illumio

EBOOK

The Breach Containment Buyer's Guide

Separating containment claims from reality and what it takes to stop the spread

Contents

When containing the breach is the only outcome that matters

In September 2025, one of the UK's most iconic manufacturers went quiet.[1]

Jaguar Land Rover's production lines stopped. Factories that normally ran around the clock shut down. Employees were told to stay home. What started as a cyber incident quickly became an operational standstill.

Days turned into weeks. Assembly lines stayed dark. Suppliers felt the impact. Logistics partners scrambled. The disruption rippled far beyond IT and deep into local economies.

The real failure wasn't the breach. It was allowing that breach to spread unchecked, turning what could have been a small security incident into a long systemwide shutdown.

Investigators later confirmed that the attackers had traveled through internal networks and production control systems that kept factories running and supply chains moving.

By the time the intrusion was fully understood, the damage was already set in motion.

In today's threat landscape, the question isn't whether you can prevent every breach. Even the most mature security programs experience intrusions. Modern enterprises are just too complex, too connected, and too heavily targeted to assume perfect prevention. Attackers need only one opening; defenders need perfection everywhere.

While breaches are inevitable, disasters don't have to be. But even as the industry moves toward containment-first thinking, a hard truth is emerging: many tools that promise containment don't actually stop breaches from spreading.

This buyer's guide breaks down why most tools that claim to "contain" breaches fall short and what it actually takes to stop lateral movement using AI-powered observability and microsegmentation.

[1] Georgia Collins (Cyber Magazine). "How JLR's Category 3 Cyber Attack Caused Production Shutdown." October 2025.

The breach containment gap

Security failures rarely start with dramatic exploits.

They start quietly, with a stolen credential, a misconfigured cloud workload, or a trusted vendor connection. Initial access is often unremarkable. What determines the outcome is what happens next.

Once inside, attackers move laterally. They map the environment, identify high-value systems, and pivot using legitimate credentials and protocols security teams often trust. By the time malicious intent is confirmed, attackers have already crossed multiple trust boundaries and expanded far beyond the initial compromise.

This is the breach containment gap.

Detection technology continues to improve. Intrusions are spotted faster, alerts fire sooner, and investigations begin earlier. Yet breaches still escalate into business crises because the most dangerous phase happens after detection, when attackers exploit implicit internal trust and move laterally without resistance.

Traditional security tools weren't built to stop this. They surface alerts and coordinate response, but they don't control how attacks spread once inside. A sensor may flag suspicious behavior on one system, but without architectural controls to block lateral movement, the blast radius keeps growing while teams investigate.

That's why organizations can detect a breach early and still suffer catastrophic impact. Detection alone doesn't limit reach. Without enforced controls on how systems communicate, every compromised asset becomes a launchpad.

In today's threat landscape, the real question isn't whether you can detect a breach. It's whether you can contain the breach before it becomes a business disaster.

Doing that requires more than faster alerts or response workflows. Containment depends on AI-powered observability to understand how an attack is spreading, paired with automated microsegmentation to enforce policy across every workload at machine speed. Together, they turn insight into control.

Containment isn't an add-on. It's the architecture that determines whether an intrusion spreads or stops where it started.

The 2025 IBM Cost of a Data Breach report found organizations that contain incidents faster consistently **reduce breach costs by *millions***, largely by limiting lateral movement and dwell time.[2]

[2] IBM. "2025 IBM Cost of a Data Breach Report." July 2025.

Where the most popular “breach containment” tools break down

The security industry knows that breach containment is the future. And many security tools claim to support containment.

Few actually do.

That’s because containment is an architectural problem, and most tools operate at the tool or workflow level. They observe activity, generate alerts, or orchestrate response steps, but they don’t control how traffic flows between systems in real time.

Here are cybersecurity tools that claim to support breach containment and why they fall short.

Endpoint detection and response (EDR)

EDR plays an important role in modern security teams. It’s great at spotting suspicious behavior on individual systems, identifying known attack techniques, and giving analysts the forensic detail they need during an investigation.

