

Multi-Tenancy In AI Clouds

Network Isolation with Netris®

Whitepaper

Abstract

This paper presents an architectural model for hard multi-tenancy — true multi-tenancy built on hard isolation (enforced on networking hardware) — in AI factories built on NVIDIA® DGX™ systems. Tenant intent is defined once and applied consistently across heterogeneous connectivity domains: East–West fabrics based on NVIDIA Spectrum-X™ or NVIDIA Quantum InfiniBand™, Ethernet-based North–South fabrics, NVL72 rack-scale NVLink® Multi-Node fabrics, and NVIDIA BlueField® DPU-augmented domains.

The paper examines aligning Ethernet, InfiniBand, and NVLink partition boundaries; coordinating lifecycle updates across fabrics; and maintaining consistency through topology-aware validation — and shows how the Netris® NAAM (Network Automation, Abstraction, and Multi-Tenancy) platform implements this model using supported NVIDIA management APIs.

This guidance is intended for CTOs, AI/cloud architects, NVIDIA Cloud Partner (NCP) operators, and enterprise infrastructure teams deploying shared DGX-based clusters, including NVIDIA Mission Control-managed environments.



4701 Patrick Henry Drive, Bldg. 25
Santa Clara, California, 95054, USA

Copyright Notice, Authorization to Use and Disclaimers

© 2026 Netris, Inc. All rights reserved

This whitepaper (the “Whitepaper”) is the confidential and proprietary information of Netris, Inc. (“Netris”). Netris authorizes you to use this whitepaper solely in connection with your use of Netris products and services. You may not use this Whitepaper to benchmark any Netris products or services, or to develop products or services that compete with those made available by Netris.

Netris may change or update this Whitepaper and/or the Netris products and services referenced therein at any time and without notice. The information in this Whitepaper is made available to you on an “AS IS” basis, without any warranties or representations, and Netris hereby disclaims all warranties and representations of any kind, whether express, implied, or statutory.

The terms with respect to this Whitepaper do not amend or modify the terms applicable with respect to Netris products and services.

Netris® is a registered trademark of Netris, Inc. Other company and product names may be trademarks of the respective companies with which they are associated.

In no event will Netris be liable for (a) any indirect, punitive, special, incidental, or consequential damages, cost of procuring substitute goods, loss of profits, revenues, use, data, or goodwill arising out of or related to this Whitepaper; or (b) aggregate liability in excess of one hundred U.S. dollars (US\$100).

Document History

| Version | Date | Authors | Description of Change |
|---------|------------|--|---|
| 1.0 | 02/20/2026 | Alex Saroyan (Netris) Andrey Khomyakov (Netris) | Initial Version |
| 2.0 | 07/01/2026 | Andrey Khomyakov (Netris) | Adds NVIDIA Mission Control context. Expands BlueField DPU coverage. Updates management network section to reference the Netris Management Network Reference Architecture. Improves readability. Clarifies terminology. |
| | | | |
| | | | |
| | | | |
| | | | |
| | | | |
| | | | |

Executive Summary

The Netris NAAM (Network Automation, Abstraction, and Multi-Tenancy) platform — the most widely deployed network automation and multi-tenancy platform for AI infrastructure — provides the coordination layer that enables hard multi-tenancy in NVIDIA DGX-based AI factories. Netris translates tenant intent into consistent isolation boundaries, with hard isolation (enforced on networking hardware), across Ethernet, InfiniBand, NVLink multi-node, and BlueField DPU-enabled networking domains.

In this architecture, tenant intent is defined once using a high-level abstraction and translated into the isolation constructs required by each fabric, including Ethernet VRFs/VXLANs, InfiniBand partitions (PKeys), and NVLink GPU partitions.

Netris implements this model by automatically generating and applying configuration across switches, DPUs, and integrated fabric management systems while continuously validating the resulting network state.

This approach allows operators to manage tenant boundaries consistently across heterogeneous fabrics — onboarding, resizing, and retiring tenants in minutes as the AI factory grows — while compute and storage orchestration systems evolve independently.

In One Sentence

Netris enables hard multi-tenancy in NVIDIA DGX-based AI factories by defining tenant intent once and automatically aligning isolation boundaries — enforced on networking hardware — across Ethernet, InfiniBand, NVLink, and BlueField DPU domains.

Architecture Principles

The architecture described in this paper is based on three principles:

1. Define tenant intent once.
Tenant isolation is expressed using a high-level abstraction representing the intended isolation domain.
2. Translate intent across fabrics.
Netris converts tenant intent into the isolation constructs required by each participating fabric, including Ethernet VRFs/VXLANs, InfiniBand partitions (PKeys), NVLink GPU partitions, and host networking policies.
3. Continuously validate the isolation state.
Netris continuously validates topology and configuration state to ensure tenant boundaries remain consistent as infrastructure evolves.

Table of Contents

| | |
|--|-----------|
| The Multi-Tenancy Model In AI Cloud | 7 |
| Fabric Domains in Multi-Tenant AI Factories | 7 |
| Tenant Isolation Across Heterogeneous Fabrics | 7 |
| Purpose of the Multi-Tenancy Model | 8 |
| Multi-tenancy Challenges in AI Infrastructure | 8 |
| Recommendations for Implementing a Multi-Tenant Environment | 10 |
| 1. Establish a Unified Tenant Abstraction | 10 |
| 2. Translate the High-Level Abstraction into Fabric Configurations | 10 |
| 3. Apply Tenant Configurations Across Fabrics Without Disruption | 10 |
| 4. Validate Topology and Isolation Continuously | 11 |
| 5. Provide an API for Upstream Platforms | 11 |
| Architecture of a Multi-Tenant AI Cloud with Netris | 12 |
| Interpreting Operator Intent and Supporting Alignment Across Fabrics | 13 |
| Upstream Platform Integration | 15 |
| Separation of Responsibilities | 16 |
| Simulation | 16 |
| An Example Network Isolation Implementation With Netris | 18 |
| Network Provisioning with Netris | 18 |
| East-West Fabric | 18 |
| If East-West is implemented using Spectrum-X Ethernet | 18 |
| If East-West is implemented using InfiniBand | 19 |
| North-South Fabric | 20 |
| NVLink Multi-Node Fabrics (NVL72) | 20 |
| Management Networks (Out-of-Band) | 21 |
| Network Services for Tenants | 23 |
| VPC-peering (Shared storage and more) | 23 |
| Shared Storage Access via VPC Peering | 23 |
| Alternative: Shared Storage Access via VLAN Subinterfaces | 24 |
| VPC Connect | 25 |
| Edge Services (NAT, Layer-4 LB) | 27 |
| Access Control Lists | 28 |
| Concurrent Multi-Tenancy / Host-Level Segmentation | 29 |
| DPU-based/Hardware-Accelerated Host Segmentation (HBN) | 29 |
| Software-Based Host Segmentation (EVPN on Host) | 32 |
| Node Provisioning with NVIDIA Mission Control | 34 |
| Shared Services Model | 34 |
| Network Validation in Multi-Tenant AI Infrastructure | 35 |

| | |
|--|-----------|
| Topology-Aware Validation | 35 |
| Detection of Partial States | 35 |
| Ongoing Validation and Drift Detection | 36 |
| Data-Plane Reachability Validation in EVPN Fabrics | 36 |
| Validation That Scales With Change | 36 |
| Summary and Architectural Outlook | 37 |
| About Netris | 38 |

The Multi-Tenancy Model In AI Cloud

AI infrastructure is increasingly being operated as a shared platform across enterprises and NVIDIA Cloud Partners (NCPs). These AI factories host many independent teams or customers on the same physical resources. Hard isolation (enforced on networking hardware), therefore, becomes a core architectural requirement: each tenant must receive predictable performance and strict isolation while participating in shared infrastructure.

Fabric Domains in Multi-Tenant AI Factories

AI factories and NCPs built with NVIDIA technologies commonly include:

- **East–West fabric (backend, scale-out), which provides a high-bandwidth, low-latency interconnect that supports RDMA-based collective operations and inter-GPU communication within the GPU cluster.** This fabric may be implemented using NVIDIA Quantum™ InfiniBand (native RDMA) or NVIDIA Spectrum-X™ Ethernet configured for lossless transport and RDMA over Ethernet (RoCE).
- **North–South fabric (frontend), which provides external connectivity for ingress, egress, storage access, and service traffic between the GPU cluster and other parts of the data center, cloud environments, and the public Internet.** It is implemented using general-purpose Ethernet. BlueField DPUs may participate in this fabric to enable concurrent multi-tenancy through host-level segmentation. The North–South domain also supports edge connectivity services, including routing, NAT, and firewalls.
- **NVIDIA NVLink™ Multi-Node rack-scale fabrics, shipping with NVIDIA GB200 NVL72 and GB300 NVL72 systems, provide the high-bandwidth, low-latency fabric that forms a unified GPU collective for tightly coupled AI workloads.** These systems can be partitioned to define one or more GPU collectives.
- **Management Network:** Often implemented as a logically isolated domain within the North–South fabric; in some deployments, provided as a separate management fabric.

Each fabric applies its own isolation mechanism. In large-scale multi-tenant AI factories, this creates a requirement for an architectural model that treats these connectivity fabrics as parallel isolation planes, together defining the tenant’s complete operational scope.

