NetFoundry

POST-MYTHOS **ROADMAP**

Mythos and LLMs changed the game,
but we can now move the playing field.

EXECUTIVE SUMMARY

The disclosure of the Claude Mythos Preview in April 2026 fundamentally shattered the traditional cybersecurity paradigm. **With expert analysis from CSA, SANS, and OWASP documenting a collapse in the mean time-to-exploit from 2.3 years to less than 24 hours, organizations can no longer rely on traditional patching to stay ahead of automated threats.** Because patches only provide temporary relief until the next zero-day vulnerability is discovered, enterprises must immediately categorize assets based on how quickly their attack surfaces can be eliminated.

The Post-Mythos Roadmap outlines an urgent, pragmatic approach: stop racing against every attacker on the Internet and start making critical workloads unreachable. By utilizing NetFoundry, organizations can transition to a proactive, zero-trust posture. NetFoundry transforms internet-facing perimeter devices and exposed enterprise APIs by denying all inbound traffic and utilizing outbound-only, identity-bound private overlays with no public listeners.

This roadmap provides a phased, 150-day timeline to eliminate exposure across all vital categories, including B2B APIs, OT/ICS networks, third-party vendor tunnels, and emerging Agentic AI workloads.

Rather than constantly racing to patch or replace infrastructure, enterprises can effectively cloak their resources, by making them unroutable from the network to secure the non-human workload connectivity that is most heavily targeted by Mythos-class models. Ultimately, this framework provides the immediate quick wins necessary to survive the post-Mythos landscape while laying an API-first, modular foundation for the future.

