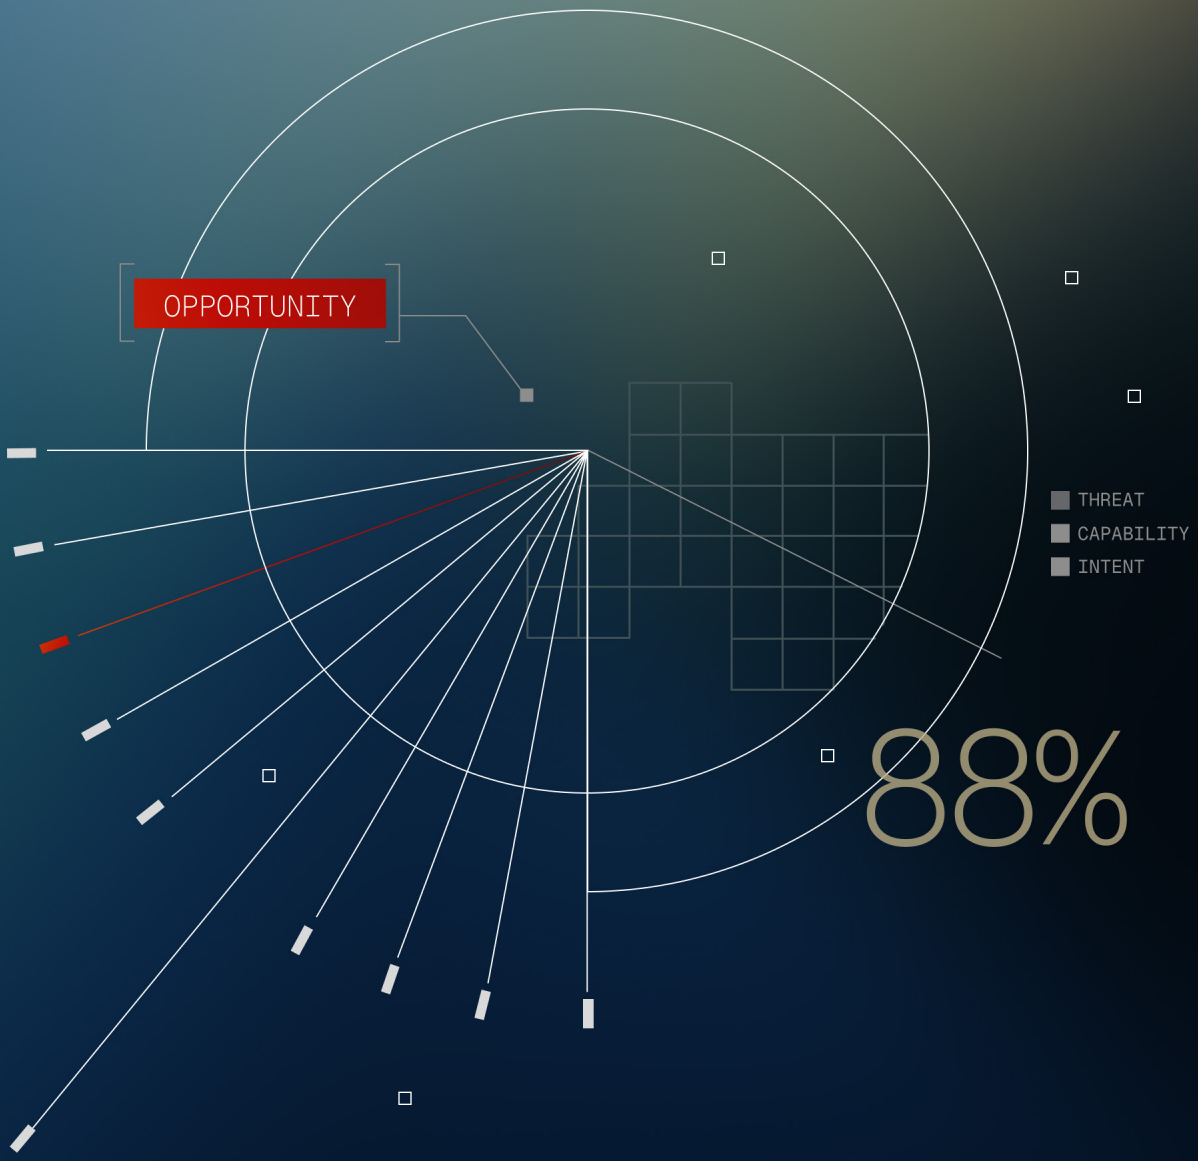
ARMADIN

Opportunity Intelligence

Reducing Exposure to Adversaries



Executive Summary

Threat is commonly assessed as a combination of an adversary's capability, intent, and opportunity, and threat intelligence has historically helped organizations understand the first two components in depth. In the agentic age, it is critical to understand all three elements. Armadin's agentic offensive operations add a view of the third component, measuring the **opportunity** an adversary would have to exploit a given environment.

