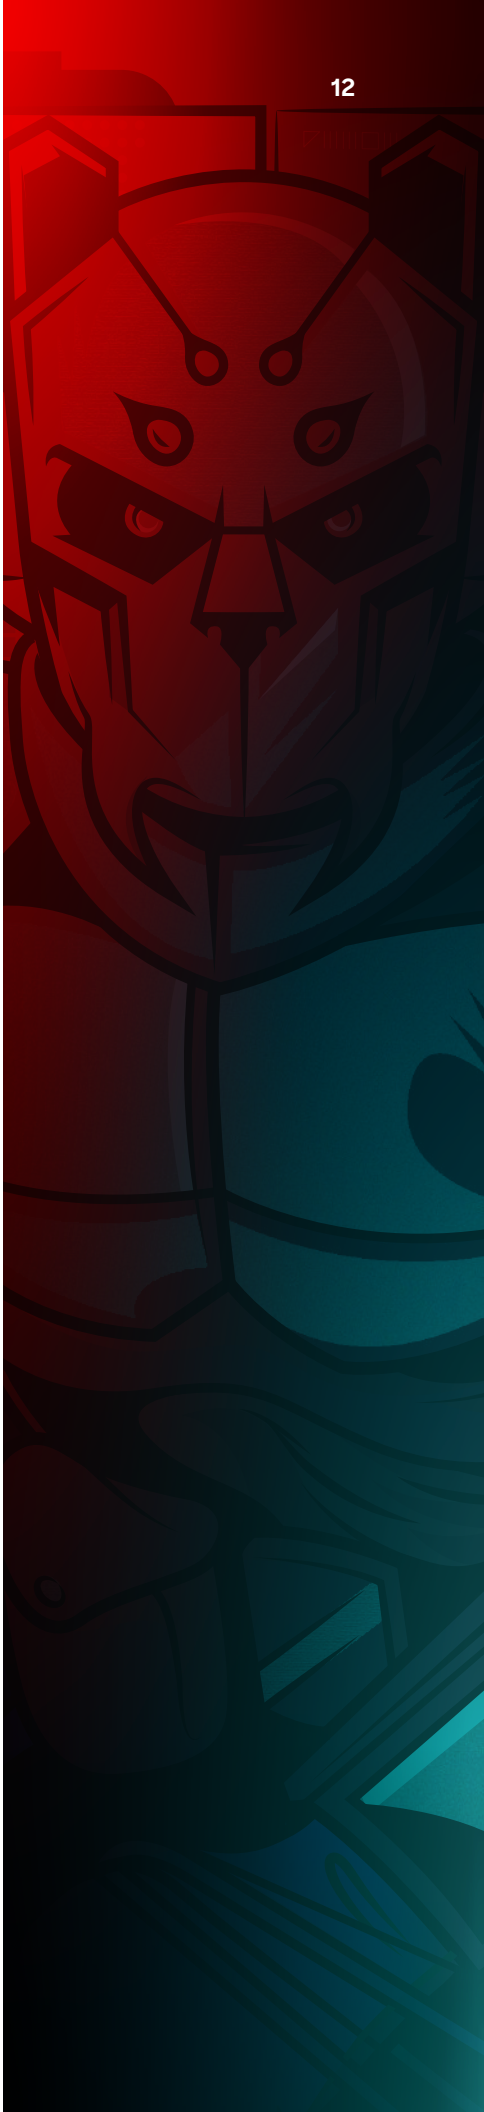

Prevention Playbook

China-nexus adversaries are among the most prolific and have grown increasingly sophisticated and stealthy. Countering their attacks requires a defense-in-depth approach.

HARDEN THE PERIMETER

Hardening against China-nexus adversaries starts with understanding which assets make up the network perimeter. As a first step, organizations should establish a clear inventory of all on-premises and cloud-based assets exposed to the internet. Which assets are internet-facing can shift through both intended and unintended administrative changes. External attack surface management (EASM) tools can help identify real-time changes in exposure. Organizations can also reduce the external attack surface by replacing traditional VPN appliances with Zero Trust network access solutions. Once the perimeter is defined, teams should review and harden configurations for perimeter devices, deploy EDR where feasible, control access between these devices and the internal network, and apply strong vulnerability management practices, including frequent scanning and off-cycle patching or remediation for critical vulnerabilities.

Given the persistent risk of zero-day exploits, organizations cannot rely on vulnerability management alone to mitigate China-nexus attacks. Organizations should implement defense-in-depth mitigations for techniques most likely to follow successful exploitation, such as credential harvesting and privilege escalation. Start with rigorous, consistent security hygiene for all internet-facing assets. For example, do not expose management interfaces to the internet, store credentials in plaintext, or use administrator accounts with excessive privileges.

