AVIATRIX®

THREAT OVERVIEW

LiteLLM Supply Chain Attack: Defending Against TeamPCP PyPI Compromise with Aviatrix®

At a Glance

TeamPCP poisoned litellm v1.82.7 and v1.82.8 on PyPI via CI/CD credential theft from a compromised Trivy scanner. The packages steal cloud credentials, SSH keys, Kubernetes secrets, and AI API keys, exfiltrating them encrypted to attacker infrastructure. v1.82.8 runs on every Python startup via a .pth file - no import needed.

Approach

Trivy v0.69.4-v0.69.6 poisoned to steal CI/CD secrets. LiteLLM ran Trivy unpinned, leaking PyPI credentials from CircleCI. Attackers registered litellm.cloud (typosquat) and published malicious packages exfiltrating credentials encrypted to models.litellm.cloud (45.148.10.212).

Related CVEs

- N/A - No CVE assigned at time of publication. Supply chain attack via compromised CI/CD credentials, not a vulnerability-based exploit.

Mitigations

- Remove litellm v1.82.7 and v1.82.8
- Find and remove litellm_init.pth
- Rotate all credentials
- Block egress to 45.148.10.212
- Pin Trivy to v0.69.3 in CI/CD

Understanding the LiteLLM Supply Chain Attack

The LiteLLM PyPI compromise represents a concerning evolution in TeamPCP's tactics: a pivot from direct cloud infrastructure attacks to software supply chain infiltration targeting the AI ecosystem. LiteLLM is a widely deployed open-source AI proxy library used by organizations to route requests to OpenAI, Anthropic, and other model providers. By poisoning it, TeamPCP gained a foothold on any host where the malicious versions were installed - including cloud-hosted AI inference services, development environments, and CI/CD runners.

So how does it work?

TeamPCP first compromised Aqua Security's Trivy vulnerability scanner (v0.69.4–v0.69.6) by force-pushing malicious tags across all 76+ trivy-action releases and publishing poisoned Docker images. The poisoned Trivy binary exfiltrated environment variables and secrets via Cloudflare Tunnels. LiteLLM's CI/CD pipeline ran Trivy via apt without version pinning; when a CircleCI run occurred during the compromise window, the pipeline's PYPI_PUBLISH_PASSWORD and a GitHub PAT were stolen. The attackers then registered litellm.cloud (typosquatting the legitimate litellm.ai) on 45.148.10.212, a VPS shared with the Trivy campaign, using models.litellm.cloud as the exfiltration endpoint. On March 24, 2026, they published litellm==1.82.7 and litellm==1.82.8 to PyPI. Version 1.82.7 embeds a credential stealer in proxy_server.py triggered on proxy import. Version 1.82.8 plants a litellm_init.pth file in Python site-packages that executes the payload on every Python interpreter startup - no import needed. The payload collects SSH keys, environment variables, cloud credentials (AWS/GCP/Azure), Kubernetes configs, AI API keys, database passwords, SSL private keys, and shell history; encrypts all of it with AES-256-CBC and a hardcoded RSA-4096 attacker public key; then POSTs the bundle to models.litellm.cloud.

By targeting the AI toolchain, TeamPCP gains access to high-value secrets - cloud IAM credentials, Kubernetes tokens, and AI provider API keys - that are routinely co-located with LiteLLM deployments. This highlights the need for network-level egress controls to block credential exfiltration, supply chain integrity verification for PyPI packages, and cloud native metadata service controls (AWS IMDSv2, GCP metadata concealment) to limit IAM credential harvesting.