NetFoundry

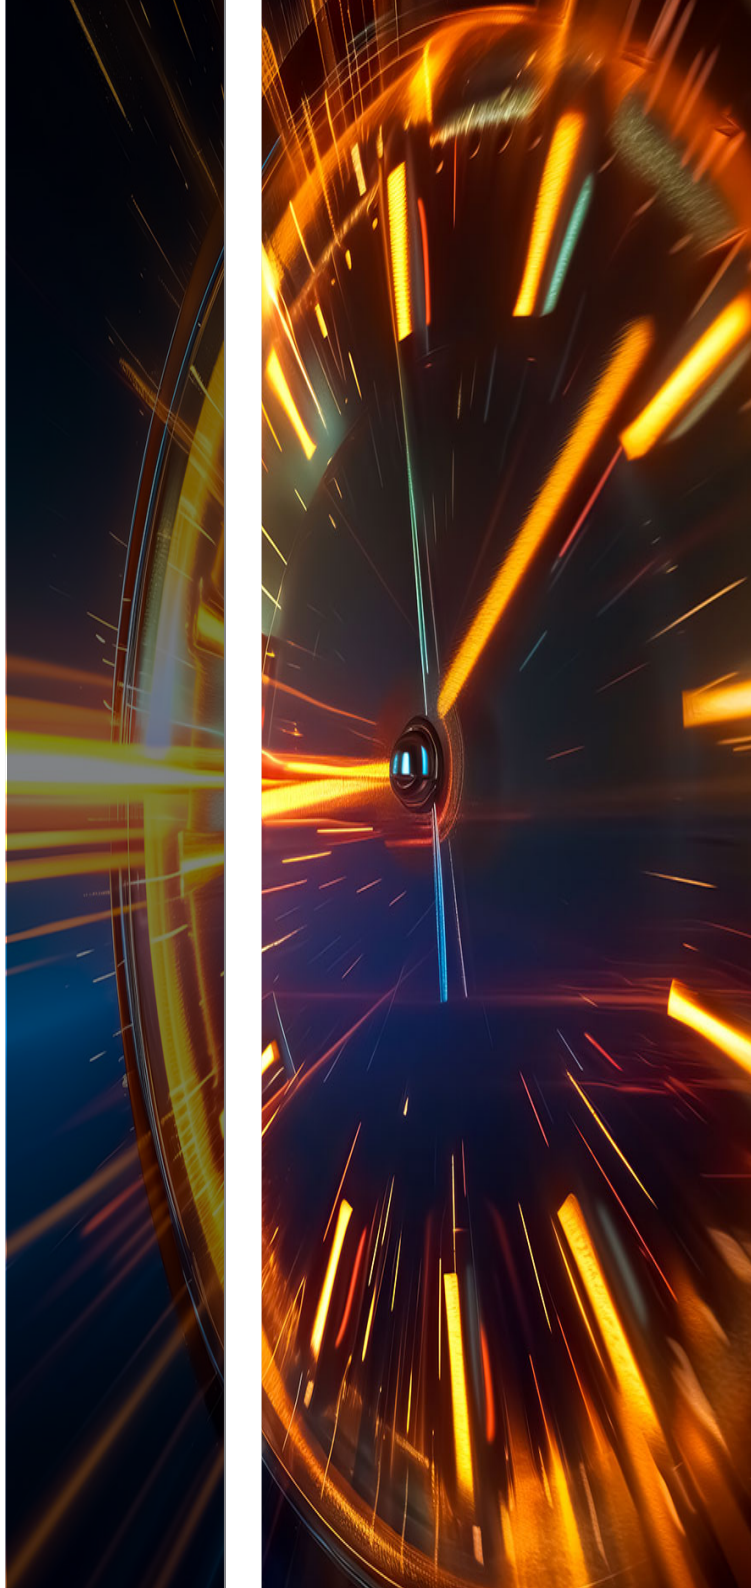

INTRODUCTION

Claude Mythos Preview, disclosed by Anthropic on April 7, 2026, broke the old cybersecurity model. Expert analysis from CSA/ SANS/OWASP documents a collapse in mean time-to-exploit from 2.3 years in 2018 to less than 24 hours. Every asset needs to be categorized by 'how quickly its attack surface can be eliminated', because:

- We will not be able to patch all of them fast enough
- The patch only lasts until the next zero day is found.
- Non-critical software serves as lateral movement to more critical software or data – we can't leave anything exposed.

By first eliminating the attack surfaces which we can eliminate, we focus patching, upgrade and replacement efforts on a viable target, and are not racing against every attacker on the Internet while we execute those plans.

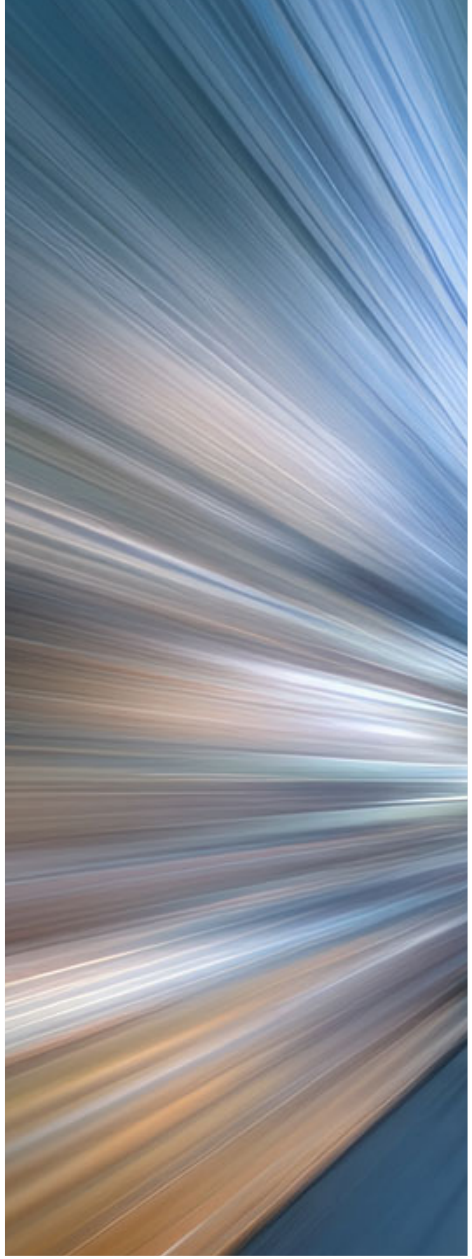

NetFoundry

SUMMARY MATRIX

PRIORITIZING BUSINESS CONTINUITY

Every enterprise is unique, but this is a map of the terrain:

#	Category	Timing	Outcomes
1	Internet-facing perimeter devices	Phase One	You can't patch or replace all the perimeter devices, but you don't have to.
2	Publicly exposed APIs (B2B, partner, mobile backends)	Phase One	Quick removal of a high-value target class.
3	OT / ICS / IoT / edge – phase one	Phase One	Prioritize critical sites and workloads
4	Agentic AI / LLM / MCP workloads – phase one	Phase One	Reduce shadow AI risk and do not add AI to the attack surface
5	Third-party / MSP / vendor – phase one	Phase Two	Start with the persistent, bidirectional connections
6	Legacy internal apps	Phase Two	Fast wins
7	Cloud / hybrid / multi-cloud workload-to-workload	Phase Three	Eliminating the network and firewall dependencies
8	Phase two of ICS, 3rd-party and AI	Phase Three	According to risk and business priorities

This assumes ZTNA or SASE is already in place for internal users accessing internal applications, and that is not the surface most impacted by Mythos – it is the non-human centric workload security model which has entered a new world.