BOLSTER IDENTITY AND NETWORK DEFENSES

Defense-in-depth should also account for common lateral movement targets, with a particular focus on virtualization infrastructure. Organizations should secure hypervisors from multiple perspectives. From an identity perspective, consider enrolling all domain-joined accounts with administrative access to hypervisors in a privileged access management solution and require MFA with factors resistant to social engineering, such as hard tokens. For enhanced account security, defenders should use a separate identity provider to further limit exposure to privilege escalation. From a network perspective, organizations should isolate virtualization management consoles on a dedicated subnet, allow access only through a hardened jump host, and block all outbound traffic from hypervisors. From a host perspective, teams should maintain rigorous patching for hypervisors and virtualization components, forward relevant logs such as authentication and configuration changes via syslog to a SIEM, and create detection rules for unexpected administrative activity.

PRACTICE AGAINST LIVE FIRE

Finally, organizations should use red team/blue team and adversary emulation exercises to not only identify how adversaries can target and exploit network perimeter assets, but also to map the most likely attack paths a China-nexus adversary could take to move laterally after successful perimeter exploitation.

How CrowdStrike Helps in a Crisis

Similar to supply chain intrusions, China-nexus investigations require a methodical, efficient approach to determine what the adversary did, how they gained access, and which persistence pathways they established in the victim environment.

In practice, this means rapidly collecting and correlating evidence across the environments China-nexus adversaries commonly access or compromise, then using that visibility to drive containment and eviction. Falcon modules such as Falcon Identity Threat Protection, Falcon Shield, Falcon Next-Gen SIEM, and Falcon Exposure Management support rapid, robust data collection across endpoint fleets, identity providers, cloud control planes, and SaaS environments to help responders understand what the adversary is doing or has already done.

Falcon Next-Gen SIEM accelerates this work by enabling rapid ingestion of real-time and historical data from a victim's existing SIEM, alongside raw, unmanaged datasets such as perimeter edge device logs. By unifying these sources, CrowdStrike responders can set alerts for known indicators of attack and quickly identify adversary activity in datasets that are typically cumbersome and reactive to triage.

> "China-nexus investigations require a methodical, efficient approach.

Supply Chain Attacks Enable Evasion of Traditional Security Controls

Threat Overview

Attackers are increasingly breaking into organizations by exploiting trust in software and cloud services rather than attacking individual companies directly. In 2025, adversaries took advantage of the assumption that software is clean, updates are safe, and development and SaaS ecosystems are secure. By abusing trusted distribution channels, integrations, and dependencies, attackers turned software providers and platforms into force multipliers — a single compromise could cascade to thousands of downstream organizations. These attacks often bypass traditional security controls, enabling widespread credential theft, data exposure, and long-term compromise at scale.

Frontline Trends and Observations

CrowdStrike observed an increase in supply chain attacks in 2025. These compromises exploit a basic reality of modern business: Organizations must trust partners, vendors, and third-party applications. That trust can provide adversaries with a privileged path into the environment, and it often contributes to longer dwell time because the compromised entity or application is already considered legitimate.

In one example, adversaries exploited a vulnerability in a victim's application to steal an API key used for managing customer remote access. They then used that access to embed additional persistence into the victim's downstream customer environment. In another instance, adversaries abused third-party identity integrations to gain access into victim environments. A notable public example is the recent Salesloft incident, where adversaries stole OAuth tokens and used them to impersonate the Drift app. This enabled access to customer data in SaaS applications where Drift was already trusted and integrated.

The common thread is trust: When a vendor, integration, or non-human identity in the supply chain is compromised, the intrusion can be widespread and harder to detect because it blends into normal business activity.

