

EIDF KUBERNETES

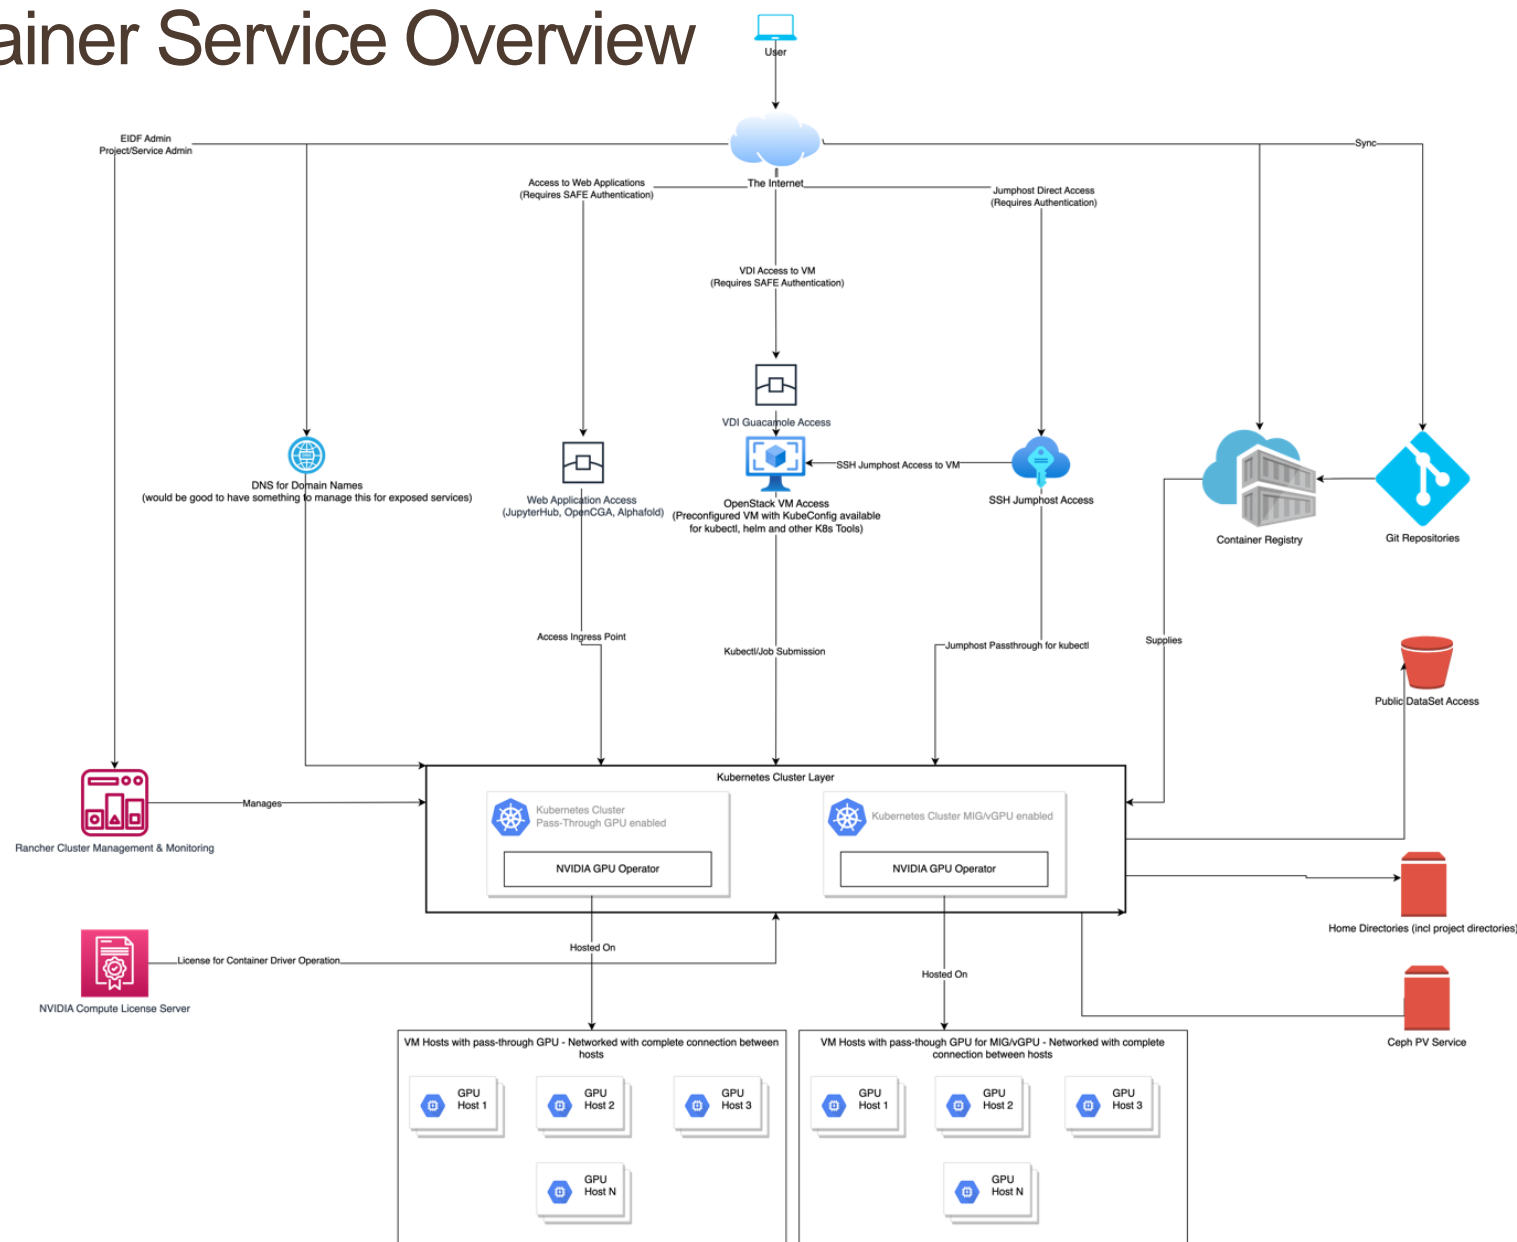
Kubernetes in EIDF

- OpenStack Kubernetes
 - Project Private Clusters
 - OpenStack VM Hosts
 - User Managed Cluster
- EIDF Container Service
 - Private Project Namespace
 - “Public” Services
 - KVM VM Hosts
 - RKE Kubernetes
 - EIDF Managed Cluster

Openstack: self managed project private clusters

- Various options:
 - Microk8s for a single node cluster
 - Very convenient; includes ingress, MinIO, storage, DNS, ...
 - Manual installation with kubernetes
 - Straightforward too but not convenient or scalable
- StackHPC Azimuth
 - Uses Cluster API to create Kubernetes clusters on demand
 - Didn't work until very recently without DNS: private Openstack API server domain could not be resolved on instances; with Openstack Designate (DNSaaS) it works as of last week!
 - Very easy to use; provides monitoring and auto-scaling
 - Integration with EIDF/EPCC infrastructure required; users don't have access to Horizon
- Use Octavia or MetalLB as load balancer
- Issues with CNI canal (known problem?), calico or weave net were fine

