WHITE PAPER

Prevention & Detection

AKIRA SonicWall Ransomware Campaign

Table of Contents

04 — Introduction

05 — AKIRA SonicWall Campaign: Overview

- 5 — Who is AKIRA?
- 5 — What did they do?
- 7 — SonicWall Response

8 — Campaign Insights and Takeaways

- 8 — High Level Precautions
  - 8 — High Level Precaution #1: Backups
  - 8 — High Level Precaution #2: Segmentation
  - 10 — High Level Precaution #3: Patch Management and Response Protocols

11 — Technical Details

14 — Conclusion

15 — Glossary

16 — Appendix A

AUTHORS

John Fromholtz

Director of Intelligence and Development

John curates BlueVoyant's ransomware intelligence and specimen collection. He has extensive knowledge of programming, code analysis, and incident response. John holds a MS in Cybersecurity from Stevens Institute of Technology and previously managed the security operations center for major financial institutions.

Izz Sacks

Threat Hunt Team Lead

Izz runs the BlueVoyant Threat Hunt team which works to proactively identify threats based on intelligence from the BlueVoyant Threat Intelligence and DFIR Teams. Izz holds a BTech of Information Technology from Kwantlen Polytechnic University.

The DFIR Team

Consisting of the Digital Forensics and Hunt Operations personnel

BlueVoyant's Digital Forensics and Incident Response team combines decades of experience, meticulously crafted techniques, and the best tools available to provide top tier service where it's needed, as soon as it's needed.

Introduction

We hope that the information, insights, and examples explored here prove useful in justifying your organization's existing controls and protocols, inspiring new projects, or both.

Starting in summer 2025 the ransomware group AKIRA initiated a spree campaign targeting SonicWall firewalls, while the overall campaign demonstrates how a zero-day vulnerability is not strictly necessary to effect substantial impact [1]. When the BlueVoyant DFIR team responded to matters stemming from this campaign, several patterns became readily apparent, both in AKIRA's typical attack chain and among the decisions made by their victims that mitigated or exacerbated the impact.

Whether insufficient segmentation or vulnerability management allowed for a major incident or well-planned backups and disaster recovery plans allowed for a swift recovery, these observations are applicable beyond this ransomware campaign. In this paper, we will examine AKIRA's SonicWall campaign through a mix of open-source intelligence and the incidents responded to by the BlueVoyant DFIR team with an eye towards lessons and ideas for those who wish to improve their security posture.

After this we will pivot to look closer at the security controls and organizational policies that can help harden an organization against malicious attention, AKIRA or otherwise. This section will review these categories at a high level before diving deeper into the technical aspects and implications for those looking for further guidance on implementing these concepts in their organization.

Readers should note that implementing the suggestions and controls discussed in this paper are likely to require some customization to best fit an organization's needs and circumstances. While the resources needed to fully implement everything discussed may be substantial, even a partial implementation can improve an organization's security posture. Selecting, prioritizing, and implementing controls and policies inspired by this paper can be done efficiently with a well-planned roadmap and diligent documentation. Wherever possible, further reading has been footnoted, and a glossary of technical terms has been appended.

1 — Detailed in the main body of this paper and in the Q3 2025 and Q4 2025 editions of Quarterly Intel Highlights

AKIRA SonicWall Campaign: Overview

Who is AKIRA?

AKIRA is a long running ransomware-as-a-service[2] group whose attacks are characterized by extremely rapid tempos[3] that lack the improvised nature of more typical "hit and run"[4] style ransomware attacks. Based on the attacks observed by BlueVoyant's DFIR team, AKIRA achieves much of this speed by templating their attacks, allowing for streamlined and efficient kill chains at the cost of difficulty overcoming unexpected obstacles.

Though not observed in this campaign, AKIRA has demonstrated a willingness to experiment with novel methods of data theft[5] as part of an overall double extortion approach[6]. AKIRA is known to favor American targets in the Business Services, Manufacturing, and Technology sectors[7] with attacks often taking less than a day after initial access. This is despite a complete kill chain resulting in data theft and encryption[8].

What did they do?

AKIRA is known for its unusually fast double extortion attacks and the cases handled by the BlueVoyant DFIR Team have been noticeably similar in their methods for initial access, lateral movement, and data exfiltration. These similarities in approach along with an observed difficulty in bypassing unexpected defenses or other obstacles lead BlueVoyant Researchers to believe that AKIRA affiliates rely on pre-planned and templated attacks to achieve their results. Despite this, the apparent skill displayed by AKIRA affiliates and encryptor deployment methods have varied substantially, this is theorized to have been caused by several rival ransomware groups shuttering operations and providing surviving groups such as AKIRA with a surplus of new affiliates[9].

THE ATTACKS

The sudden ramp up of AKIRA attacks in late July, an unclear initial access means, and the observation that SonicWall SSLVPN products were a common denominator initially fit the pattern of a new zero-day attack[10]. Much of the early guidance was based on this belief and recommendations to disable the SSLVPN service until the attacks are better understood[11] are best practice when a poorly understood vulnerability is being actively abused.