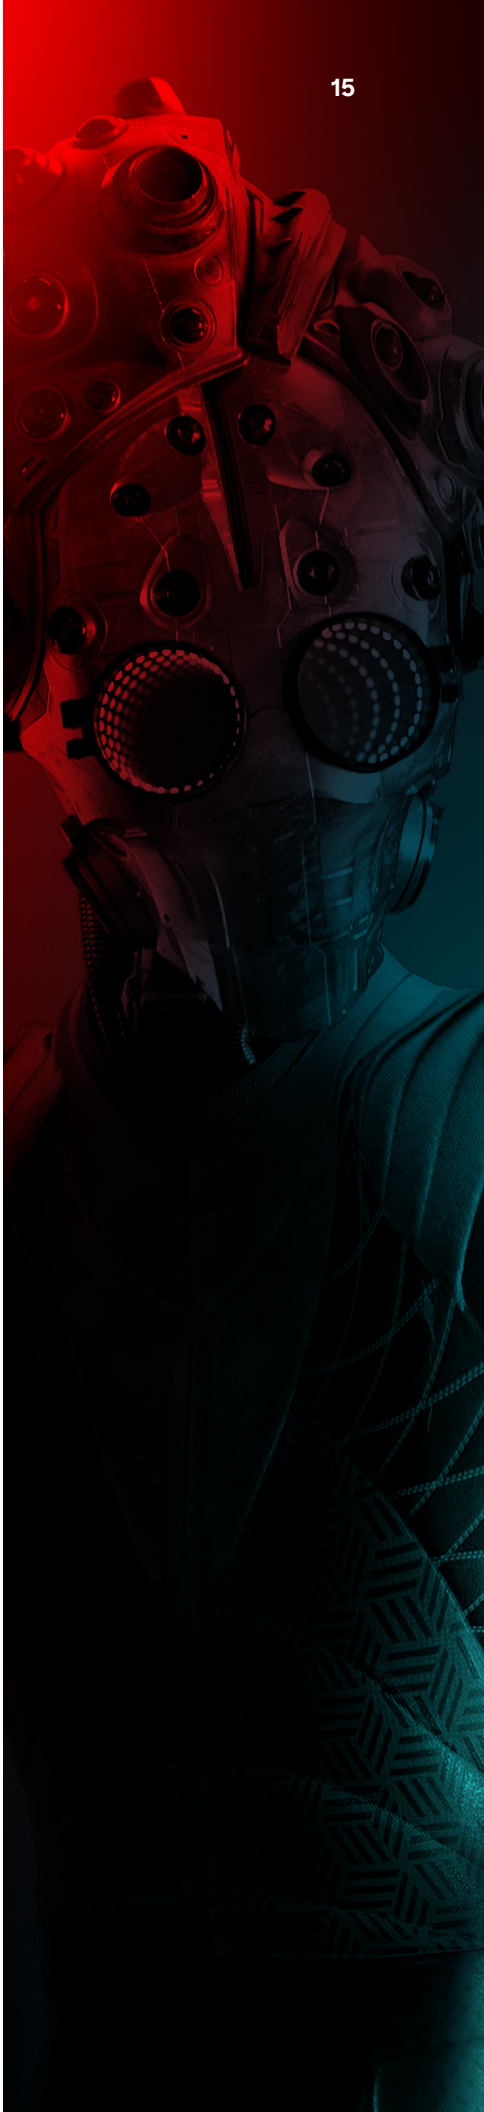

Prevention Playbook

Supply chain attacks often succeed because defenders lack visibility into the full software attack surface and because trust in third-party components is assumed rather than continuously validated. Effective prevention starts with tighter insight into software components and dependencies, then applies controls that limit what can be introduced into production, reduce credential and access risk, and ensure tooling and response plans are prepared for supply chain scenarios.

MAINTAIN ASSET VISIBILITY

Just as asset inventory is vital to securing traditional IT infrastructure, it is also vital for software components. Building a comprehensive asset inventory alongside a detailed software bill of materials (SBOM), using both manual and automated recursive dependency scanning, extends visibility into software assets, including nested dependencies that may otherwise remain hidden.

IMPLEMENT ARTIFACT CONTROL AND VULNERABILITY MANAGEMENT

Once an organization has a clear view of its software attack surface, it can begin to control the use of third-party dependencies. Organizations should implement a software composition analysis (SCA) tool as part of their code deployment pipeline to identify vulnerable dependencies as an integrated step in the build process.

While SCA tooling checks for known vulnerable dependencies, organizations can exert further control by leveraging a centralized artifact repository from which developers source external dependencies. Caching external components allows an organization to enforce policies requiring the use of approved libraries.

SECURE IDENTITY AND ACCESS

Limiting the impact of supply chain compromises requires organizations to move away from long-lived and hardcoded secrets that attackers can leverage to move across an environment. OpenID Connect (OIDC) provides secure identity verification using short-lived tokens to reduce the risk of credential theft. Organizations can also leverage PAM solutions to provision just-in-time access for third-party access to the environment and limit the availability of external access.

REVIEW TOOLING CAPABILITIES AND MATURITY

With some prevention solutions existing as SaaS applications, ensuring proper tool configuration is a critical step in implementation. Falcon Shield helps organizations harden SaaS applications through built-in and custom security checks that detect misconfigurations.

Organizations should also regularly evaluate their risk posture against current security controls to guide roadmap development and future tooling acquisition. They can leverage the CrowdStrike Cybersecurity Maturity Assessment to evaluate their security posture and identify gaps across several domains, including third-party risk management and application services.