Tenant Isolation Across Heterogeneous Fabrics

AI factories and NCPs may simultaneously incorporate several types of connectivity fabrics, each defining tenant boundaries using its own isolation construct:

- **VRF / VXLAN** on Ethernet fabrics,
- **Partition Keys (PKeys)** on InfiniBand fabrics (the InfiniBand equivalent of a VRF: defines which adapters can communicate within a partition),
- **NVL partitions** on systems equipped with an NVLink Multi-Node fabric, such as NVL72.

In shared AI factories and NCPs, properly aligned tenant boundaries across all participating connectivity fabrics reduce operational and security risks to the business stemming from error-prone, manual configuration and coordination processes or inadvertent exposure of one tenant's data to another.

Purpose of the Multi-Tenancy Model

This model establishes the foundation for multi-tenancy in DGX-based AI factories and helps operators to:

- Maximize GPU utilization by dynamically reallocating capacity across tenants
- Provision tenants immediately with hard isolation (enforced on networking hardware) configured automatically
- Eliminate misconfigurations that disrupt customer workloads
- Support rapid tenant onboarding and lifecycle at AI-fabric scale.

The next section examines the technical challenges involved in implementing this model across heterogeneous fabrics.

Multi-tenancy Challenges in AI Infrastructure

Running a multi-tenant AI factory means keeping tenant isolation consistent across three completely independent fabrics — Ethernet, InfiniBand, and NVLink — each with its own control plane, its own configuration syntax, and no native awareness of the others. Every operator building on this infrastructure will encounter the following challenges.

1. Coordinating Many Independent Configurations Across Multiple Fabrics

AI factories and NCPs routinely create, expand, shrink, or retire tenant environments (e.g., adding or removing servers from a tenant). Each lifecycle event must be reflected across all participating fabrics to maintain alignment which can include changes to VRFs, VXLANs, IP ranges, switch ports, RoCE, InfiniBand PKeys and HCA GUIDs, NVLink partitions and GPU UIDs, DPUs, edge policies, and many other highly specialized parameters. Manually applying configuration changes across Ethernet, InfiniBand, and NVLink fabrics introduces complex, time-consuming, and error-prone workflows.

Dynamic hard isolation (enforced on networking hardware) — the only viable foundation for securely sharing GPU resources across multiple tenants — is not operationally achievable through manual processes.

2. Security and Operational Challenges

Misaligned tenant boundaries can expose one tenant’s traffic to another, disrupt workloads, or violate expected isolation. These issues create operational, security, and compliance risks, including higher remediation costs, reduced customer confidence, regulatory scrutiny, financial penalties, and loss of usable GPU capacity.

3. Limited Visibility Across Fabrics

Operators cannot easily see how Ethernet, InfiniBand, and NVLink tenant boundaries relate to one another, making misconfigurations difficult to detect until they cause an incident.

4. Inconsistent Reconfiguration Timing

Updates often apply at different speeds across fabrics, creating short windows where tenant isolation is incomplete or misleading.

5. Rollback Difficulty After Failures

Undoing a failed tenant change requires coordinated reversions across all fabrics, and any mismatch can prolong outages or expose traffic.

6. Validation Complexity Before Changes

It is difficult to determine, in advance, whether a change to one fabric will preserve alignment with the others.

7. Configuration Drift Across Toolchains

Independent automation systems or manual workflows can introduce subtle, compounding differences across fabrics over time.

8. Staffing and Knowledge Retention

The required domain expertise across Ethernet, InfiniBand, and NVLink fabrics is difficult to find and expensive to retain. Building in-house cross-fabric automation requires years of time and significant capital, and maintaining it pulls engineering focus away from the GPU cloud business itself. When the engineers who built or operate these workflows leave, they take a significant portion of institutional knowledge with them. Encoding operational intent in a platform — rather than in individual expertise — is a meaningful form of business risk mitigation.

9. Limited Automation Standards

There is no shared, cross-fabric representation of tenant intent at the hardware or

fabric-management layer, leading operators to rely on custom tooling when higher-level orchestration is required.

10. Scaling Pressure From Rapid Tenant Turnover

Fast-changing or short-lived tenant environments increase the frequency and volume of multi-fabric updates.

The following section describes the architectural practices that address each of these challenges.

Recommendations for Implementing a Multi-Tenant Environment

At the AI-factory scale, operators benefit from a unified tenant abstraction, such as a Virtual Private Cloud (VPC), to help reason consistently about isolation and membership across heterogeneous fabrics. The following architectural practices support this objective.

1. Establish a Unified Tenant Abstraction

Define the tenant as a high-level abstraction that includes an isolated communication boundary along with the set of workloads assigned to it. The nature of the workloads may evolve, but the tenant abstraction remains constant, providing a stable way to reason about isolation and membership regardless of the connectivity fabrics in the AI factory.

2. Translate the High-Level Abstraction into Fabric Configurations

Programmatically translate this tenant abstraction into the appropriate configurations for each fabric, including:

- VRFs and VXLANs on Ethernet,
- PKeys on InfiniBand,
- and NVLink partition membership when NVLink Multi-Node fabrics are present.

Deriving these configurations from a high-level abstraction reduces the need for operators to define and manage fabric-specific and switch-specific configurations.

3. Apply Tenant Configurations Across Fabrics Without Disruption

Tenant creation, resizing, and retirement actions must be applied across all fabrics without causing configuration conflicts and preserving the operational and isolation boundaries of all existing tenants.

4. Validate Topology and Isolation Continuously

Continuous validation helps detect divergences between operator intent and the running state of all fabrics. This validation should confirm that partition membership, routing boundaries, and GPU collectives conform to the tenant abstraction. Proactive validation helps prevent outages and reduce the risk of data loss.

5. Provide an API for Upstream Platforms

AI factories and NCPs frequently integrate networking with higher-level orchestration layers such as IaaS and PaaS systems and AI factory management platforms such as NVIDIA Mission Control. Exposing a consistent high-level abstract networking API allows these systems to integrate with the networking layer without requiring specialized fabric-specific knowledge across multiple control planes.

The remainder of this paper describes one implementation of these practices: the Netris NAAM (Network Automation, Abstraction, and Multi-Tenancy) platform.

Architecture of a Multi-Tenant AI Cloud with Netris

The Netris NAAM (Network Automation, Abstraction, and Multi-Tenancy) platform implements each of the practices described earlier. Operators define tenant isolation in high-level terms — which servers belong to which tenant — and Netris translates that into the configurations required by every fabric involved: Ethernet VRFs, InfiniBand PKeys, NVLink partitions, and DPU policies.

In this role, Netris serves as a foundational coordination layer of AI networking infrastructure, aligning tenant-related configuration across the participating fabrics while integrating with NVIDIA fabric management systems.

In the most basic sense, the IaaS or PaaS operator requests an action of creating, modifying, or retiring a tenant using a Netris API endpoint, and the Netris algorithm translates this request into the appropriate configurations for the Ethernet, InfiniBand, and NVLink Multi-Node fabrics and applies them consistently and securely across the environment.

Figure 1 shows that Netris sits between the upstream platform and the individual fabric control planes — the upstream platform interacts with Netris alone, while Netris translates tenant intent to UFM, NMX-C, and the Ethernet fabric independently.

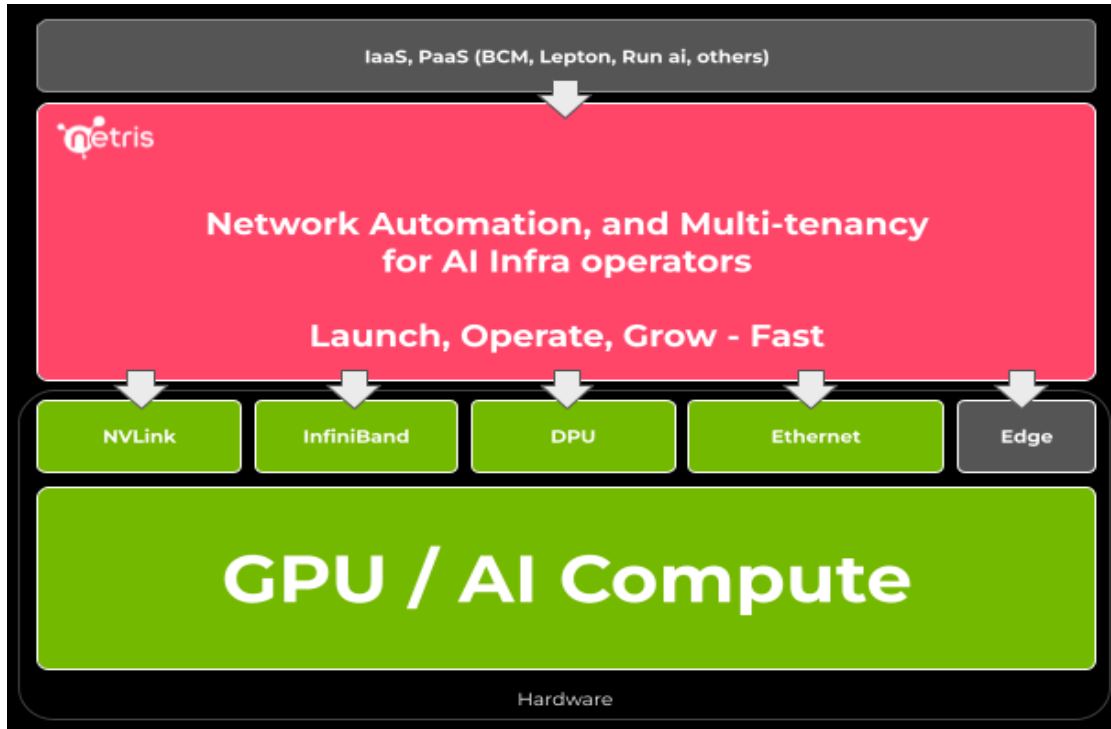


Figure 1. Integration of Infrastructure and Platform Layers Using Netris within the NVIDIA AI Ecosystem