Across operations conducted against production systems throughout 2026, 88% of the severe findings Armadin identified fell outside existing intelligence sources, and 93% could be remediated by the affected organization without third-party action.

> Armadin refers to this evidence of where organizations can be exploited as opportunity data, which informs **opportunity intelligence**.

Activity data, the focus of most threat intelligence programs, records adversary behavior and the harm it caused in specific environments, while opportunity data shows whether other organizations are similarly exposed, regardless of which adversary would pursue the path.

Actors without established tradecraft can now conduct operations, while experienced attackers are no longer bound by the limits of manual execution, raising both operational tempo and targeting breadth. As attacks increase in quantity and quality, and the cost of frontier models decreases, Armadin assesses that defending only against adversary behavior observed in other environments will likely become less economically viable.

Measured exposure becomes the practical basis for deciding what organizations must fix, especially as they attempt to harden their environment.

Because Armadin's testing is agentic, our opportunity intelligence offers some of the first insights into what AI-driven attackers will be able to exploit as the barrier to entry for cyber operations continues to shrink.

AT A GLANCE

88% of the severe findings Armadin identified fell outside existing intelligence sources, and 93% could be remediated by the affected organization without third-party action.

How a Threat is Defined



Key Findings

FINDING

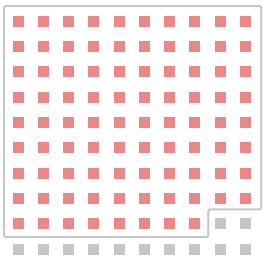
96%



96% of clients had at least one critical finding. Critical findings represent a direct and immediate threat to the organization and require emergency remediation.

FINDING

88%



88% of the severe findings were not identified through external intelligence sources, including vendor advisories, public security research, or CISA's KEV list. Severe findings refer to those that encompass criticals and highs (further defined in the appendix). The majority of vulnerabilities were either in the clients' own software, a product of deployment configuration, or had not yet been discovered.

FINDING

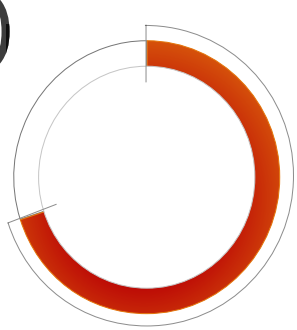
71%

71% of credentials used in kill chains were non-human identities. The majority of these were stored in readable form, with organizations disclosing them unintentionally.

FINDING

~70%

Nearly 70% of severe findings could be reached by someone without an account, and the majority of the documented kill chains started from an unauthenticated position. More than half of those kill chains ended with Armadin holding sensitive data, administrative control, or a combination of both.



FINDING



93%

93% of severe findings could be mitigated without third-party action and 49% did not require modification to application code. This means nearly all identified exposure could be mitigated directly by the organizations, and time to remediate depends primarily on the organizations' capabilities.

Our Top Recommendations

01 Prioritize by Measured Exposure

Base remediation priorities on the measured risk of your own environment, including internally built software, which holds exposure no external source will describe.

02 Close Unauthenticated Access

Require authentication for functions and data stores, change default settings, and remove administrative consoles and installers from internet exposure, starting with internet-facing systems.

03 Govern Service Account Credentials

Locate, remove, and rotate credentials stored in readable form, starting with service accounts and client-side code, and enforce policies for how teams provision, store, and rotate them.

04 Restrict Lateral Movement & Privilege Escalation

Restrict user and service account permissions to what their function requires, and use micro-segmentation to isolate network segments and contain potential breaches.

05 Fast-Track Configuration Changes

Give setting, permission, and credential changes an approval path that moves faster than a code release.

06 Patch Management & Compensating Controls

Patch internet-facing vendor products first, and add a compensating control as a stopgap to prevent abuse where no patch currently exists.

Operations and Methodology

The risks and vulnerabilities Armadin found are **not hypothetical**. Each finding and attack narrative was discovered during active tests against live systems across external, internal, and web application attack surfaces.

Since February 2026, we identified critical findings in 96% of our client base. At nearly three quarters of these clients, testing either confirmed code execution or identified a technically supported path to it that was not executed because of engagement restrictions.

We also accessed sensitive personal, financial, health, or confidential business data in the vast majority of our client base, demonstrating the impact of these operations.

In the vast majority of these cases, Armadin needed **only three or fewer steps** to produce a critical outcome without valid credentials.

These kill chains often did not require substantial effort. Approximately half of the documented kill chains did not require an exploit to be used.

Armadin defines a finding as a single recorded weakness, and a kill chain is a documented multistep path to achieve an objective. A step is one recorded action within a chain. Severe findings refer to combined scores for findings rated Critical or High.