PLAN FOR RESPONSE EFFORTS

Part of prevention is being prepared to respond to supply chain incidents, which often require specialized response procedures. Organizations should develop incident response plans that explicitly address supply chain compromise scenarios. These plans should include clear procedures to identify and isolate affected dependencies to prevent the spread of compromise. CrowdStrike Tabletop Exercises can be tailored to an organization's needs, allowing teams to prepare specifically for a supply chain compromise.

In the event a third-party dependency is compromised, organizations should determine the extent of impact within their environment. CrowdStrike Compromise Assessments work to uncover activity and threats and can target specific threats, vulnerabilities, and exploits.

Finally, CrowdStrike Response Readiness Exercises provide workshops that cover a variety of topics, including software supply chain and vulnerability exploitation readiness. These sessions connect security teams with CrowdStrike experts to strengthen defensive planning and apply insights from frontline experience.

How CrowdStrike Helps in a Crisis

Supply chain incidents require a rapid response, but also a highly thorough one, given the likelihood of nation-state involvement and the associated risk of prolonged dwell time, far-reaching access, and durable persistence. CrowdStrike responders establish an initial foundation for the investigation by collecting high-fidelity artifacts across domains, combining Falcon Identity Threat Protection for identity sources, Falcon Shield for SaaS applications, Falcon Next-Gen SIEM for historical and streaming datasets from existing sources, and traditional endpoint triage using CrowdStrike Falcon® Insight XDR and the Falcon Forensic Collector. This cross-domain collection allows for a thorough understanding of the attack scope to inform containment and eradication.

As forensic data is collected, it is analyzed in parallel, and responders work with the victim organization and supporting security partners to develop a containment and eviction plan. When the intrusion involves a highly sophisticated adversary, including nation-state actors exploiting the supply chain, CrowdStrike responders can lead a coordinated remediation event designed to reduce the likelihood of persistence or recurring access, so the business can continue operations with greater confidence.

Adversary Objectives Shape Zero-Day Exploit Selection

Threat Overview

Attackers are not using zero-day exploits randomly. In 2025, adversaries deliberately chose zero-day vulnerabilities based on what they were trying to achieve, such as stealing data, gaining long-term access, or launching large-scale attacks. Nation-state threat actors focused on breaking into exposed edge systems like VPNs and mail servers to stay hidden and persistent, while cybercriminals prioritized vulnerabilities in common enterprise apps and operations systems that were easy to exploit at scale. Some adversaries repeatedly target the same products, maximizing specialized knowledge. As a result, zero-day attacks increased sharply and are being used more strategically, making them harder to defend against and more likely to lead to serious business impact.

Frontline Trends and Observations

Adversaries typically follow the path of least resistance, especially when it comes to resource- and cost-intensive zero-day exploits. Many will avoid "burning" a zero-day unless it is necessary due to time pressure or operational complexity. In other cases, adversaries conduct their own form of return-on-investment analysis, deciding an exploit is worth using if it can yield high-value information, meaningful leverage, or direct financial gain.

CrowdStrike responders observed multiple investigations where adversaries leveraged zero-days and n-days. Notably, nation-states are not the only threat actors using zero-days. eCrime adversaries such as GRACEFUL SPIDER leveraged a campaign involving a prominent zero-day exploit, CVE-2025-61882. The exploit enabled authentication bypass to execute malicious code, ultimately granting full remote control without requiring a password.

In other nation-state investigations, China-nexus adversaries leveraged multiple zero-day SQL injection chains to issue unauthenticated, API-driven SQL commands against victim databases. China-nexus adversaries also leveraged an n-day (CVE-2025-5777) that enabled session token theft and MFA bypass on remote access gateways and application delivery infrastructure.

Prevention Playbook

While the inherent nature of zero-day threats implies a sense of inevitability, there are actions organizations can take to harden systems and processes and limit the potential impact of novel exploits. As adversaries increase investment in cultivating new vulnerabilities and exploits, organizations should review foundational security capabilities and controls to determine whether they are equipped to meet the challenge.

DETECT AND PREVENT MALICIOUS BEHAVIOR

For targeted intrusions where adversaries leverage novel zero-day exploits, behavior-based prevention is more reliable than signature-based controls because it is tuned to attack classes and post-exploitation activity. Focusing on common attacker motives and tactics, rather than specific tools, malware, or exploit code, allows organizations to detect and potentially block both known and unknown threats. Understanding industry- or company-specific targeting trends and the associated tactics, techniques, and procedures (TTPs) can help further tune behavior-based detections. In the best of cases, well-tuned behavior detections can enable proactive blocking of select high-risk processes, such as likely malicious file system access. Security teams should consider a phased approach to deploying these blocks to balance prevention with business operational needs.