Interpreting Operator Intent and Supporting Alignment Across Fabrics

Operators or upstream platforms define a tenant using a high-level abstraction, such as a Netris Server Cluster—essentially a set of servers allocated to a tenant. The Netris algorithm automatically translates this high-level definition into the fabric-specific configurations required by the participating connectivity fabrics:

- Spectrum-X East-West Ethernet:** Netris functions as the fabric manager for this Ethernet fabric. Netris automatically generates tenant VRFs and L3VPN VXLAN configurations compliant with the appropriate version of the NVIDIA Spectrum-X Deployment Guide, and then assigns switch ports to the appropriate VXLANs. At the operator’s discretion, the optional Netris Host Networking (NHN) plugin can also apply the required BlueField-3 SuperNIC configurations, such as RoCEv2, to each compute node in compliance with the appropriate version of the NVIDIA Spectrum-X Deployment Guide for lossless inter-GPU communication.
- InfiniBand East-West:** Netris integrates with the NVIDIA Unified Fabric Manager (UFM)—NVIDIA’s own controller for managing InfiniBand fabrics—to discover and

synchronize server inventory and HCA GUIDs for each registered server node. Netris, through the NVIDIA UFM API, creates, modifies, and removes InfiniBand partitions (PKeys) based on the high-level tenant definitions (Server Cluster), then associates the relevant HCA GUIDs with the appropriate PKeys.

- **NVLink Multi-Node Fabrics:** Netris integrates with the NVIDIA NMX Controller (NMX-C)—the NVLink control plane application for managing NVLink Multi-Node fabric switches. When NVLink Multi-Node fabrics, such as NVL72, are present, Netris discovers, creates, modifies, and deletes NVLink partitions as appropriate and assigns the relevant GPU GUIDs to these NVLink partitions using the NVIDIA NMX-C gRPC API.
- **North-South Ethernet Fabric:** Netris functions as the fabric manager for this Ethernet fabric. Netris configures the tenant VRF and VXLAN-based routing domains for ingress, egress, storage access, and service communication. DPU configurations inherit the same tenant definition. Netris applies the relevant configurations consistently across switches and DPUs, and provides the automation that aligns North-South connectivity with the rest of the fabrics in the AI factory.
- **Edge Connectivity:** When deployed, Netris SoftGate, a multi-tenant edge gateway, provides tenant-aware NAT, Layer-4 Load Balancing, and other edge services automatically aligned within the tenant boundaries.

Figure 2 shows that Netris operates as the network control layer between the upstream platform and the underlying fabric systems — UFM, NMX-C, Ethernet switches, and DPUs — handling cross-fabric consistency so the platform layer above never needs to interact with individual fabric controllers directly.

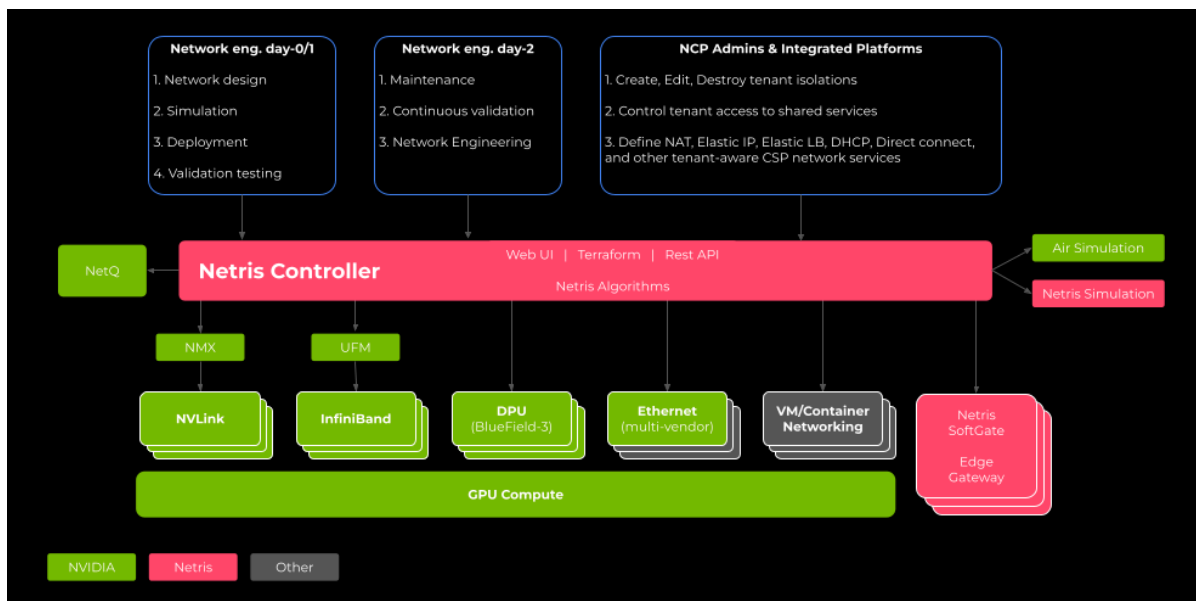


Figure 2. Operational Architecture of Netris in an NVIDIA-Based AI Infrastructure

This unified intent-driven model supports the implementation of dynamic, multi-tenant clouds on NVIDIA DGX-based architectures by interpreting the intent expressed by upstream platforms in one place and automatically and dynamically applying the relevant configurations consistently across all connectivity fabrics participating in the AI factory.

Upstream Platform Integration

AI factories typically operate within a broader platform stack that includes user portals, billing systems, PaaS layers, and administrative tooling. These upstream systems are responsible for user interaction, policy, and workload orchestration, and are often custom-built by the operator.

Netris integrates as one of several infrastructure management components consumed by the platform. It focuses specifically on network-related tenant isolation and connectivity across the underlying fabrics, while remaining independent from compute and storage management systems.

Figure 3 illustrates this integration model. Upstream platforms coordinate tenant lifecycle actions and interact independently with networking, compute, and storage management components. Netris operates alongside compute and storage managers, translating tenant intent into network configuration and ensuring that network isolation remains aligned with the overall platform architecture.

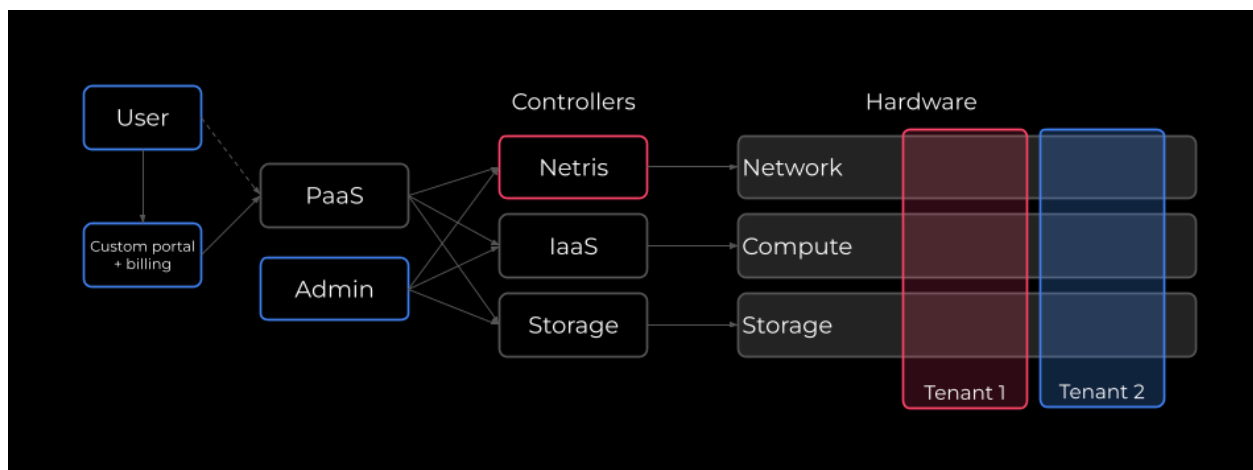


Figure 3. Netris platform integration architecture

Upstream platforms, such as PaaS layers or administrative tools, typically interact with Netris through a stable API to express tenant lifecycle intent—for example, creating, expanding, or retiring a tenant. This interaction does not require the upstream platforms to

directly manage fabric-specific details or communicate with individual network control planes.

In parallel, the same upstream platforms interact with infrastructure services responsible for compute and storage. Each of these systems—networking, compute, and storage—retains responsibility for its own domains. Netris does not replace upstream orchestration platforms or infrastructure managers; it provides a consistent integration point for applying tenant-related network isolation across Ethernet, InfiniBand, and NVLink fabrics.