Deeper examination would disprove this initial theory[12] and an understanding of the initial attack began to take shape. Use of CVE-2024-40766 or publicly accessible administrative portals allowed AKIRA to log in with default or stolen credentials. Because the default LDAP group applied to all local accounts in the firewall, these accounts often had permissions beyond what the maintainers of these devices expected. Once the firewall was accessed, AKIRA could create or modify accounts as they wished to grant themselves a VPN account.

It was observed in several of the cases where the firewalls were patched, but not before AKIRA or an initial access broker[13] could establish themselves that a common pattern was followed for privilege escalation. When logs were available, the BlueVoyant DFIR team noted that initial access was gained through a default account local to the SonicWall that the victim did not expect to have access to the firewall, had been unaware of, or was otherwise forgotten. With this initial foothold the intruder would seek to pivot to an account with sufficient privileges, often the LDAP service account.

2 — https://www.cloudflare.com/learning/security/ransomware/ransomware-as-a-service/
3 — https://blog.qualys.com/vulnerabilities-threat-research/2024/10/02/threat-brief-understanding-akira-ransomware
4 — https://www.imperva.com/blog/know-your-enemy-the-four-types-of-cyber-attackers-trying-to-breach-your-security-today
5 — In some cases, AKIRA has been observed using https://github.com/peak/s5cmd to effect extremely rapid exfil
6 — https://gca.isa.org/blog/double-extortion-ransomware-what-it-is-and-how-to-respond
7 — https://www.ransomware.live/groupstats/akira
8 — https://www.darkreading.com/endpoint-security/akira-ransomware-lightning-fast-data-exfiltration-2-hours
9 — https://www.theregister.com/2026/01/08/ransomware_2025_emsisoft/
10 — https://www.securityweek.com/sonicwall-hunts-for-zero-day-amid-surge-in-firewall-exploitation/
11 — https://www.bleepingcomputer.com/news/security/sonicwall-urges-admins-to-disable-sslvpn-amid-rising-attacks/
12 — https://thecyberexpress.com/sonicwall-ssl-vpn-flaw-cve-2024-40766/
13 — https://www.packetlabs.net/posts/what-is-an-initial-access-broker/

INITIAL ACCESS, RECONNAISSANCE, AND LATERAL MOVEMENT

While not a zero-day, the collection of vulnerabilities leveraged in this campaign yielded comparable results and allowed AKIRA to establish access and persistence to their victim's network. By creating or modifying a VPN enabled account for themselves, AKIRA affiliates could directly connect their own device to their victim's network and begin in depth reconnaissance while avoiding significant activity on any monitored endpoints. Once reconnoitered, AKIRA affiliates frequently focused singularly on file servers and critical infrastructure to affect data theft and prepare to launch their encryptor.

Observed data exfiltration during the SonicWall campaign was typically limited in size and scope. While AKIRA affiliates have been observed using advanced data exfiltration techniques, the attacks in this campaign instead relied on more common tools[14] and showed little indication that file selection was customized for specific victims.

Encryption was frequently restricted to just hypervisor servers and adjacent critical infrastructure. This is common among swifter ransomware attacks and is very common for AKIRA attacks. By detonating an encryptor on an ESXI or Hyper-V server, AKIRA can encrypt the drives of all hosted virtual machines and achieve exceptional impact in a short time frame. In a small subset of cases, a script was used to run the encryptor from a central location against system devices via smb shares[15]. While requiring more time and skill to prepare, this style of encryption offers a far greater blast radius while limiting the easily observed and blocked activity of the encryptor to a single device that has been prepared for this purpose.

CVE-2024-40766 or publicly accessible administrative portals allowed AKIRA to log in with default or stolen credentials. No zero-days used, but just as effective.

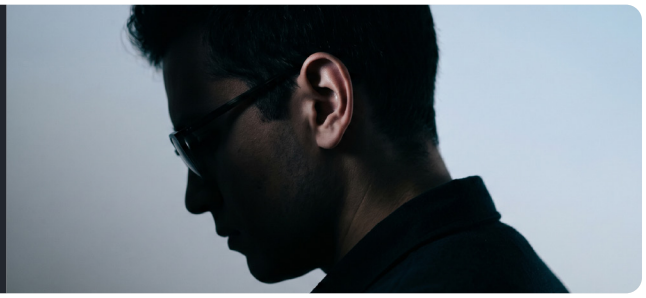

THE VULNERABILITIES

Despite its efficacy, the initial access attack that facilitated this campaign leverages previously discovered vulnerabilities and design flaws in a novel way and do not fit the definition of a zero-day attack[16]. CVE-2024-40766[17] is the oldest vulnerability involved and consists of flaws in the management access and SSLVPN components of SonicWall firewalls prior to version 7.0.1 that could allow access regardless of configuration. At the time of its announcement this vulnerability AKIRA being among the first[18] ransomware groups to use it and a patch was released by SonicWall shortly after[19]. As with many vulnerabilities that allow unauthorized access or remote code execution, patching the vulnerability does nothing to remove any accounts that may have been compromised or created already, nor would it undo any configuration changes that would permit continued access. As a result, BlueVoyant's DFIR Team continues to see a common pattern of incidents among clients who may have been swift to apply the patch but failed to completely remediate the vulnerability by auditing accounts and settings.