REDUCE TARGETS OF OPPORTUNITY INSIDE THE PERIMETER

In addition to behavioral prevention, organizations should examine their environment as an adversary would and reduce opportunities for lateral movement, privilege escalation, and data exfiltration. Network segmentation, data loss prevention, and adherence to least privilege are critical in limiting what an adversary can do after initial exploitation.

DEVELOP RAPID PATCHING PROTOCOLS

For zero-day exploits learned through public disclosure, these same principles apply. However, organizations often have more opportunity to remediate or mitigate the underlying vulnerability before opportunistic attackers adopt the exploit. To do this effectively, organizations need responsive threat intelligence and vulnerability management programs.

Threat intelligence teams should collect key details about the vulnerability, validate community reporting, and assess mitigations, particularly when vendors do not immediately provide fixes. Membership in intelligence sharing organizations, such as information sharing and analysis centers (ISACs) and information sharing and analysis organizations (ISAOs), can support peer collaboration and the exchange of indicators and mitigations through trusted channels.

Once the intelligence is analyzed, organizations should be able to activate emergency vulnerability management procedures that go beyond patching. This includes configuration changes and containment measures. In many instances, it may be necessary to disable vulnerable functionality or remove network access for affected devices.

Teams responsible for emergency vulnerability management should have the authority to implement these measures, or rapid access to those who do. They also need a clear understanding of business impact. By mapping system dependencies, particularly for critical applications and services, organizations can make informed decisions and apply containment measures with greater confidence.

Effectively preparing for and responding to the increasing use of zero-day exploits depends on foundational security capabilities and processes that are appropriately tuned and well managed. Security teams need a clear understanding of their environment, operations, and threats to reduce disruption following the next major zero-day disclosure.

How CrowdStrike Helps in a Crisis

When a suspected zero-day or n-day exploit is involved, CrowdStrike responders begin with forensic timeline analysis to pinpoint the earliest evidence of compromise. In many investigations, this timeline narrows initial access to a specific endpoint, edge device, or application. From there, responders assess whether the activity aligns with a known exploit pattern or a recently identified vulnerability leveraged by the adversary through available logging and close collaboration with the victim organization to recreate the exploit scenario.

Determining how initial access occurred depends on reliable telemetry and log coverage in the right places. Falcon Shield supports investigations involving SaaS applications, while Falcon Next-Gen SIEM enables rapid analysis of network edge device logs. If the suspected source is offline, responders use legacy data collection as a fallback to preserve investigative continuity and reconstruct the initial access path, which is often where exploitation occurred.

As the initial point becomes clear, the focus shifts to preventing re-entry and accelerating containment. CrowdStrike responders leverage Falcon Exposure Management to identify unpatched vulnerabilities and misconfigurations across the device fleet, including conditions an adversary could reuse to regain access. Falcon Exposure Management also surfaces attack paths, showing which high-value targets are reachable from the initial breach point. This helps responders prioritize emergency patching or isolation.

Identifying zero-day or n-day exploit use provides high confidence in how the intrusion succeeded and what must be fixed to stop it from recurring. That clarity improves containment and eradication outcomes and supports a safer return to normal business operations once the vulnerability is patched or mitigated.

Adversaries Subvert Trust in Cloud Platforms and Services

Threat Overview

Adversaries are increasingly gaining access by abusing trust in cloud and SaaS platforms rather than exploiting software flaws directly. In 2025, attackers frequently authenticated using legitimate credentials and trusted access paths, then moved through cloud environments by abusing identities, authentication flows, and partner relationships. Hybrid identity abuse became a primary enabler of this tradecraft, allowing adversaries to pivot between on-premises and cloud environments. Because this activity relies on valid access and native cloud services, it often blends into normal user behavior and is difficult to detect.

Frontline Trends and Observations

CrowdStrike responders continue to observe adversaries subverting trust in cloud environments by exploiting legitimate identities. These identities are commonly obtained through infostealers, social engineering, or access brokers, then used to operate as authorized users and blend into normal activity.

This risk is compounded in organizations with hybrid identity setups rather than cloud-only accounts. In these environments, an adversary who compromises an identity in the cloud control plane may also be able to access on-premises directory services, expanding both reach and impact. Responders have also observed adversaries exploit trusted third-party relationships, including compromises of managed service providers (MSPs) and non-human service principals. These access paths can grant persistent, high-level tenant access that appears fully authorized, and depending on the identity architecture, may also enable traversal into on-premises systems.

CrowdStrike responders are also increasingly seeing attackers take advantage of cloud control plane complexity. Rather than relying on a sophisticated software exploit, adversaries weaponize common misconfigurations and overly privileged roles. In effect, they exploit the fact that the cloud trusts what has been configured. This includes abuse of hardcoded secrets in code, long-lived access keys, and overly permissive or implicitly trusted authentication conditions. Together, these weaknesses enable persistence, extend dwell time, and support follow-on objectives across cloud platforms.