NetFoundry

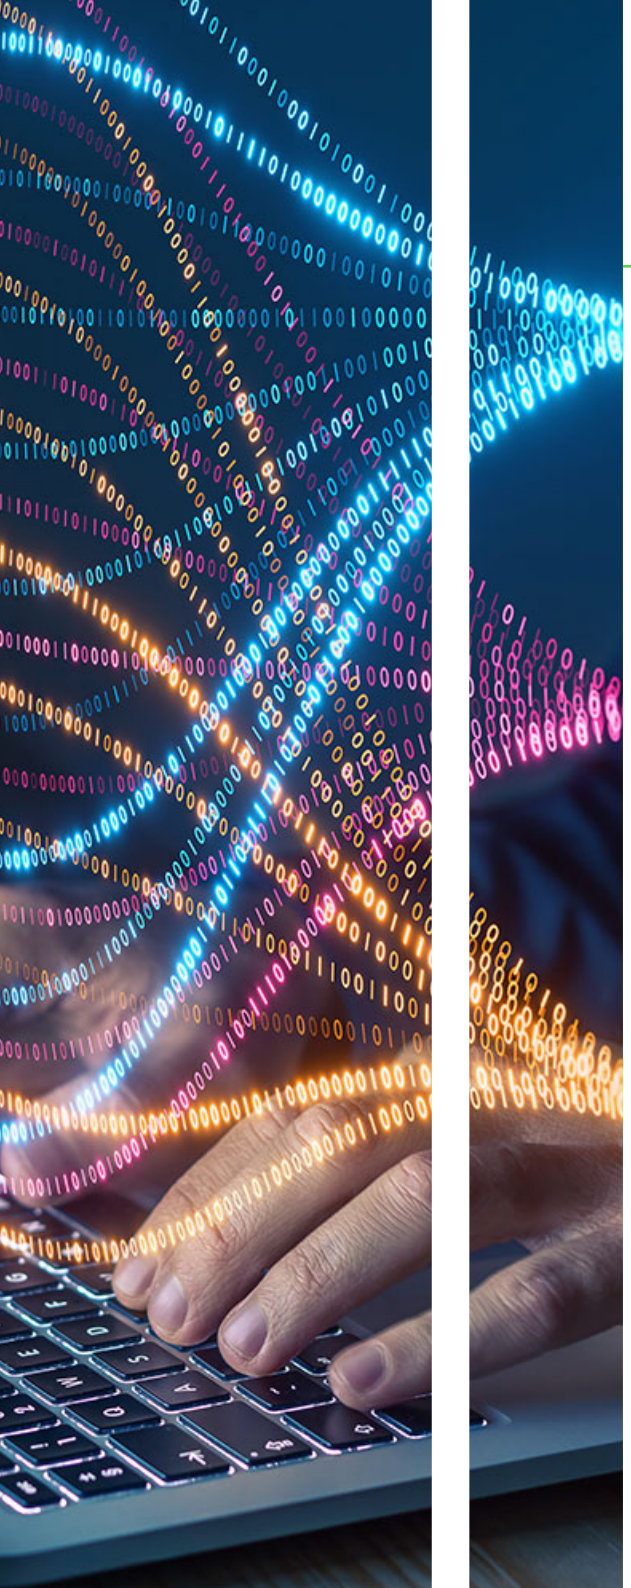

CATEGORY-BY-CATEGORY ANALYSIS

1. INTERNET-FACING PERIMETER DEVICES

Includes VPN concentrators, corporate firewalls with management UIs exposed, API gateways, enterprise NAS, perimeter routers, and load balancers with admin interfaces reachable from the internet. These devices are **a high-priority target class in a 20-hour MTTE world**: they are always on, always listening, often not patchable within vendor cadence, and carry trusted network identities.

Per Forescout (2025) and post-Mythos analyses from the CSA and ISACA, routers alone account for over half of the most critically vulnerable Internet-exposed systems. The April 2026 Synology SSL-VPN advisory (CVE-2021-47960/47961) is a textbook case: a legacy bug becomes a zero-day again in a Mythos world.

