





# AI solution overview

For organizations implementing agentic solutions



# AI challenges

Organizations need **both** governance and speed. Both are very difficult to get.



## Fragmentation

Different solutions for different clouds, LLMs and applications, each with lock-in risk and governance dilemmas.



## Intellectual property risk

Using public LLM providers, exposing remote MCPs, shadow AI and novel AI cybersecurity threats.



## Uptime and resiliency risk

Uptime is threatened by firewall and network surgery, even if there are no breaches. Resiliency threatened by disparate solutions.



## Operations and governance complexity

Of adding non-deterministic workflows without robust visibility and controls. Often forces orgs to slow down.



## Delays

From the extensive network, firewall, VPN, DNS, integration and security work required by most AI solutions.



## Lack of visibility and attribution

Agents using human identities across fleets of sub-agents and processes. Hurts operations, speed and security.



## Security risk

AI breaks most inspection-based tooling, add threats and have unprecedented lateral movement.



## Cost concerns

Of foundation models, but there are quality and security concerns of using multiple LLMs. Impact is often to hit the brakes.



# NetFoundry's AI mission

NetFoundry securely delivers billions of sessions per month for hundreds of organizations, including 2 of the largest 5 in the US, and innovators around the world.

NetFoundry is SOC2 Type II compliant and helps customers achieve FIPS, CRA, NIST, NIS2, IEC 62443, NERC/CIP, CJIS, HITRUST, HIPAA



NetFoundry's mission is to help organizations:



## End the tug-of-war between governance and speed

Organizations need both. NetFoundry enables AI speed, within the governance framework.



## Get AI sovereignty, control and resiliency

Keep control of your AI everywhere - every cloud, every model, every site, including on-prem. Single approach avoids resiliency risks.



## Secure AI without disruption

Zero trust, identity-microsegmented AI with no changes to firewalls, networks or tooling.

# Governance and speed **both** start with identity



*“NetFoundry provides the secure network foundation Rhapsody needs to support private, policy-based access across healthcare, including applications, APIs, workloads, and emerging AI-enabled workflows.”*

**Kevin Day**  
**Rhapsody Health CTO**

## AI speed, sovereignty and resiliency, within governance framework



### 1. Identity

- Every agent proves who it is before it gets access and inventory is auto-built as they spin up and down.
- Shadow AI can't get access to enterprise resources and credentials.
- Accelerate – identity-first overlays eliminate the need to perform firewall, VPN, network and DNS surgery.



### 2. Zero trust & sovereignty

- Deploy anywhere, including on-prem. Get one solution across all sites and LLM models.
- Agents can't even see tools, API keys, resources, data paths which policy does not explicitly allow.
- Make MCP servers unreachable to unauthorized agents.



### 3. Governance & mitigation

- Central visibility and controls across cloud and network borders.
- Both broad and narrow kill switches to optimize uptime and business continuity.
- Every session tagged with the identity and policy which allowed it.

# NetFoundry AI solution overview

Use separately or use together as one platform

-  **Agentic identity and visibility**  
Sovereign identities, bootstrapped as agents spin up, and options to use other identity and authentication factors.
-  **Agent sandboxing**  
Kernel-level sandboxing, with agents born into governance framework.
-  **Zero trust gateways: MCP and LLM**  
Enable remote MCP, unreachable gateways, model routing.
-  **Mint: AI as code**  
Spin up AI templates on demand, born into your controls and governance.
-  **Enclave: secure AI without the cloud tax or lock-in**  
Zero trust AI without touching firewalls and networks. One solution across all clouds and LLMs, including on-prem AI.
-  **Universal zero trust**  
The first zero trust solution which secures every type of workload, at every type of site, from legacy apps to AI agents.

# Solution #1: agents get verifiable passports

Sovereign, cryptographic identity for every agent, MCP server and model endpoint, bootstrapped at birth



## Sovereign agent identity

A compromised token or API key no longer grants access. NetFoundry builds in X.509 and enables 3<sup>rd</sup>-party identity (SPIFFE/SPIRE, Keyfactor, OAuth, etc.) to gate network access and auto-build inventory.




## Agent-level visibility

See and control every connection, attributed to named agents, not IPs, independent of agents, LLMs, sites.



## No integration required

Works out of the box; optionally chain your enterprise IdP, CA or 3<sup>rd</sup> party identity as additional or primary factors.



# Solution #2: AI agents, born into sandboxes

Kernel-level eBPF containment around every agent container, for every agent.



## Default-deny egress

Enforced in the kernel by CIDR, port, protocol, executable at two independent layers (packet + socket)




## No keys in the sandbox

Model traffic exits only via the LLM Gateway, which holds all provider credentials; unagreed destinations are unroutable.



## Every attempt on record

Allowed and denied connections logged with process, container and destination  
Streamed to your SIEM / AI-SOC.



# Solution #3: govern access with zero trust MCP & LLM gateways

This means shadow AI agents can't reach enterprise resources...



## Route any model

Multi-provider routing across dozens of models. OpenAI-compatible, with a 3-layer semantic cascade.




## Budgets, costs, guardrails

Identity-based access and rate limits, per-team/project cost tracking, PII and content-safety guardrails. Built-in zero trust.



## Solve MCP sprawl and security

Tool aggregation and namespacing, structural permission filtering, per-client isolation. MCP servers made unreachable from shadow AI. Move from local MCP servers to centralized.



# Solution #4: ship governed agent fleets in one command

Pull, verify, run: signed agent packages that arrive with their governance built in.



## Signed registry

Pull from and publish to a NetFoundry-operated registry; every entry signed and vouched, mirrored into your own sealed, private fabric (even air-gapped).



## Recipes

Versioned, immutable, reproducible packages of your agents, model access and tools. Governance: approvals, spend ceilings, run bounds.



## Harness

A lightweight runtime and CLI. Model and tool paths route via the gateways, so governance is unconditional. All on your infrastructure, including on-prem.



```
$ mint run acme/support-agent
```

*governance travels with the package, unconditionally*

# Solution #5: zero trust AI enclaves

Spin up enclaves in minutes for one app or everything, without touching underlay network and firewalls.



## Invisible until authorized

No inbound ports, public IPs or VPNs; resources do not exist on the network.




## Identity-based microsegmentation

Each agent, MCP server, model and data source reaches exactly what policy allows; lateral movement dies by default.



## Connect anything without disrupting network

SDK in the app, connectors on hosts or at edges.



# Solution #6: universal zero trust