When something goes wrong, EDR is usually the first place teams look to understand what happened on a compromised endpoint.

But EDR lives at the endpoint. Breaches don’t.

Once attackers get in, they don’t sit still. They move. They reuse stolen credentials, abuse trusted relationships, and pivot across workloads using everyday protocols like RDP, SMB, WinRM, SSH, and APIs.

EDR can see what’s happening on one system, but it can’t see or control how an attack spreads across the environment. Its visibility stops at the edge of the host.

Even when EDR quickly detects malicious activity, containment tends to be narrow and reactive. Isolating an endpoint can help, but it often happens after credentials have already been stolen and used elsewhere.

By then, the attacker no longer needs that original machine. The blast radius is already growing.

That’s the core limitation. EDR is very good at explaining what happened and how it happened. It’s far less effective at stopping where the attacker goes next. And during a live breach, that difference is decisive.

EDR detects compromise at the asset level. True breach containment requires controlling communication and trust relationships across the entire environment.

Identity and access management (IAM)

Identity controls are foundational to modern security. Strong authentication, authorization, and access governance reduce risk and shut down a lot of common attack paths. That's why so many organizations now call identity the new perimeter.

The problem is what happens after access is granted.

Most identity tools are built on a simple assumption: if a user or service is authenticated and authorized, communication is allowed. During a breach, that assumption becomes a liability.

Attackers know this. Instead of breaking in with noisy exploits, they steal credentials. Once they have a valid identity, they move laterally using access paths that IAM systems are designed to permit.

In fact, in the 2025 Verizon Data Breach Investigations Report, stolen credentials were used as the initial access vector in roughly 22% of confirmed breaches. That makes credential abuse one of the most common ways attackers gain entry before moving laterally within a network.[3]

From the identity platform's point of view, nothing looks wrong. The login was successful, so it approved permissions. But now, the attacker blends in as a legitimate user.

IAM has other blind spots. Application-to-application traffic, service accounts, legacy systems, and machine-to-machine communication often bypass user authentication entirely.

These paths are common in hybrid and cloud environments and are exactly where attackers move once they're inside.

So while IAM can tell you who accessed a resource, it can't control how far an attacker can go after that access is granted.

Identity tools gate access. They don't restrict lateral movement between workloads or stop attackers from abusing trusted internal communication paths.

The limits of identity-only security

In 2023, MGM Resorts was hit by a ransomware attack that began with stolen credentials.[4] Attackers used social engineering to impersonate an employee, convinced the help desk to reset access, and logged in using valid identity credentials. From there, they moved laterally across systems that trusted that identity.

The result was a massive business shutdown. Casino floors stalled. Hotel systems went offline. Operations across multiple properties were disrupted for days.

The attack succeeded because trusted identity was treated as trusted movement. Once access was granted, there were few controls to limit where that identity could go or what systems it could reach.

[3] Verizon. "2025 Data Breach Investigations Report." April 2025

[4] Tom Singleton and Joe Tidy (BBC). "MGM Resorts: Slot machines go down in cyber-attack on firm." September 2023.

SIEM, SOAR, and incident response platforms

Security information and event management (SIEM) and security orchestration, automation, and response (SOAR) sit at the center of many security operations teams.

They pull in telemetry from across the environment, correlate signals from multiple tools, and help teams coordinate investigations and response. During an incident, they provide valuable context and keep everyone working from the same playbook.

When it comes to breach containment, the problem is what they're built (or not built) to do.

SIEM and SOAR are coordination engines, not enforcement engines. They tell teams what's happening and help manage the response. But they don't actually stop anything on their own.

To take action, these platforms rely on predefined playbooks, integrations, and often human approval. Teams must review alerts. They must make critical decisions. And they must coordinate downstream tools to respond correctly — all while the breach is still unfolding.

Even with automation in place, all of this takes time. In fast-moving attacks, especially ransomware and cloud-native intrusions, that delay gives attackers the space they need to spread.

Just as important, SIEM and SOAR don't control traffic or trust relationships directly. They depend on other tools to enforce decisions.