Why use NetFoundry? Replacing these devices or constantly patching them is not realistic. Patches need to be hygiene – not critical path. NetFoundry changes the role of these devices – they deny all inbound traffic, and open a single outbound port to a private overlay, only available to authorized endpoints, replicated by an outbound-only zero-trust overlay with no public listener. They become true firewalls instead of part of the attack surface.

Crawl, walk, run. These devices usually secure server-side resources. The client side requires the application of a NetFoundry-provided identity, authentication, authorization solution. There are many NetFoundry options, including clientless for certain endpoints and clients which provide continuous authentication and proof of possession-based identity for other endpoints, but some endpoints will be more difficult than others.

The full stack. The firewalls don't go away. They become more effective and simpler to operate (eliminate the firewall ACL maintenance nightmare). Likewise, downstream SIEM and SOC are made cheaper and more effective – no longer filtering the Internet and storing all that data.

2. PUBLICLY EXPOSED ENTERPRISE APIs

Exposed enterprise APIs have survived because they are often snowflakes which cost attackers too much time and money too breach. **That is no longer the case - Mythos-class models enumerate, fuzz, and chain API flaws against B2B APIs, mobile backends, partner webhooks, and AI-service endpoints.**

Why use NetFoundry? Quick wins to make B2B APIs and internal APIs dark – literally serve them from private IP address space API gateways. NetFoundry provides the easy button to make the API endpoint dark — no public listener, IdP-gated session setup, short-lived token based overlay access.

Crawl, walk, run. More robust solutions include end-to-end mTLS and proof of possession identity, as well as developing new APIs as natively zero trust.

Big picture: NetFoundry does not solve truly public, anonymous APIs (e.g., unauthenticated search, marketing site lookups, content CDNs) - those require a hardened public edge with a modern WAF, bot management, and API discovery.

3. OT / ICS / IOT / EDGE

This will not be an immediate priority for every enterprise, but it is growing. **There are 21.1B connected IoT and OT devices and many run firmware images that have not been patched and may never be.** Mythos finds new zero-days across the 150,000+ firmware stacks in industrial and embedded fleets – even the rare ones.

Why use NetFoundry? Immediately take care of the most important problems – deny all inbound on OT firewalls (especially at remote, unmanned sites), limit machine to IT/edge/cloud and secure access to per-session, per-device, identity-bound access. CISA's ICS remote-access guidance and NIST align with this pattern.

Crawl, walk, run. For greenfield, NetFoundry SDKs provide zero-listener telemetry.

Big picture: OT protocol inspection (Modbus/TCP, DNP3, OPC-UA payload analysis), ICS-specific anomaly detection, and safety-system isolation remain but become simpler and cheaper as NetFoundry changes the risk profile to internal-attacks, as opposed to Internet-based attacks.

4. AGENTIC AI / LLM / MCP WORKLOADS – PHASE ONE

40% of enterprise apps will have task-specific AI agents by end-2026 (Gartner). Agents are non-human identities that talk to APIs on behalf of humans - exactly the shape of target Mythos-armed attackers love – but also the area where NetFoundry tilts the playing field against them.

How to use NetFoundry: Focus phase one on greenfield and reducing shadow AI. With NetFoundry, MCP servers and tool endpoints are reachable only over identity-bound, scoped sessions. You flip MCP gateways from an attack surface to a zero trust firewall – use NetFoundry to put critical enterprise APIs, tools, data and resources behind the gateways such that even if someone runs a shadow AI project that project can't reach the critical resources. NetFoundry provides identity, control and visibility for AI agents – separate or built into the harness.

Big picture: NetFoundry is made to integrate with semantic-layer controls such as prompt injection defense, tool-call authorization policy, agent attestation, content safety. Pair NetFoundry with these solutions and MCP-aware authorization (emerging from Anthropic, Google A2A, and the IETF WIMSE).

NetFoundry

THIRD-PARTY / MSP / VENDOR / EXTERNAL CONNECTIVITY

Kaseya, ScreenConnect, Change Healthcare, MOVEit were all the same fatal flaw: a trusted inbound tunnel or remote-admin channel becomes the entry vector. In a Mythos world, these tunnels are one vendor-CVE away from being the front door for any attacker.

This multiplication is problematic because the average supply-chain breach already costs: $4.91M with 267 days to contain.

How to use NetFoundry: Start with the most dangerous – the persistent, two-way, site-to-site tunnels and firewall ACLs which enable vendors to initiate inbound traffic into your network. Replace those with outbound-initiated, identity-bound, per-session overlays. Make controls, policy and logging available to both the enterprise and the vendor. Historically, this was an 'it takes two to tango' problem which resulted in the firewall/VPN problem persisting – NetFoundry transforms it to a win for both, and it can be implemented by either the enterprise or the supplier.

