

THE CONTAINMENT PLATFORM

How Cloud Native Security Fabric Closes
the Architectural Divide

An Executive Perspective



Executive Summary

The Containment Era demands a platform that does four things: enforces default-deny communication at every workload, operates across all compute models without agents, propagates a single policy universally in subseconds, and measures success by blast radius rather than detection speed. This paper introduces the architecture that delivers all four.

The Cloud Native Security Fabric (CNSF) is Aviatrix's answer to the Architectural Divide. As documented in Paper 1, the threat model outgrew the architecture. Three forces are converging in what we call the Toxic Combination: attackers industrializing, AI putting frontier offensive capability into millions of hands at consumer cost, and a cloud that is insecure by default by design. Three operational gaps compound one another.

Chokepoint Security cannot close them. A new architectural category is required. CNSF is that category. It is not another tool to bolt on at the edge. It is a new foundation for security embedded directly within the cloud fabric itself. It delivers a real-time, policy-driven enforcement layer that inspects, segments, and secures communication between every cloud workload, with particular leverage on the AI workloads that the toxic combination concentrates risk on.

CNSF is not an evolution of existing security architecture. It is the architecture the Containment Era requires: embedded in the fabric, enforcing at every workload, governing every communication path.

Defining Containment

The Containment Era demands a platform that delivers containment. But what, precisely, is containment?

Across the industry, the term is used broadly and without rigor. Every vendor claims to deliver it. NGFW vendors say their firewalls contain threats. Microsegmentation vendors say their agents contain lateral movement. CNAPP vendors say their posture tools contain risk. Without a formal definition, these claims cannot be evaluated—and the buyer has no rubric for distinguishing marketing from architecture.

This paper provides that definition.

“When prevention fails and detection is too slow, containment decides whether the incident becomes a breach.”

That is the executive formulation. The formal definition is precise:

Containment is the architectural enforcement of explicit communication policy at every workload—governing what it can reach and what can reach it, at the granularity of workload identity and protocol—on every path available to it, independent of whether a compromise has been detected.

Every clause is load-bearing. “Architectural enforcement” means this is not a dashboard finding or a recommendation—it is enforcement in the data plane. “At every workload” means enforcement at the workload, not at a centralized appliance. “At the granularity of workload identity and protocol” means L7 identity—the boundary is “myproject@github.com,” not “github.com.”

“On every path available to it” means path-complete: if a communication path exists, policy governs it. “Independent of whether a compromise has been detected” means containment is always on—it is the architectural state that limits what incident response needs to clean up.

Five Testable Properties

Any architecture claiming to deliver containment must demonstrate all five:

- 1. Path-complete.** Enforcement governs every communication path available to a workload, including those that bypass centralized inspection points.
- 2. Identity-aware at L7.** Policy operates at the granularity of workload identity and application protocol—not IP addresses and ports.
- 3. Detection-independent.** Enforcement holds before, during, and after a breach, without requiring that the breach first be detected.
- 4. Compute-model agnostic.** Enforcement reaches every workload type—VMs, containers, serverless functions, managed services, partner VPCs—without requiring agent installation on each.
- 5. Universally propagated.** A single policy change enforces across providers, regions, and clusters within subseconds.

These five properties are not aspirational. They are the minimum threshold for containment.

Identity governs who can act. Containment governs what they can reach. The two are complementary, not competing: identity without containment leaves the blast radius unbounded when credentials are compromised, and containment without identity cannot distinguish between legitimate and illegitimate actors. Both are required.

The Evaluation Framework

Defining the five properties is necessary. But a definition without application is academic. The value of these properties is that they are testable—any architecture can be measured against them, and the results are binary. A property is either demonstrated or it is not.

Before applying them, each property in plain language:

- 1. Path-complete.** Every communication path governed. No bypasses.
- 2. Identity-aware at L7.** Workload identity, not IP addresses.
- 3. Detection-independent.** Enforced before, during, and after breach.
- 4. Compute-model agnostic.** VMs, containers, serverless. No agents.
- 5. Universally propagated.** One policy. Every cloud. Subseconds.

