AVIATRIX®

Architecture Won. A Fortune Global 500 Enterprise Has the Block Logs to Prove It.

Security architecture is a bet you make before the threat has a name. You design the controls, build the policy, and trust that when something real comes, the foundation holds.

A $46 billion Fortune Global 500 enterprise–580,000 employees, 3,000+ locations, 30 countries–made that bet. Last week, it paid off.

When the LiteLLM supply chain attack surfaced, they did not need to build a new control, stand up a new tool, or wait for an emergency change window. The architecture was already there. The policy was already enforced. All they had to do was add four IPs to a rule that was already running in production.

That is what architecture winning looks like. Not a scramble. A decision made months ago, doing exactly what it was built to do.

What Happened

When the LiteLLM supply chain attack surfaced, the security team at this Fortune Global 500 enterprise identified suspicious IP addresses associated with the campaign and reached out to their cloud infrastructure team: can we block these?

The answer was immediate. They already had an Aviatrix Distributed Cloud Firewall with a Global IP blocklist policy in place, following standard deployment best practices. The infrastructure engineer added four IPs to the existing rule:
46.151.182.203 83.142.209.11 83.142.209.203 45.148.10.212

That last one is the confirmed LiteLLM credential exfiltration server documented in our supply chain post.

He made the change without knowing what any of them were connected to. That detail matters more than anything else in this story.

When we followed up after publishing the LiteLLM story, what we heard back was worth sharing: not only had they added the block, they had actual hits on the rule. Traffic had reached those addresses. The block was not hypothetical. And the infrastructure engineer had no idea what he had been blocking until we connected it to the LiteLLM campaign.

Why That Last Detail Matters

The infrastructure engineer did his job without needing to understand the threat. The security team identified the IPs and escalated. He blocked them using a policy framework that was already operational. The architecture handled the rest.

This is what good security posture looks like in practice. It is not a team that reads every threat intelligence report and manually responds to each one. It is a team that has built the right controls so that response is operationally fast and does not require deep attacker context to execute.

The alternative plays out like this: no enforcement architecture in place, security team identifies the suspicious IPs, requests a firewall change, that change goes through a ticketing process, gets reviewed, gets approved, gets deployed. The credentials may already be gone. The window that matters in a supply chain exfiltration is minutes, not days.

This enterprise closed that window. Not because they had the fastest SOC or the best threat intelligence. Because they had the right architecture already running.

Even Partial Deployment Stopped a Live Attack

This enterprise is building toward where every security team wants to be: workloads fully mapped, an explicit allow list in place, and everything else denied by default. They are not there yet. But here is what matters: they did not need to be. A blocklist policy running in production–one piece of a broader enforcement architecture–was enough to stop a live supply chain exfiltration. That is the power of getting the architecture right early.

The full zero-trust destination–deny-by-default egress, microsegmentation between workloads, east-west traffic governance–eliminates the exfiltration path entirely, regardless of what the attacker knows or what lands on the host. This enterprise is building toward that. But the foundation they already have in place stopped a real attack from a real threat actor. That is not a starting point to apologize for. That is architecture winning.

The Architectural Divide

Our LiteLLM and TeamPCP ransomware posts make the case that the right response to supply chain attacks is architectural: control what your workloads can reach, so that even a fully compromised host cannot exfiltrate credentials. TeamPCP has already cascaded from Trivy to LiteLLM to Telnyx. The campaign is ongoing and the pattern is the same every time: compromise a trusted dependency, exfiltrate credentials, move laterally. The organizations that survive this pattern are the ones with enforcement built into the cloud network fabric–not bolted on at the endpoint or the application layer. This enterprise did not set out to stop the LiteLLM attack. They built the right architecture, their security team identified suspicious indicators, and the block landed before anyone connected the dots.

That is the model. The control does not need to know the attacker's name. It needs to be in place before the attack arrives.

If you want to know whether your environment would have stopped this exfiltration–or if your workloads have paths to the internet that they should not have–the free Workload Attack Path Assessment is the right starting point.