In this most recent SonicWall campaign, CVE-2024-40766 would be used in tandem with the leveraging of a flawed design in ldap group provisioning that could cause certain accounts to have much broader permissions on the firewall than intended[20]. These weaknesses in combination with default or leaked credentials set the stage for AKIRA's campaign.

14 — https://www.security.com/threat-intelligence/ransomware-data-exfiltration
15 — Appendix A.
16 — https://www.ibm.com/think/topics/zero-day
17 — https://nvd.nist.gov/vuln/detail/cve-2024-40766
18 — https://www.bleepingcomputer.com/news/security/critical-sonicwall-sslvpn-bug-exploited-in-ransomware-attacks/
19 — https://psirt.global.sonicwall.com/vuln-detail/SNWLID-2024-0015
20 — https://www.sonicwall.com/support/knowledge-base/ldap-configuration-sslvpn-default-user-groups-security-risk/250813061722917

PERSISTENCE, EXFILTRATION, AND ENCRYPTION

In line with known AKIRA attacks[21], the cases handled by the BlueVoyant DFIR team were typically swift and focused in their tempo and scope. The use of valid VPN credentials allowed AKIRA to directly connect their own device to a victim's network even after patches closed the vulnerabilities involved with initial access. With this persistent access, AKIRA affiliates could then openly use tools and techniques without concern that they would be detected by endpoint monitoring solutions to identify critical systems such as domain controllers and gather the information needed to continue their attack. This technique often meant that AKIRA's presence in a network was only detected by endpoint monitoring solutions when the affiliate connected to a system to steal data or begin encryption.

While BlueVoyant has observed advanced data exfiltration techniques in other AKIRA cases, this campaign seemed to favor the use of filezilla, rclone, or winscp to upload finance, accounting, HR, and similar folders. Based on the consistent nature of stolen files observed, BlueVoyant Researchers speculate file and folder selection to have been keyword driven with little effort given to actively selecting files based on the victim's business model or sector. Whether this indicates a preference of the AKIRA affiliates conducting the actual attacks, or a choice by the AKIRA group's leadership to avoid using their best tools in a sure to be highly studied campaign is unknown.

THE AFTERSHOCKS

Though most observed attacks occurred in latter half of 2025, BlueVoyant Investigators continue to respond to a steady trickle of related cases well into 2026.

In these 'aftershock' attacks initial access was likely sold by an initial access broker (IAB)[22] or an affiliate[23] seeded initial access to use. In these cases, a commonly observed pattern was that a victim had vulnerable firewalls during the height of AKIRA's campaign, these firewalls were patched, however they were not audited for new or modified user accounts, unusual authentication patterns, or modified access rules. This combined with the short lifespan of local log storage frequently left the precise details of initial access unrecoverable.

In the cases where relevant logs could be recovered, the BlueVoyant DFIR team observed that the local account used to facilitate LDAP connectivity was a popular target, likely due to its expected permissions and ubiquity simplifying the threat actor's efforts to locate and obtain administrative credentials across their targets. From here the attacks reconverged onto the attack patterns typically seen throughout the SonicWall campaign[24], though often at a noticeably slower tempo. Based on the cases where relevant logs were available, times between initial access and encryption were 1-3 weeks, with most of the observable activity happening in the days prior to encryption. This wide range and the continuance of similar attacks[25] further supports multiple sources of initial access.

SonicWall Response

SonicWall has since released additional patches addressing management of local and default user accounts and ldap groups, though the need to audit for maliciously created or modified accounts remains good practice after updating. SonicWall devices running patch versions 7.3[26] or lower since the summer of 2025 should be considered potentially compromised and in need of appropriate remediation.

21 — https://www.darkreading.com/endpoint-security/akira-ransomware-lightning-fast-data-exfiltration-2-hours
22 — https://www.cisecurity.org/insights/blog/initial-access-brokers-how-theyre-changing-cybercrime
23 — https://www.halcyon.ai/faqs/what-is-a-ransomware-affiliate
24 — https://www.darktrace.com/blog/inside-akiras-sonicwall-campaign-darktraces-detection-and-response
25 — https://reliaquest.com/blog/threat-spotlight-vpn-exploitation-when-patched-doesnt-mean-protected
26 — https://software.sonicwall.com/Firmware/Documentation/232-006386-00_RevA_SonicOS_7.3_ReleaseNotes.pdf

Campaign Insights and Takeaways

Despite the speed and scale of this campaign, the victims who found themselves having to consider a ransom payment to restore their business were in the minority among the cases handled by the DFIR team. Some common infrastructure and policy choices were observed among the victims who fared better than. Conversely, the DFIR team also noted a series of common weaknesses and architectural flaws that resulted AKIRA traversing their victim's network faster and reaching further than they otherwise could have. These observations have been collected into 3 high level precautions for ease of reading.