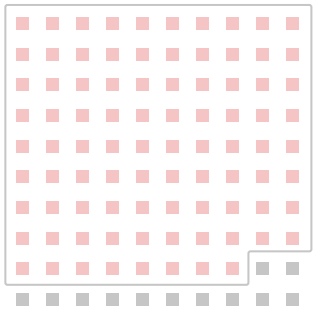
RISK	IMPACT AND REMEDIATION
CRITICAL	These findings represent a direct and immediate threat, and remediation should be treated as an emergency.
HIGH	These findings pose a significant threat to operations, and require urgent attention.
MEDIUM	These findings indicate a tangible security weakness that a motivated attacker could exploit, and should be addressed as part of regular patching cycles.
LOW	These findings represent minor security weaknesses, and remediation should be performed as resources allow.

See details on methodolgy >>

FINDING

88% of Armadin's Severe Findings Had No Matching Coverage in Published Sources

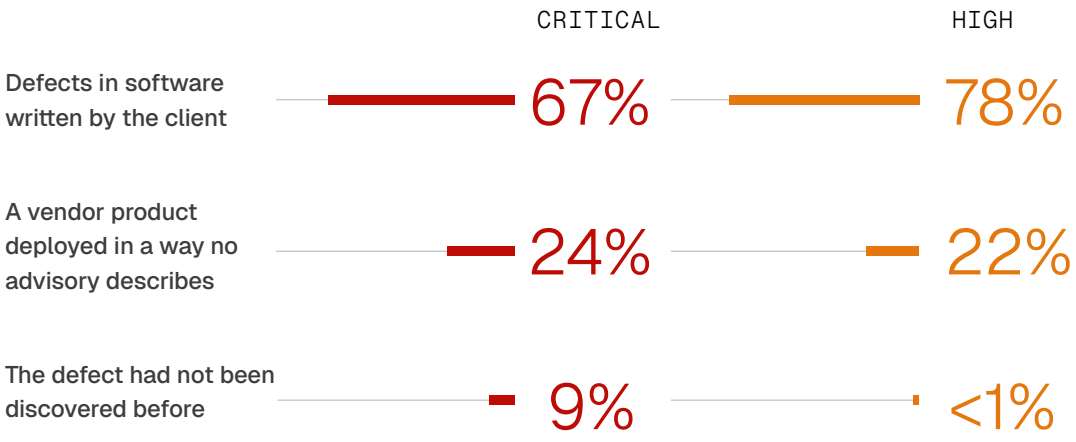
88%



Existing intelligence sources described only a fraction of the severe findings we identified.

88% of severe findings had no matching coverage in the external published sources reviewed, including CVE records, CISA KEV entries, vendor advisories, and public security research. These were instances that outside sources had no ability or reason to document. An organization managing exposure from external reporting alone would have lacked insight into the majority of routes that would matter to real-world attackers.

Why Findings Fell Outside of External Coverage



45%

of Critical findings related to third-party software

94%

of severe findings were not matched in a public advisory

Most of this exposure was in software the client built, which placed it beyond the reach of external threat intelligence.

Key Statistics

Across all severe findings, approximately two-thirds sat in software the client built and one-third were in third-party products.

Third-party software carried a larger share of the most serious findings, accounting for 45% of critical findings against 29% of high findings. Only about 6% of severe findings matched a published advisory for a vendor product, which means 94% were not. Within third-party software, roughly one in six findings matched a published advisory, while more than half arose from how the product was deployed and configured.

Much of the remainder was from a vendor product operating as designed. Because the code is not defective, advisories would not address how threat actors could take advantage of it. The exposure originates in a deployment decision that only someone inside the organization, or someone targeting it, would observe. This includes administrative consoles published to the internet, authentication available but switched off, and defaults that could allow entry by an attacker.

Nearly two-thirds of the findings that gave Armadin code execution on a system had no published CVE or advisory when we found them. Code execution findings were more likely than other severe findings to sit in vendor products, and about a sixth of them were zero-days Armadin discovered.

A separate portion of findings without advisories were zero-days that Armadin discovered in vendor software. We found them across an identity governance platform, an enterprise backup product, a network monitoring suite, a web application firewall, a content management platform, and other areas before reporting each to its vendor. Vendors frequently take months to publish these vulnerabilities, so their customers often remain exposed for prolonged periods.

