


The containment stall and how to get past it



Introduction

Ideas about effective security strategies have always evolved with the threat landscape. This has been a reactive, back-and-forth process, as new controls and approaches become available and attack strategies, tools and techniques adapt to break them. The emergence of increasingly rapid, autonomous, agent-driven attacks has necessitated that defenders assume workloads will be compromised as the starting point to build an effective defense.

The critical question to answer today is: What additional assets can a compromised workload reach? The past “janitorial” model — cleaning up after compromise — and the enhanced perimeter defenses that succeeded it were built for simpler attacks. Organizations must move beyond these legacy detection and response patterns to maintain a reasonable security posture. Isolation and segmentation are useful tools, albeit underutilized. However, they, too, leave large gaps. To provide meaningful protection in a world where compromise must be assumed, security professionals need to focus on containment.

Containment bounds what a compromised workload can reach. Detection remains critical, but with attack timelines approaching zero, containment is the strategy that bounds how far a successful attack can reach. This represents a change in thinking, shifting to a preemptive posture that can constrain any successful attack by understanding the potential reach and impacts of a compromise before it happens. This approach puts isolation and segmentation controls to work in new ways that can overcome the perceived risks that have held back deployment strategies such as microsegmentation.

The goal in the current threat environment must be to achieve true containment of successful attacks. This is a deceptively simple concept, but it requires a sophisticated approach to build effectively — and it needs to happen today.

The Take

For too long, security teams have been asking the wrong questions about security protections and using the wrong metrics to generate answers. In an AI-driven world, the time from vulnerability disclosure to exploit is no longer meaningful. The question of how long it takes for an organization to implement a patch loses relevance. Organizations have effective controls at their disposal, and yet they are not set to enforce the isolation and segmentation policies that would contain a successful attack. To deliver security efficacy, teams need to ask what is holding them back from using those controls not just as passive monitors, but to enforce protections before attacks are launched. The key question today is not whether containment is the goal, but why so few organizations are implementing the enforcement policies to achieve it. Enterprises must establish capabilities that give them confidence to address risks and achieve containment without breaking production.

Today’s hesitant security reality

Modern security teams do not lack tools. Most organizations have a wealth of tools and an extensive set of controls that can address a range of risks. The challenge most face is a lack of clarity about the potential impacts of full enforcement on the applications that run in their environments. As a result, proof-of-concept deployments drag on, and security products are kept in “monitor only” mode in perpetuity. In the typical enterprise, various organizational dynamics make security teams hesitant to lock down their infrastructure.

In the modern enterprise, technology teams are often partitioned by specialty. While that can help to develop depth of skill and scale service delivery, it also creates silos across whose boundaries context and visibility may be lost. When security teams do not understand changes in application behaviors, they risk blocking key business processes. When application developers lack the latest security insights, they can build with vulnerable code. These dynamics create motivations that do not serve organizational security. Rather than assessing security posture through the most effective means, teams resort to metrics that are easy to measure but have limited impact.


A typical example is vulnerability remediation processes. It is easy to measure the rate of vulnerability discovery and remediation. Vulnerabilities become the point of exchange between security and development teams. While there is a metric to track the process, it is disconnected from the enterprise’s effective security posture. A fundamental problem is that it is difficult to measure the attacks that don’t happen. Teams are seldom rewarded for successful defense, but are penalized for breaking an application or delaying a new service.

Organizational dynamic	Challenge	Outcome
Silos limit context and visibility	Controls enforcement is limited	Security protections are limited
Security metrics are misaligned	Easily tracked metrics are improved	Security effectiveness is not improved
Vulnerability remediation becomes the focus	Identification and prioritization face scaling challenges	Security is working backward while attackers move forward

In the historical threat landscape, this misalignment could be tolerated. There were improvements in security posture over time, and the time horizon of most attacks was not tremendously far ahead of remediation efforts. The assumption that most environments would be patched prior to being attacked often held true. However, the shift to autonomous attacks in a post-Mythos world has changed that. The first instinct in many organizations has been to reinvigorate the vulnerability management program — in other words, to apply the old measure with more effort.

AI-driven vulnerability discovery has multiple implications that render untenable the expectation that organizations can keep up with vulnerability management. Not only are vulnerabilities being generated at a dramatically greater rate, the ability of agents to coordinate and chain vulnerabilities into a successful attack increases the ease of exploitation. It is as though every vulnerability becomes a zero-day attack.