Building on this, we will then dive into a broader and more technical discussion regarding the implementations and implications that surround these choices. Wherever possible, observations from BlueVoyant's Managed Detection and Response have been included to provide the reader with a broad set of options to inspire further refinement of their security posture.

Organizations with isolated or immutable backups often recovered without ransomware payment consideration.

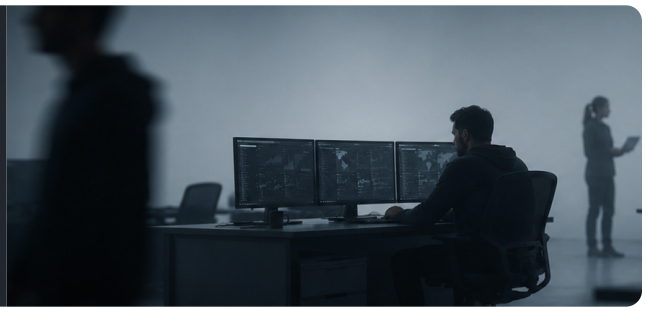

High Level Precautions

High-Level Precaution #1 – Backups

It is no secret that ransomware groups will make a point of seeking out and destroying backups prior to encryption[27]. The prevalence of immutable backup services[28] arose as a response to this and has proven effective in many of the incidents handled by the DFIR team. However immutable backups are not a perfect solution as they can be costly to provision, complex to implement and maintain, and slow to download for small and medium businesses.

Amongst the victims aided by the DFIR team who could restore their systems from backups, nearly all had immutable backups as a long-term backup storage solution to complement local backups in some capacity. Several took the additional step of hardening their local backup solution by deliberately not joining it to their primary domain. How the personnel responsible for these backup solutions managed access varied, but in these cases the threat actor was unable to use domain administrator credentials to access and destroy local backups[29].

Hardened backup solutions saved these organizations substantial time and cost as rebuilding from a recent checkpoint was chosen over considering the ransom demands or rebuilding from nothing. Not having to resort to the retrieval of offsite backups for much of the rebuilding efforts dramatically sped the recovery efforts and offered another layer of protection.

Care must be taken when restoring backups that any persistence mechanisms left by a threat actor are not also restored. When restoring data and files this risk can be mitigated by vetting the files for malicious signatures and suitable security monitoring on the new devices using this data. However safe restoration of operating systems typically requires the restore point predates initial access, which requires forensic examination to determine with any confidence.

27 — https://techwireasia.com/2025/09/google-security-chief-ransomware-now-targets-backups/
28 — https://www.commvault.com/explore/immutable-backup
29 — In one case, the client reported that they had found a device with the web UI of their backup solution open and locked after the domain administrator account had been attempted too many times.

High-Level Precaution #2 – Segmentation

Defining authentication and network segments are foundational to organizing and securing an organization's infrastructure. In AKIRA's SonicWall campaign, it was common for the threat actor to connect their own device to their victim's VPN and pivot to critical systems as soon as they gain suitable credentials. The DFIR team has observed that victims with well-planned and implemented segmentation schema were able to mitigate the damage caused by the encryption attack, either by limiting what systems the threat actor could quickly access or by limiting the utility of compromised credentials.

From these cases and similar examples[30] it is apparent that a lack of network or authentication segmentation was frequently a factor in the cybercriminal's favor. In many cases, the device that AKIRA connected to their victim's VPN was able to directly access domain controllers, database servers, and other critical infrastructure, using an administrative account found on the firewall. Appropriate barriers between perimeter, intermediate, and core regions of a network could have greatly complicated this attack chain and presented additional opportunities to detect and evict AKIRA before they completed their plans.

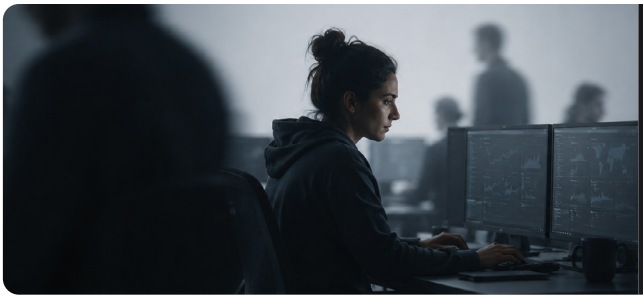

Victims with proper network and authentication segmentation limited AKIRA's blast radius.

NETWORK SEGMENTATION

While a well secured and maintained perimeter is vital to an organization's security[31], no single security control can be expected to mitigate all risks. Among the layers of defense[32] commonly placed behind the perimeter are a series of strategically placed firewalls, access rules, and subnets[33]. Defining the segments of a network according to department and/or purpose naturally lends itself to a segmentation plan that restricts access[34] and simplifies internal security monitoring by specifying accepted gateways between sections.

Segmenting according to system criticality requires additional planning but is commonly done to further protect critical infrastructure from malicious attention and monitoring entry points. This often goes hand in hand with identifying and isolating entry points into a network and into highly restricted and monitored segments[35]. Typically, these layers are a low criticality/highly exposed layer at the "outside" edges of a network (VPN and IOT subnets), a high criticality/low exposure "core" layer (domain controllers, backup systems, other critical infrastructure), and one or more intermediate layers (various workstation subnets or noncritical special purpose subnets). While most common infrastructure readily lends itself to one of these layers, some items may require careful consideration and a judgement call based on an organization's needs and risk tolerances.