Once tenant isolation is established at the network level, compute and storage systems operate within those boundaries using their existing mechanisms. In this way, tenant environments span network, compute, and storage domains, while each domain continues to be managed by the systems designed for that purpose.

In NVIDIA Mission Control-managed environments, Netris operates in the same way: Mission Control components manage cluster and workload lifecycle, while Netris provides the network automation and multi-tenancy layer, integrating with UFM and BCM through the same interfaces described in this paper.

Separation of Responsibilities

Netris architecture helps operators maintain a clear division of responsibilities:

- **Network engineering teams:** design and validate connectivity fabrics using Netris, which implements the orchestration model, integrates with NVIDIA management systems, and serves as the Ethernet fabric manager.
- **Platform operators:** manage tenant lifecycle and express intent through Netris northbound API.
- **NVIDIA management systems (UFM, NMX-C):** maintain inter-GPU network fabric health and GPU operational behavior. In Mission Control-managed environments, UFM is delivered as a Mission Control component.

This architectural approach supports multi-tenancy at scale, consistent tenant security boundaries, and complex distributed operations.

Simulation

Many operators use simulation environments to explore network designs, tenant boundaries, and planned changes before touching a live AI factory.

NVIDIA DSX Air is NVIDIA's SaaS platform for simulating AI factory infrastructure. Introduced at GTC 2026, it creates high-fidelity digital twins of the full hardware stack — GPUs, SuperNICs, DPUs, and switches — and integrates with partner solutions for storage, orchestration, security, and networking through open API-based connectivity. A complete AI factory environment can be constructed, configured, and validated in DSX Air before a single server is unboxed.

Netris integrates with DSX Air, allowing operators to include network automation and multi-tenant configurations in the simulation alongside compute and storage systems. Tenant isolation constructs, VPC boundaries, and fabric configurations built in Netris can be tested in DSX Air exactly as they would be deployed in production.

Operators use simulation across three phases: pre-deployment validation, where designs are tested before hardware arrives; proof-of-concept work, where new architectures or tenant models are explored without risk to live environments; and ongoing change management, where upgrades, tenant changes, and maintenance procedures are rehearsed before being applied to production. By shifting integration and troubleshooting into simulation, operators reduce the time from hardware delivery to first operational workload and lower the probability of configuration errors reaching the live environment.

An Example Network Isolation Implementation With Netris

This section illustrates how Netris orchestrates network provisioning and tenant isolation across the different fabric domains of an NVIDIA-based AI infrastructure. The examples highlight how multi-fabric configuration and tenant alignment are automated through the Netris Server Cluster construct, supporting consistent behavior across InfiniBand, Ethernet, NVLink Multi-Node, and management networks.

Network Provisioning with Netris

A Server Cluster is the primary mechanism Netris uses to define tenant boundaries; in its most basic form, it is a list of servers assigned to a tenant. The Netris Controller automatically creates a VPC and all other relevant isolation primitives such as V-Nets, subnets, etc. in the Netris Controller database based on a predefined Server Cluster Template. Then, using this list of servers along the topology information stored in Netris Topology Manager, discovered GPU UIDs, HCA GUIDs, and other metadata, the Netris algorithm generates the following:

- which switch ports participate in the tenant's VRF and VXLAN,
- which HCAs belong to the tenant's InfiniBand partition (PKey),
- and which GPU identifiers (GPU ID) are included in the tenant's NVLink partition.

Netris then distributes the appropriate configuration to NVIDIA UFM, NMX-C, and directly to the Ethernet components, including SuperNICs and DPUs.

This approach aligns tenant isolation across Ethernet, InfiniBand, and NVLink Multi-Node domains without requiring operators to configure each fabric directly.

Tenant "Alpha" is used as a symbolic example; vrf-alpha, vni-alpha, switchport-set-alpha, pkey-alpha, hca-set-alpha, nvl-alpha, and gpu-set-alpha represent abstract configuration constructs derived from the Server Cluster definition.

East-West Fabric

The East-West fabric carries inter-GPU communication and collective operations traffic within a GPU cluster.

If East-West is implemented using Spectrum-X Ethernet

Tenant Alpha's East-West Ethernet boundary is represented through:

- **VRF:** vrf-alpha
- **VNI:** vni-alpha
- **Switch Port Set:** switchport-set-alpha (list of switch ports)

These constructs follow the isolation and transport requirements described in the NVIDIA Spectrum-X deployment guide. Netris configures these VRFs, VNIs, and inserts the appropriate switch ports into the correct VXLANs (VNI) on every Ethernet switch in the East-West fabric, so that the routing domain for inter-GPU communication remains consistent with the intended tenant boundary.

GPU collective operations — all-reduce, all-gather, and similar patterns used in distributed training — are exceptionally sensitive to packet loss. RDMA-based collectives stall when any participant drops a packet, because the operation cannot complete until all endpoints respond. NVIDIA Spectrum-X achieves the required lossless behavior through Priority Flow Control (PFC), which applies per-hop backpressure to prevent drops at congested switch ports, and Data Center Quantized Congestion Notification (DCQCN), which signals endpoints to reduce injection rates before drops occur. Spectrum-X reference architectures also use a rail-optimized topology, in which GPUs that participate in the same collective operations connect through a common leaf switch, keeping most collective traffic within a single switch tier. Netris applies the required RoCEv2 settings on BlueField-3 SuperNICs through the NHN plugin, ensuring each tenant's compute nodes comply with the lossless transport model defined in the NVIDIA Spectrum-X Deployment Guide.

If East-West is implemented using InfiniBand

For InfiniBand-based East-West fabrics, Tenant Alpha's communication boundary is defined by an InfiniBand partition and its associated host channel adapters (HCAs):

- **InfiniBand Partition Key (PKey):** pkey-alpha
- **HCA membership set:** hca-set-alpha (list of HCA GUIDs — the hardware identifiers for InfiniBand adapters, analogous to MAC addresses)

Netris integrates with NVIDIA UFM to collect and maintain an inventory of servers and their corresponding HCA GUIDs. This data is stored in the Netris Controller and used to derive InfiniBand partition membership from the Server Cluster definition.

When a tenant is created, Netris creates the corresponding InfiniBand PKey in UFM and associates the required HCA GUIDs with that partition. When the tenant is resized, Netris updates the membership of the existing PKey it previously created. When the tenant is

deleted, Netris removes the associated PKey from UFM. In this way, the lifecycle of the InfiniBand partition is aligned with the tenant lifecycle defined in Netris.

The assigned PKey defines the communication domain for HCAs belonging to the tenant. Netris keeps in-network computing services, such as SHArP-based collective acceleration (an InfiniBand in-network computing capability that offloads collective operations to network switches, reducing host CPU involvement), aligned with the tenant's InfiniBand partition boundary.

North-South Fabric

The North-South fabric supports ingress, egress, storage access, and service connectivity for each tenant. Tenant Alpha's routing boundary is configured using:

- **VRF:** vrf-alpha
- **VNI:** vni-alpha
- **Switch Port Set:** switchport-set-alpha (list of switch ports)

This domain remains consistent whether East-West is implemented via Spectrum-X Ethernet or InfiniBand. Netris derives routing, segmentation, and forwarding behavior using vrf-alpha, vni-alpha, and switchport-set-alpha, and applies these semantics across the North-South connectivity layer.

BlueField DPUs on the North-South fabric interpret the same constructs to align tenant routing and service policy at the host edge.

At the North-South fabric perimeter, SoftGate can provide routing, NAT, and related ingress/egress services, promoting consistent isolation across ingress and service paths.

NVLink Multi-Node Fabrics (NVL72)

When NVLink Multi-Node fabrics are present, Tenant Alpha's GPU grouping is represented as:

- **NVLink Partition Identifier:** nvl-alpha
- **GPU Identifier Set:** gpu-set-alpha (list of GPU IDs)

Netris automatically discovers and maintains an inventory of GPU IDs for all servers registered in the Netris inventory by synchronizing with NVIDIA NMX-C instances in each NVL72 rack. From this inventory Netris builds a list of GPU IDs based on the list of member

servers in the Server Cluster and then communicates this grouping to NVIDIA NMX-C using the supported NMX-C gRPC API.

Management Networks (Out-of-Band)

AI infrastructure typically includes a dedicated management network to provision, monitor, and operate infrastructure components independently of the tenants' data-plane traffic. This network provides access to server management interfaces, DPU management interfaces, and other control endpoints and is intentionally isolated from tenant workloads.

In some architectures, this management network is referred to as out-of-band (OOB), a term historically used to allude to management connectivity that is isolated from tenant networks and often implemented as a physically separate network.

In practice, management connectivity may be deployed either as a logically isolated domain within the North–South Ethernet fabric or as a physically separate Ethernet fabric dedicated to management traffic. Both deployment models are common in DGX-based AI factories and are supported by Netris.

Regardless of physical topology, management connectivity is carried over Ethernet and uses the same isolation primitives applied elsewhere in the architecture. A management network is configured using:

- a dedicated routing context (VRF),
- an associated VXLAN overlay identifier (VNI), and
- a defined set of switch ports that provide connectivity to management interfaces.

These primitives are used consistently whether the management network is implemented on the shared North–South fabric or on a standalone management fabric.

