Modern Multi-tenancy Solutions for HPC and AI Clusters supporting Kubernetes and MLOps

Ron Bosworth, Solutions Architect
XENON Systems

Outline

- What does Multi-Tenancy mean?

- Typical Solutions

- Trade-offs

- From Hard to Soft Multi-Tenancy

- Modern Implementations

- Storage Multi-tenancy

- Challenges

What does Multi-Tenancy mean?

- **Multiple User:**
  - **within the same team** and organisation
  - collaboratively running workloads, which might need to communicate with each other
  - direct access to cluster

- **Multiple Teams:**
  - separate groups **within the same organisation**
  - direct access to cluster

- **Multiple Customers** ("SaaS" multi-tenancy)
  - **strongly isolating** different workloads
  - no access to customer
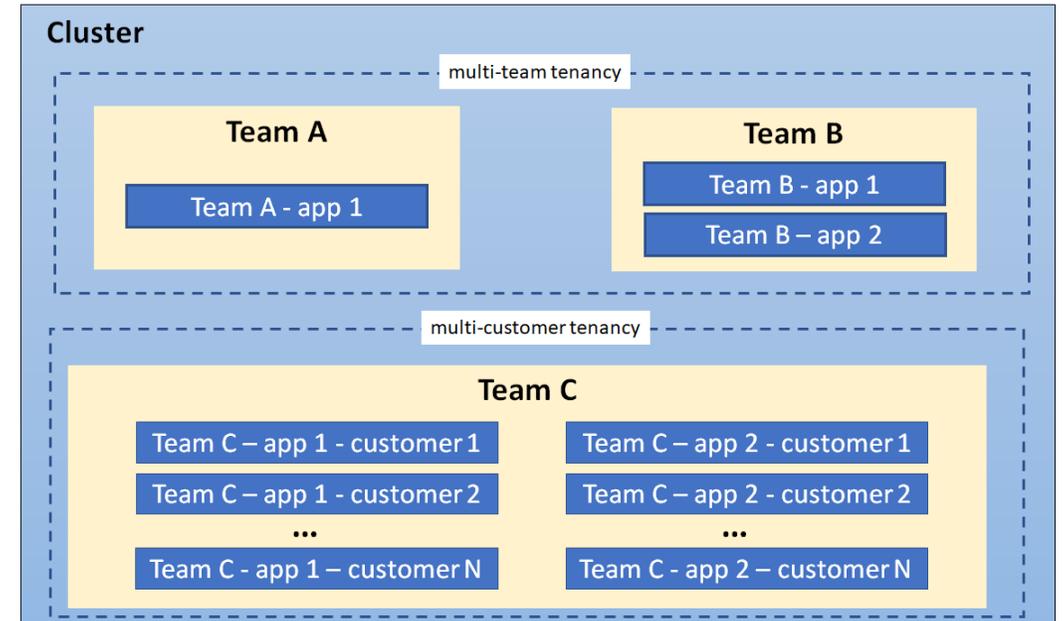
Typical Solutions