Figure 1: Top strategic security objectives for 2026



Q. What are your top strategic security objectives for 2026? Please select all that apply. (Top five responses.)
Base: All respondents (n=252).
Voice of the Enterprise: Information Security, Budgets & Outlook 2026.
Source: 451 Research by S&P Global.

At the same time, agentic applications within the enterprise are scaling up, leaving less time to evaluate and secure them. Traditional security processes are being overwhelmed, as the volume of agents and their ephemeral nature outstrips security teams’ ability to secure them. With executive-level support for agentic applications, the penalties for holding them back can be significant.

It is not that enterprises are not reacting to address these changes. The complication is that most organizations are focusing on improving remediation, rather than on enforcing controls. Teams are taking steps to push patches out faster and accepting more downtime to apply them, but in doing so, they are racing toward the wrong goal. They are stalling in efforts to achieve containment because the risks of enforcement and the potential to break production applications appear greater than the risks associated with patching.

What is holding back better security?

Security teams often understand what they could be doing better, but are held back by operational constraints and organizational behavior. Organizational friction points that stall effective security implementations can be broken down into four major types:

- **Operational blindness:** The exponential expansion of interconnections between applications can effectively blind security teams. Demands for new network paths are growing much faster than the ability to understand the implications. Hybrid infrastructure has created separate islands of telemetry in different environments. Operational blindness stops efforts to implement effective controls before they can begin, as it prevents teams from identifying a starting point or tracing likely downstream effects. The resulting idea that change must occur across all applications and environments simultaneously can be an insurmountable obstacle.
- **Control chaos:** Having too many uncoordinated tools may be just as bad as not having enough. This is another friction point that has increased with the adoption of hybrid infrastructure. There are differing controls across various cloud and on-premises venues, and creating matching policies can be a nightmare that requires deep specialization in all environments. This chaos can prevent the implementation of effective security as teams face the complexity of applying the right controls in the right locations.
- **Blocking blocks:** Fear of breaking production applications often prevents strong controls from being enforced. One wrong rule can result in a production outage and assignment of blame. Additionally, concerns about blocking innovation can keep environments open to attackers. This is particularly true of egress controls. Attempts to avoid blocking can stop security efforts before they become truly effective: when controls are in place but frozen in monitoring mode.
- **Maintenance malaise:** Even when controls are implemented, there is often resistance to change as environments and applications evolve. When it is no longer clear who owns a permission or how it is being used, making a change can be risky. The rapid rate of change in agentic applications raises the risk of granting too many permissions for too long, as security teams try to avoid blocking innovation. Agents may accumulate standing reach that no one remembers granting. This inertia keeps security from adapting effectively once permissions are in place.

Friction point	Stops security at:	Owner
Operational blindness	Initial evaluation	Security architect
Control chaos	Control insertion	Platform teams
Blocking blocks	Control enforcement	Security operations
Maintenance malaise	Operational change	Compliance and audit

Unlocking security velocity and depth

To build effective security capabilities for today’s threat landscape, enterprises must do two things: Shift to a containment mindset, and overcome the implementation frictions that have held them back. Achieving containment requires that controls be in place before an attack. Visibility into the current state of enterprise-wide infrastructure is essential to reduce the operational risk of enforcement. Security policies must extend across all parts of the infrastructure without requiring translation. And those policies must be written to the workload, rather than the address, because addresses change by the minute in cloud and agentic environments. Additionally, controls need to achieve comprehensive isolation, which means strong blocking of both ingress and egress traffic must be applied to every workload. To be effective, containment measures must also govern what partners or software-as-a-service connections can reach once inside.

Overcoming implementation friction requires both better technology and, by extension, the establishment of trust based on the data that technology generates. At the heart of required technology improvements is the need for integrated visibility that spans infrastructure boundaries. Visibility must integrate cloud and on-premises data while maintaining sufficient context about the workloads being secured. To be maximally effective, controls cannot carry the baggage of agent-based approaches, as environments often present too many barriers to deploying additional agents. Strong blocking policies must therefore be driven by data from visibility capabilities; this affords sufficient trust in the policies to support implementation of comprehensive blocking to limit an attacker’s lateral movement. Importantly, a policy should be able to show clearly what it would block before it blocks anything.


