AVIATRIX

SOLUTION BRIEF

Contain Azure AI Foundry Agents

Powered by Aviatrix + Microsoft

Aviatrix Validated Containment Architectures are lab-tested containment deployment blueprints for the AI platforms enterprises are actually running. Ship-ready, policy-included, validated before they arrive. This Validated Containment Architecture covers Azure AI Foundry Agents.

The Threat

Azure AI Foundry BYOVNet mode solves the network isolation problem – agents run in subnets you control. But subnet delegation for hosted container creates a no-control egress path that is invisible to enterprise security policy.

Three problems compound this exposure:

Threat	What It Means
Unrestricted tool-call egress	Every tool call, MCP server connection, and Code Interpreter outbound request routes through the delegated agent subnet to the open internet. Prompt injection or a compromised MCP server can exfiltrate data through that path with no network-layer barrier.
Indiscriminate TLS decryption breaks the runtime	Applying TLS inspection broadly to the hosted agent subnet path breaks Azure hosted agent control-plane traffic – AAD token acquisition, container image pulls, ACA infrastructure calls. The security solution optionally decrypts tool-call traffic while explicitly bypassing decryption for Azure control-plane destinations.
Per-stack firewall sprawl	Azure Firewall on a stick solves egress for one Foundry deployment on one stack. Enterprises running Foundry alongside AWS Bedrock AgentCore, Kubernetes-hosted agents, and self-hosted frameworks end up with incompatible policy engines, separate audit logs, and no unified enforcement primitive across the AI agent estate.

The Architecture

The VCA uses two enforcement layers, both transparent to the agent container:

Layer	What It Does
Egress enforcement with selective TLS decryption	The Aviatrix Spoke Gateway is deployed in the Foundry spoke VNet. A User-Defined Route (UDR) on the hosted agent subnet routes all egress through the Spoke Gateway. DCF selectively decrypts tool-call traffic (external MCP servers, Code Interpreter, REST API calls) while explicitly bypassing decryption for Azure control-plane FQDNs and Service Tags – AzureActiveDirectory, AzureContainerAppsService, and ACA runtime requirements – via DECRYPT_NOT_ALLOWED rules. The ACA runtime never sees a certificate it does not trust.
East-West lateral movement containment	SmartGroup East-West policy denies traffic from the Foundry agent subnet to adjacent workload spokes. A compromised agent cannot reach databases, internal APIs, or other workloads in connected VNet spokes – even if those resources are reachable by other workloads in the same enterprise transit.
Two deployment paths	Enterprises already running Aviatrix Enterprise Transit attach the Foundry spoke to the existing transit – same control plane, same console, same policy model as every other spoke in the fabric. Organizations deploying Foundry for the first time with no Transit in place use a standalone Spoke Gateway inside the Foundry VNet – one UDR on the agent subnet, no hub peering, no on-premises connectivity required. The standalone gateway connects to Enterprise Transit later if the network footprint grows.
Policy as code	The DCF WebGroup allowlist and SmartGroup definitions are Terraform-native. They live in the same repository as the agent configuration and update in the same pipeline run. When an agent adds a new MCP server or changes its tool set, the egress policy updates in the same commit – no out-of-band network team coordination required.