Four architectural approaches dominate the market today. Each was designed to solve a legitimate problem. None was designed to deliver containment as defined in this paper. The table below applies all five properties to each.

Five Testable Properties: Architectural Comparison

Applied to four dominant security architecture models

Property	Centralized NGFW	CNAPP / Posture	Agent-Based Microseg	Cloud Native Security Fabric
1. Path-complete	✗	✗	✗	✓
2. Identity-aware at L7	✓	✗	✓	✓
3. Detection-independent	✓	✗	✓	✓
4. Compute-model agnostic	✗	✗	✗	✓
5. Universally propagated Properties demonstrated	2 of 5	0 of 5	2 of 5	5 of 5

Figure: The Five Testable Properties applied to four architectural approaches. No alternative delivers more than two.

The results are architectural, not competitive. Each failure traces directly to the definitions above. A centralized NGFW is not path-complete because it governs only traffic that routes through it—every bypass path is ungoverned. CNAPP tools deliver none of the five because they observe, score, and recommend; observation is not enforcement. Agent-based microsegmentation enforces only where the agent runs—serverless functions, managed services, and pods exiting through node NAT are unreachable. No current alternative demonstrates more than two of the five properties.

The diagnostic is simple: ask every vendor in your environment to demonstrate all five properties. Document what they cannot. That gap is the Architectural Divide in your organization—and it is the space a Containment Platform must close.

The sections that follow demonstrate how the Cloud Native Security Fabric delivers each one.

What Cloud Native Security Fabric Is

Aviatrix is defining the Cloud Native Security Fabric: a new category of security designed to enforce trust dynamically across the entirety of an organization's cloud workloads. CNSF delivers what existing architectures lack: a real-time, frictionless, intent-based policy enforcement layer embedded inside the cloud fabric itself.

It instantiates the missing posture of Zero Trust by inspecting, segmenting, and securing cloud workloads as they communicate. Encryption, segmentation, and identity-aware controls all operate inline. With CNSF, security becomes embedded, not bolted on. Dynamic and distributed, moving with the workload. Agentless and inline, requiring no agents to deploy. Real-time, enforcing as workloads connect. Pervasive, spanning cloud, data center, edge, and AI workload landscapes. And developer-transparent, requiring no application code changes.

How CNSF Closes the Architectural Divide

The Architectural Divide manifests as three gaps. CNSF closes each one architecturally.

AVIATRIX

Chokepoint Security vs. Containment Architecture

A chokepoint governs the traffic that routes through it. Communication Governance governs every path.

X CHOKEPOINT SECURITY

✓ CONTAINMENT ARCHITECTURE



<p>X K8s Pod Egress Exits via Node NAT</p> <p style="text-align: right;">UNGOVERNED</p>	<p>✓ K8s Pod Egress Enforced at Pod</p> <p style="text-align: right;">GOVERNED</p>
<p>X Serverless Functions Exits via Provider NAT</p> <p style="text-align: right;">UNGOVERNED</p>	<p>✓ Serverless Functions Enforced at Function</p> <p style="text-align: right;">GOVERNED</p>
<p>X East-West VPC Traffic Direct Peering</p> <p style="text-align: right;">UNGOVERNED</p>	<p>✓ East-West VPC Traffic Enforced at Workload</p> <p style="text-align: right;">GOVERNED</p>
<p>X New VPC / Policy Gap No Routing Configured</p> <p style="text-align: right;">UNGOVERNED</p>	<p>✓ New VPC / Policy Gap Auto-Propagated</p> <p style="text-align: right;">GOVERNED</p>
<p>Governs only traffic that routes through it</p>	<p>Governs every workload, every path, every region</p>

Detailed Comparison

Dimension	Chokepoint Security	Containment Architecture
Enforcement Point	Central transit firewall	Every workload
K8s Pod Egress	Invisible	Governed
Serverless Functions	Invisible	Governed
East-West Traffic	Depends on routing	Governed
Policy Propagation	Hours/days per device	Subsecond, universal
Blast Radius	Network-wide	Single workload

The distinction is not "egress filtering vs. no egress filtering." The distinction is where the enforcement lives.