Given the rapid pace of change in IT environments, the operational aspects of a containment posture must continually adapt to shifting demands while remaining able to identify active protections and permissions and their purposes. That can be achieved by using the same work model that applications already use: Policy is declared as code, versioned, tested and deployed through the same pipelines that ship the application. In this way, a policy change moves at the speed of a deployment rather than a change ticket. The process must also satisfy audit and compliance requirements by enabling reporting on the current state of protection and the supporting rationale. The system must express risk exposure in a form a board can understand and track — i.e., what a compromised workload can reach today and how that changes over time.

Friction point	Addressed by:	Outcome
Operational blindness	Extensive visibility, including mapping which workloads communicate with which destinations, not just a list of assets	Automated policy and reduced risk
Control chaos	Control insertion without disruption or reliance on agents	Coordinated controls with minimal disruption
Blocking blocks	Policies that specify what they would block before they block	Blocking implementation with low risk
Maintenance malaise	Policy as code, with data-driven validation and adaptation	Continuous compliance with provable protections

Implementing an effective containment strategy also requires thoughtful transition planning. Overcoming existing organizational dynamics requires building trust, making rip-and-replace deployment unfeasible. An initial deployment phase is needed to collect operational data to inform policy and establish context for the environment. Ideally, this is when application relationships and workload characteristics are mapped. This stage provides an opportunity to build trust across teams for whom a well-built visibility capability can act as a single source of truth. It is also useful to plan for team integration — extending responsibilities across previous silo boundaries.

When built in an integrated fashion, a containment approach can mitigate some of the more complex challenges of security operations, such as managing data ownership and access policies. This is an area where the demand for agentic access can be a strong motivator. Proper controls can enable innovation while ensuring appropriate data protection.

Figure 2: Top security tool categories by planned increase in spending



Q. For each of the following vendor-based security tools, how will your organization's spending change in 2026 — if at all — compared to 2025?
Base: All respondents.
Voice of the Enterprise: Information Security, Budgets & Outlook 2026.
Source: 451 Research by S&P Global.

Conclusion

The step change in the speed of attacks demands that organizations fundamentally shift the focus of security operations. To defend effectively, they need to have strong, active controls that can contain compromises when they happen. Isolation and segmentation must function at the workload level, and they must be implemented via systems dynamic enough to respond at the pace of enterprise innovation. To reduce the risk of breaking business-critical applications, telemetry needs to extend across an organization's entire infrastructure — on-premises, in the cloud and everywhere in between. Policy changes must be enacted with confidence that they will not cause damage. Containment is not a one-and-done action. It is a continuous process that requires integrating controls with context to improve security footing and maintain operational velocity.



Aviatrix is the network containment company. Aviatrix Clearance gives every cloud workload clearance for the paths it needs and nothing else. It is two controls and an autonomous defense layer, Outbound Containment, Lateral Containment and Aviatrix Active Defense, purchased together or in any order. Every policy shows what it would have blocked before it blocks anything, enforcement is one setting, and backing out is one action. Start with the free Containment Assessment. In five minutes, it shows what a compromised workload in your cloud can reach today: aviatrix.ai/containment-assessment.

About the author



Eric Hanselman

Chief Analyst

Eric is the chief analyst at 451 Research by S&P Global. He coordinates industry analysis across the broad portfolio of research disciplines, with an extensive, hands-on understanding of subject areas including information security, networks and semiconductors, and their intersection in areas such as AI, 5G and quantum computing.

Eric helps 451 Research's clients navigate these turbulent waters and capitalize on potential outcomes. He is a lifetime member of the Institute of Electrical and Electronics Engineers (IEEE) and a Certified Information Systems Security Professional (CISSP). He is also a frequent speaker at leading industry conferences and hosts the Next in Tech technology podcast.

About 451 Research

451 Research by S&P Global provides essential insight into key trends driving digital transformation across the entire technology landscape. By offering a combination of expert analyst insight and differentiated data, 451 Research provides client organizations with the critical information and perspectives they require to make more effective business decisions.

451 Research offerings are organized around the prevailing themes driving digital transformation today, including data centers and energy, AI, cybersecurity, digital industries, and resilient digital infrastructure.

For more information about 451 Research, please go to: spglobal.com/451research.