Netris uses topology information and metadata (e.g., interface labels) to ensure that interfaces connected to management endpoints—such as server management ports—are consistently placed into the management network. This allows management connectivity to remain stable and isolated as servers are added, removed, or reassigned across tenant environments, without requiring manual reconfiguration.

One management-plane consideration is specific to automation platforms: the system that manages the fabric cannot depend on that same fabric for its own reachability. Netris addresses this with a Controller Management Network (CMN) — a small, statically configured, non-Netris-managed Ethernet network that connects the Netris Controller nodes. Management connectivity for all other devices is aggregated by the OOB network, which connects to the CMN either directly (in smaller deployments) or through a set of

seed switches provisioned before the rest of the fabric (in larger deployments). In both cases, tenant data-plane forwarding does not depend on the availability of either management network.

The Netris Management Network Reference Architecture describes this design in greater detail and is available by contacting sales@netris.io.

Network Services for Tenants

Multi-tenant AI environments often require network services at the tenant boundary, interconnections between tenant VPCs, and reachability to and from external networks. These services are distinct from the core fabric routing and isolation functions and are commonly shared across tenants.

VPC-peering (Shared storage and more)

VPC peering provides controlled Layer-3 connectivity between two or more VPCs within the same Netris-managed fabric. A peering relationship allows selected IP prefixes to be exchanged between VPCs while preserving each VPC's independent routing domain and isolation boundaries.

Conceptually, a VPC peering relationship can be viewed as a routed boundary between two tenant networks. Each VPC retains its own routing context, and only explicitly permitted IP prefixes are advertised across the peering relationship. Prefix-lists define which routes are exchanged. Layer-2 domains are never extended across the peering relationship. The operator should ensure that the IP address space of the peered VPCs does not overlap to avoid reachability conflicts.

Netris implements VPC peering using the VRF route leaking feature.

Although VPC peering is technically feasible wherever routed connectivity exists, it is primarily intended for use on the North–South fabric. In practice, VPC peering is most commonly used to provide controlled access to shared services—such as storage systems or platform components—without collapsing tenant isolation or introducing direct connectivity between tenant VPCs.

Shared Storage Access via VPC Peering

Storage systems in AI factories are frequently shared across multiple tenants. Unlike compute resources, storage endpoints are often not dedicated on a per-tenant basis and may have practical limits on the number of logical interfaces they can expose. As a result, storage isolation and sharing are typically achieved through a combination of network segmentation and storage-platform controls rather than dedicated hardware per tenant.

A common and scalable pattern is the Shared Services VPC model, in which storage endpoints are placed into a dedicated VPC. Tenant VPCs establish VPC peering relationships with this shared services VPC, allowing tenants to access storage through controlled, routed connectivity.

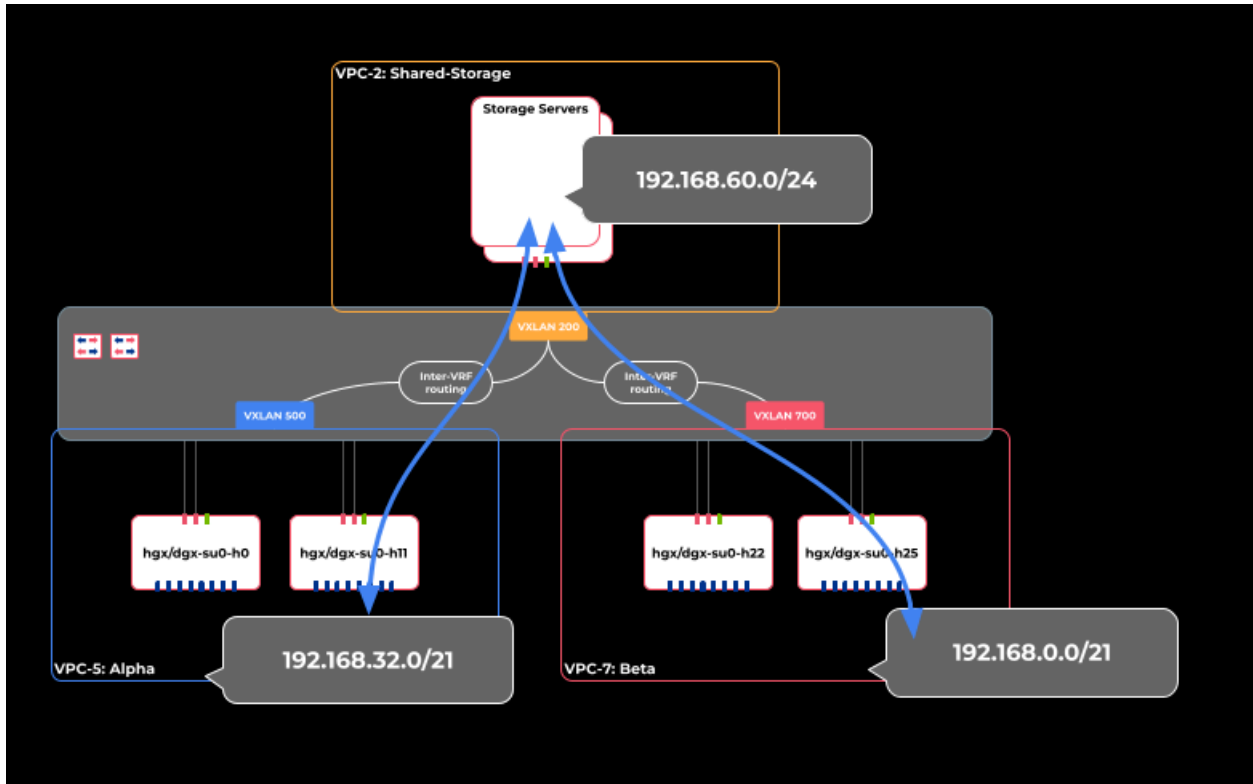


Figure 4. Shared storage access with VPC peering

This approach:

- Preserves strict isolation between tenant VPCs
- Avoids per-tenant scaling limits on storage-side interfaces
- Centralizes storage access control and lifecycle management
- Simplifies onboarding and offboarding of tenants

Each tenant’s access to storage is governed by the VPC peering relationship and the associated prefix-lists, ensuring that only explicitly permitted prefixes are reachable.

Alternative: Shared Storage Access via VLAN Subinterfaces

Some storage platforms support exposing logical interfaces—such as subinterfaces or virtual ports—directly into each tenant’s VPC. In this model, the storage system participates directly in each tenant’s North–South routing domain, eliminating the need for VPC peering to access storage.

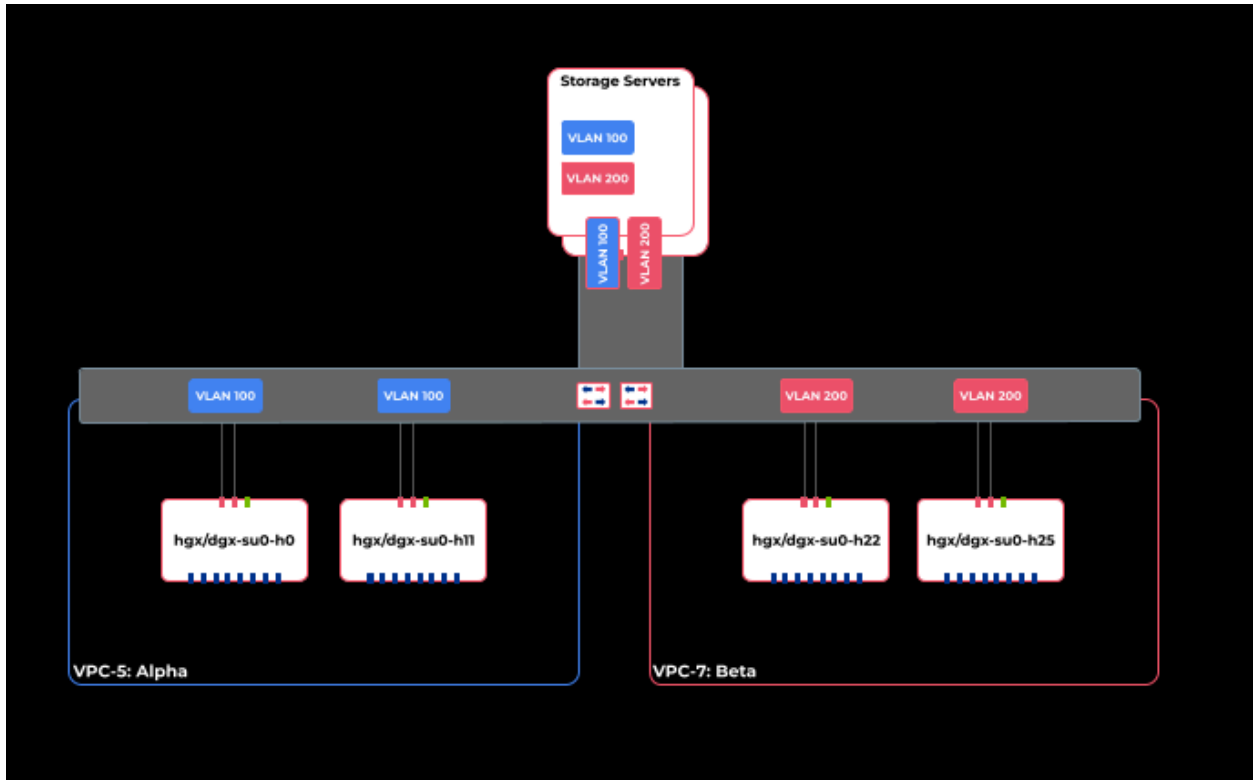


Figure 5. Shared storage access with VLAN subinterfaces

While this approach can provide clear per-tenant visibility and straightforward routing semantics, its scalability is constrained by the storage platform’s ability to support a large number of logical interfaces. Operators should account for these limits when designing environments.