Key Threats and Vulnerabilities and How Aviatrix Solves Them

CI/CD Supply Chain Compromise	TeamPCP poisoned Trivy security scanner releases (v0.69.4-v0.69.6) to harvest CI/CD secrets from any pipeline running Trivy without version pinning, stealing PyPI publishing credentials from LiteLLM's CircleCI environment.	Aviatrix DCF egress controls limit what CI/CD runners can reach, reducing the blast radius of compromised pipeline tools. CoPilot flow logs provide visibility into unusual outbound connections from build infrastructure, enabling early detection of credential exfiltration.
PyPI Package Poisoning	Attackers used stolen credentials to publish malicious litellm versions (v1.82.7, v1.82.8) to PyPI. Version 1.82.8 plants a .pth file that executes the payload on every Python startup, affecting all Python processes on the host.	Aviatrix DCF blocks the network-layer exfiltration path: even if a compromised package is installed, outbound connections to the C2 IP (45.148.10.212) and typosquat domain (models.litellm.cloud) are denied before stolen credentials leave the environment.
Cloud Credential Theft	The malicious payload harvests AWS, GCP, and Azure credentials from environment variables, config files, and the cloud metadata service (169.254.169.254), enabling attacker access to cloud accounts and resources.	Metadata service traffic (169.254.169.254) is routed at the hypervisor layer, not the network - use AWS IMDSv2 and GCP metadata concealment. Aviatrix DCF blocks the exfiltration POST to 45.148.10.212 and models.litellm.cloud, stopping stolen credentials at the network layer.
Kubernetes Secret Exposure	The payload collects kubeconfig files and in-cluster service account tokens, enabling lateral movement across Kubernetes environments and multicloud accounts.	Aviatrix SmartGroups isolate Kubernetes workload VPCs from sensitive control planes. DCF policies enforce zero trust boundaries, limiting attacker movement after credential theft. CoPilot provides unified multicloud visibility.
AI API Key Theft	LiteLLM deployments routinely co-locate AI provider API keys (OpenAI, Anthropic, etc.) with the proxy process. The payload enumerates all environment variables, exfiltrating these high-value keys along with cloud credentials.	Aviatrix blocks the exfiltration path at the network layer before stolen keys reach attacker infrastructure. DCF rules denying outbound traffic to 45.148.10.212 and ThreatGroup feeds ensure that even a fully compromised host cannot successfully complete the exfiltration POST request.
Encrypted Data Exfiltration	Stolen credentials are encrypted with AES-256-CBC and a hardcoded RSA-4096 attacker public key, then POST'd via curl to models.litellm.cloud. Encryption means only the attacker can decrypt the bundle - detection must happen at the network layer.	Aviatrix DCF with ThreatGroup automatically blocks egress to known malicious IPs. Explicit deny rules for 45.148.10.212 block the exfil endpoint. CoPilot flow logs capture every attempted connection as forensic evidence, even for blocked traffic.
Persistent Execution (v1.82.8)	Version 1.82.8 drops litellm_init.pth in Python site-packages, causing the credential-stealing payload to execute on every Python interpreter startup - including unrelated scripts, tools, and scheduled jobs running on the same host.	Aviatrix DCF rules apply at the network layer regardless of how the malicious process is triggered. Blocking egress to 45.148.10.212 stops exfiltration even when the payload runs from unexpected Python processes. CoPilot anomaly detection alerts on unusual outbound patterns from application VPCs.

Aviatrix: Comprehensive Defense Against the LiteLLM Supply Chain Attack

The LiteLLM supply chain compromise demonstrates that AI infrastructure is now an active target for sophisticated threat actors. TeamPCP has proven capable of multi-stage attacks - poisoning upstream security tooling to compromise downstream software publishers - representing a new class of threat to organizations running AI workloads in the cloud. Defending against this requires network-layer controls that operate independently of host integrity: even a fully compromised host cannot exfiltrate credentials if the egress path is blocked. By partnering with Aviatrix, organizations deploying AI infrastructure can take control of their cloud security with solutions that deliver:

✓ Egress Control and Exfiltration Prevention

Aviatrix DCF blocks outbound connections to TeamPCP C2 infrastructure (45.148.10.212) and automatically blocks new malicious IPs via ThreatGroup feeds - stopping credential exfiltration at the network layer even when a compromised package is installed on a host.

✓ Egress Blocking: Defense-in-Depth Against Credential Theft

Cloud metadata service traffic (169.254.169.254) is routed at the hypervisor layer and never traverses spoke gateways - DCF cannot intercept it directly. The correct control is cloud native: enforce AWS IMDSv2 to require session-oriented credential requests, and GCP metadata concealment to restrict access. Aviatrix DCF's critical role is blocking the exfiltration path after any credential harvest - denying outbound connections to 45.148.10.212 and models.litellm.cloud regardless of how credentials were obtained.

✓ AI Workload Segmentation and Zero Trust

Aviatrix SmartGroups isolate AI/ML workloads from sensitive control planes and data stores, limiting blast radius if a host is compromised. DCF policies enforce zero trust network segmentation, ensuring that stolen credentials cannot be leveraged for lateral movement across cloud environments.

✓ Real-Time Visibility and Anomaly Detection

Aviatrix CoPilot continuously monitors network traffic, detecting anomalous outbound patterns from AI workload VPCs and providing flow logs for forensic analysis. Blocked connection alerts give security teams immediate notification of exfiltration attempts.

✓ Supply Chain Defense in Depth

Aviatrix provides the network-layer safety net that catches supply chain compromises that bypass host-level controls. Even when a malicious package evades SCA scanning or package integrity checks, DCF rules ensure the exfiltration path is blocked before stolen data leaves your environment.