- **Dedicated Hardware:**
  - Physically isolated: all components separate and isolated
- **Logically separated networks:**
  - Dedicated compute and storage resources
  - Networks configured for logical separation
- **Virtualisation:**
  - KVM (OpenStack, Open Nebula, ProxMox)
  - Xen (XenServer, XCP-NG)
  - VMWare
  - Hyper-V
- **Containers:**
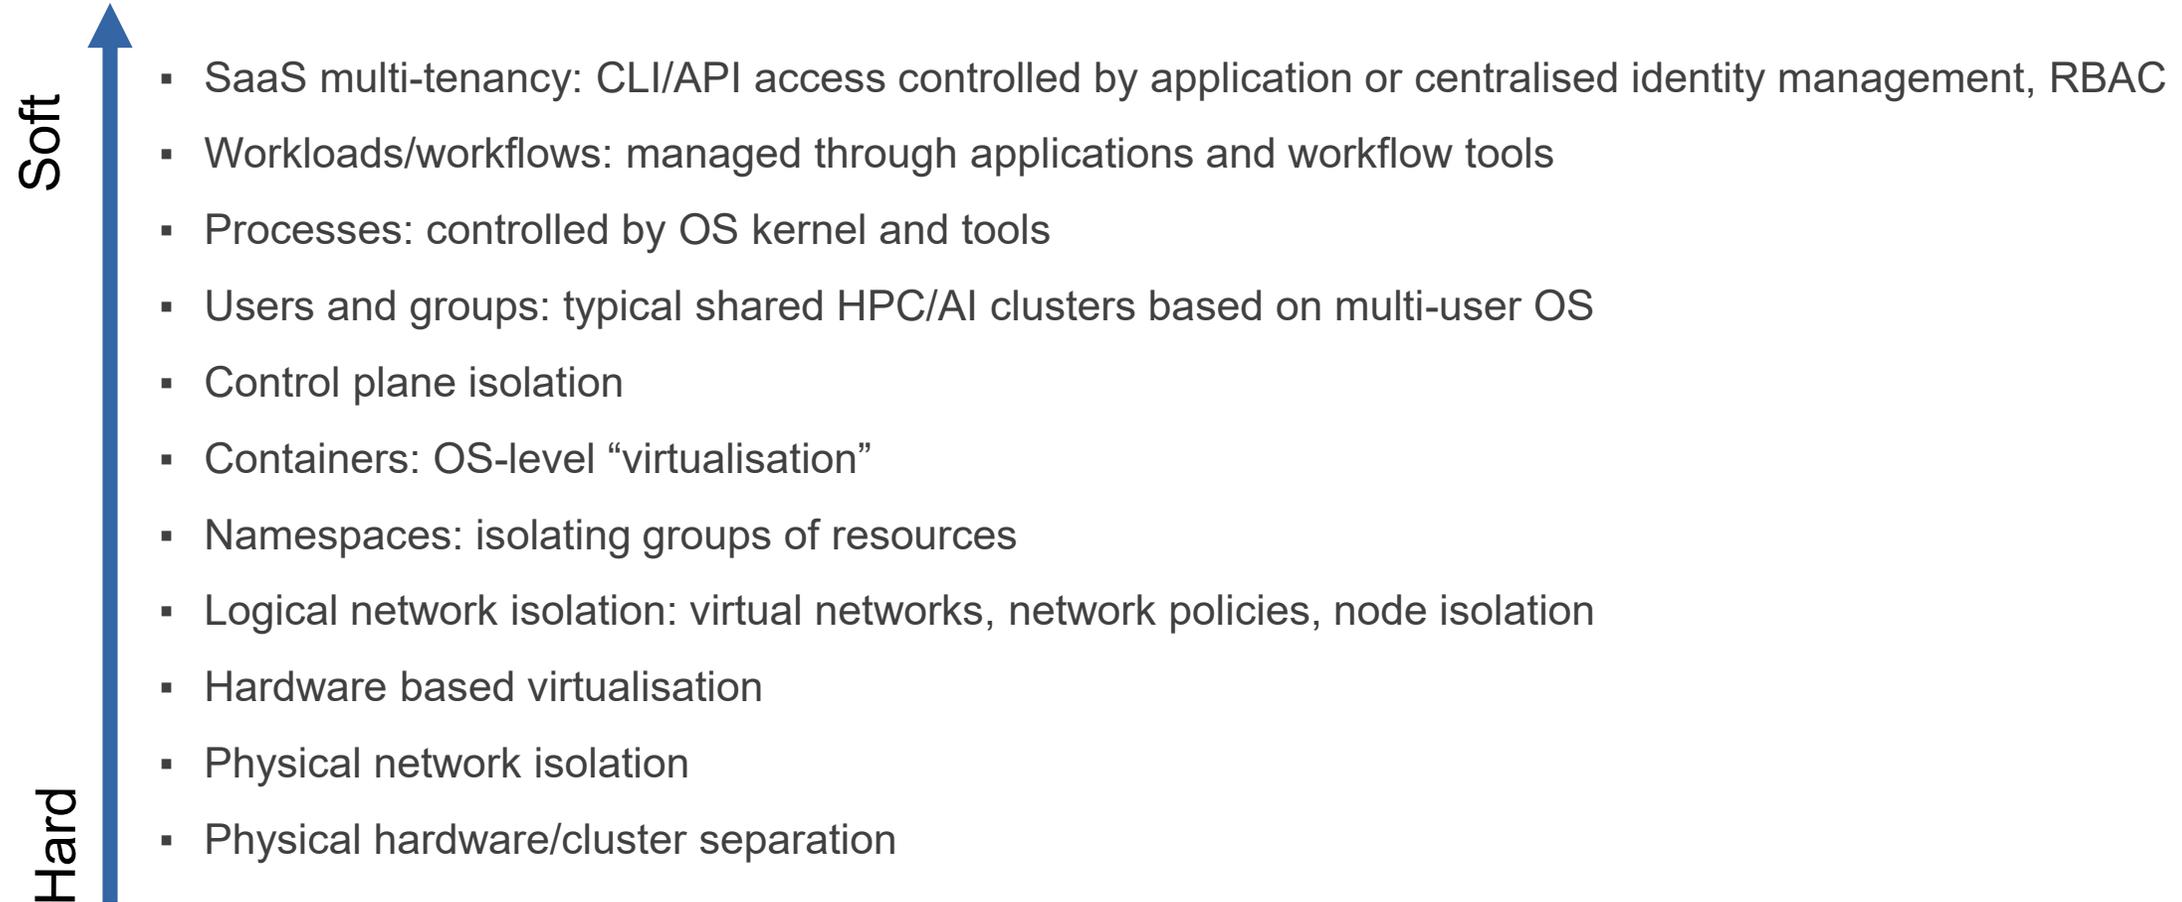 - Docker, Podman
  - Kubernetes

Typical Solutions

Trade-offs:

- Hardware management and utilisation

- Scheduling, Fairness, QoS

- Cluster management

- Security

- Operating system management, updates

- Performance: bare metal vs overheads (abstractions)