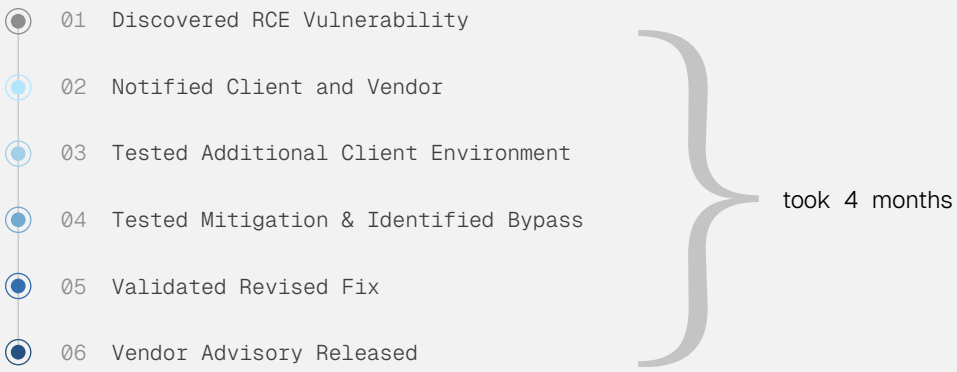
CASE STUDY

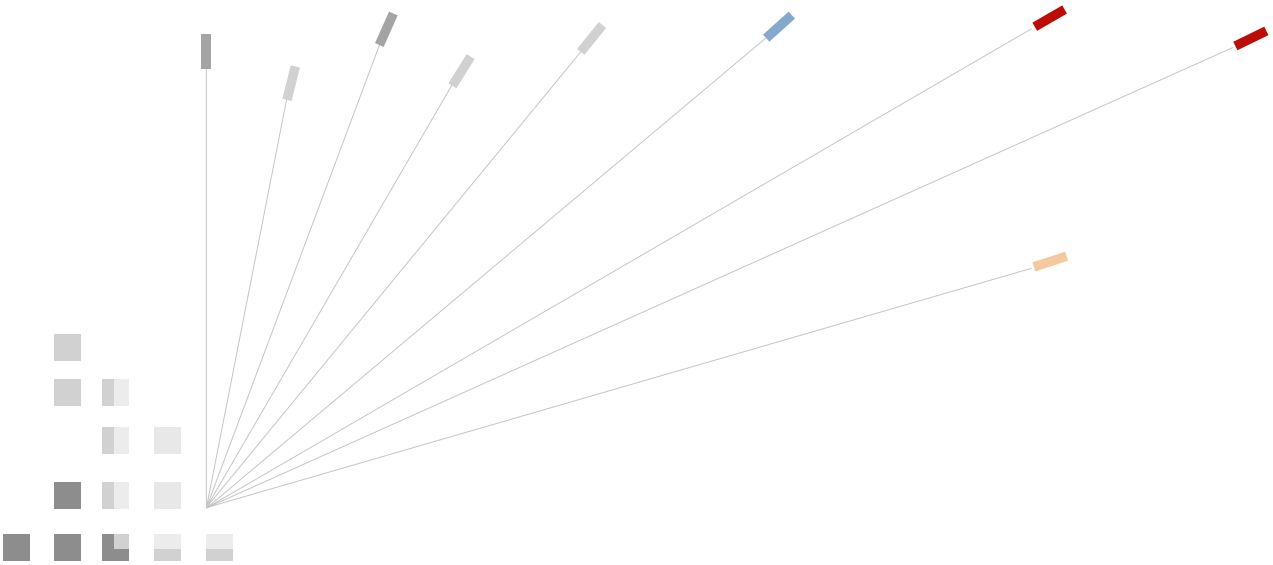
A Critical Defect in an Identity Governance Platform

The affected product was a platform the Armadin client had purchased in order to govern access across its environment. Armadin found unauthenticated remote code execution (RCE) in it. The platform exposed an endpoint to unauthenticated callers by design, along with the evaluated submitted values, allowing Armadin to remotely execute code on the server running the application.

Armadin notified the vendor, then tested the same defect at a second client running the same platform as part of its ongoing testing. With no patch available, the goal was protection before the patch, or measuring what an attacker holding the same zero-day could reach, mapping the kill chains from that foothold, and helping the client sever them while a fix remained in development. That client blocked the vulnerable path and Armadin established that the block could be trivially circumvented. The client's revised block deployed the following day withstood the original payloads, the circumvention, and the additional variants Armadin attempted.

It took the vendor nearly four months to publish an associated advisory. Between the time we identified it and that advisory, the only organizations aware of it were the ones Armadin tested.





GUIDANCE FOR DEFENDERS

When a zero-day vulnerability leaves an organization exposed before a vendor patch is available, perimeter defenses and tracking published vulnerabilities alone are insufficient. During this period, organizations can reduce the potential impact of successful exploitation by strengthening controls that limit what an attacker can accomplish afterward. Enforcing least-privilege for service accounts and implementing network segmentation can restrict lateral movement, limit access to sensitive systems, and reduce the likelihood that exploitation of a single vulnerability results in broader network compromise or data exfiltration.

Armadin assesses that it is likely becoming less viable for organizations to prioritize risk mitigation based on external vulnerability data. We recommend organizations shift the basis for deciding what to fix from what has been reported about other organizations to the measured risk of their individual environments.