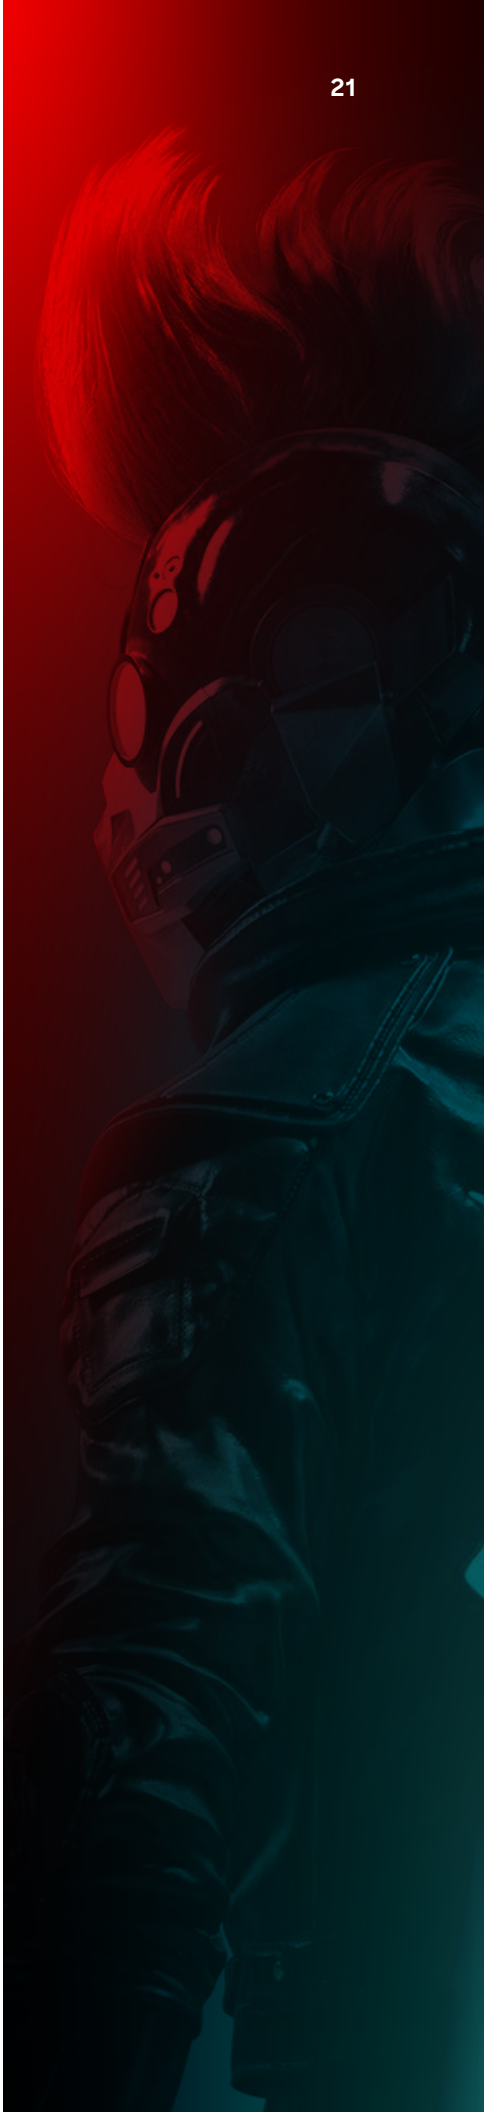

Prevention Playbook

Adversaries subvert trust in cloud platforms by abusing credentials, permissions, and configurations that were designed to enable speed and scale. Effective prevention reduces credential exposure, limits what compromised identities can do, and increases visibility into the control plane changes adversaries rely on for persistence.

PROTECT CLOUD CREDENTIALS

Countering the abuse of trust in cloud platforms starts with limiting the exposure of identities and credentials. Organizations should prohibit users, including third-party providers, and applications from storing access keys, API keys, and other cloud credentials in insecure locations such as public code repositories, infrastructure-as-code templates, and local devices. Instead, provide users and applications with access to a secrets manager, key vault, or PAM solution to store and automatically rotate sensitive cloud credentials. Security teams can use tools such as Falcon Cloud Security to conduct periodic scans that identify cloud credentials stored in unapproved locations.