Regardless of the exact layout of layer and departmental segments, traffic and authentication between segments should be restricted to only what is necessary and monitored closely for deviations from authorized patterns.

30 — https://zeronetworks.com/blog/mgm-resorts-microsegmentation-thwarts-ransomware
31 — https://fidelissecurity.com/threatgeek/network-security/perimeter-security-and-defense/
32 — https://www.fortinet.com/resources/cyberglossary/defense-in-depth
33 — https://www.cisa.gov/sites/default/files/2023-01/layering-network-security-segmentation_infographic_508_0.pdf
34 — i.e., the manufacturing/scada segment of the network should not be able to access devices in the finance and hr segments, and vice versa.
35 — i.e., devices connected to the vpn should not be able to directly log into a domain controller.

AUTHENTICATION SEGMENTATION

Separating everyday user accounts from administrator accounts[36] helps to limit the damage that can be caused by compromised credentials and is closely tied to the principle of least privilege[37]. Improvements to this policy can be made by further separating administrative accounts into purpose specific ones with narrow scopes of authority and privilege[38]. The narrower the intended use case for an account, the easier it becomes to detect potentially concerning behavior and the harder it is to abuse[39] compromised credentials. A deep dive into this matter is explored in the Technical Details section.

Every vulnerability allowing unauthorized access or remote code execution requires post-patch auditing for persistence artifacts, new accounts, and modified rules.

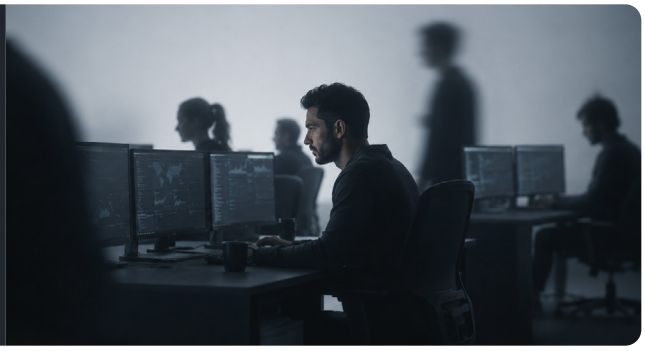

High-level Precaution #3 – Patch Management and Response Protocols

Pivoting to preventative measures, many of the attacks in AKIRA's campaign could have been prevented outright had the vulnerabilities enabling them been remediated entirely. An unfortunately common pattern was observed by BlueVoyant's DFIR team in which the victim had thought the application of a patch to be sufficient remediation, only to discover the error of this assumption weeks later. Had an audit of the affected firewalls been performed, the victims could have noticed that AKIRA had already granted themselves the credentials needed to continue their attack.

Whether it is a zero day or not, applying the patch is often not the final step when remediating a detected vulnerability. Vulnerabilities that permit remote code execution, unauthorized access, or other modifications to a system will require that these systems be audited for signs of compromise and either cleared for continued use or taken offline until there is reason to believe they are clean. In AKIRA's campaign, a sweep for unexpected or modified accounts and access rules would be sufficient to detect and remove any access that the ransomware group provisioned itself. In the wake of similar campaigns, the DFIR team has also seen webshells[40] or specific indicators in available logging[41] while responding to incidents or performing compromise assessments.

In some situations, a clear understanding of how to detect traces of an exploited vulnerability may not be immediately available. These circumstances require that a decision must be made whether to leave the device online and monitored until suitable checks can be made, or if the disruption of temporarily replacing the device with an unaffected equivalent is acceptable relative to the understood risks.

36 — i.e., John.Smith@acme.corp could be the user's account for accessing their workstation and email, but Admin.John.Smith@acme.corp would be what this person uses to perform administrative tasks such as configuring and maintaining company infrastructure.

37 — https://www.fortinet.com/resources/cyberglossary/principle-of-least-privilege

38 — To continue the example, Admin.John.Smith@acme.corp could be broken up as DCAdmin.John.Smith@acme.corp for accessing and administrating the domain controller and DBAdmin.John.Smith@acme.corp for the database product

39 — DCAdmin.John.Smith@acme.corp should only be able to access the domain controller(s), attempts to access other systems can be flagged as cause for concern.

40 — https://cloud.google.com/blog/topics/threat-intelligence/pst-want-shell-proxyshell-exploiting-microsoft-exchange-servers

41 — https://www.splunk.com/en_us/blog/security/detecting-cve-2020-0601-exploitation-attempts-with-wire-log-data.html

Technical Details

Expanding beyond our study of AKIRA's SonicWall campaign, further examples can be found to help an organization improve their security posture and mitigate risk exposure. The topics we will explore can be organized around four broad areas: patching policy, network hardening, auxiliary actions, and account management. A reoccurring thread through these sections is the use of a risk register and Business Impact Analysis (BIA) to prioritize resources where it can have the best. In a similar vein to the previously explored topics, fully implementing everything discussed here can require a significant investment of resources. However even partial implementation of these concepts and controls can substantially improve an organization's security posture and a well-planned roadmap can see these improvements integrated efficiently over a time span and cost that fits an organization's needs.