Figure: Chokepoint Security vs. Containment Architecture

Closing the Fragmentation Gap

The Fragmentation Gap exists because security ownership is distributed across tools that were never designed to share a common operating model. CNSF provides the unifying layer. A single control plane spans AWS, Azure, GCP, and OCI. Policy is defined once, in business-relevant language, and enforced everywhere, without rewriting rules for every provider, without translating security intent into provider-specific constructs, without the drift that inevitably occurs when the same intent must be expressed five different ways.

This is accomplished through SmartGroups: dynamic constructs that automatically classify workloads based on cloud-native metadata, tags, and resource attributes. As infrastructure changes, SmartGroups adapt automatically. No manual rule updates. No stale policies. No gaps. When a CISO asks whether their policy covers the new VPC their platform team spun up yesterday, the answer is yes, because the policy is defined by workload identity, not by infrastructure location.

Closing the Runtime Enforcement Gap

The Runtime Enforcement Gap exists because the industry invested in visibility without enforcement. CNSF sits in the data path. It does not advise. It does not alert and wait. It enforces, in real time, at the workload, before lateral movement occurs, before data leaves the environment, before the alert reaches a human console.

Aviatrix operates inline, not as a proxy, not as a sidecar, not as a centralized inspection point that traffic must be redirected through. Policy enforcement happens directly in the data path at the workload level, enabling real-time, agentless enforcement across multi-cloud and hybrid environments without the latency, bottlenecks, or single points of failure inherent in centralized approaches.

In the data path, by design. Not observing from the edge. Not advising after the fact. Enforcing at the workload, in real time, on every communication path.

Closing the Ownership Gap

The Ownership Gap exists because security ownership is distributed across teams that were never designed to share a common enforcement plane. CNSF provides that plane. Platform engineering, application teams, security teams, and cloud operations all operate on a shared architectural layer with a common source of truth for policy. When an incident occurs, the question of which team owns the workload, which tool has the telemetry, and which console to use is answered by the fabric itself. Every DENY log entry is a forensic record: which workload, which destination, what time. The architecture provides both the enforcement and the evidence.

Why CNSF is the Architecture for AI Workloads

Paper 1 in this series identifies AI workloads as the bullseye of the toxic combination. AI workloads are ephemeral, highly privileged, and rapidly shipped. The architecture that secures them must match their shape. CNSF was built for exactly this surface.

Ephemeral workloads cannot host agents. CNSF does not require them. Enforcement is inline at the fabric, not on the workload, so a serverless function that lives for three seconds is governed identically to a virtual machine that lives for three years. Highly privileged workloads concentrate non-human identity risk. CNSF expresses policy in terms of workload identity, not IP address, so an AI agent that needs to reach a specific service can be granted exactly that path and nothing more. SmartGroups make this expressible in business terms rather than network primitives.

Rapidly shipped workloads outrun traditional security review. CNSF's policy-as-code model lets developers declare connectivity requirements alongside their application manifests, validated in CI/CD before deployment. Security becomes a stage in the pipeline, not a ticket queue that slows shipping.

The same properties that make CNSF the right architecture for AI workloads make it the right architecture for cloud workloads generally. AI workloads are the most acute case. They are not the only case.

The Lineage and the Leap

The concept of workload-level security enforcement is not new. Previous platforms proved that lateral-movement containment at the workload layer was valuable. That lineage matters, and it would be dishonest to pretend it does not exist.