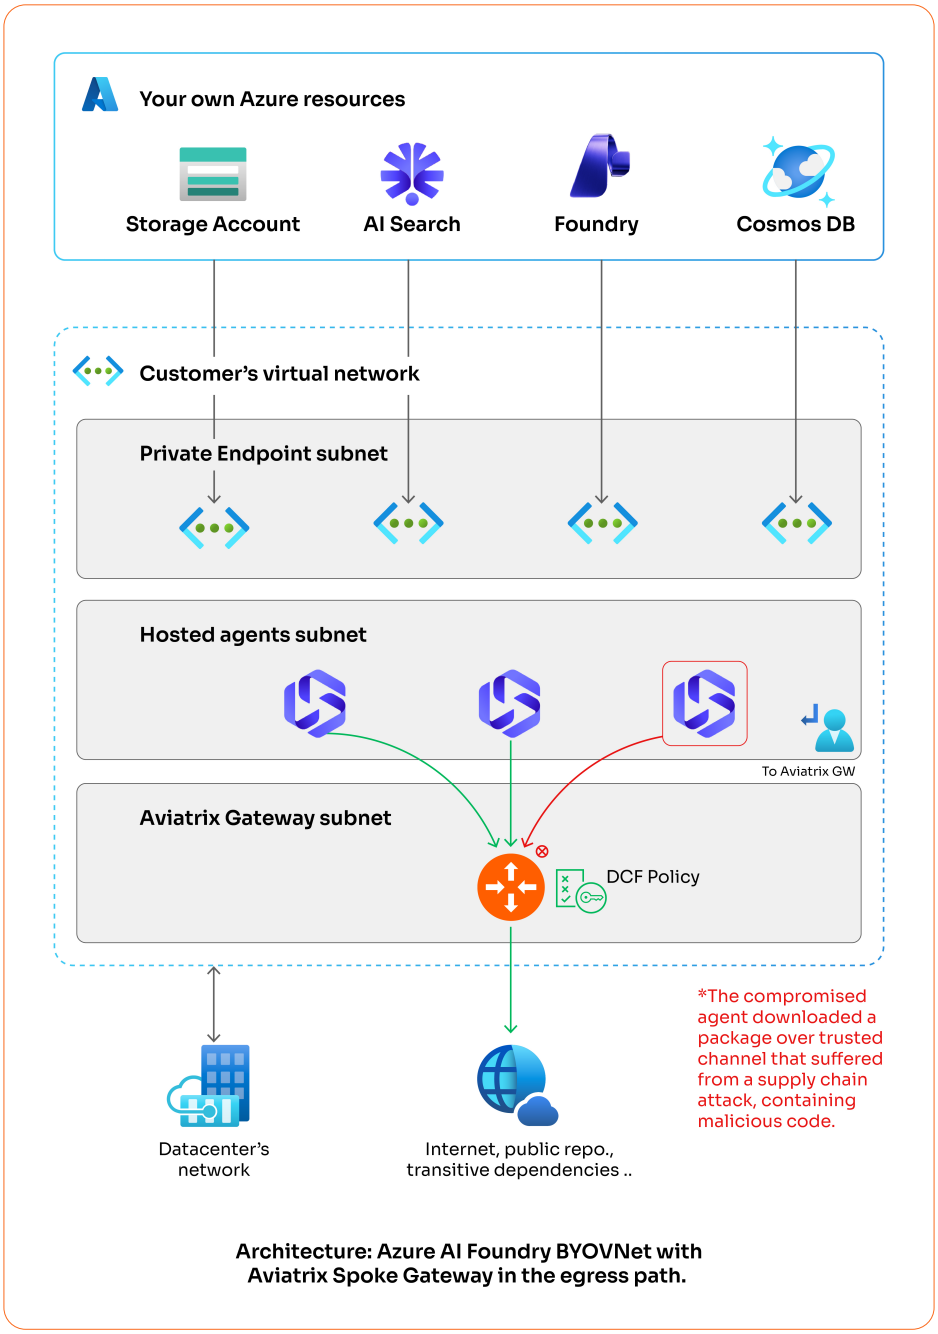

Architecture: Azure AI Foundry BYOVNet with Aviatrix Spoke Gateway in the egress path.

Three Things Your Current Stack Can't Do

01 Inspect Foundry agent tool calls without breaking the runtime

Azure Firewall cannot apply TLS inspection selectively – it either decrypts everything, which breaks ACA control-plane traffic, or nothing. Aviatrix Distributed Cloud Firewall uses per-rule decryption scoping: tool-call egress is inspected, Azure control-plane traffic is bypassed automatically via Service Tags. No manual IP maintenance. No runtime failures.

---

02 Block lateral movement from a compromised agent to adjacent spokes

Azure Firewall governs internet egress. Foundry Guardrails govern content. Neither controls East-West traffic inside the VNet. Aviatrix SmartGroup policy enforces identity-aware deny rules between the Foundry agent subnet and adjacent workload spokes – before the agent ever tries to reach them.

03 Update egress policy in the same pipeline as the agent

Azure Firewall is managed outside the agent's release pipeline. Every new MCP server or tool change requires a separate network team ticket – pressure to over-permit rather than stay least-privilege. Aviatrix policy is Terraform-native: the allowlist lives in the same repo as the agent config and deploys in the same pipeline run.

Compliance Evidence

For HIPAA, PCI-DSS 4.0, SOC 2, EU AI Act, DORA, and FedRAMP environments, auditors require architectural proof of enforcement, not a policy document. CoPilot, Distributed Cloud Firewall, and SmartGroup logs as well as infrastructure-as-code evidence prove that this Validated Containment Architecture provides continuous enforcement. The proof of enforcement is the architecture itself – not a statement about the architecture.

Questions Worth Asking

- How many AI agents are running in your environment right now – and how many are fully inventoried?
- If a compromised agent attempted to exfiltrate data to an external domain, what in your stack would catch it?
- How are you currently proving to auditors that AI workload egress is controlled – not just described in policy?

What's Included

Deliverable	Detail
Insertion pattern	Aviatrix Spoke Gateway deployed in the Foundry spoke VNet with a User-Defined Route on the hosted agent subnet – transparent to the agent container, no code changes required
SmartGroup and WebGroup model	SmartGroup and WebGroup objects keyed to Foundry's identity primitives: hosted agent subnet, hosted agent platform FQDNs, Azure Service Tag SmartGroups (auto-updating), approved tool-call destinations, and adjacent spoke resources for East-West containment
Default-deny DCF policy pack	Prioritized rules covering threat intelligence deny, hosted agent runtime platform permits, tool-call egress with optional selective TLS decryption, and East-West lateral movement deny – all logging-enabled to CoPilot and DCF logs.
Deployment guide	Lab-validated reference architecture tested against prompt injection exfiltration, lateral movement to adjacent spoke private endpoints, and hosted agent runtime startup with Aviatrix in-path
Terraform blueprint	Infrastructure as code for the full deployment – ships in the GitHub repository, deployable against existing Aviatrix Enterprise Transit or as a standalone Spoke Gateway with no Transit required

Get Started

The Validated Containment Architecture for Azure AI Foundry Agents is available today. The Terraform blueprint and deployment guide ship together. Aviatrix Controller 8.2 or later is required.

- Existing Aviatrix Enterprise Transit customers: Attach the Foundry spoke to your existing transit and apply the SmartGroups and WebGroups from the blueprint.
- New customers: Deploy a standalone Aviatrix Spoke Gateway inside the Foundry VNet – one UDR on the agent subnet, no hub peering required.
- Ask your Aviatrix account team for a guided deployment: terraform apply, the scenario cards, and the back-out path – in under an hour.

Request a 30-minute architecture review

Walk through the enforcement model in your environment, map your current Foundry agent inventory, and identify the egress paths you currently cannot see.

About Aviatrix

Aviatrix® is pioneering the Cloud Native Security Fabric – the architecture the Containment Era requires. The Cloud Native Security Fabric 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.