PATCHING POLICY

Organizations should update or create their internal patching policy to address important software updates in a structured, timely manner and aim to cover the full life cycle from the patch announcement to final implementation. An organization will need to find the right balance of speed while minimizing disruption to existing infrastructure, using a risk register and business impact analysis to inform and track patching targets.

A risk register can be used to determine which systems carry the highest potential business impact, establish a clear priority order for which software, services, and devices are patched. Maintaining a vulnerability feed tied to a priority and triage system can further support this and alert pertinent stakeholders when a relevant vulnerability is identified. Several free and paid services are available for vulnerability feeds including but not limited to the OWASP Dependency Check[42] project, Microsoft Defender Vulnerability Management[43] and Exploit Database[44].

Reducing time between patch announcement and full implementation limits the window for a threat actor to exploit a vulnerability. Automated vulnerability reporting and news feeds, such as those leveraging STIX/TAXII or RSS (footnote on best practices for implementation) can help accelerate this cycle. When deploying patches, organizations can leverage their risk register and/or BIA to determine prioritization order, determine their maximum tolerable downtime, and determine a path to apply patches quickly and efficiently. Fallback devices or a phased deployment strategy can be used to maintain business critical functions during patching windows. In virtual machine dependent environments, maintaining an up to date golden image of virtual machine baselines can enable the rapid deployment of fully patched instances when needed.

Before broad rollout, patches should be tested on a dedicated device to surface any operational pitfalls or compatibility issues introduced by the update. The 2024 CrowdStrike update which caused widespread device failure is a recent and highly visible example of an issue that could have been caught through appropriate patch testing and rollout strategies prior to full deployment[45]. This process may also reveal compatibility issues with internal or industry specific applications in a controlled and low impact environment which may require additional review and validation.

Organizations that implement flexible patch management systems with the ability to address updated vulnerabilities can benefit from a more thorough remediation system.

42 — https://owasp.org/www-project-dependency-check/
43 — https://learn.microsoft.com/en-us/defender-vulnerability-management/defender-vulnerability-management?tabs=preview-customers
44 — https://www.exploit-db.com/
45 — https://en.wikipedia.org/wiki/2024_CrowdStrike-related_IT_outages

NETWORK HARDENING

As companies grow, more devices, services, and software are frequently added to the network. Over time, these assets may change ownership, be deprioritized during role transition or lose their responsible owner entirely as personnel and infrastructure shift and grow effectively becoming 'ghost' assets. Though forgotten, these devices and services can be a significant entry point for a threat actor. An old firewall appliance that was deprecated but never decommissioned could serve as an initial access point for a threat and has been the root cause of many incidents handled by BlueVoyant's MSS and DFIR teams.

Diligent asset management program can establish a clear inventory of all devices in an environment: where each device lives, what function it serves, and who is responsible for it. This enables the visibility needed to identify and remediate assets that have fallen out of active management before they can be exploited.

Using the risk register as a guide, organizations can assess which devices are directly accessible from the internet or reachable from perimeter devices. With this information, understanding what ports and services need to be exposed (if any) can be paired with the risk register to craft and prioritize new security controls. Systems and services that need substantial external exposure can be placed in a demilitarized zone (DMZ) to mitigate the risks of their potential compromise.

Organizations should review their BIA to evaluate whether multiple critical systems share the same network segment and ensure a network containment strategy is in place should one of those systems become compromised. In many ways network segmentation can be compared with passive mitigation systems in the physical world, such as the passive fire protection systems common to any large structure. Placing barriers and establishing dedicated and controlled gateways between critical regions of an environment provides a wealth of opportunities to mitigate damage or preemptively detect an issue before it becomes an incident.

AUXILIARY ACTIONS

Continuous education within the security and patch management teams are an important complement to patching and hardening efforts. When reviewing vulnerability disclosures, key phrases such as "Unauthorized Access" or "Remote Code Execution (RCE)" are indicators that accounts or credentials on the affected system may have been at risk. When identifying these vulnerabilities, follow on actions may include cycling the credentials of all local accounts and any accounts that have accessed the affected appliance or software or examining relevant systems and s.

These risks are the core reason why patching alone is not always sufficient remediation. The AKIRA campaigns observed over the past year have demonstrated that threat actors left persistence mechanisms in place that cannot be resolved through patching alone. If a vulnerability was exploited prior to being patched that resulted in new account creation or credential exfiltration, the risk to the environment persists even after the patch is in place. A complete remediation process must account for post exploitation artifacts not just the underlying vulnerability.

ACCOUNT MANAGEMENT

Just as a risk register and BIA are used to prioritize patching and reduce device exposure, the same framework can be applied to account management. Mapping user accounts and service principles that have access to high-priority systems identified in the risk register with particular attention to those at or near the network perimeter can permit an organization the visibility needed to assess access to high value resources and infrastructure. Careful evaluation can then be given to whether access is necessary for the intended use. In a similar vein to monitoring for vulnerabilities, this provides the foundation for a defensible account management posture and ensures that privileged access is controlled and monitored for concerning events.