VPC Connect

VPC Connect provides a mechanism for connecting tenant VPCs with external networks, including non-Netris-managed data center networks, WAN routers, and security appliances.

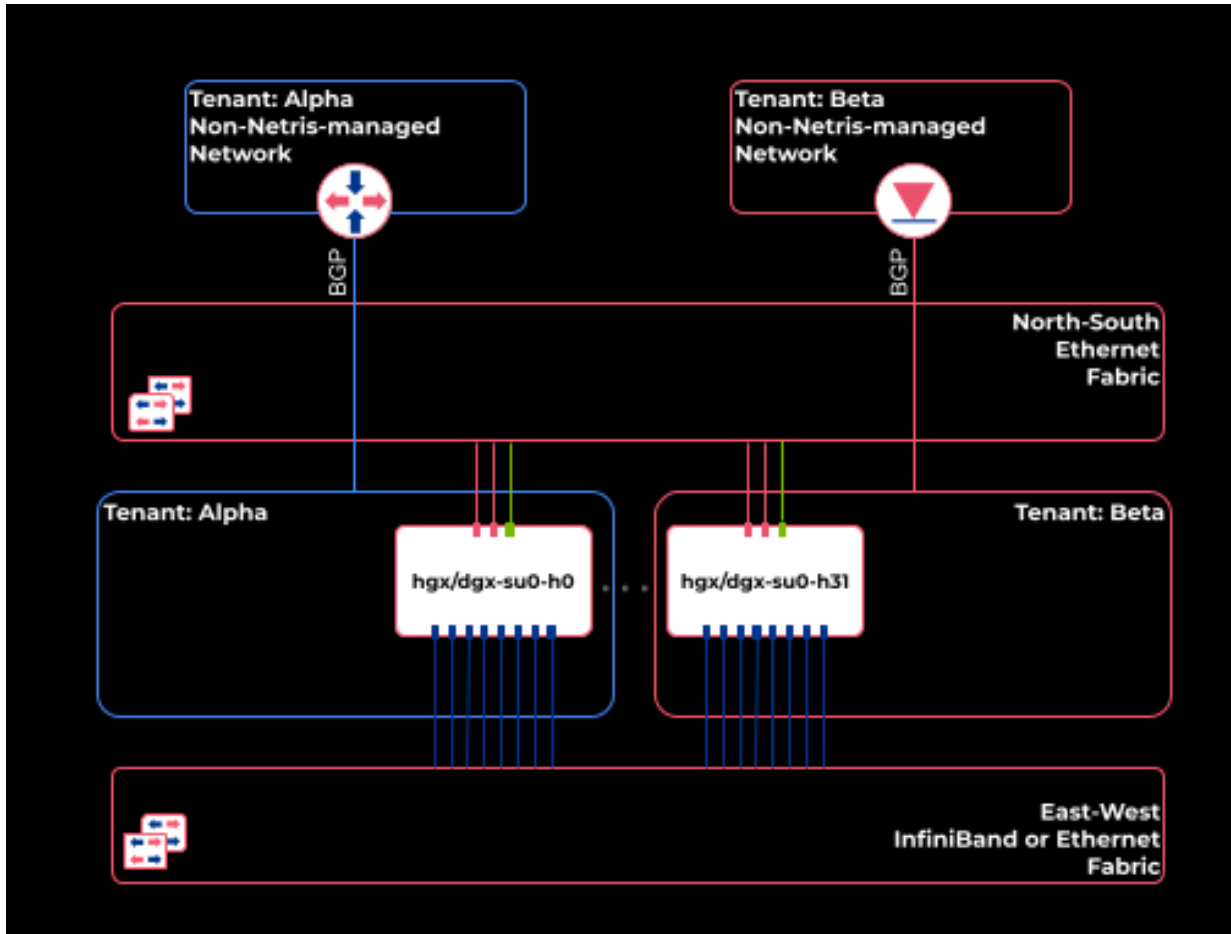


Figure 6. Netris VPC Connect

In this model, Netris exposes a switch port on the North-South fabric. When a single-tenant VPC is in scope, the port is configured as a routed port, and the eBGP neighbor relationship terminates directly on that switch port. When multiple tenant VPCs are in scope, the port is configured as IEEE 802.1Q tagged, with each tenant mapped to a dedicated VLAN, and eBGP neighbor relationships terminate on corresponding SVIs (Switch Virtual Interfaces) established on the switch or elsewhere in the fabric, as determined by Netris.

Route exchange between tenant VPCs and the external network is explicitly controlled using route-maps and prefix-lists, allowing operators to define which prefixes are advertised in each direction while preserving clear routing boundaries.

The external endpoint—such as a router, firewall, or other network transport device—is typically owned and operated by the operator’s networking team.

Edge Services (NAT, Layer-4 LB)

Typically, AI workloads need Internet access. Inference workloads need to be accessed from the Internet. Additionally, some workloads need access to the Internet for integrations with third-party services and other administrative use cases (downloading packages, performing DNS lookups, etc.).

In practice, two classes of services are foundational at the VPC boundary.

- Tenant-aware (VPC-aware) Network Address Translation (NAT) enables controlled ingress and egress, allowing workloads inside a VPC to initiate outbound connections and, where required, be accessed from external networks without exposing internal addressing.
- Tenant-aware (VPC-aware) Layer-4 load balancing provides a stable service endpoint (frontend public IP address) for high-volume inference and application traffic, distributing TCP and UDP flows across backend workloads within a tenant's routing domain.

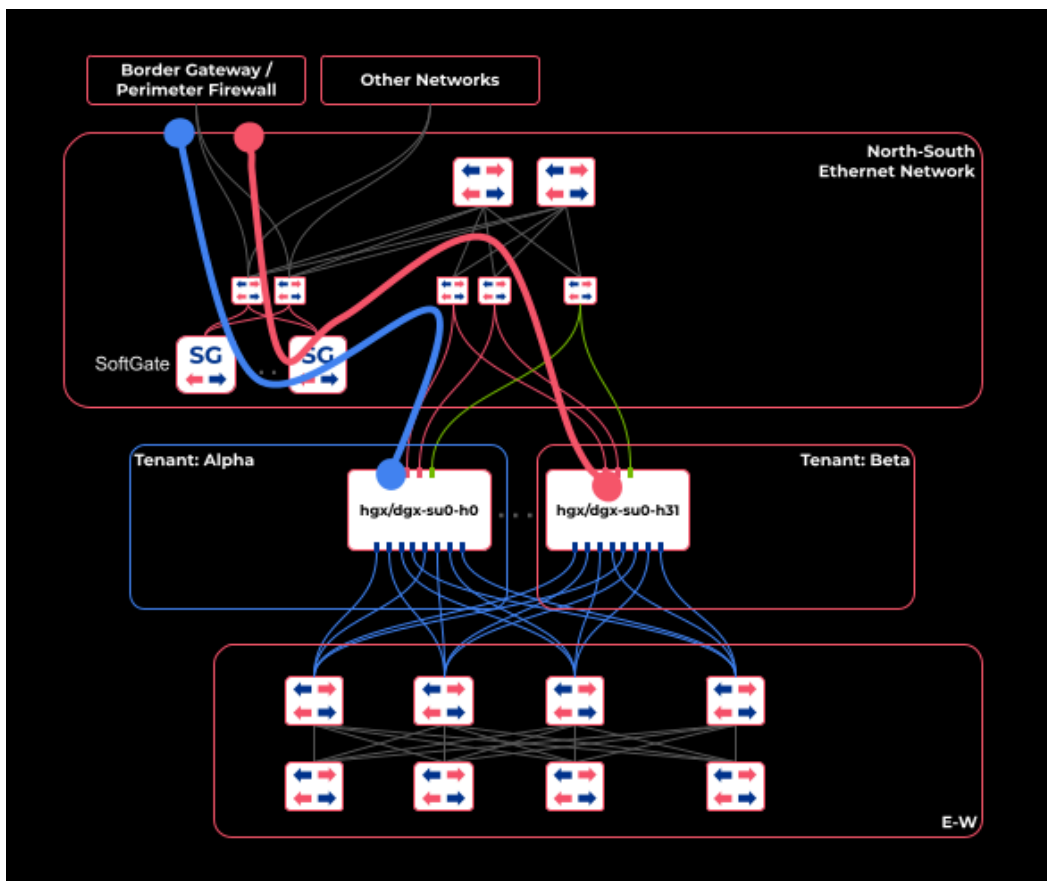


Figure 7. VPC edge services provided by Netris SoftGate

Netris SoftGate is an optional, horizontally scalable edge services layer that operates as an integrated extension of the Netris-managed North-South fabric. Deployed as a cluster of dedicated bare-metal nodes running the SoftGate dataplane and agent, it delivers VPC-aware network edge services for controlled ingress and egress while remaining consistent with tenant isolation boundaries defined in Netris.

