

Infrastructure as a Service with Strong Tenant Separation on a Supercomputer

Riccardo di Maria
Chris Gamboni
Manuel Sopena Ballesteros
Hussein Harake
Mark Klein
Marco Passerini
Miguel Gila
Maxime Martinasso
Thomas C. Schulthess
ETH Zurich, Swiss National Supercomputing Centre
(CSCS)
Lugano, Switzerland

Alun Ashton
Derek Feichtinger
Marc Caubet
Elsa Germann
Hans-Nikolai Viessmann
Achim Gsell
Krisztian Pozsa
PSI – Paul Scherrer Institute
Villingen, Switzerland

ABSTRACT

This paper explores the innovative implementation of Infrastructure-as-a-Service (IaaS) on a HPE Cray Shasta EX supercomputer. In cloud environments, IaaS offers scalable on-demand access to virtualized resources. However, applying IaaS principles to high-performance computing (HPC) systems without relying on virtualization technologies poses some challenges, since they typically have a tightly coupled software stack. We address these challenges in a co-design partnership between an HPC provider, CSCS, and an end-user institution, PSI, by developing a suite of technologies for the HPE Cray Shasta EX system architecture that supports resource isolation and granular control. This approach not only provides the IaaS model in supercomputing environments but also enables dynamic resource management.

Our contributions include a detailed exploration of the technological advancements necessary for integrating IaaS into HPC, together with the lessons learned from our collaborative efforts. By extending IaaS capabilities to supercomputers, we aim to provide scientific institutions with unprecedented flexibility and control over their computational resources.

CCS CONCEPTS

• **Computing methodologies** → *Massively parallel and high-performance simulations*; • **Computer systems organization** → *Interconnection architectures*; Multicore architectures; • **Software and its engineering** → *Cloud computing*.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

CUG, May 04–08, 2025, New Jersey

© 2025 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 978-1-4503-XXXX-X/2025/05
<https://doi.org/XXXXXXX.XXXXXXX>

KEYWORDS

IaaS, Network, Multi tenancy, Supercomputer

ACM Reference Format:

Riccardo di Maria, Chris Gamboni, Manuel Sopena Ballesteros, Hussein Harake, Mark Klein, Marco Passerini, Miguel Gila, Maxime Martinasso, Thomas C. Schulthess, Alun Ashton, Derek Feichtinger, Marc Caubet, Elsa Germann, Hans-Nikolai Viessmann, Achim Gsell, and Krisztian Pozsa. 2025. Infrastructure as a Service with Strong Tenant Separation on a Supercomputer. In *Proceedings of Cray User Group (CUG)*. ACM, New York, NY, USA, 10 pages. <https://doi.org/XXXXXXX.XXXXXXX>

1 INTRODUCTION

Infrastructure-as-a-Service (IaaS) [12] has fundamentally transformed the landscape of computing resources allocation, offering scalable and flexible virtual computing resources over the Internet. Originating in the cloud, IaaS allows users to provision processing, storage, and networking resources on-demand, thereby eliminating the need for significant upfront capital expenditures on hardware and data center space, together with reducing the operational costs associated with managing physical servers. Cloud-based IaaS platforms, such as Amazon Web Services, Microsoft Azure, and Google Cloud Platform, have democratized access to computing resources, enabling organizations of all sizes to deploy and scale enterprise applications rapidly in response to changing demands.

In the cloud, IaaS is characterized by several key elements that distinguish it from other service models. At its core, IaaS provides virtualized computing resources, encompassing virtual machines (VMs) with customizable configurations of CPU, memory, and storage. Networking is another critical component, offering scalable and secure connections between VMs, and between VMs and the Internet, often through virtual private networks (VPNs) or dedicated connections. Storage solutions within IaaS are also highly flexible, ranging from block storage for general-purpose data, such as file storage, to object storage for scalable, unstructured data. In addition, IaaS platforms provide a comprehensive suite of management tools and APIs for automating and orchestrating resource

provisioning, scaling, and monitoring, enabling users to manage their infrastructure programmatically and efficiently.

Despite the widespread adoption of IaaS in cloud computing environments, extending IaaS principles to High-Performance Computing (HPC) and supercomputers introduces challenges and opportunities. HPC systems, exemplified by the HPE Cray Shasta EX [11], are engineered for peak performance, traditionally employing a vertically integrated software stack and tightly integrated hardware to execute compute-intensive tasks through batch processing. These systems are pivotal for addressing complex problems in the scientific, engineering, and data analysis domains, leveraging parallel programming paradigms and running scientific applications at unprecedented scales. A key characteristic of HPC systems is their avoidance of virtualization technologies, such as VMs, or the usage of VPNs, to ensure direct access to hardware performance. Moreover, HPC systems are single-tenant. In this context, the tenant—either a data centre or HPC provider—offers computing services to various scientific communities that share batch queues. The lack of virtualization layers significantly complicates the implementation of IaaS on HPC platforms, especially in achieving robust network-level separation among multiple tenants, a critical aspect that the IaaS model inherently depends on.