Big picture: Human remote access is less dangerous and easier to police but still needs a modern solution. NetFoundry can provide it, but also partners with the leaders of human access solutions.

5

LEGACY INTERNAL APPS

The “VPN in, flat network out” pattern is still all too common. It is difficult to add identity, authentication and encryption to older applications. These may not all be high risk, but start with the ones which may enable lateral movement.

How to use NetFoundry: Without changing the app, make it reachable only by authorized identities and segment the server side so it can’t egress to unauthorized destinations and can’t facilitate lateral movement. Rollout is fast because there are clientless, client-based and gateway-based solutions, and none of them require networking, firewall, DNS and NAT changes.

Big picture: Some apps will be out of scope, such as apps which may require broadcast/multicast, peer-to-peer discovery protocols, or protocols that embed IP addresses at L7. The longer term solution for them is often to use NetFoundry SDKs, but these apps may initially need a separate solution or compensating controls.

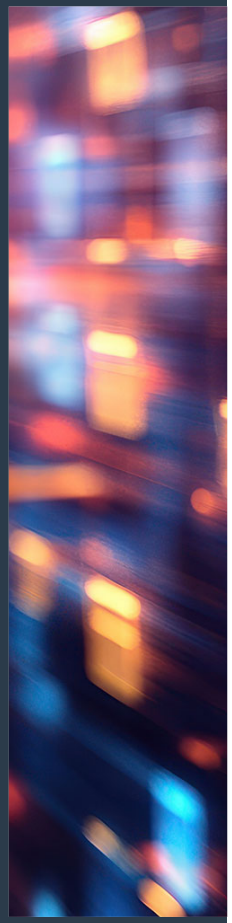

7. CLOUD / HYBRID / MULTI-CLOUD SITE-TO-SITE AND WORKLOAD-TO-WORKLOAD

70% of enterprises run hybrid cloud (Flexera 2025) and misconfiguration is the #1 cloud threat (CSA 2024). The CSA "Mythos-ready" guidance makes this explicit: zero trust must extend beyond user-to-app and cover workload-to-workload traffic.

How to use NetFoundry: Leverage NetFoundry as an intra-cloud and cross-cloud, zero trust native fabric. Workloads in AWS, Azure, GCP, on-prem, and edge only reach each other over identity-bound, outbound-only overlays – even PrivateLink endpoints. Eliminate VPC and VNet-inbound exposure such that the next LLM-found zero-day bug can't be exploited from the outside. Segment the inside without dependencies on firewalls and IP addresses to restrict lateral movement and data exfiltration.

Big picture: It is still ok to carry east-west traffic inside a single Kubernetes cluster via service mesh or similar. NetFoundry extends the mTLS / zero trust outside of the cluster, across administrative and network boundaries. NetFoundry can also leverage the same identities if desired, e.g. SPIFFE/SPIRE.

8. AGENTIC AI / LLM / MCP WORKLOADS

Importance: 4 | NF Triage Speed: 4

40% of enterprise apps will have task-specific AI agents by end-2026 (Gartner). 25% of enterprise breaches will trace to AI-agent abuse by 2028 (Gartner). OWASP Top 10 for Agentic Applications 2026 calls out Goal Hijack, Tool Misuse, Identity & Privilege Abuse, Agentic Supply Chain, and Rogue Agents. Agents are non-human identities that talk to APIs on behalf of humans — exactly the shape of target Mythos-class adversaries love.