The primary services provided by SoftGate include:

- Network Address Translation (NAT) for controlled ingress and egress
- Layer-4 load balancing for TCP and UDP traffic, including health checks
- DHCP services, where required

Access Control Lists

Access Control Lists (ACLs) in Netris are expressed as a high-level policy associated with a tenant VPC, rather than as a configuration bound to a specific device or service. Operators define the desired access-control intent without specifying where that policy must be enforced within the infrastructure.

Netris determines the enforcement point automatically. Depending on the deployment, access control policies may be enforced on fabric switches, DPUs, host-based enforcement points, or edge service components.

This approach reduces cognitive load and misconfiguration risk while ensuring consistent enforcement aligned with tenant network boundaries. Access Control Lists operate independently of NAT, load balancing, and DHCP services and do not require deployment of Netris SoftGate.

Concurrent Multi-Tenancy / Host-Level Segmentation

So far, this document has described a multi-tenancy use case in which the atomic unit of tenancy is the whole bare-metal server. More specifically, this means that any given bare-metal server in the deployment, and thus all of its GPUs, is assigned to only one tenant at any given time.

Operators may choose to enable concurrent multi-tenancy (scenarios in which bare-metal servers are shared across multiple tenants) in their environments to improve hardware utilization across some or all of their AI clusters. In this model, isolation boundaries extend beyond the network fabric and into the hosts themselves. Host-level segmentation enables a single physical server to be partitioned so that multiple tenants can safely share GPUs, connectivity, and services simultaneously.

This approach changes how tenant isolation must be enforced. Network boundaries can no longer rely solely on server membership; they must extend inside the server and apply consistently to GPUs, host networking, and shared services.

There are two mechanisms available to operators for deploying concurrent multi-tenancy in NVIDIA-based AI infrastructures:

- DPU/hardware-accelerated segmentation.
- EVPN-on-Host software-based segmentation.

DPU-based/Hardware-Accelerated Host Segmentation (HBN)

A BlueField DPU (Data Processing Unit) is a system-on-a-chip, an Ethernet-switch-like device installed inside the server, intended to offload data center networking, storage, and security functions from the general-purpose host CPU to dedicated hardware.

NVIDIA BlueField DPUs can operate in one of 3 modes: regular NIC mode, DPU mode, and Zero-Trust mode.

- In NIC mode, the BlueField card behaves like a traditional network adapter from the host's perspective.
- In DPU mode, the embedded ARM subsystem owns and controls the NIC resources and data path. It means that the on-chip ARM cores manage the network controller functions and traffic before anything reaches the parent host.

- Zero-Trust mode is an enhanced security variant of the DPU mode. In Zero-Trust mode, the host admin can't use typical host paths to access the DPU hardware for administration. DPU management is performed via the embedded ARM and out-of-band BMC interfaces instead.

In the concurrent multi-tenancy use case, the DPU is typically configured in Zero-Trust mode. Host-Based Networking (HBN) is an NVIDIA DOCA service (Data Center Infrastructure-on-a-Chip Architecture—NVIDIA's software framework for DPU application development). HBN runs on the DPU in Zero-Trust mode and implements server-side routing and network services using standard Layer-3 protocols.

Zero-Trust DPU mode with HBN enables hardware-based tenant isolation to be securely enforced outside the host operating system, even when multiple tenants share a single physical host.

Each DPU typically includes:

- A dedicated management port used for access to the DPU operating system and the DPU's BMC
- One or more high-speed uplink ports connected to the North-South Ethernet fabric

When operating in DPU mode, the DPU presents Virtual Functions (VFs) to the host. These VFs appear to the host operating system as network interfaces and can be independently configured. From the host's perspective, they behave like regular NICs. From the platform's perspective, each VF is functionally equivalent to a physical switch port — an isolated endpoint with hard isolation (enforced on networking hardware).

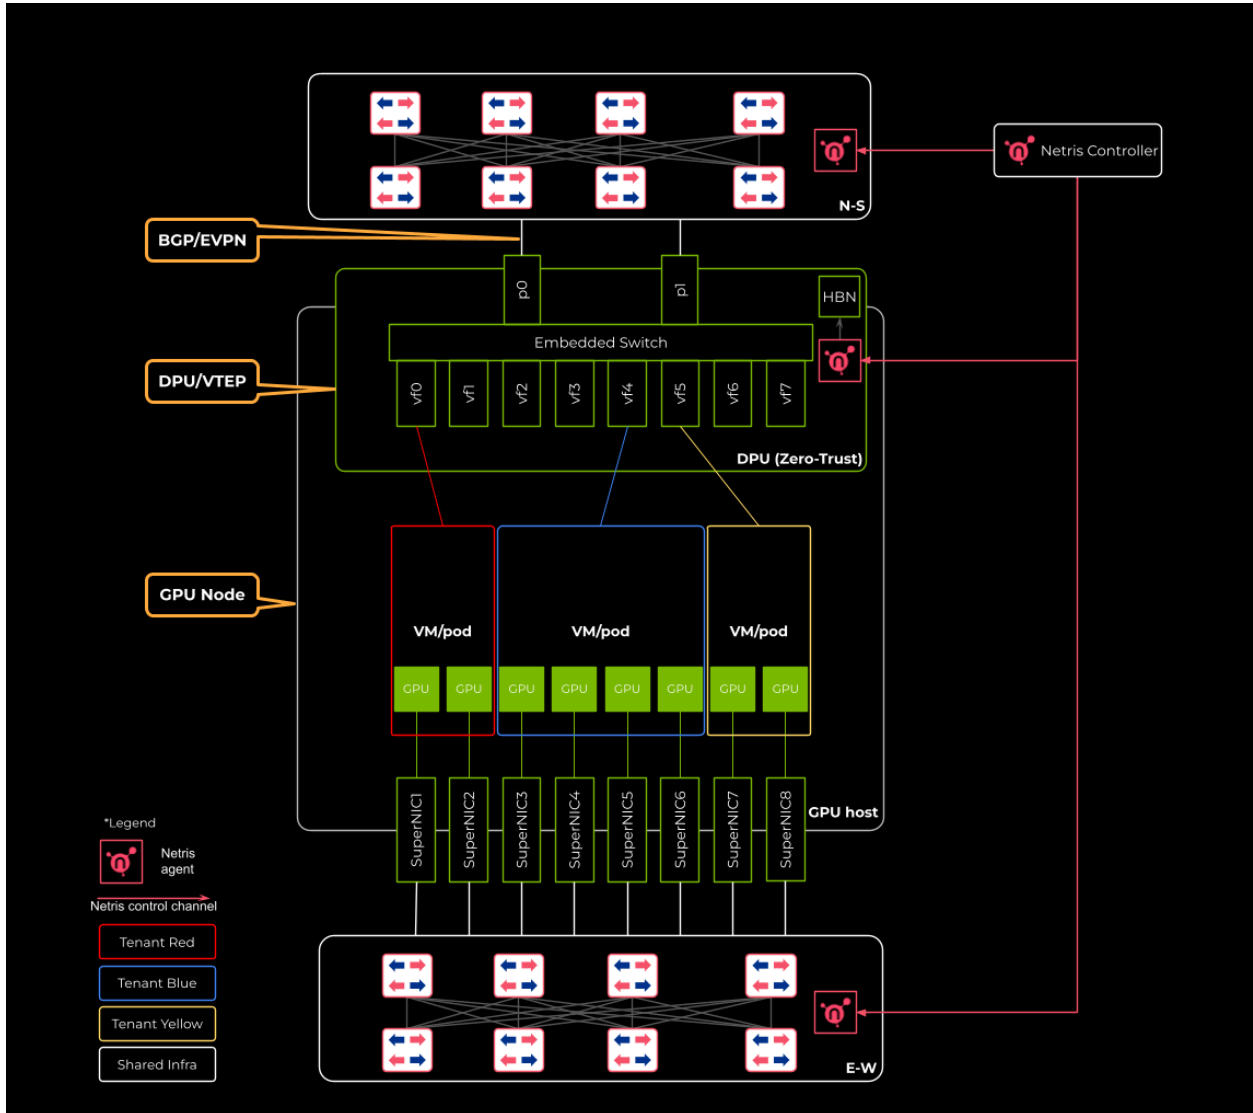


Figure 8. BlueField-3 DPU in Zero-Trust mode with Netris agent

Netris manages BlueField-3 DPUs as network devices in the server inventory. A registered DPU participates in the North–South Ethernet fabric alongside physical switches: it establishes BGP EVPN adjacencies, functions as a VXLAN Tunnel Endpoint (VTEP), and exposes its VFs as assignable endpoints through the same workflow used for switch ports. Because switch ports and DPU VFs participate in the same EVPN control plane, operators can place workloads on DPU-equipped servers and directly-attached non-DPU servers into the same isolation domain, with consistent isolation across both device types.

This model supports two deployment patterns:

- **Single-tenant server with North–South connectivity via DPU.** One DPU per server, with all North–South traffic passing through the DPU, which handles BGP and

VXLAN on behalf of the host. For single-tenant servers, a Physical Function (PF) may be exposed directly to the host instead of a VF, simplifying the data path.