That means containment hinges on how quickly people and integrations can react, not on controls that operate continuously in the background.

SIEM and SOAR tools may help teams coordinate response. But they can't stop lateral movement in real time while an attack is spreading.

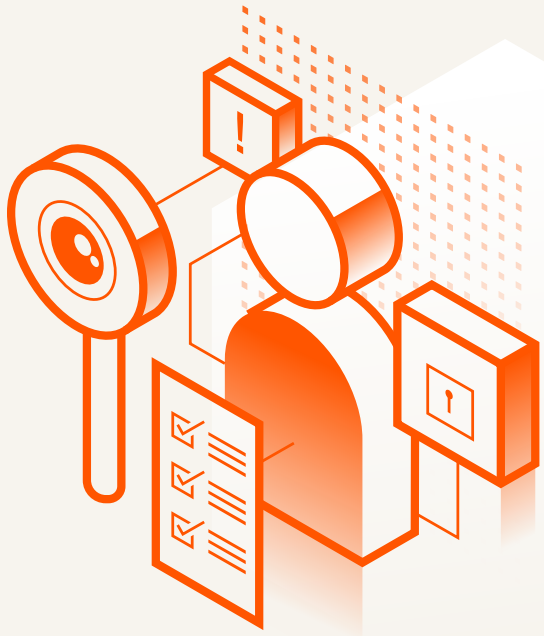

Security service edge (SSE) and secure access service edge (SASE)

SSE and SASE platforms combine networking and security controls at the edges of your network.

They're often positioned as all-in-one security architectures. And they do a solid job securing user access to applications and protecting traffic that flows through them.

The problem is where breach containment actually happens.

Most lateral movement doesn't occur at the network edge. It happens inside the environment, between workloads, services, and systems that never pass through SSE or SASE enforcement points.

Once traffic starts moving east-west in cloud environments or data centers, edge controls are no longer in the path of attack.

That creates a false sense of containment. The perimeter looks locked down, but internal communication is still implicitly trusted and largely unrestricted.

If an attacker gets inside, they can move freely without ever hitting an SSE or SASE control again.

Edge security protects entry points, but it doesn't stop internal spread.

Network firewalls and Zero Trust network access (ZTNA)

Firewalls and ZTNA are often described as containment tools because they control access to applications and resources. And they do play an important role. They reduce exposure and help block unauthorized access at key entry points.

The limitation with breach containment is where and how they operate.

Traditional firewalls are built for north-south traffic. ZTNA focuses on user-to-application access. Neither is designed to control east-west communication between workloads, services, and internal systems.

Once an attacker gets inside the environment, these controls are often no longer in the path.

In modern hybrid and cloud environments, lateral movement doesn't follow clean network boundaries. It follows application dependencies and service relationships that firewalls and ZTNA were never built to manage at scale.

Firewalls and ZTNA protect access points. But they don't control the internal pathways attackers use to spread.

Network detection and response (NDR)

NDR tools add valuable visibility by watching network traffic for suspicious behavior and known attack patterns.

They often spot lateral movement that endpoint tools miss, which is why they're sometimes positioned as a containment layer.

The problem is what happens next.

NDR is built to observe, not to enforce. Most platforms raise alerts and rely on integrations or manual action to block traffic.

In cloud-first, highly distributed environments, that handoff introduces delay and complexity. Even when blocking is possible, it's often too coarse to safely stop only the malicious paths without risking disruption.

Seeing lateral movement is useful. But during an active breach, insight without immediate control gives attackers time to keep moving while defenders investigate.

Visibility shows you where the attack is going. Containment is what stops it.

Data loss prevention (DLP)

DLP tools are built to stop sensitive data from leaving the environment. Because of that, some organizations treat DLP as a form of breach containment.

But in practice, DLP shows up very late in the attack.

Industry data shows why DLP arrives too late. The Verizon report consistently finds that data exfiltration happens after attackers have already moved laterally and established control — long after the breach has spread internally.[5]

By the time data exfiltration is detected, attackers are already disrupting systems. DLP doesn't stop ransomware from spreading, take back control of compromised systems, or prevent internal damage.