CONTACTS

www.spglobal.com
www.spglobal.com/en/enterprise/about/contact-us.html

Copyright © 2026 S&P Global Inc. All rights reserved.

These materials, including any software, data, processing technology, index data, ratings, credit-related analysis, research, model, software or other application or output described herein, or any part thereof (collectively the **“Property”**) constitute the proprietary and confidential information of S&P Global Inc its affiliates (each and together **“S&P Global”**) and/or its third party provider licensors. S&P Global on behalf of itself and its third-party licensors reserves all rights in and to the Property. These materials have been prepared solely for information purposes based upon information generally available to the public and from sources believed to be reliable.

Any copying, reproduction, reverse-engineering, modification, distribution, transmission or disclosure of the Property, in any form or by any means, is strictly prohibited without the prior written consent of S&P Global. The Property shall not be used for any unauthorized or unlawful purposes. S&P Global's opinions, statements, estimates, projections, quotes and credit-related and other analyses are statements of opinion as of the date they are expressed and not statements of fact or recommendations to purchase, hold, or sell any securities or to make any investment decisions, and do not address the suitability of any security, and there is no obligation on S&P Global to update the foregoing or any other element of the Property. S&P Global may provide index data. Direct investment in an index is not possible. Exposure to an asset class represented by an index is available through investable instruments based on that index. The Property and its composition and content are subject to change without notice.

THE PROPERTY IS PROVIDED ON AN “AS IS” BASIS. NEITHER S&P GLOBAL NOR ANY THIRD PARTY PROVIDERS (TOGETHER, **“S&P GLOBAL PARTIES”**) MAKE ANY WARRANTY, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE, FREEDOM FROM BUGS, SOFTWARE ERRORS OR DEFECTS, THAT THE PROPERTY'S FUNCTIONING WILL BE UNINTERRUPTED OR THAT THE PROPERTY WILL OPERATE IN ANY SOFTWARE OR HARDWARE CONFIGURATION, NOR ANY WARRANTIES, EXPRESS OR IMPLIED, AS TO ITS ACCURACY, AVAILABILITY, COMPLETENESS OR TIMELINESS, OR TO THE RESULTS TO BE OBTAINED FROM THE USE OF THE PROPERTY. S&P GLOBAL PARTIES SHALL NOT IN ANY WAY BE LIABLE TO ANY RECIPIENT FOR ANY INACCURACIES, ERRORS OR OMISSIONS REGARDLESS OF THE CAUSE. Without limiting the foregoing, S&P Global Parties shall have no liability whatsoever to any recipient, whether in contract, in tort (including negligence), under warranty, under statute or otherwise, in respect of any loss or damage suffered by any recipient as a result of or in connection with the Property, or any course of action determined, by it or any third party, whether or not based on or relating to the Property. In no event shall S&P Global be liable to any party for any direct, indirect, incidental, exemplary, compensatory, punitive, special or consequential damages, costs, expenses, legal fees or losses (including without limitation lost income or lost profits and opportunity costs or losses caused by negligence) in connection with any use of the Property even if advised of the possibility of such damages. The Property should not be relied on and is not a substitute for the skill, judgment and experience of the user, its management, employees, advisors and/or clients when making investment and other business decisions.

The S&P Global logo is a registered trademark of S&P Global, and the trademarks of S&P Global used within this document or materials are protected by international laws. Any other names may be trademarks of their respective owners.

The inclusion of a link to an external website by S&P Global should not be understood to be an endorsement of that website or the website's owners (or their products/services). S&P Global is not responsible for either the content or output of external websites. S&P Global keeps certain activities of its divisions separate from each other in order to preserve the independence and objectivity of their respective activities. As a result, certain divisions of S&P Global may have information that is not available to other S&P Global divisions. S&P Global has established policies and procedures to maintain the confidentiality of certain nonpublic information received in connection with each analytical process. S&P Global may receive compensation for its ratings and certain analyses, normally from issuers or underwriters of securities or from obligors. S&P Global reserves the right to disseminate its opinions and analyses. S&P Global Ratings' public ratings and analyses are made available on its sites, www.spglobal.com/ratings (free of charge) and www.capitaliq.com (subscription), and may be distributed through other means, including via S&P Global publications and third party redistributors.

S&P Global uses generative AI to create content in accordance with our [Terms of Use](#).