Organizations should maintain a clear and current inventory of all privileged accounts, including local administrator accounts and default appliance accounts. Typically tracked information for each of these accounts includes:

- Where does this account exist?
- Who has access to it?
- What can it access?

Shared admin accounts, or admin privileges attached as secondary roles to standard user accounts increase the difficulty of attribution during an incident and expand the blast radius of a credential compromise. Individual named accounts should replace shared credentials or default accounts, and privilege levels should be limited to the minimum required for the role. Distinguishing between roles such as global admin and global reader in environments like Microsoft Entra ID can significantly limit the damage of a compromised account. Solutions such as Microsoft Local Administrator Password Solution (LAPS) paired with specific-use administrator accounts can automate the management and rotation of local administrator credentials across endpoints, reducing the risk of lateral movement through reused local admin passwords.

Service principles, non-human accounts used by applications, scripts and automated processes to authenticate and interact with resources are a frequently overlooked attack surface. As with ghost devices on a network, these accounts can accumulate over time as applications are deprecated or projects end, leaving behind accounts with standing access that no longer have an active owner. By maintaining an inventory of all service principles, auditing their permissions regularly, and decommissioning any that are no longer in active use an organization can limit the number and duration of these orphaned accounts. Service principles should follow the principle of least privilege, with access scoped only to the resources required for their function. Where possible, time limited or just in time (JIT) access models, for example Microsoft Entra Privileged Identity Management (PIM), can be used to reduce the duration of standing elevated access.

These untracked service accounts pose a similar risk as untracked user accounts and devices, especially once their intended purpose has ended and they are left adrift in an environment. User accounts that are not properly offboarded, service accounts tied to departed employees, or credentials created during a specific third-party project and never removed can all become entry points for a threat actor. These ghost accounts carry the same fundamental risk as ghost devices: they are accessible, often under monitored, and may retain the access they held at the time they were last actively managed. A formal account lifecycle process, from creation through to deprovisioning, can be enforced with periodic reviews to identify and remove stale accounts or devices. This review cadence can use the same process as a risk register to prioritize patching and network hardening. Pairing this with a BIA to identify high priority systems and the accounts that can access can further refine prioritization and planning.

For standard user accounts Group Policy Objects (GPOs) and role-based access controls can be used to ensure accounts are scoped to their function. To give some examples, a standard user in the HR group should not have access to infrastructure documentation, and a member of the helpdesk group should not hold the ability to modify firewall rules or access domain controllers. Account creation policies should be clearly defined to prevent privilege accumulation over time and aid in tracking currently active accounts and their permissions. A common issue is where access is granted reactively and rarely audited after the event, such occasions are inevitable in a live environment, but can be mitigated with automatic time limits to this additional access, follow up reviews for the access at a more practical time, and/or regularly scheduled reviews of the environment to ensure that documented access matches reality.

Conclusion

Through the course of this paper we have delved into the security controls and organizational policies that harden businesses against threat actors such as AKIRA through both open-source intelligence and information gained through the DFIR team since last summer. We hope to have imparted deeper understanding on the importance of strong security policies such as good segmentation and thorough vulnerability management and their positive effect on an organization's security posture.

Incidents such as the ones that inspired this paper underscore the need for vulnerability management protocols that account for remediation steps beyond simply updating to the newest version. The traces left by the leveraging of a vulnerability can guide detection and remediation strategies. Through understanding potential detection and remediation steps insight into threat actor objectives is gained and relevant security controls can be refined. If the vulnerability has seen active use by cybercriminals, then remediation should be regarded as an urgent matter.

Thank you for reading this paper, we hope this work provides inspiration for new defenses or further justifies the infrastructure and policy decisions currently in place. As impactful as a zero-day based campaign can be, AKIRA's latest spree demonstrates that zero-days are not a necessary component of a major campaign. Thankfully the precautions to mitigate the risks of experiencing such a campaign firsthand also provide protection against a myriad of other risks to business continuity and integrity.

Glossary

Business Impact Analysis (BIA): is a systematic process to identify and evaluate the potential effects of an interruption to critical business operations because of a disaster, accident, or emergency. As opposed to risk assessment, which focuses on threats, BIA focuses on the consequences of disruption. If a critical process stops, what happens after one hour, one day, or one week? It assesses financial, operational, regulatory, and reputational impacts. BIA results define recovery objectives: Recovery Time Objective (RTO), how quickly processes must resume, and Recovery Point Objective (RPO), how much data loss is acceptable. Maximum Tolerable Downtime (MTD) is the longest period a system or process can be unavailable before the impact becomes unacceptable; RTO and RPO are set within that boundary. BIA is the foundation for business continuity and disaster recovery planning. Combined with a risk register, the BIA ensures that hardening actions are aligned with business priority: high impact systems receive faster patching, tighter network controls, and stricter account management. These combined efforts directly reduce the value of a compromise to groups like Akira that seek to disrupt or extort critical operations.