At best, it limits some types of data loss after the damage is already done.

DLP deals with the aftermath. It doesn't stop the breach from spreading.

Emerging AI security tools

AI has made security tools smarter. It has improved signal analysis, helped teams prioritize alerts, and sped up investigations. Because of that, many vendors now claim AI can also speed containment.

The catch is enforcement. Most AI tools act like a high-tech smoke detector. They can tell you there is a fire and even tell you which room it is in. But they don't always have the power to turn on the sprinklers.

They surface risk scores, summaries, or recommended next steps, then hand those decisions off to analysts or other tools. Without enforcement built directly into the infrastructure, AI insight stays advisory.

To actually contain the breach, a person or another piece of software usually has to "flip the switch." If the AI isn't directly plugged into the company's network controls, its "insight" is just a suggestion while the attack continues to spread.

Containment doesn't work on suggestions. It works when decisions are executed automatically and safely in real time.

AI can help you understand an attack faster. But without built-in enforcement, it can't stop that attack from spreading.

[5] Verizon. "2025 Data Breach Investigations Report." April 2025.

Comparing breach containment capabilities across popular security tools

Tools	Observes activity?	Understands attack paths?	Actively blocks lateral movement?
EDR	Yes	Limited (endpoint only)	No
IAM	Limited (authentication events)	No	Partial (access gating only)
SIEM / SOAR	Yes	Limited (post-correlation)	No
Incident response platforms	Limited	No	No
SSE / SASE	Limited	No	Partial (edge traffic only)
Firewalls & ZTNA	Limited	No	Partial (north-south only)
NDR	Yes	Partial (network-only view)	No
DLP	Limited (exfiltration-focused)	No	No
AI point tools	Yes	Partial	No

What modern breach containment actually requires

Containment isn't something you switch on when an incident is declared. By then, it's already too late.

Real containment is the result of how your environment is designed and run every day.

Modern attacks don't wait for clean handoffs between detection and response. Initial access, lateral movement, and impact often overlap. That means containment has to work continuously, not as a last-ditch action.

That gap is why containment so often stays theoretical.

This is where AI-powered observability and microsegmentation come together. Observability shows how an attack can move. Segmentation makes it possible to stop that movement, everywhere it matters.

With both, containment becomes something teams can actually do.

From guesswork to guardrails: three pillars of real-time containment

To contain a breach, security teams need clear, real-time answers to three questions.

- **What is communicating with what, right now?** Teams need an accurate view of how workloads and services are actually talking to each other, not how diagrams say they should. During an incident, yesterday's assumptions fall apart fast.
- **Which of those paths actually matter?** Not every connection is risky. The challenge is knowing which paths an attacker can use to spread and which ones the business depends on. Without that context, teams either block too much or hesitate to block anything at all.
- **Can we stop those paths immediately, without breaking the business**? This is where most containment plans fail. Blocking the wrong thing causes outages. Moving too slowly gives attackers time to pivot again. Containment only works when enforcement is fast, precise, and safe, even while the investigation is still ongoing.

Most security tools can help answer one of these questions. Very few answer all three together.

How Illumio delivers containment by design

Most security tools treat containment as something you attempt after an incident is confirmed. In this process, an alert fires, a ticket opens, a playbook runs, and teams scramble to slow the damage.

Illumio takes a different approach.

Illumio is built on the idea that breach containment should be a continuous operating state, not an emergency response.

Instead of waiting for humans to react, Illumio assumes attackers will eventually get in and designs controls that limit how far they can go from the start.

The goal is simple and practical: reduce blast radius automatically, even while an incident is still unfolding.

That is why the Illumio platform pairs two capabilities that are rarely designed to work together. AI-powered observability with Illumio Insights to understand how attacks move, and microsegmentation with Illumio Segmentation to stop that movement everywhere it matters.

