

The First Tenant-Managed ALPS vCluster - a presentation by CSCS and PSI

HPC-CH

Hussein Harake, CSCS

05.10.2023



CSCS


Centro Svizzero di Calcolo Scientifico
Swiss National Supercomputing Centre

ETH zürich

The First Tenant-Managed ALPS vCluster - a presentation by CSCS and PSI

The First Tenant-Managed ALPS vCluster - a presentation by CSCS and PSI


- ALPS
- vCluster
- Multitenancy
- Manta



The First Tenant-Managed ALPS vCluster - a presentation by CSCS and PSI

Is there any HPC machine that satisfies all user requirements and scientific applications?

- Diversity of Scientific Needs
- Technological Specializations
- Resource Constraints
- Software and Programming Models
- Usability and Accessibility
- Evolving Technologies



ALPS

Alps is the HPE Cray EX Supercomputing

- Alps stands out as our latest flagship HPE Cray EX supercomputer.

Deployment in Phases

- Installation commenced in a structured multi-phase approach, beginning in 2020.




The First Tenant-Managed ALPS vCluster - a presentation by CSCS and PSI

Networking and Accessibility

- **Network:** Utilizes the sophisticated Slingshot network.
- **Availability Zones:** Configured with two availability zones to ensure High Availability (HA) and offer a non-HA option.

Cooling Efficiency

- **Liquid Cooling:** Boasts 100% liquid cooling, extending directly to the chips themselves.






CSCS

Centro Svizzero di Calcolo Scientifico
Swiss National Supercomputing Centre

ETH zürich

Versatile software-defined cluster (vCluster)

Technology



User environments management

Platforms and services management

Infrastructure as Code


Versatile software-defined cluster (vCluster)

- Custom user environments
- Manage platform services, provisioning of clusters
- Tenant: network and storage isolation

vCluster Management and Technology Utilization

- **Dedicated Administration for Every vCluster**
 - Each vCluster is managed by its own team of platform administrators, equipped with a distinctive set of users, as well as its unique configuration and policies.
- **Efficiency through Unified Infrastructure**
 - Utilizing a shared infrastructure for hardware operation, and adopting innovative IaaS and PaaS alongside DevOps technologies, accelerates the path to finding solutions.
- **Enhanced Flexibility in Technological Choices**
 - This approach provides augmented flexibility concerning choices in hardware and software, facilitating optimized selections for specific needs.
- **Elasticity and Scalability of Resources**
 - Ensure resource adaptability and extensibility: the capability to easily expand or contract vClusters enhances operational agility.

vClusters currently in operation on Alps



Platform	vCluster name	Scope
MCH	Balfrin	R&D
	Tasna	Production
HPC Platform	Eiger	Production
	Pilatus	Staging
	Rigi	R&D
WLCG Platform	Fort	Production
	Gele	R&D
	Noir	Production (CTA)
AI/ML Platform	Clariden	AI/ML
Testbeds	Rosa	K8s testing
	Adula	HW validation
	Zinal	Internal development
	Bar	Internal development
	Hohgant	Internal development
PSI	PSI-dev	R&D
	PSI-tds	TDS for PSI-dev



CSCS

Centro Svizzero di Calcolo Scientifico
Swiss National Supercomputing Centre

ETH zürich

Multitenancy

Multitenancy Overview:

•Definition:

- Enabling multiple tenants to share common computing resources.
- Ensuring isolated, secure, and private environments for each tenant.

•Bare Metal Cluster Context:

- **Physical Resource Sharing:** Allocating bare-metal resources among tenants.
- **OS-Level Isolation:** Employing operating system capabilities for securing tenant spaces.
- **Direct Resource Allocation:** Handling straightforward physical resource distribution.



Other Contexts

•Kubernetes (K8s) Context:

- **Namespace Utilization:** Separating tenants via namespaces.
- **Resource Management:** Implementing quotas and limit ranges per namespace.
- **Security Assurance:** Applying RBAC and Network Policies for safe multi-tenancy.

•General Aspects:

- **Isolation:** Essential for maintaining tenant privacy and security.
- **Scalability:** Ensuring the system can adjust to varying tenant demands.
- **Resource Optimization:** Efficient utilization and management of shared resources.
- **Access Control:** Strictly regulating who has access to specific resources and data.
- **Tenant Independence:** Ensuring one tenant's activities do not impact others.