FINDING

71% of Credentials Used in Kill Chains Were Non-Human Identities

71%

71% of credentials used in kill chains were non-human identities. The majority of these were stored in readable form, with organizations disclosing them unintentionally.

Software held most of the credentials that enabled the kill chains we found. Exposed credentials enabled just under half of the kill chains that began from an unauthenticated position. In 71% of those kill chains, the credential belonged to an application or service account, with the majority in readable form.

In many cases, Armadin encountered cleartext service account credentials that were stored in several locations accessible to unauthenticated or low-privilege attackers.

How the Credential Was Obtained	Share of Kill Chains That Used a Credential
Stored and accessible in readable form	46%
No secret required because the application issued a session/accepted a token without checking it	24%
Default, weak, or breach-sourced password	19%
Recovered from memory or cracked offline	11%

90%

of these credentials were accessible without extraction

33%

of the kill chains that reached administrative rights included an exposed secret

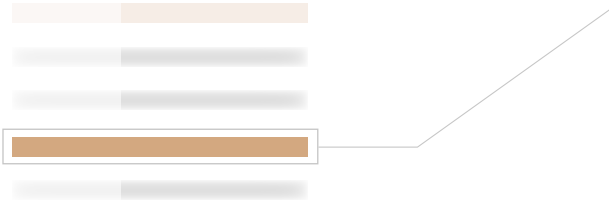
Missing authentication was our most common finding, and authentication gaps also appeared at the login itself.

Nearly 90% of the credentials used in kill chains did not require active extraction or cracking; they were simply exposed in plain text. Of the credentials stored in readable form, the largest share were in client-side code, followed by configuration and environment files, as well as application responses.

Software-owned accounts are frequently vulnerable, given they commonly lack the protective measures applied to human accounts. They often hold privileges granted at installation without amendment, infrequently have credentials rotated, and are rarely offboarded. In most cases, they are not covered by multifactor authentication.

Exposed credentials were also frequently responsible for us reaching administrative control. More than a third of the kill chains that reached administrative rights included an exposed secret, and about a quarter relied on a weak or default password.

These credentials can regularly be extracted with a single request. At one Armadin client, a server answered an anonymous upload and returned its own stored settings inside a page response, including the access keys to the production cloud account, which carried administrative rights across the resources in it.



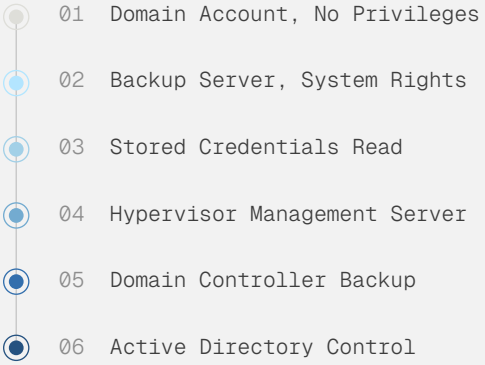
CASE STUDY

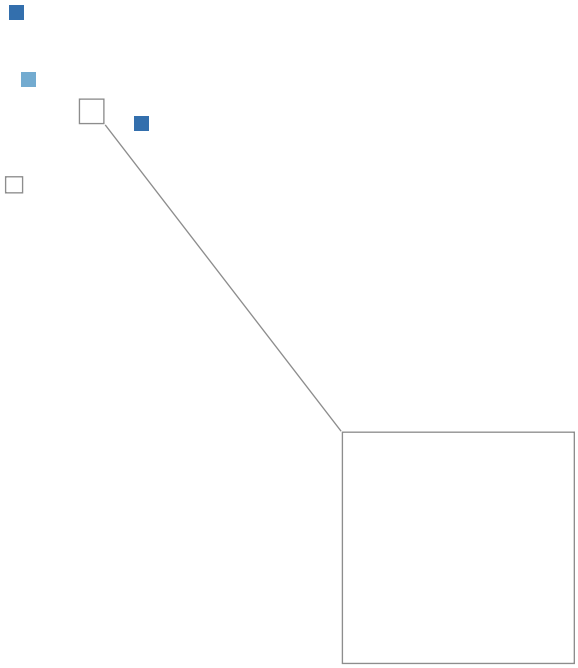
Backup Platform Holding Credentials

Armadin began with an ordinary domain account holding no privileged group memberships and used a weakness in a privileged component of the backup software to execute commands on a backup server with full system privileges. The same service was present on eight backup servers across the network.

We then read the credentials the backup platform had stored for the systems it protects. These covered the virtualization infrastructure, cloud storage keys, and privileged operating system accounts. One of the accounts was granted administrative rights to a virtualization system. That allowed us to identify hypervisor management servers hosted elsewhere on the network that were accessible. Armadin logged in to a server, which served as the management plane for 65 virtual machines across three hosts and a shared datastore.