- Application management: dependencies, software supply chain/asset management

From Hard to Soft Multi-Tenancy

Multi-Tenancy is based on a range of concepts and thus provides a continuum solutions

Soft ↑

- SaaS multi-tenancy: CLI/API access controlled by application or centralised identity management, RBAC
- Workloads/workflows: managed through applications and workflow tools
- Processes: controlled by OS kernel and tools
- Users and groups: typical shared HPC/AI clusters based on multi-user OS
- Control plane isolation
- Containers: OS-level "virtualisation"
- Namespaces: isolating groups of resources
- Logical network isolation: virtual networks, network policies, node isolation
- Hardware based virtualisation
- Physical network isolation
- Physical hardware/cluster separation

Hard

XENON
High Performance Computing

Implementing Network Multi-Tenancy

Multi-tenancy at network level:

- Ethernet (high speed 'North-South', 'East-West', and out-of-band networks)
  - Manual network configuration on switches using VLANs
  - Manual configuration using EVPN/VXLAN
  - Semi-automated with configuration management tools (e.g. Ansible)
  - Automated software defined network configuration using e.g. Netris, Aarna, Rafay
- Infiniband: Partition Keys

Manual configuration:
1. Configure base configuration on switches (OOB, MLAG, interface type/speed, NTP, etc)
2. Create a VLAN (or VXLAN) configuration for the new tenant in NVIDIA Air, eg:
3. Copy the config file to /home/cumulus/tenant-vlan.yaml (or paste it in via VI or NANO)
4. nv config patch /home/cumulus/tenant-vlan.yaml
5. nv config apply
6. Confirm the config is running with: nv config show
7. Repeat for all switches as required.

Semi-automated with Ansible:
1. Setup OOB Management and disable ZTP on the switches
2. Enable NVUE API on the switches
3. Add nvidia.nvue module to Ansible environment.
4. Prepare Production Ready Automation
5. Run playbook to load config to switches