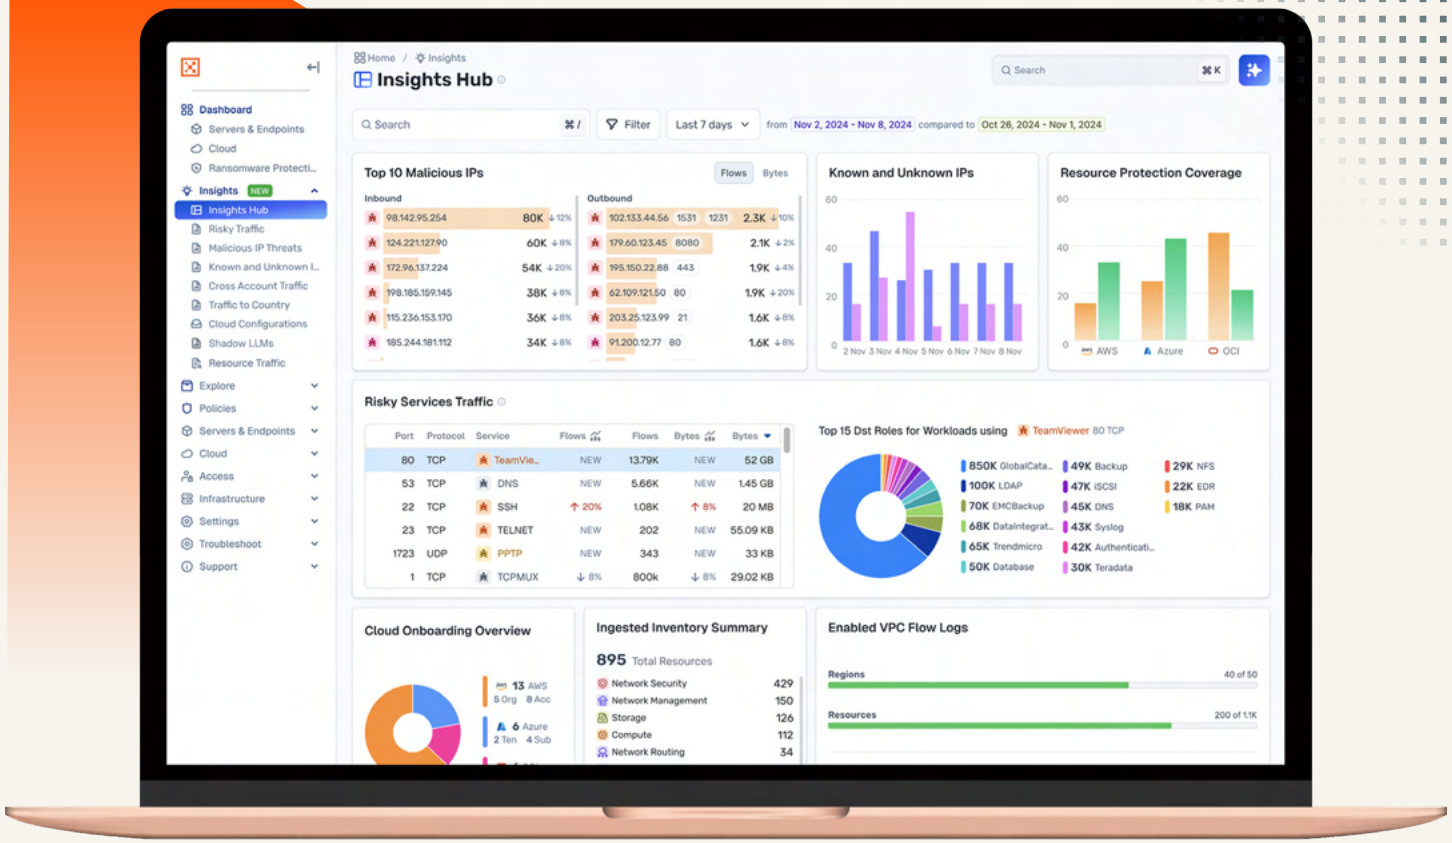

Illumio Insights: AI observability focused on lateral risk

Illumio Insights delivers cloud-native, AI-powered observability built specifically to answer one question most tools struggle with: how does risk move through the environment?