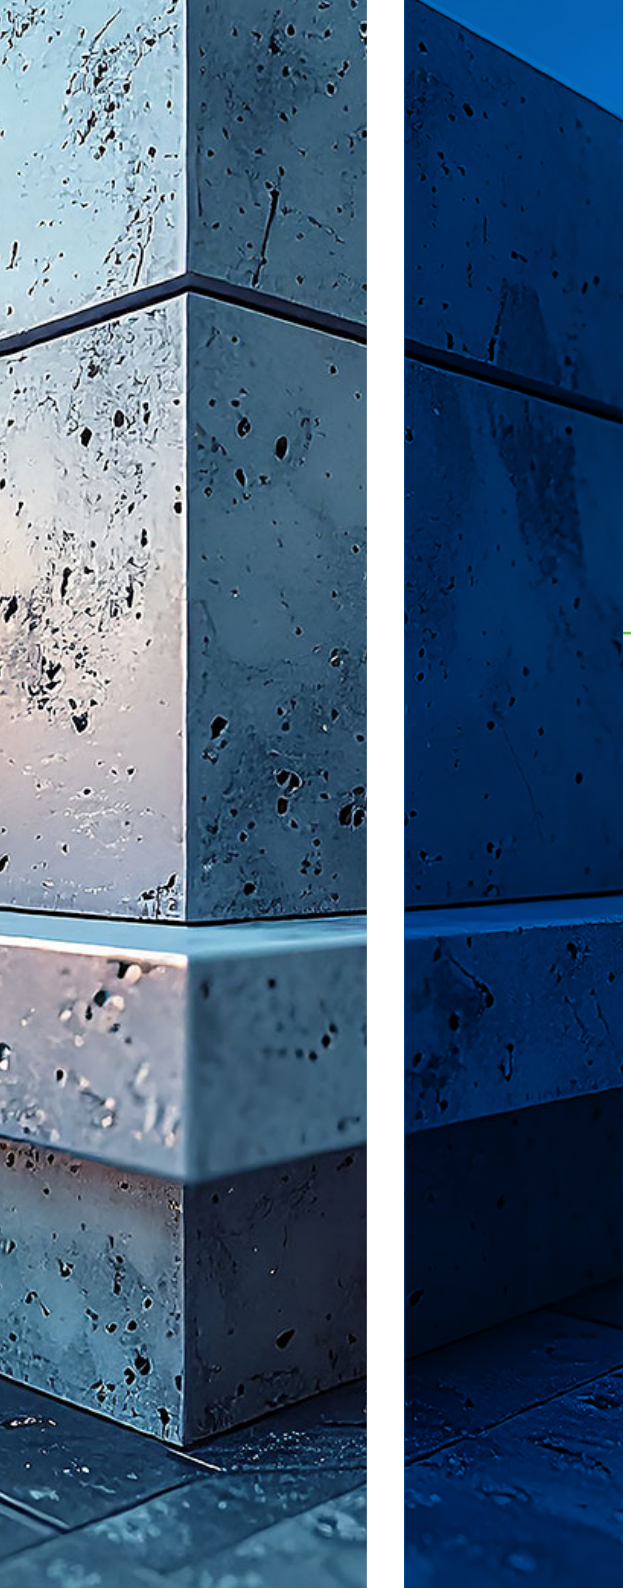

Where NetFoundry wins: SDK-embedded agent carriers. The agent binds to the fabric, not to an OS socket; MCP servers and tool endpoints are reachable only over identity-bound, scoped sessions. This is greenfield work — crawl/walk/run largely collapses to "run" because the code is new. NetFoundry's OpenZiti AI-workload enclave (March 2026) is explicitly aimed here.

Where it does not: Semantic-layer controls — prompt injection defense, tool-call authorization policy, agent attestation, content safety — are not network-layer problems. Pair NetFoundry with MCP-aware authorization (emerging from Anthropic, Google A2A, and the IETF WIMSE drafts) and agent-identity platforms. NetFoundry carries; it does not adjudicate.

PHASE TWO AND BEYOND FOR THE ABOVE CATEGORIES

This is of course a broad and deep category. Not many generalities apply – each business will need to make their own prioritization decisions.

However, NetFoundry is a horizontal platform – providing identity, authentication, authorization, controls, visibility and audit for any workload.

NetFoundry is also an API-first, modular architecture, designed to fit well in a heterogenous vendor, hybrid, quickly evolving world. Hence, future phases have a solid starting point, and expert SIs, VARs and MSPs are enabled to help define and implement them.

NetFoundry

SUMMARY

Mythos and LLMs changed the game, but we can now move the playing field.

Use NetFoundry to change the first question from ‘how long will it take to patch, upgrade or replace’ to ‘which workloads do I want to make unreachable today’. The two go together – removing the workload from the Internet gives you time to properly patch, upgrade or replace, without racing against the entire Internet.

Although this document discusses quick wins, your enterprise is putting in a full foundation to support the future roadmap.

LET’S TALK

Contact us today to learn more about how to secure your AI deployments without slowing them down. We’ll show you how to:

- Deploy a secure, zero trust enclave for your users, AI agents, MCP servers, and LLMs.
- Avoid wasting time on never-ending firewall rule changes.
- Centralize AI cost accounting and optimization.

NetFoundry

www.NetFoundry.io | Info@NetFoundry.io