The LiteLLM attack underscores that supply chain threats targeting AI infrastructure are not hypothetical - TeamPCP has demonstrated the ability to compromise widely used AI libraries through upstream toolchain infiltration. Organizations running LiteLLM or similar AI proxy infrastructure should treat network-layer egress controls as a critical defense layer, not an optional enhancement.

Aviatrix delivers a scalable, cloud-agnostic security solution that blocks credential exfiltration through DCF, enforces zero trust workload segmentation with SmartGroups, provides unified multicloud visibility through CoPilot, and enables TLS inspection to detect domain fronting - ensuring that even a fully compromised AI workload host cannot successfully exfiltrate credentials to TeamPCP infrastructure.

Aviatrix Protection Strategy

1. Network Segmentation

- Isolate AI/ML workloads from control planes and sensitive data stores
- Separate CI/CD infrastructure from production cloud environments
- Create network boundaries between cloud providers
- Enforce cloud native metadata service controls (AWS IMDSv2, GCP metadata concealment)

2. Distributed Cloud Firewall (DCF) Rules

Navigate to: CoPilot → Security → Distributed Cloud Firewall → Policies

✓ View step-by-step guide for creating DCF rules

✓ Learn about ExternalGroups and ThreatGroup for C2 blocking

✓ **Rule 1:** Block TeamPCP Exfiltration Server (IP + Domain)

Source: Any → Destination: 45.148.10.212/32 + models.litellm.cloud → Action: **DENY** → Logging: ON

Blocks outbound connections to the C2 IP and typosquat exfil domain. Apply to all Spoke gateways.

✓ **Rule 2:** Block Known Malicious Infrastructure via ThreatGroup

Source: Any → Destination: ExternalGroup (ThreatGroup) → Protocol: Any → Action: **DENY** → Logging: ON

Automatically blocks egress to known C2 infrastructure. Provides ongoing coverage as TeamPCP infra evolves.

✓ **Rule 3:** Default-Deny Egress from AI Workloads to Untrusted Destinations

Source: SmartGroup (AI-ML-Workloads) → Destination: Any → Action: **DENY** → Logging: ON

Use CoPilot live flow logs to baseline required egress endpoints, then add explicit PERMIT rules. Default deny stops unknown exfil paths. R

3. Advanced Configurations

- Enable TLS inspection on all egress from AI workloads - domain fronting can hide malicious POST paths behind a legitimate SNI; inspection exposes the actual request path
- Use CoPilot live flow logs to baseline authorized egress endpoints before enforcing default-deny policies - identifying required domains is the key operational challenge
- Enable CoPilot NetFlow analytics and alerting for anomalous egress from AI workload VPCs
- Configure SmartGroups for AI/ML workload-based segmentation
- Implement PyPI dependency pinning with hash verification across all Python projects

Resources & References

Aviatrix Threat Intelligence:

Aviatrix Threat Research Center

Source Research:

BerriAI/litellm GitHub Issue #24518 | Trivy Supply Chain Attack Analysis (ramimac.me)
Aqua Security Blog | Socket.dev: Trivy Under Attack

Prior Aviatrix Advisory:

TeamPCP Cloud Worm (AVX-SEC-2026-002) - same threat actor's prior campaign targeting Docker APIs and Kubernetes

Aviatrix Documentation:

Creating DCF Rules | ExternalGroups & ThreatGroup | Kubernetes Security | CoPilot

Supply Chain Security Resources:

LiteLLM Official GitHub Releases | LiteLLM Trivy Pin Fix PR #24525

Need Assistance?

Under Active Attack or Suspect Compromise?
Get immediate incident response assistance: Aviatrix Breach Lock

Assess Your Security Posture
Run a free Workload Attack Path Assessment: Start Assessment

Additional Support:

- **Support Portal: support.aviatrix.com**
- **Professional Services: Contact your account team**
- **Security Incidents: Email with "SECURITY INCIDENT" in subject - ref. AVX-SEC-2026-003**

ABOUT AVIATRIX®

Aviatrix® is the cloud network security company. We're on a mission to simplify cloud networking so companies stay agile. Trusted by more than 500 of the world's leading enterprises, our cloud networking platform creates the visibility, security, and control needed to adapt with ease and move ahead at speed. Combined with the Aviatrix Certified Engineer (ACE) Program, the industry's leading multicloud networking and security certification, Aviatrix empowers the cloud networking community to stay at the forefront of digital transformation. Learn more at www.aviatrix.com.