Netris

Cloud provider-style network automation and abstraction suitable for both multi-tenant public cloud providers and private cloud environments
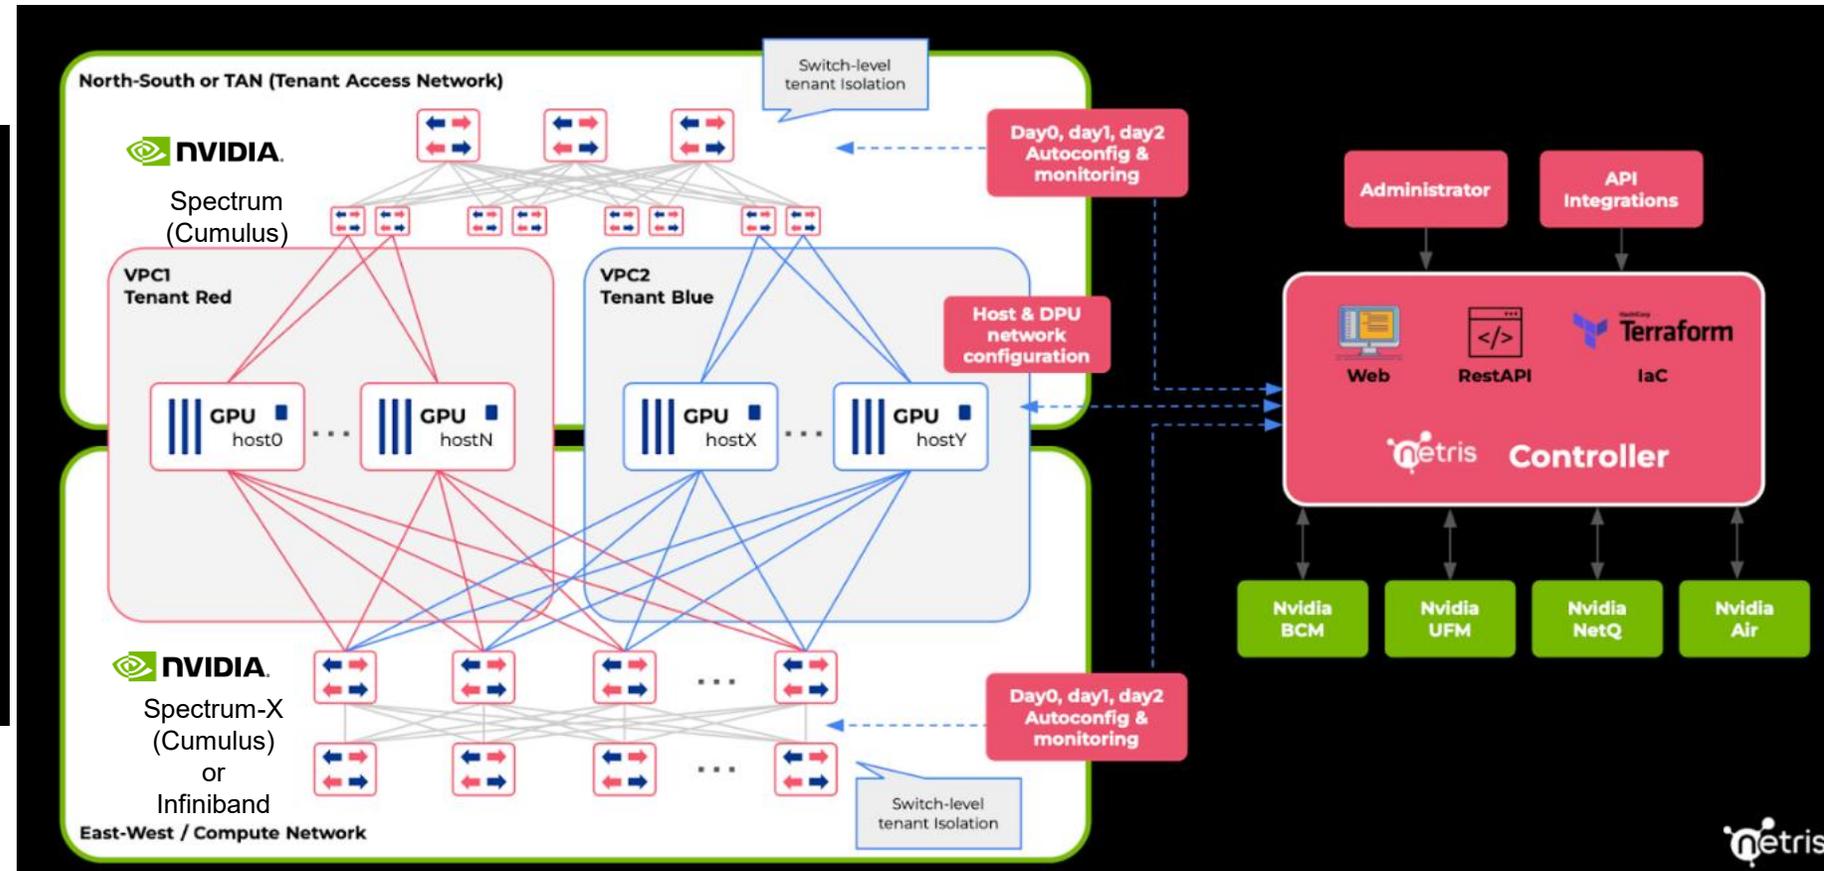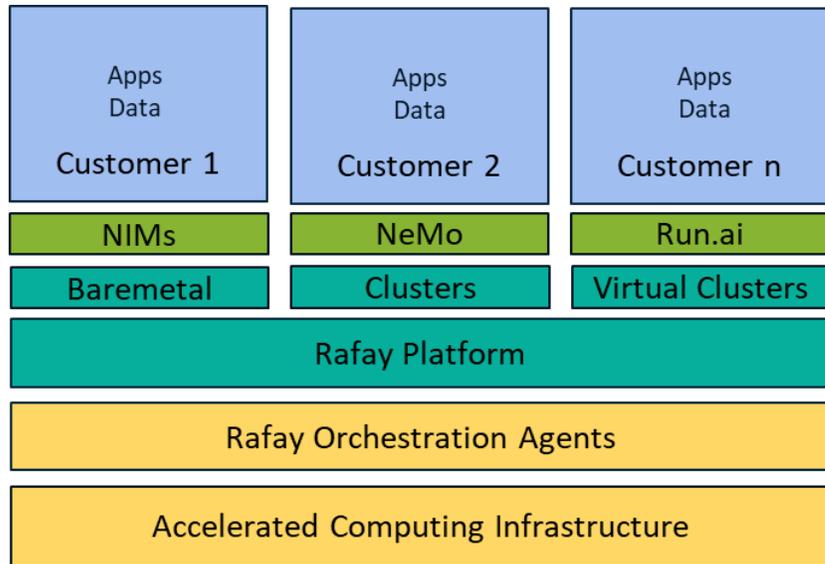
Netris