Armadin then identified the primary domain controller through the hypervisor management server access and reconnected to the backup server with the compromised administrator account. We used the backup system access to recover a domain controller backup, extracting hashes for domain administrators and enterprise administrators. Holding this credential material enabled control of the Active Directory environment, all without code running on a domain controller or data written to production storage.





GUIDANCE FOR DEFENDERS

The accounts that carried these attacks fall outside the routines most organizations use to govern credentials since they belong to software and no joiner or leaver event brings them into review. A platform that backs up a domain controller holds a copy of the credentials inside it, and whoever controls the restore process can reach the network without touching production.

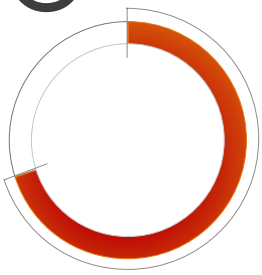
Armadin assesses that this exposure will likely widen over the next 12 months as automation adds software identities faster than governance extends to them and as AI-enabled collection lowers the cost of finding the credentials they depend on.

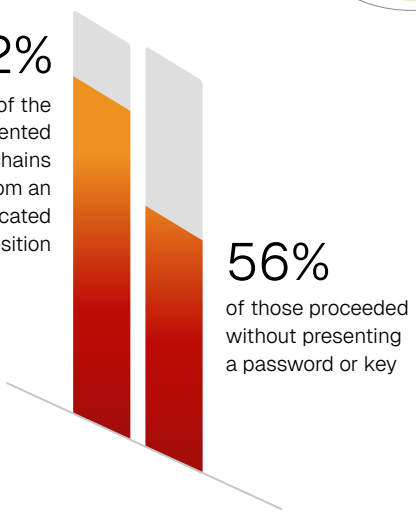
We recommend that organizations define and enforce standardized policies for how teams provision, store, and rotate service account credentials. We also recommend that organizations proactively evaluate their client-side code and configuration files to ensure that any potentially exposed secrets are remediated prior to exploitation. Finally, we recommend that organizations limit service accounts to the minimum privileges their function requires, and use micro-segmentation to isolate environments and the systems that connect them, so a single compromised credential or host cannot carry an attacker from one network to the next.

FINDING

Nearly 70% of Severe Findings Could Be Reached without an Account

~70%





82%
of the documented kill chains started from an unauthenticated position

56%
of those proceeded without presenting a password or key

The majority of our findings could be reached by someone without an account. This is because they sat on an unauthenticated attack surface, with no credentials needed to reach them. Full kill chains reflected this same trend. In fact, 82% of the documented kill chains began from an unauthenticated position, and 56% of the kill chains progressed without using a password or key.

More than half of these routes ended with Armadin holding sensitive data from the client, administrative control, or both. The defects that enabled access without accounts were mostly ordinary application flaws.

SQL injection was as common as it was consequential in our operations. More than 90% of SQL injection findings were rated severe, and they appeared in 60% of our client environments. Because each flaw sat in front of the point where a credential would be checked, strengthening authentication controls alone would not eliminate the exposure.

CASE STUDY

Unauthenticated Endpoint Returning Full Personnel Database

During an external engagement, Armadin identified a web application that disclosed information about a backend API. Upon testing that API, Armadin found that it responded to an unauthenticated request with the client's complete technician database in a single response, holding full names, national identity numbers, email addresses, phone numbers, dates of birth, and links to photographs of identity documents and signatures.

The application enforced authentication correctly across its other endpoints, so the system as a whole was not left unprotected. However, the one reachable route provided information that could enable identity fraud, regulatory exposure, social engineering, and other risks.

GUIDANCE FOR DEFENDERS

In our experience, organizations that maintain exposure inventories at the system level often record the type of exposure we identified in these cases as secured. The majority of this exposure sits in front of the point at which a credential would be evaluated, placing it outside the reach of identity and single sign-on investment. Placing authentication in front of exposed endpoints was the remediation that most often addressed the first exploited step of a kill chain, applying to approximately a quarter of them.



Armadin assesses that automated reconnaissance is likely compressing the interval for actors to identify and exploit unauthenticated vulnerabilities. We recommend that organizations focus their security investments on *both* unauthenticated and authenticated attack surfaces.