Instead of flooding teams with raw alerts or isolated signals, Insights focuses on communication behavior. It analyzes traffic patterns, relationships between workloads, and changes in behavior to surface the paths attackers are most likely to exploit.

The result is not more noise, but clearer context.

During an active incident, Insights helps teams move beyond detection and into understanding. It highlights risky communication paths, exposed dependencies, and potential blast radius so teams can see how an attack could spread before it does.

Insights answers the questions security teams struggle to answer under pressure:

- Where can this attack spread next?
- Which assets are actually exposed, not just noisy?
- Which communication paths matter most right now?

This clarity changes how teams respond. Instead of chasing alerts or guessing where to act, they can focus on the paths that truly drive risk.

That shift from reactive investigation to informed decision-making is what makes containment possible.

Illumio Segmentation: enforcing containment at machine speed

Understanding risk is only half the problem. Stopping it requires enforcement that works everywhere, all the time.

Illumio Segmentation enforces Zero Trust principles directly at the workload level. Policy is distributed across the environment rather than centralized in a single control point.

Enforcement happens close to the workloads themselves, without requiring network redesigns, brittle appliances, or privileged access into systems.

Because enforcement is built into the architecture, containment doesn't depend on tickets, playbooks, or manual approval. Lateral movement gets blocked instantly, even while investigations are still ongoing.

This is where Insights and Segmentation come together.

When Insights identifies risky paths or expanding blast radius, Segmentation ensures those paths can be shut down safely and precisely. Communication is restricted to what is required for the business to operate and nothing more.

The blast radius shrinks before attackers can pivot again.

The result is containment that is fast, targeted, and repeatable. Not a one-time response, but a continuous control that holds even when everything else is under pressure.

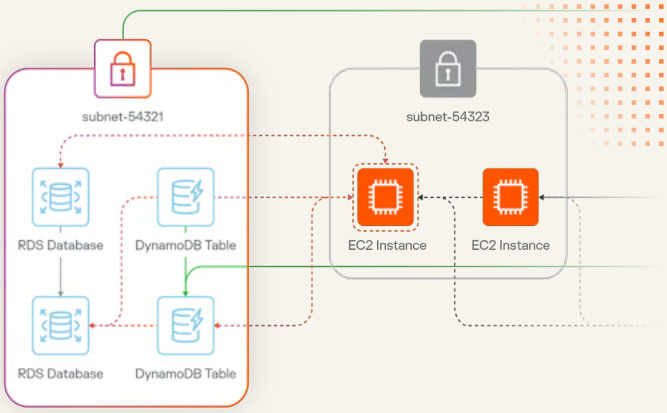

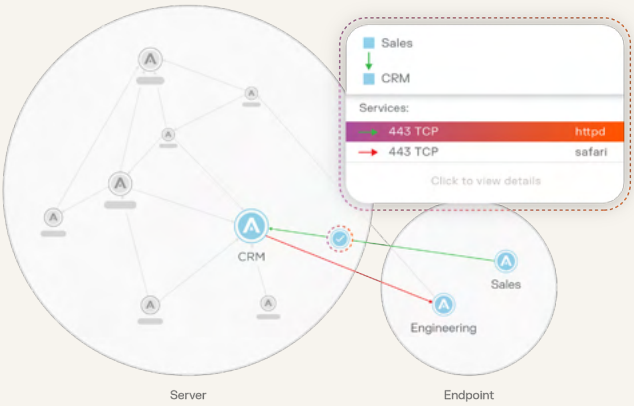

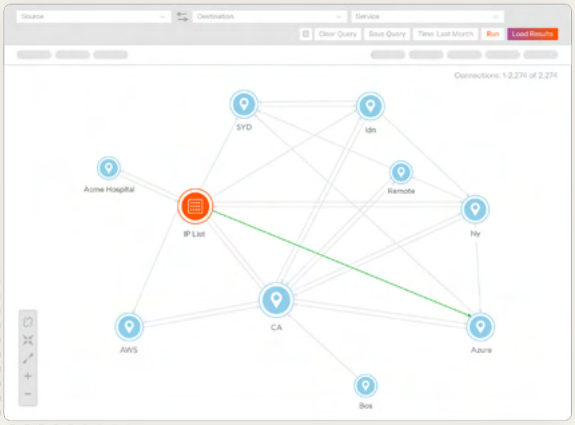

Containment is the new security baseline

In a world of inevitable breaches, security success can't be measured in prevention alone. Breaches are inevitable, and modern cybersecurity is all about reducing their effect.