- Dynamically assigns infrastructure components to various tenants
- Managed through Netris web console or by integrating Netris APIs.
- Automatically configures the appropriate VRFs, VXLANs, and all other low-level networking details across every network switch and managed gateway
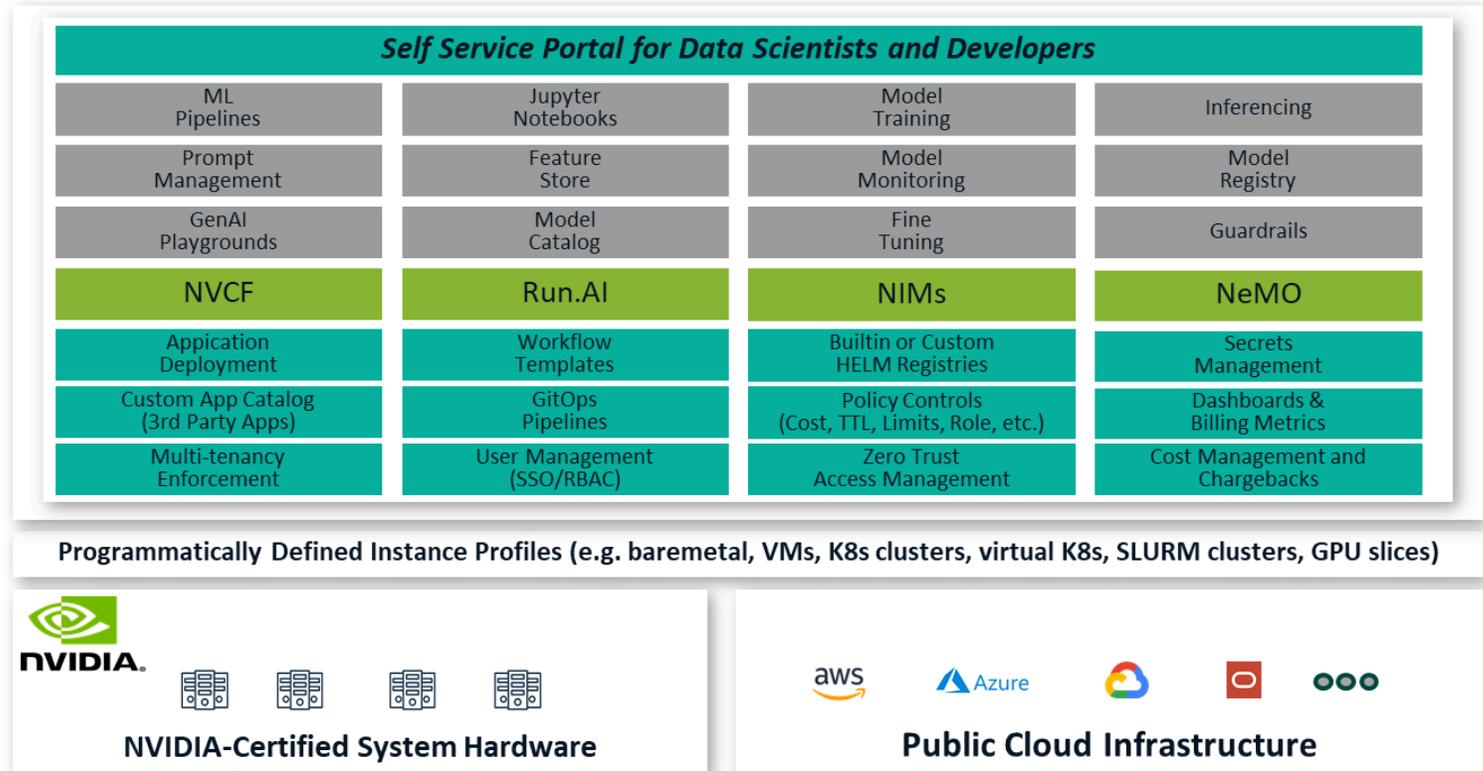- Delivers highly scalable multi-tenancy and network-level isolation.
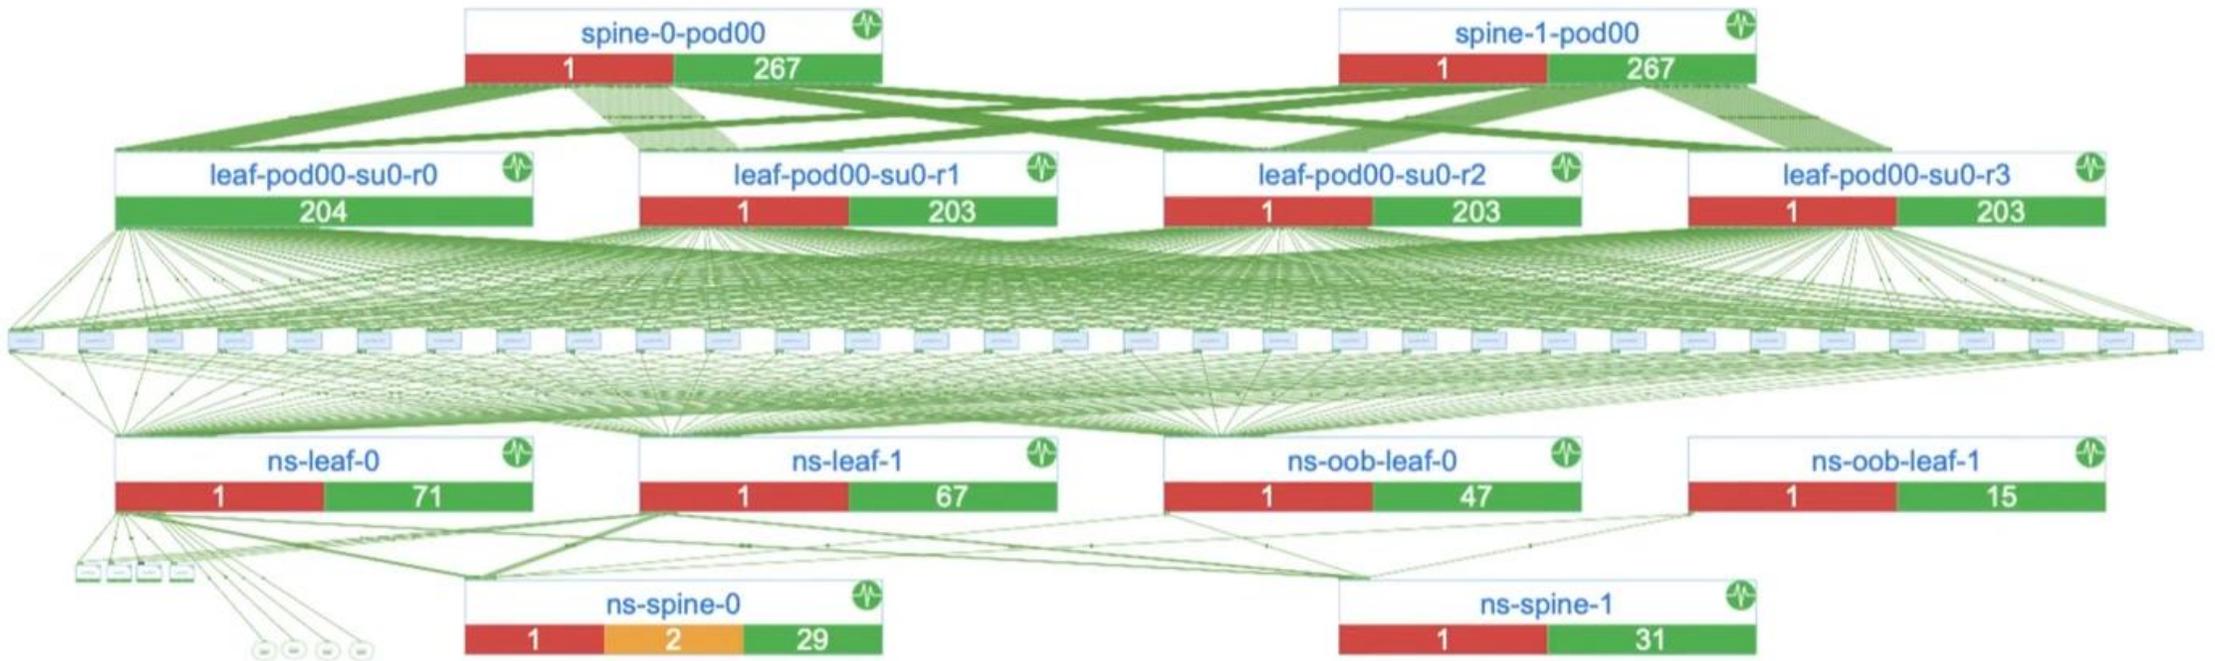
Rafay