- **Concurrent multi-tenancy on a shared server.** Each tenant's workload binds to one or more VFs. Netris assigns each VF to the appropriate tenant network, and the DPU enforces isolation between tenants in hardware, even though the workloads execute on the same physical host.

Responsibilities are divided between NVIDIA DPF (DOCA Platform Framework) and Netris. DPF manages the DPU device itself — OS installation, firmware, mode configuration — and deploys the Netris agent onto the DPU. The Netris agent configures the DPU's DOCA HBN service as a VTEP, including BGP/EVPN adjacencies, and applies tenant-specific networking configuration, mapping tenant definitions directly to DPU-managed VFs.

DPU-accelerated segmentation is one mechanism for enabling concurrent multi-tenancy. In environments without DPUs, similar isolation semantics can be achieved using host-based techniques described in the following section.

Software-Based Host Segmentation (EVPN on Host)

In software-based host segmentation, tenant isolation is enforced with the host operating system itself, allowing a shared server to participate directly in tenant-specific network domains without dedicated DPU hardware.

This model is commonly used for shared CPU-based infrastructure, such as control-plane nodes for managed Kubernetes services, shared service hosts, or other environments where multiple tenants' resources must coexist on the same server while maintaining strict network isolation.

In this architecture, the host becomes a VTEP in the North-South EVPN/VXLAN fabric. Netris supports this model by deploying the Netris EVPN-on-Host agent to servers in scope for concurrent multi-tenancy. The agent configures the host networking stack to implement the same network constructs used across the rest of the North-South fabric, including:

- Tenant-specific VRFs
- VXLAN bridges and VNI mappings
- Routing tables aligned with tenant VPCs

From the platform's perspective, a host configured with EVPN-on-Host behaves like any other EVPN endpoint in the fabric. Virtual machines or containers running on the server can be connected directly to tenant VPCs (i.e., VXLANs) by attaching them to the

appropriate Linux bridges established by the Netris agent. This approach enables secure communication between containers, virtual machines, GPU resources (whether dedicated or shared), and external infrastructure endpoints—such as storage systems or proprietary appliances—while preserving network-enforced isolation between tenants, implemented with the same EVPN/VXLAN constructs used across the rest of the fabric.

Node Provisioning with NVIDIA Mission Control

NVIDIA Mission Control is NVIDIA's integrated AI factory management platform, combining NVIDIA's operational best practices and AI cluster automation into a single control plane. Mission Control includes NVIDIA Base Command Manager (BCM) for cluster management and node provisioning, NVIDIA Run:ai for workload orchestration, NVIDIA UFM for InfiniBand fabric management, and NVIDIA NetQ for network observability. NVIDIA GB200 and GB300 NVL72 systems require Mission Control for installation and management.

Within this stack, BCM remains responsible for bootstrapping and lifecycle management of DGX nodes. To continuously manage a node, BCM must maintain network connectivity to that system. This requirement influences how DGX nodes are integrated into a multi-tenant network architecture.

Shared Services Model

In practice, operators commonly deploy BCM within a dedicated management VPC. This VPC contains the BCM instance and associated management components and is isolated from tenant workload VPCs.

Tenant VPCs are then selectively peered with the BCM VPC using controlled Layer-3 connectivity—VPC peering. This peering establishes a constrained communication path that allows DGX nodes allocated to a tenant to reach BCM for provisioning and lifecycle operations. At the same time, routing policies prevent lateral connectivity between tenant VPCs and prevent BCM from serving as a transit path between tenants.

In this model, DGX nodes remain associated with tenant isolation constructs at the networking layer (e.g., VPC boundaries, InfiniBand partitions, NVLink partitions), while maintaining the required management-plane reachability to BCM. The result is a separation of concerns: BCM retains lifecycle control of the nodes, and tenant isolation is preserved through network-level segmentation and controlled route exchange.

Network Validation in Multi-Tenant AI Infrastructure

Operating a multi-tenant AI infrastructure requires coordinating multiple independent systems—network fabrics, accelerators, and control planes—that were not designed to collectively validate tenant intent. Each system may appear healthy in isolation, while the overall tenant configuration is incorrect.

Network validation must therefore be system-aware, topology-aware, and explicitly tied to tenant intent rather than to individual device state.

Netris maintains a canonical tenant definition using high-level constructs such as VPCs and Server Clusters. Based on this definition, Netris algorithms automatically generate the appropriate configuration for Ethernet fabrics, and the appropriate partitioning information for InfiniBand and NVLink fabrics.

Topology-Aware Validation

Validation must account for network topology. It must be clear how switches and servers are connected and whether this connectivity matches the operator's declared design.

Netris acts as the fabric manager for Ethernet fabrics. It maintains the expected Ethernet topology as defined by the operator's declared configuration.

Netris continuously verifies that the live Ethernet fabric topology matches the declared topology maintained in the Netris topology manager. This validation covers both switch-to-switch links and switch-to-server links. For each declared connection, Netris confirms that the corresponding physical interfaces are connected to the expected peer device and port. When the observed physical connectivity does not match the declared topology, Netris raises a notification event to indicate the mismatch.

Detection of Partial States

Tenant lifecycle changes are applied across multiple systems. Validation must show whether these changes were fully applied.

Netris tracks the status of the configuration application during tenant operations. It surfaces cases where the configuration was not applied, failed to apply, or was only applied to part of the environment.

This allows operators to identify incomplete changes directly.

Ongoing Validation and Drift Detection

Over time, the deployed configuration can diverge from the tenant definition due to manual changes or system updates. Validation must therefore be ongoing.

Netris treats the declared tenant intent stored in the controller as authoritative and continuously reconciles the deployed configuration against it. When Netris detects drift, or when repeated configuration changes indicate that the applied configuration is not stable, it raises a notification event to alert the operator.

Data-Plane Reachability Validation in EVPN Fabrics

In EVPN VXLAN-based Ethernet fabrics, correct operation requires VTEP-to-VTEP reachability across the fabric.

Netris tracks reachability between all VTEPs in each managed EVPN fabric. Netris verifies that every VTEP can reach every other VTEP and surfaces issues if reachability is lost.

Validation That Scales With Change

As tenant onboarding and the rate of changes increase, validation must scale with them.

Netris centralizes tenant definitions, exposes cross-fabric membership through APIs, maintains Ethernet topology consistency, tracks configuration status, and monitors EVPN reachability — supporting repeatable validation as environments grow.

Summary and Architectural Outlook

As AI factories mature into long-lived, shared platforms, the networking layer increasingly defines the operational limits of multi-tenancy. In DGX-based environments, tenant isolation results from the coordinated behavior of multiple independent fabrics—Ethernet, InfiniBand, NVLink Multi-Node, and management networks—each with distinct control planes, operational semantics, and failure modes.

This document described an architectural model in which tenant intent is defined once and consistently expressed across all participating connectivity domains. Aligning Ethernet routing domains, InfiniBand partitions, and NVLink GPU collectives lets operators maintain predictable isolation while supporting tenant onboarding, resizing, and retirement without relying on tightly coupled, fabric-specific workflows.

Netris serves as the authoritative source of the desired state for fabric-wide Ethernet configuration and coordinates the appropriate partitioning across InfiniBand and NVLink Multi-Node fabrics. As NVIDIA continues to evolve its multi-tenancy guidance across connectivity domains, Netris, as an Independent Software Vendor (ISV) within the NVIDIA ecosystem, incorporates this guidance into a consistent, multi-tenant network architecture using supported integration interfaces. This includes NVIDIA Mission Control-managed environments, where Netris integrates with UFM and BCM and provides the tenant isolation and network automation functions that sit outside Mission Control's scope.

Netris maintains a consistent tenant intent representation and consistent Ethernet fabric behavior, enabling upstream platforms (IaaS/PaaS) to evolve independently of the underlying network implementations while preserving a clear separation of responsibilities.

As AI infrastructure continues to incorporate new compute and networking platforms—including DPUs, evolving GPU architectures, new interconnect technologies, and deeper platform integrations—tenant abstractions are also expected to evolve—from server-based allocation toward finer-grained GPU allocation models. Architectures that centralize tenant intent while decoupling it from individual fabric control planes offer greater durability over time. A unified, intent-driven networking layer—combined with continuous, cross-fabric validation—establishes a foundation that can accommodate these shifts without requiring fundamental changes to the tenant model or operational workflow.

About Netris

Netris is the leading provider of network automation and multi-tenancy for AI infrastructure.

The Netris NAAM (Network Automation, Abstraction, and Multi-Tenancy) platform is the most widely deployed platform — trusted by high-growth neoclouds, NVIDIA Cloud Partners, AI factories, and leading AI platform providers. Netris provides a unified control plane and native integrations across the complete AI infrastructure networking stack — Ethernet (including NVIDIA Spectrum-X), NVIDIA Quantum InfiniBand, NVL72, NVIDIA BlueField DPUs, and virtual and edge networking.

Netris enables operators to deploy and operate shared GPU infrastructure in weeks instead of years, provision tenants immediately with hard isolation (enforced on networking hardware) configured automatically, maximize GPU utilization by dynamically reallocating capacity across tenants, ensure network stability, and future-proof AI infrastructure.

Learn more at netris.io.