But the Containment Era demands more than lateral-movement governance between workloads inside the network. It demands Communication Governance at every workload, governing not just which workloads can talk to each other internally, but which workloads can reach any destination on the internet. It demands coverage across compute models that agent-based approaches cannot reach: serverless functions, managed Kubernetes pods, PaaS services. It demands a single policy that propagates across every cloud provider in subseconds, not per-host deployments that take weeks to roll out.

The distinction is architectural scope. Previous solutions answered the question: can this workload communicate with that workload inside our network? CNSF answers a broader question: for every workload in our environment, what can it communicate with, internally and externally, on every path available to it? That includes the Kubernetes pod path that bypasses the transit inspection point. The serverless function that never touches customer infrastructure. The east-west traffic between VPCs that peers directly. These are the paths the Cascade exploited, and the paths that the previous generation of tools was never designed to govern.

Intent-Based Policy: Security in Business Terms

Traditional network security policy is built on IP addresses, port numbers, and CIDR blocks, constructs that made sense when infrastructure was static. In cloud environments, workloads are ephemeral, IP addresses are dynamically assigned, and infrastructure scales in minutes. Policies anchored to network primitives are perpetually stale.

Aviatrix takes a fundamentally different approach. With Intent-Based Policy, security teams define what should be allowed based on what a workload is, its application, environment, compliance requirements, and business function, rather than where it sits on the network.

This intent is expressed through SmartGroups. Because Aviatrix policies are expressible as code, they integrate directly into CI/CD pipelines, version-controlled, peer-reviewed, and deployed alongside the infrastructure they protect. Security becomes part of the workflow, not a gate that slows it down.

The result is Communication Governance: not a one-time policy, but continuous verification. Given current identity and context, should this workload be allowed to communicate with that destination? That question, answered at every workload, on every path, at all times, is what distinguishes a Containment Platform from Chokepoint Security.

From Default-Deny to Operational Reality

The most common response to default-deny communication is also the most legitimate: it will break production. Every network operator who has attempted to move from permit-all to deny-all at scale knows the operational risk. CNSF was designed with this reality as a first principle.

Intent Replaces IP Complexity

The operational brutality of traditional default-deny comes from the policy model, not from the concept itself. When enforcement requires manually maintaining thousands of IP-based allow rules across dynamic cloud infrastructure, default-deny is unmanageable. SmartGroups eliminate this entirely. Policy is expressed in business terms, “Production-App communicates with Database on HTTPS,” and the platform translates that into enforcement. When infrastructure changes, policies adapt automatically.

Three stakeholder groups experience this differently. Cloud architects gain operational clarity, no more updating rules every time an application team scales a service. Security operations teams gain consistent Zero Trust posture without slowing development. DevOps teams gain velocity, as long as infrastructure is tagged correctly, the network automatically grants the connectivity required.

Policy-as-Code in the Developer Workflow

CNSF embeds security into the workflows developers already use. Developers declare their connectivity requirements using custom resource definitions, the same Kubernetes manifests they already write for their applications. Before deployment, guardrails validate every policy against organizational standards. If a policy fails validation, the pipeline blocks the merge. Security becomes a CI/CD stage, not a ticket queue.

At cloud runtime, the fabric processes the validated policies and enforces them at the gateway level. Two layers, two enforcement points, defense in depth.

Observe Before Enforce

New policies deploy in observe mode, logging what would be denied without actually blocking traffic. Teams validate their SmartGroup definitions and connectivity declarations against real production traffic patterns. When confidence is established, enforcement is activated surgically, per workload group, not as a big-bang cutover.

This is the safety mechanism that makes the difference between “we turned on default-deny and broke production” and “we observed for two weeks, tuned our policies, and then enforced with confidence.”

The Contain-Detect-Eliminate Progression

The Containment Era does not render the Detection Era irrelevant. It completes it.

For fifteen years, the detection industry has been solving the right problem with the wrong precondition. Detection tools are built to find anomalies. But in an unbounded environment, every communication path is potentially legitimate and potentially malicious. The signal-to-noise ratio is crushing. SOC analysts drown in alerts because there is no architectural boundary separating normal from abnormal communication.