FINDING

93% of Severe Findings Could Be Mitigated without Third-Party Action

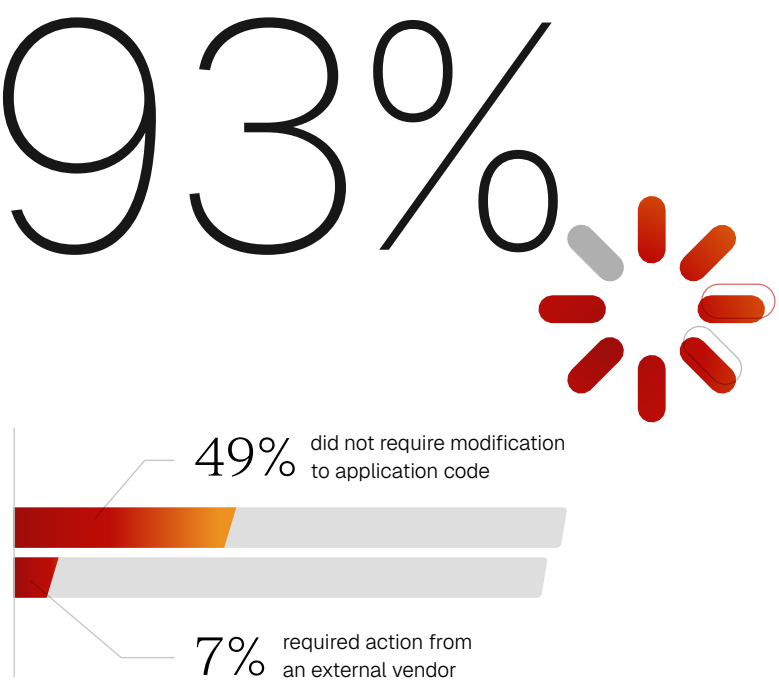
Conventional threat exposure models often assume that remediation is constrained by third-party vendor disclosure and patch timelines, but new opportunity data indicates that 93% of severe findings had at least one mitigation or remediation path that did not require a vendor or third party to act.

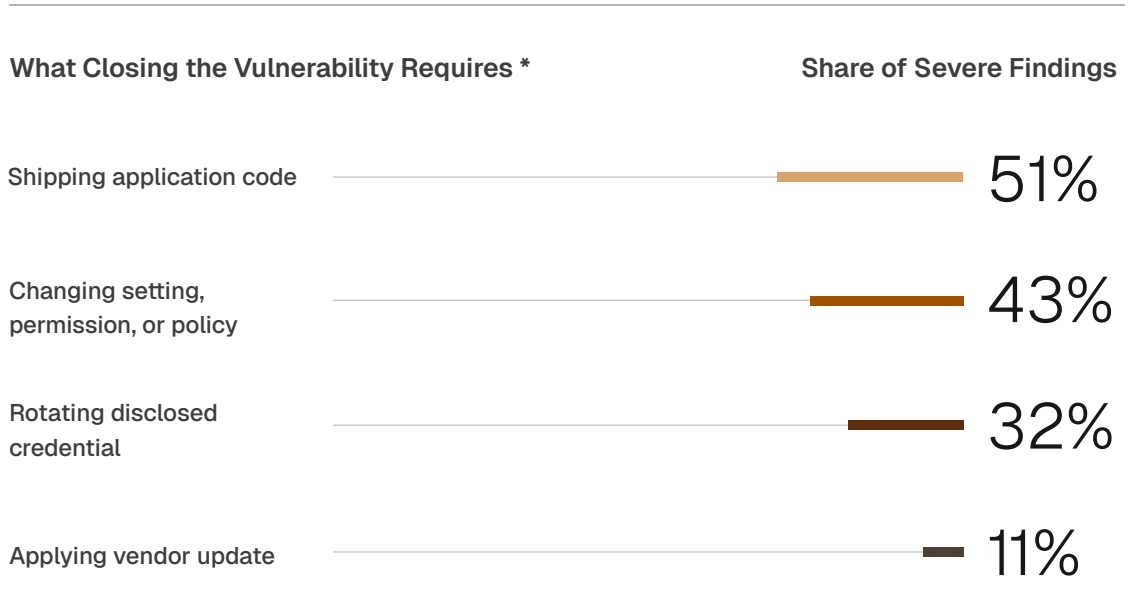
Only 7% of the severe findings we recorded could be closed only through a third party. For the rest, organizations could close the exposure another way, usually by rotating a credential or changing a configuration.

Code execution was the exception. 30% of the findings that gave Armadin code execution could only be closed by a vendor, against 7% of severe findings overall, because code execution concentrates in vendor products and zero-days.

Almost half of all severe findings were addressable through routine security controls and administrative policy enforcement. Addressing them also carries low engineering overhead, which may allow security teams to respond faster.

Armadin found that most kill chains combine technology and development flaws, including coding errors, permissions, and architecture, with process gaps in development, patching, and quality assurance. No single category of change addressed the first exploited step of most kill chains.





*Many findings needed more than one action to close, such as rotating an exposed credential and also fixing the code that exposed it, so the shares add up to more than 100%. The 11% includes every finding a vendor update would close, including those that could also be closed another way, while the 7% includes only findings with no setting, code, or credential route available to the organization.