Extreme Multitenancy: Network-Level Separation on ALPS

Overview

- Tenant-specific network environments
- VLAN-based isolation

•VLAN Network Isolation

- Assign PSI VLAN to vCluster
- Ensuring strict network segmentation
- Managing PSI IP spaces

•vCluster Management

- Autonomous network oversight
- Assurance of optimized network functionality

•Operational Considerations

- Attentive allocation of network resources
- Ensuring scalability and stringent compliance

•Benefits and Challenges

- Enhanced security vs. VLAN management complexity



CSCS

Centro Svizzero di Calcolo Scientifico
Swiss National Supercomputing Centre

ETH zürich

Manta

Manta - Enabling Autonomous vCluster Management

Manta: Developed at CSCS

- A specialized, binary-less tool for centralized vCluster management
- Utilizes an opinionated Shasta CLI and ergonomic design
- Aggregates information across various systems

Facilitating Admin Autonomy

- Allows PSI admins to independently manage clusters
- Minimizes the need for ALPS infrastructure / Platform admin intervention



Versatile & Empowering Functionalities of Manta

Centralized & Simplified Cluster Operations


- Replaces standard k8s and CSM tools
- Enables easy management via a single, centralized tool

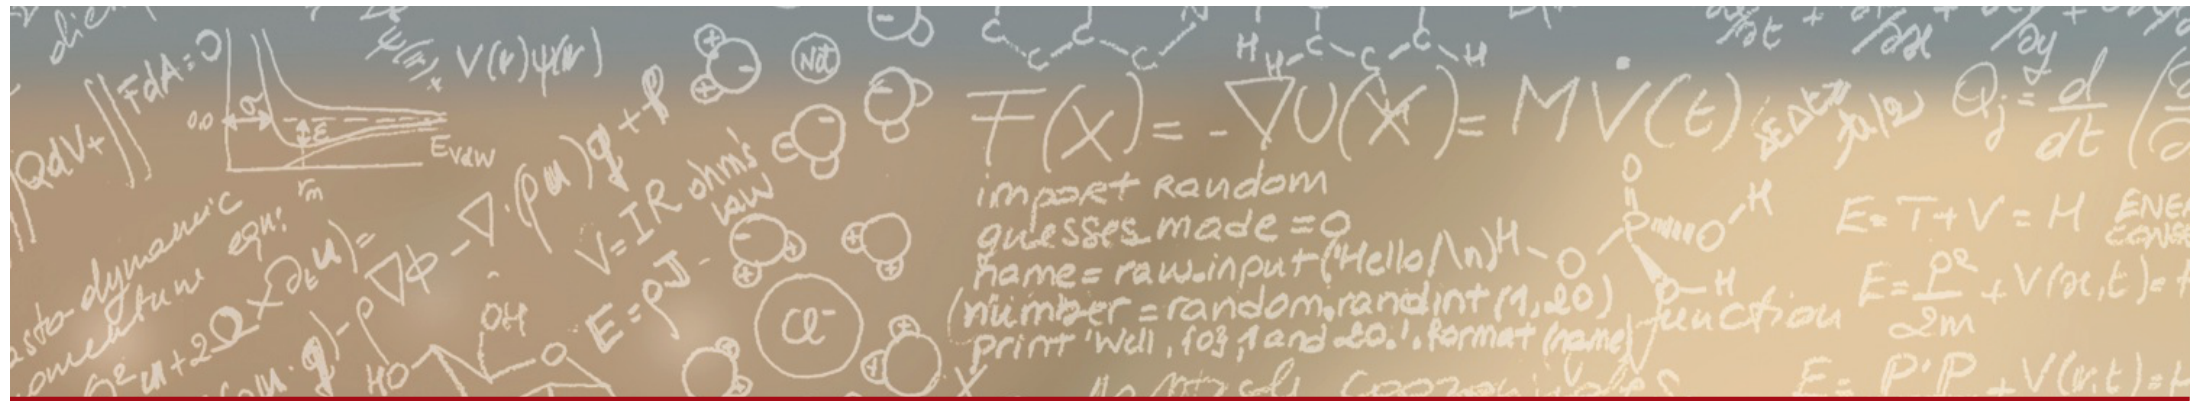
Versatile Cluster Management Capabilities

- Recreate OS images, modify Ansible layer, and redeploy a cluster
- Reboot a cluster
- No need for infrastructure admin intervention

Secure Cluster Management

- Ensures a secure, flexible, and autonomous management environment
- Customizes cluster settings and environments as per specific needs





Thank you for your attention.