PRACTICE GOOD ACCOUNT HYGIENE

A defense-in-depth approach to mitigating cloud trust abuse should enforce least privilege to minimize what a threat actor can do with a compromised identity or credential. This requires configuring fine-grained permissions across cloud identities, including users, roles, service principals, groups, and related constructs. Organizations should have users assume roles with task-specific permissions rather than assigning standing privileges to named user accounts. Just-in-time access features can also enforce approval workflows and time-based restrictions for highly sensitive cloud administrator activities. To prevent lateral movement between cloud and on-premises environments, organizations should use cloud-native accounts for cloud administration, not accounts synchronized with an on-premises identity store; for example, organizations can use native Entra ID accounts rather than hybrid accounts synced to on-premises Active Directory.

MONITOR FOR CREDENTIAL ABUSE

To mitigate abuse of non-human identities such as service principals, security teams should establish a baseline inventory of non-human identities, including the permissions and secrets assigned to them. With this baseline, implement rules that detect when unexpected non-human identities are added to a cloud account and when new permissions or secrets are provisioned to non-human identities.

CrowdStrike has observed incidents where adversaries modified federation settings to establish persistent access to cloud environments. For example, adversaries such as SCATTERED SPIDER have used the AADInternals tool to modify a victim's domain federation settings to trust a fictitious identity provider using a signing certificate the adversary controls. Organizations should monitor directory administration and federation changes in Entra ID and other identity platforms to detect attempts to establish persistence.

> “ Countering the abuse of trust in cloud platforms starts with limiting the exposure of identities and credentials.

MAINTAIN SECURE CONFIGURATIONS

Adversaries frequently target exposed and misconfigured cloud resources such as publicly accessible storage containers and unsecured virtual machines. Organizations should extend traditional security practices to cloud resources rather than assuming cloud environments have built-in protection. In many incidents, organizations inadvertently exposed back-end servers, key vaults, and other sensitive resources, allowing adversaries to directly access them. Organizations should apply layered network controls by separating public and private resources and carefully controlling access between them. Defenders should also restrict outbound access to help limit data exfiltration, block command-and-control traffic, and prevent adversaries from downloading tools onto compromised endpoints.

At the resource level, organizations often struggle to determine and apply the appropriate controls to secure each resource, such as enforcing encryption at rest, controlling access to storage resources using identities instead of shared keys, and enabling versioning or immutability. Cloud security posture management (CSPM) tools such as Falcon Cloud Security can help establish baselines and alert on risky configurations, potential exploitation, and suspicious modifications, such as publicly exposing a virtual machine or storage bucket. Regularly reviewing and addressing CSPM findings can help minimize exploitable misconfigurations and reduce the likelihood and impact of an incident. Periodic third-party assessments, such as CrowdStrike Cloud Security Assessments or CrowdStrike Cloud Red Team Blue Team Exercises, can also help organizations uncover, understand, and prioritize the vulnerabilities and misconfigurations that cloud-focused adversaries are most likely to exploit.

How CrowdStrike Helps in a Crisis

When adversaries subvert trust in cloud environments, CrowdStrike responders prioritize two objectives in parallel: stopping identity-based spread and removing the control-plane conditions that enable persistence.

First, responders focus on identity, because legitimate credentials often determine how quickly an intrusion can pivot between cloud and on-premises environments. Falcon Identity Threat Protection monitors authentications across hybrid and cloud-only setups and helps detect the identity hopping common in hybrid environments. It also surfaces anomalies such as a cloud-based service principal suddenly accessing on-premises resources. When credentials appear stolen, Falcon Identity Threat Protection can enforce MFA or block sessions to disrupt active access and limit lateral movement between cloud and local domains.

At the same time, responders leverage Falcon Cloud Security to gain real-time visibility into attacks on workloads that typically lead to subsequent persistence at the identity layer. Responders also use Falcon Cloud Security, alongside native control plane access, to detect and mitigate cloud control plane risk and help achieve containment.

By combining rapid triage with deeper analysis, responders can identify persistence mechanisms such as over-privileged accounts, hardcoded secrets, and long-lived access keys. Pairing identity monitoring with active cloud threat detection and posture insights helps responders contain the active intrusion, then close the implicit-trust paths that enabled it, reducing the likelihood of re-entry after containment is complete.

Outcome-Driven MDR for Modern Attacks

Across the threats covered in this guide, one pattern holds: Adversaries move fast, operate across domains, and exploit the gaps between teams and tools. Even strong prevention can be tested at the worst possible time, whether after-hours, during a holiday, or when malicious activity blends into routine noise. That is why many organizations pair proactive hardening with continuous detection and response, providing unified visibility across endpoint, identity, cloud, and third-party data, then acting quickly when preventative controls fall short.

CrowdStrike Falcon® Complete managed detection and response (MDR) is built for that moment, delivering expert-led, AI-powered detection and response across the domains where modern attacks unfold. Whether adversaries exploit identity gaps, target perimeter infrastructure, or leverage emerging techniques such as AI-enabled social engineering, CrowdStrike's expert analysts operate 24/7/365 within globally distributed fire teams to intervene when prevention alone is not enough. Through real-time detection, agentic workflows, full-cycle remediation, and expert oversight, they rapidly determine scope and severity, contain threats, evict adversaries, and prevent intrusion from escalating into business impact.