EIDF Container Service Overview



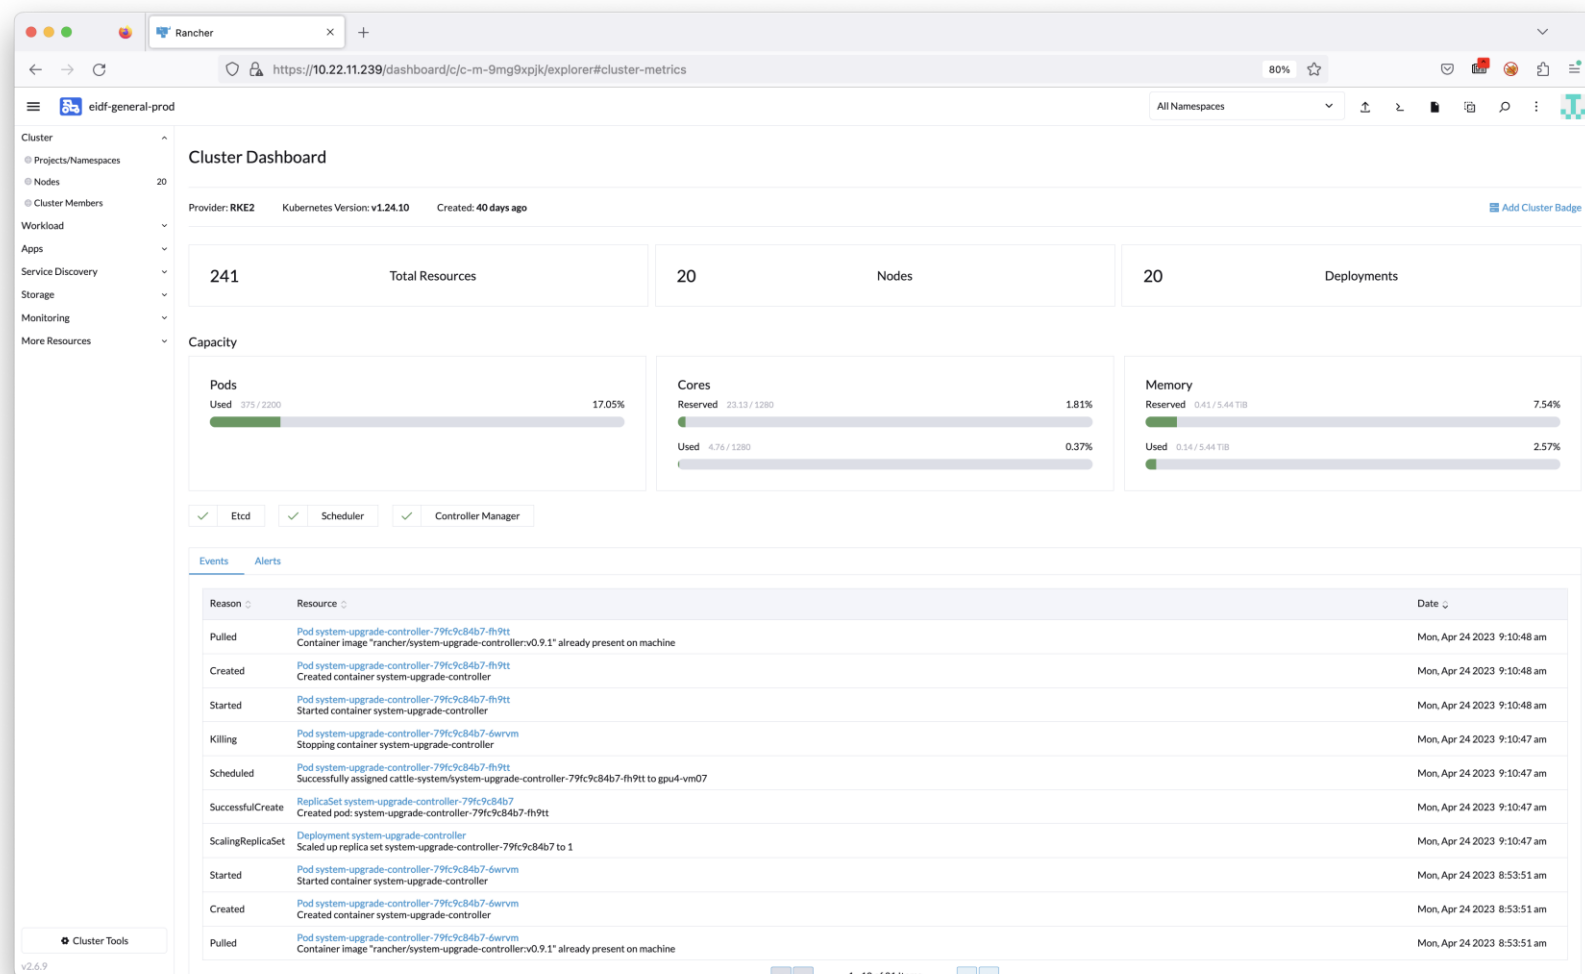
Current Status

- 11 Physical 8 GPU Nodes - 36 Virtual Nodes

Allocatable CPU (millicores)	1536000
CPU Capacity (millicores)	1536000
Current CPU Requests (millicores)	22533
Allocatable GPU	222
GPU Capacity	222
Current GPU Requests	4
GPU Type Capacity	
NVIDIA-A100-SXM4-40GB	54
NVIDIA-A100-SXM4-40GB-MIG-1g.5gb	140
NVIDIA-A100-SXM4-40GB-MIG-3g.20gb	28
GPU Type Requests	
NVIDIA-A100-SXM4-40GB	4
NVIDIA-A100-SXM4-40GB-MIG-1g.5gb	0
NVIDIA-A100-SXM4-40GB-MIG-3g.20gb	0
Allocatable Memory (MiB)	6732536
Memory Capacity (MiB)	6732536
Memory Requests (MiB)	19507

Cluster Management: Rancher

- Managed Kubernetes and Clusters
- Multi-Cluster Management Tool



Setup

- Ansible
 - Hypervisor Prep
 - Network and VM Deployment
 - VM Node Prep – NVIDIA and upgrades
 - Cluster Registration
- Manual Steps
 - Host Inventories
 - Variable Files
 - Token Creation

Projects

- Users
- Namespaces - including default tolerations
- Quotas
- Limits

- Access:
 - Ceph
 - GPU (MIG/Passthrough)
 - Specific Node Restrictions - Taints

Future

- A100 80GB GPU Nodes
- Internal Registry
- Improved Queue and Scheduling Management
- Policy Agent Additions