Containment changes that equation. When every workload has explicit communication boundaries, every communication outside those boundaries is anomalous by definition. Detection tools no longer need to distinguish malicious from legitimate across an unbounded surface. They need to investigate the bounded set of DENY events and policy violations that the fabric surfaces automatically.

The progression is structural. Contain first: bound the blast radius so that any compromise is limited to a governed space. Detect within the contained space: use the enforcement telemetry as a high-fidelity signal for investigation. Eliminate with precision: because the compromised workload is contained, response is targeted rather than reactive.

Containment does not replace detection. It is the precondition that makes detection effective and elimination possible. Without containment, detection searches an unbounded environment. With containment, every DENY log is a lead.

Why Containment Holds Against the Toxic Combination

Containment is the only control that holds equally against all three legs of the toxic combination. It does not care how the workload was compromised. It cares what the workload can reach.

Against attacker industrialization, containment bounds the blast radius regardless of the sophistication of the supply chain attack. A compromised package that runs as expected still cannot reach a destination it has no policy to reach.

Against AI-democratized exploitation, containment is independent of the volume or speed of attack attempts. The wall does not care how many actors are trying the door, or how fast they are. The wall holds.

Against the cloud's permit-all defaults, containment replaces the default. Every workload is governed by explicit policy. The tangled web of interconnected workloads with direct internet egress becomes a deliberate communication graph the defender controls.

And against the credential vector, the 82% of intrusions that ride valid credentials through legitimate channels, containment is the only control that operates without needing to distinguish legitimate from malicious. A valid credential used against a workload that cannot reach the data store is a contained event regardless of whether anyone detects the credential abuse.

The From/To

FROM: Detection Era	TO: Containment Era
Keep bad things out (perimeter frame)	Limit what bad things can reach once inside (containment frame)
Detect attacks fast (detection metric)	Minimize blast radius (containment metric)
Chokepoint Security covers some paths	Communication Governance covers all paths
Trust perimeter + trust adjacent = security	Every workload has explicit communication boundaries
One breach = network compromised	One breach = one workload contained
Contain and detect as parallel efforts	Contain first, then detect within governed space, then eliminate

Building the Fabric Control Layer

Aviatrix is building the Cloud Native Security Fabric as both a platform and an ecosystem, anchoring the real-time control layer that enforces and observes trust across clouds, workloads, and services. The platform delivers core cloud runtime enforcement, while the ecosystem integrates identity, posture, and AI signal providers through extensible APIs.

The Containment Era is not a prediction. It is a structural shift already underway. The toxic combination of industrialized attackers, AI-democratized offensive capability, and an insecure-by-default cloud is the present landscape. The Architectural Divide is the structural condition. CNSF is the architecture that closes the divide and delivers containment at cloud scale.

The space between every workload is the battlefield. CNSF is the architecture that governs it.

Aviatrix is pioneering the Cloud Native Security Fabric, a new architectural category purpose-built to deliver Containment Platform capabilities at cloud scale.

This is Paper 2 of 4 in the Containment Era series.

[Paper 1: The Containment Era – Why the Threat Model Outgrew the Architecture.](#)

[Paper 3: 144 to 1 – Why Every Workload in Your Cloud Is Already Exposed.](#)

[Paper 4: The Priority Inversion – Why the SANS Mythos Report Has the Order Wrong.](#)

About Aviatrix

Aviatrix® is pioneering the Cloud Native Security Fabric (CNSF) – the architecture the Containment Era requires. CNSF governs every workload communication path across every cloud, every VPC, every Kubernetes cluster, and every serverless function, from a single policy plane. One rule. Universal propagation. Enforced at the workload, not at a chokepoint. Trusted by more than 500 of the world's leading enterprises. For more information, visit aviatrix.ai.