Platform-as-a-Service for GPU & CPU Deployments
Infrastructure Orchestration and Workflow Automation for Cloud-Native and AI Innovations
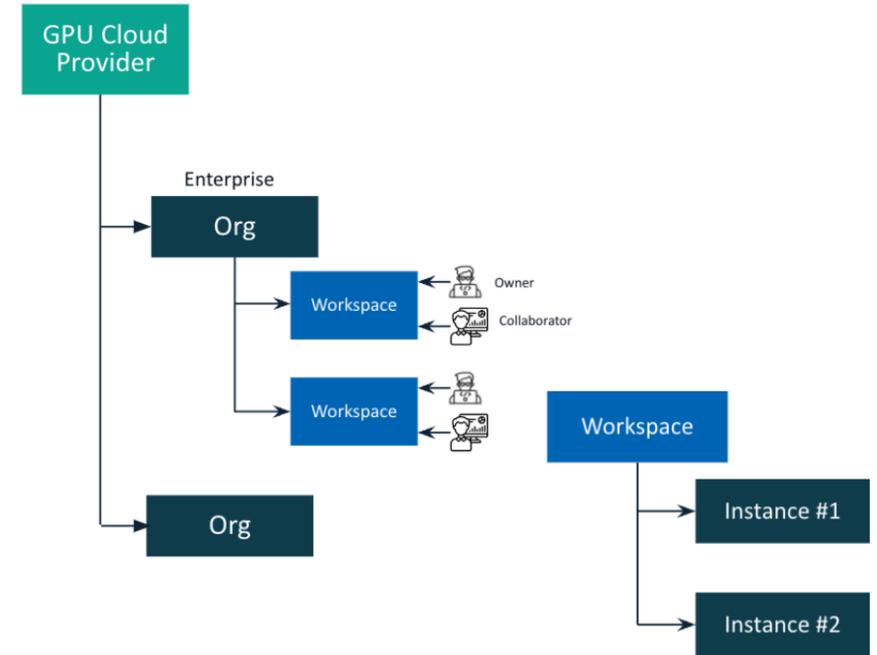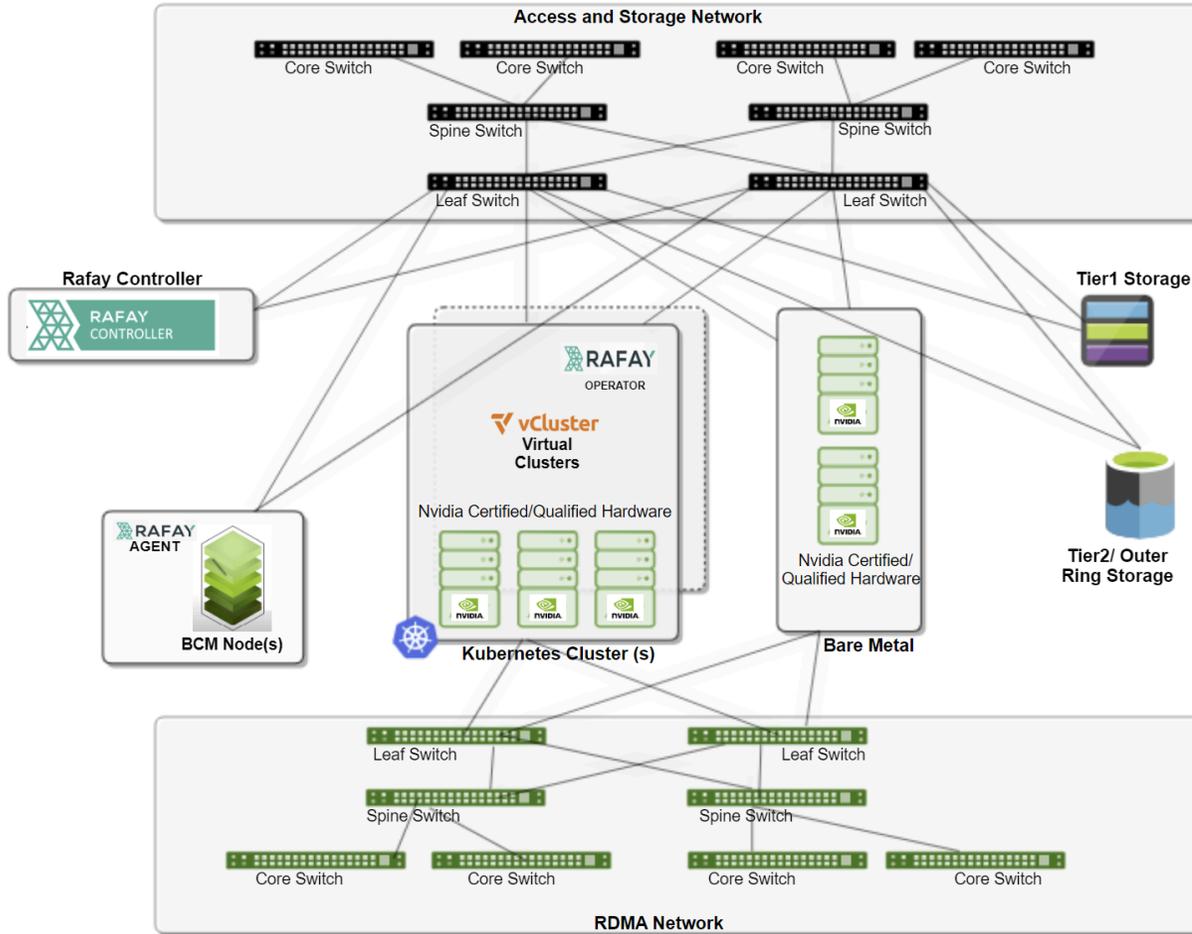