Behind every decisive response is a detection strategy forged from real-world adversary engagement. The Falcon Complete detection engineering team continuously refines detection logic based on real-world adversary behavior observed in the field, unifying cross-domain telemetry and enriching signals with threat intelligence to accelerate confident, precise response actions.

By owning the full lifecycle from detection through remediation, Falcon Complete reduces operational burden, strengthens resilience, and enables organizations to stay focused on executing the strategic prevention playbooks outlined throughout this guide.

STOPPING RANSOMWARE ACROSS DOMAINS

Ransomware operators increasingly exploit cross-domain paths, including pivots from compromised SSO credentials into VMware ESXi, and abuse of hybrid identity solutions through stolen AD Sync Agent configurations. Falcon Complete is designed to detect these pivots early and disrupt lateral movement before encryption.

Falcon Next-Gen SIEM correlates identity anomalies with endpoint activity and suspicious cloud access to surface high-fidelity alerts. Analysts then leverage this technology to execute coordinated response actions through CrowdStrike Falcon® Fusion SOAR workflows and Falcon Real Time Response, including isolating compromised endpoints and applying cloud and identity countermeasures such as password resets, session revocation, and stronger MFA enforcement and credential restrictions.

DETECTING NATION-STATE PERSISTENCE

China-nexus adversaries often begin at the perimeter, targeting web servers, VPNs, and firewalls, then pivot into virtual infrastructure to deploy webshells and establish ghost VMs for long-term persistence. Falcon Complete ingests telemetry from edge devices and cloud infrastructure that often lack endpoint visibility, and applies detection rules enriched with CrowdStrike threat intelligence.

By correlating signals across endpoint, identity, and cloud, these detections surface rapid exploitation-to-virtualization patterns and help accelerate investigation when adversaries steal credentials, abuse cloud service principals for persistence, or establish command-and-control through compromised identities. In parallel, Falcon Adversary OverWatch continuously hunts for indications of adversary activity across customer data sets.

UNIFIED DEFENSE AGAINST ADVANCED THREATS

Modern threats evolve faster than point defenses. Cloud and SaaS abuse, supply chain compromise, and zero-day exploitation can slip past traditional controls by operating through trusted paths. Falcon Complete addresses these threats with cross-domain visibility through Falcon Next-Gen SIEM and the ability to take decisive action.

When a zero-day achieves initial access, when adversaries pivot through cloud infrastructure with stolen credentials, or when trusted relationships are abused through a supply chain compromise, agentic workflows support rapid triage and response actions. This reduces investigation and response load on internal teams and helps them stay focused on the hardening initiatives outlined in the prevention playbooks.

Conclusion

Modern adversaries move quickly, operate across domains, and exploit trust paths that often resemble normal business activity. Whether an intrusion begins with identity compromise, a perimeter weakness, a trusted integration, or a newly disclosed vulnerability, the result is the same when gaps between systems go unaddressed.

The prevention playbooks are designed to translate these frontline lessons into measurable progress. Use them to prioritize the controls that reliably disrupt attacker momentum: Reduce unnecessary privilege, strengthen access and trust boundaries, harden systems that provide outsized leverage, and regularly test detection and response against realistic attack paths. The objective is not perfection. It is disciplined, repeatable hardening that makes the next intrusion slower to execute, easier to detect, and faster to contain.

If you need support prioritizing these actions, validating them in your environment, or responding to an incident, CrowdStrike is available 24/7/365 to help.

About CrowdStrike

CrowdStrike (Nasdaq: CRWD), a global cybersecurity leader, has redefined modern security with the world's most advanced cloud-native platform for protecting critical areas of enterprise risk — endpoints and cloud workloads, identity and data.

Powered by the CrowdStrike Security Cloud and world-class AI, the CrowdStrike Falcon® platform leverages real-time indicators of attack, threat intelligence, evolving adversary tradecraft and enriched telemetry from across the enterprise to deliver hyper-accurate detections, automated protection and remediation, elite threat hunting and prioritized observability of vulnerabilities.

Purpose-built in the cloud with a single lightweight-agent architecture, the Falcon platform delivers rapid and scalable deployment, superior protection and performance, reduced complexity and immediate time-to-value.

CrowdStrike: We stop breaches.

Learn more: www.crowdstrike.com

Follow us: Blog | X | LinkedIn | Facebook | Instagram | YouTube

Start a free trial today: www.crowdstrike.com/free-trial-guide