Security leaders need to prove that when something goes wrong, it stays small. There needs to be a way to keep systems available, protect critical data, and ensure the business can keep running.

The industry may be waking up to containment-first security, but not all "containment" is created equal. In fact, some approaches fail precisely when they matter most.

Illumio delivers it by design.

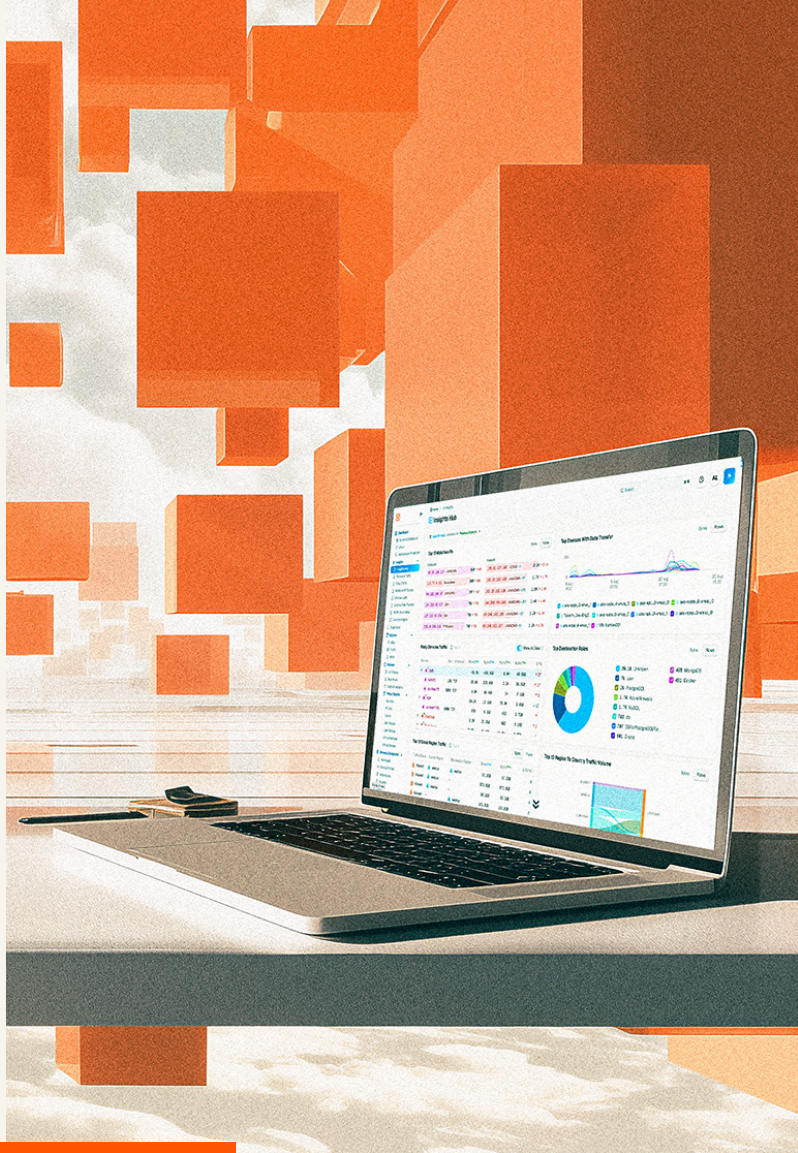

See breach containment in action

Get started with Illumio Insights free for 14 days.

illumio.com/insights-free-trial