Rafay Operating Model

Apps Data	Apps Data	Apps Data
Customer 1	Customer 2	Customer n
NIMs	NeMo	Run.ai
Baremetal	Clusters	Virtual Clusters

Rafay Platform

Rafay Orchestration Agents

Accelerated Computing Infrastructure

Self Service Portal for Data Scientists and Developers

ML Pipelines	Jupyter Notebooks	Model Training	Inferencing
Prompt Management	Feature Store	Model Monitoring	Model Registry
GenAI Playgrounds	Model Catalog	Fine Tuning	Guardrails
NVCF	Run.AI	NIMs	NeMO
Appication Deployment	Workflow Templates	Builtin or Custom HELM Registries	Secrets Management
Custom App Catalog (3rd Party Apps)	GitOps Pipelines	Policy Controls (Cost, TTL, Limits, Role, etc.)	Dashboards & Billing Metrics
Multi-tenancy Enforcement	User Management (SSO/RBAC)	Zero Trust Access Management	Cost Management and Chargebacks

GPU PaaS

Programmatically Defined Instance Profiles (e.g. baremetal, VMs, K8s clusters, virtual K8s, SLURM clusters, GPU slices)

NVIDIA

NVIDIA-Certified System Hardware

aws   Azure

Public Cloud Infrastructure

XENON
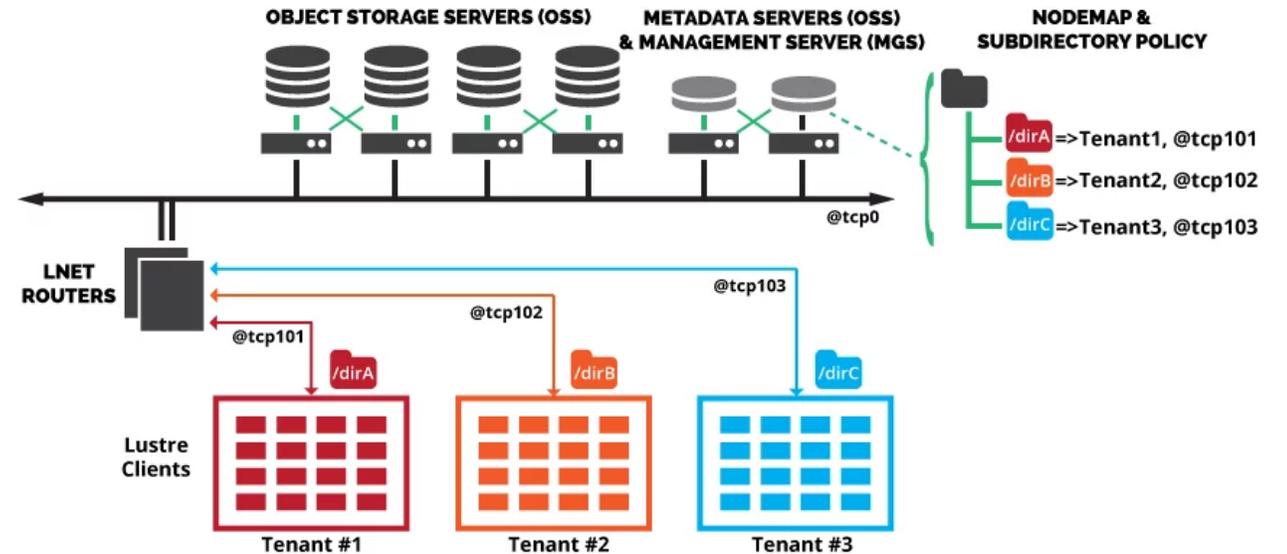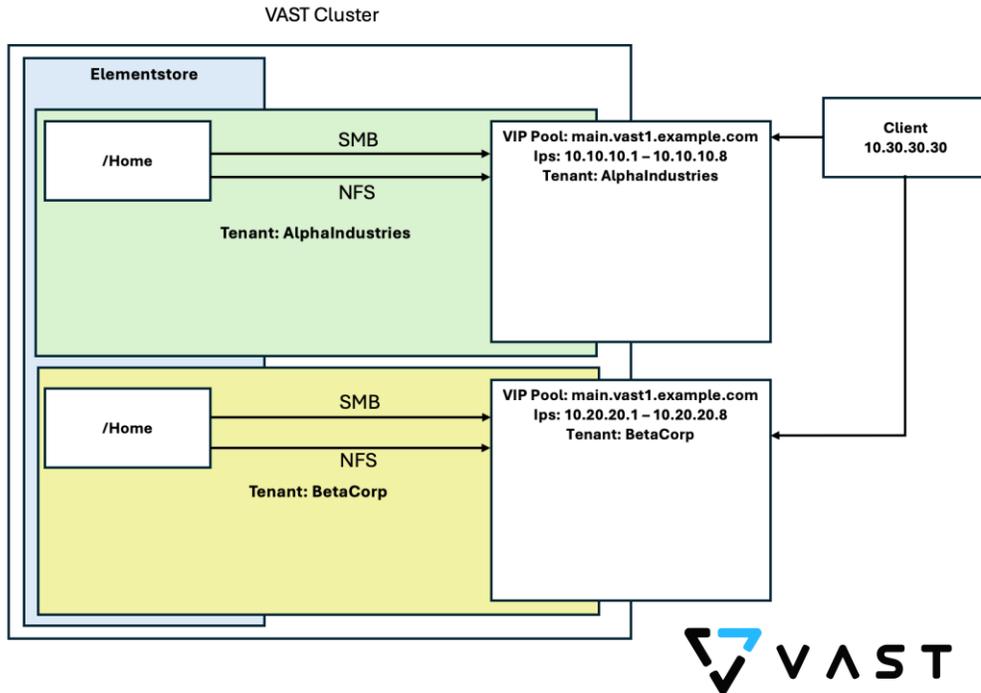High Performance Computing

Rafay

Platform-as-a-Service for GPU & CPU Deployments

Storage Multi-Tenancy

Storage systems also need to support multi-tenancy:

- Separate storage areas for different tenants

- Integration with (different) identity providers: local, AD, LDAP, or NIS

- Unique encryption keys, quotas, and quality-of-service (QoS) settings

Solutions available from various vendors (incl. NVIDIA certifications): VAST, DDN, Weka, IBM, etc.

Conclusion

Challenges

- Multi-tenancy has a large management overhead

- Manual network configuration becomes unsustainable very quickly

- Switching to a network automation solution will likely require a clean slate

Learnings

- Hard vs soft multi-tenancy is a continuum of solutions

- Carefully assess requirements against this continuum

- Plan for future requirements (many customers moving towards "harder" multi-tenancy)

- Implement automation from the start

- Pick the right solution (and partner) for the implementation

Thank You

XENON Systems Pty Ltd
10 Westall Road, Springvale, Victoria 3171, Australia

www.xenon.com.au

P +61 3 9549 1111
E info@xenon.com.au

A member of the XENON Technology Group
www.xtg.com.au