GUIDANCE FOR DEFENDERS

Organizations that treat remediation as a function of vendor disclosure are likely to misjudge what is actually slowing remediation in their systems. Half of these findings can be closed without application code modification, and 93% had at least one path that required no third party to act, which places the options for closure within the organization's authority.

For organizations already evaluating their own exposure, Armadin assesses that the limiting factor on remediation is likely shifting toward change authorization over the next 12 months. 43% of severe findings can be closed through a setting, permission, or policy change, and 32% through rotating a disclosed credential. Remediation work that moves through the same approval path as a code release can be unnecessarily delayed.

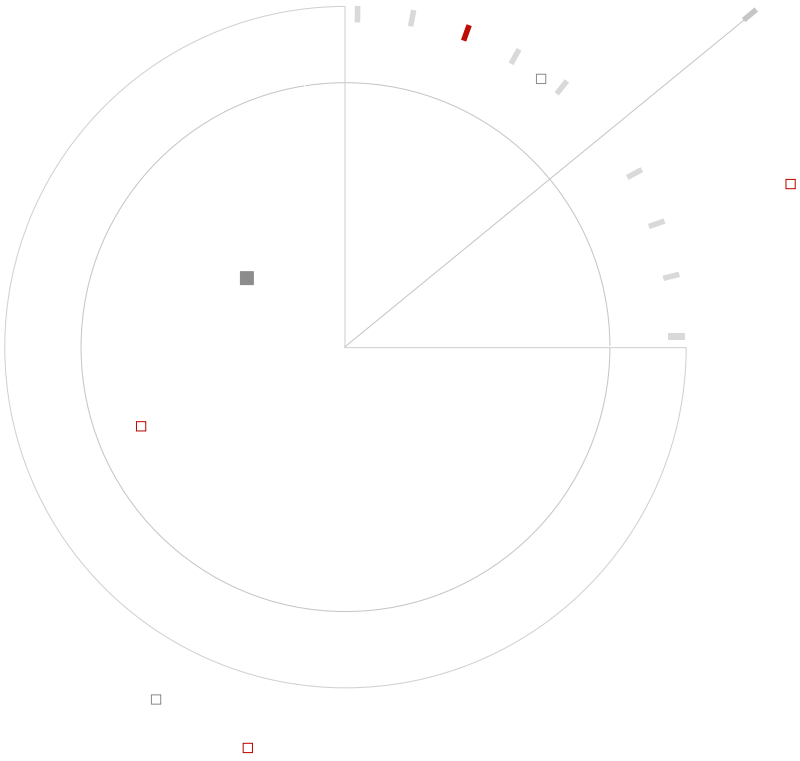
Strategic Outlook

NOW — +12 MONTHS

Armadin assesses with high confidence that the amount of available exposure converted into intrusions will continue to rise over the next 12 months, given the increasing adversarial use of AI and the decreasing cost of using AI tools for nefarious purposes.

AI tooling lowers the barrier to entry, so operators without established tradecraft can reach outcomes that previously required a specialist, while actors who already had the tradecraft are no longer bound by how fast an individual can work. For defenders the same AI tooling lowers the cost of measuring their own environment.

The exposure recorded here was measurable before an adversary arrived, and most severe findings had a mitigation or remediation path that did not require waiting for a vendor. What determines the outcome is how much exposure organizations choose to look for and how quickly they can fix it.



As discovery becomes less expensive due to agentic AI, the variable separating organizations likely becomes the speed and scale at which they act on what they find.

Appendix: Impact and Remediation Guide

Risk	Impact and Remediation
Critical	• These findings represent a direct and immediate threat. • They are characterized by a combination of high to catastrophic business impact and moderate to high exploitability. If exploited, these vulnerabilities could lead to a systemic compromise, mass data exposure, or significant disruption of critical business functions. • Remediation should be treated as an emergency and addressed immediately.
High	• These findings pose a significant threat to operations. • While they might be more difficult to exploit than critical findings, they could still lead to major compromises of sensitive data or disruption of key services. • They require urgent attention and should be remediated as a high priority.
Medium	• These findings indicate a tangible security weakness that a motivated attacker could exploit to gain unauthorized access or disrupt services. • They typically have a moderate impact or are more difficult to exploit. • They should be addressed as part of regular patching and development cycles to reduce the attack surface.
Low	• These findings represent minor security weaknesses or deviations from best practices. • While they do not pose an immediate or significant threat on their own, they could potentially contribute to an attack chain. • Remediation is recommended to improve the overall security posture and should be performed as resources allow.

ARMADIN

OPPORTUNITY INTELLIGENCE: REDUCING EXPOSURE TO ADVERSARIES

Stay Ahead of Whatever Threats May Come

AI-driven threats require a new approach and the right platform to keep environments secure.
Discover more at armadin.com

©2026 ARMADIN, INC. ALL RIGHTS RESERVED.