Digital Forensics and Incident Response (DFIR): a cybersecurity discipline focused on identifying, investigating, and containing cyberattacks. The digital forensics side collects, preserves, and examines data to understand how a breach occurred. The incident response side focuses on quickly containing active threats to minimize damage and business downtime. Together, these two functions help organizations understand the root cause of an attack and safely recover compromised systems.

DMZ: a isolated subnetwork that exposes an organization's external-facing services to the untrusted internet. It acts as a buffer zone between a secure private network and an insecure public network.

GPO: collections of registry-based settings that administrators use to manage users and computers in a Microsoft Active Directory network. They allow IT teams to centrally configure desktop environments, security policies, and software settings across thousands of machines at once.

MSS: a collection of cybersecurity tasks and capabilities outsourced to a third-party vendor. These providers offer remote monitoring and management of an organization's entire IT infrastructure to protect digital assets from threat actors.

Open Worldwide Application Security Project (OWASP): a nonprofit organization dedicated to improving software security through free, vendor-neutral, and community-driven resources. Its most notable project is the OWASP Top 10, an industry-standard, recognized report identifying critical web application security risks.

Risk register: is a record of information about identified risks (ISO 73:2009). It is used to identify, assess, prioritize, and monitor risks over time. A risk register typically includes for each risk: a description or name, the likelihood of occurrence, the impact or consequence if it occurs, a risk score or rating (usually likelihood * risk), risk category, mitigation measures, contingency plans, and risk triggers. It is a living document to support prioritization of remediation, allocation of resources, and communication of risk to leadership. In the context of this section, the risk register helps determine which systems, applications, and devices should receive highest priority for patching, network segmentation, and access review to ensure defenses are concentrated on the assets most likely to be targeted or to enable lateral movement.

Role based access control (RBAC): a security method that restricts system access based on a user's specific job role. Instead of assigning permissions to individuals, administrators assign them directly to predefined roles. Users then inherit the exact permissions tied to their assigned roles. This system simplifies user management, improves security, and ensures compliance with data protection regulations.

Structured Threat Information eXpression (STIX) / Trusted Automated eXchange of Intelligence Information (TAXII): Industry standard formats for programmatically sharing and organizing cyber threat intelligence.

Appendix A

An excerpt from a deployment script meant to run the AKIRA encryptor against a series of targeted devices via network file sharing protocols.

lock_srv.bat

```
start win_locker.exe -remote -n=3 -p=\\[redacted].10\C$
start win_locker.exe -remote -n=3 -p=\\[redacted].10\D$
start win_locker.exe -remote -n=3 -p=\\[redacted].10\F$
start win_locker.exe -remote -n=3 -p=\\[redacted].10\G$
start win_locker.exe -remote -n=3 -p=\\[redacted].11\C$
start win_locker.exe -remote -n=3 -p=\\[redacted].11\D$
start win_locker.exe -remote -n=3 -p=\\[redacted].11\F$
start win_locker.exe -remote -n=3 -p=\\[redacted].11\G$
start win_locker.exe -remote -n=3 -p=\\[redacted].15\Accounting2
start win_locker.exe -remote -n=3 -p=\\[redacted].15\apps$
start win_locker.exe -remote -n=3 -p=\\[redacted].15\[redacted]$
start win_locker.exe -remote -n=3 -p=\\[redacted].15\C$
start win_locker.exe -remote -n=3 -p=\\[redacted].15\[redacted]
start win_locker.exe -remote -n=3 -p=\\[redacted].15\[redacted]
start win_locker.exe -remote -n=3 -p=\\[redacted].15\[redacted]
start win_locker.exe -remote -n=3 -p=\\[redacted].15\[redacted]
start win_locker.exe -remote -n=3 -p=\\[redacted].15\[redacted]
start win_locker.exe -remote -n=3 -p=\\[redacted].15\[redacted]
start win_locker.exe -remote -n=3 -p=\\[redacted].15\[redacted]
start win_locker.exe -remote -n=3 -p=\\[redacted].15\[redacted]
start win_locker.exe -remote -n=3 -p=\\[redacted].15\CustServ
start win_locker.exe -remote -n=3 -p=\\[redacted].15\D$
start win_locker.exe -remote -n=3 -p=\\[redacted].15\[redacted]
start win_locker.exe -remote -n=3 -p=\\[redacted].15\[redacted]
start win_locker.exe -remote -n=3 -p=\\[redacted].15\[redacted]
start win_locker.exe -remote -n=3 -p=\\[redacted].15\F$
start win_locker.exe -remote -n=3 -p=\\[redacted].15\Fin$
start win_locker.exe -remote -n=3 -p=\\[redacted].15\GPO$
start win_locker.exe -remote -n=3 -p=\\[redacted].15\[redacted]
start win_locker.exe -remote -n=3 -p=\\[redacted].15\HR
start win_locker.exe -remote -n=3 -p=\\[redacted].15\Insurance
start win_locker.exe -remote -n=3 -p=\\[redacted].15\[redacted]
```