The primary beneficiaries of implementing IaaS and becoming a tenant on supercomputers are scientific institutions and communities, for whom this model offers flexibility and control over computational resources, without the need to invest themselves in necessary HPC-capable data center facilities and know-how. With tenant separation, these institutions gain the ability to connect their own network directly into the supercomputer network and to manage the IaaS supercomputer cluster as if the cluster is an integral part of their own data center. This enables the institution to provide and offer custom services specifically designed to meet their research requirements from their community of users. By managing multiple tenants, the HPC provider enables the dynamic nature of IaaS for efficient resource lifecycle management, allowing seamless addition or removal of computational elements to tenants in response to changing demands. Furthermore, it empowers institutions to integrate new technologies at their own pace, ensuring that their infrastructure remains at the cutting edge without disrupting ongoing projects and enabling comprehensive investment planning.

This paper addresses the novel endeavor of implementing strong tenant separation using IaaS concepts on the HPE Cray Shasta EX supercomputer, without virtualization and compromising performance. We detail the technical challenges encountered and the innovative solutions developed through a collaborative co-design approach with a partner institution.

Our contributions in this work include the following key elements:

- We introduce a co-design partnership for implementing tenant separation and IaaS services on supercomputers, emphasizing collaborative efforts between the HPC infrastructure provider and the end-user institution. This approach ensures that the IaaS model is optimally aligned with the partner’s requirements and the capabilities of supercomputing environments.

- We propose and detail a comprehensive suite of technologies tailored for HPE Cray Shasta EX systems. These technologies facilitate resource isolation and offer granular resource control mechanisms, enabling partner institutions to manage their computational resources effectively. This suite represents a pivotal advancement in supercomputing, allowing dynamic allocation and management of HPC resources in a manner akin to traditional cloud environments.
- Through a reflective analysis of our project, we share valuable lessons learned, highlight effective solutions, and outline trends for future work in the domain of tenant separation for supercomputers. Our insights aim to guide and inform ongoing and future efforts to integrate strong multitenancy capabilities into HPC systems, outside of the HPE offering, introducing both future technical solutions and the opportunities that multitenancy on HPC systems brings for innovation.

2 PARTNERSHIP IN CO-DESIGNING SOLUTIONS

The implementation of IaaS on supercomputers is based on a co-design partnership between CSCS (the HPC provider), PSI (a research institution), and HPE (the HPE Cray Shasta EX manufacturer). This collaboration ensures that the multitenant IaaS model aligns with PSI’s operational needs, drives new technology development at HPE, and integrates seamlessly into CSCS’s Alps infrastructure [23].

2.1 Initial requirements and use case for an IaaS tenant at the Partner Institution

PSI, a leading research center in Villigen, Switzerland, conducts advanced research in fields like materials science and energy. PSI aims to address some of the most important challenges facing society today, including the development of sustainable energy sources, understanding the properties of materials for new technologies, and investigating the complexities of living systems. PSI includes a third-generation synchrotron light source, alongside a computational framework designed to support both real-time (online) and post-experimental (offline) data analysis. This computational setup is crucial for processing the voluminous data generated by experiments, which is subject to strict requirements regarding availability, readiness, and latency. PSI provides its extensive research community, which spans fields from particle physics to microbiology, with access to its state-of-the-art synchrotron facility.

During live experiments, the beamlines are capable of producing several terabytes of data that require immediate post-processing and storage. The decision to upgrade the beamline equipment to a newer technology that offers significantly increased brightness—and consequently, a performance boost by up to a factor of 40—necessitates a corresponding enhancement of the computational infrastructure [20]. However, the data center currently operated is at its maximum capacity, and expanding its computational resources capacity to meet the demands of the new beamline technology presents a significant challenge. The institute is faced with two options: constructing a new data center or decentralizing its computational capabilities.

CSCS and PSI have engaged in a series of collaborative projects that have progressively aimed at coupling both research infrastructures. This evolving partnership has led to significant advancements, including the adaptation of workflow to a RESTful interface [7] for streamlined access to HPC resources [14], supporting both real-time and analytical workflows of experimental activities. Furthermore, a concerted effort has been made to implement data geo-redundancy through tape technology [24], ensuring robust and secure storage of experimental data across both locations. The synergy between the two institutions is consolidated and challenged by shared financial support from a common funding agency, reinforcing their joint commitment to pushing the boundaries of scientific research. The cumulative success of these collaborations has led to another strategic stepping stone by adopting an IaaS tenant model for a research cluster serving about 80 groups of users. Both institutions recognized that PSI's requirements for direct control over HPC resources and the ability to deploy custom services extends beyond the capabilities of mere RESTful API access. The idea of physically housing PSI's own computational resources within CSCS's facilities was evaluated but ultimately dismissed in favor of the tenant model on a shared, software defined infrastructure, which offers superior scalability and flexibility to meet PSI's future computational demands, in particular, as CSCS is deploying a large heterogeneous HPC infrastructure [23]. For the same set of resources, the cost to do IaaS on the Cloud would have been prohibitive.

PSI's requirements emphasize autonomous management of their computational infrastructure, enabling the deployment of specialized services tailored to their needs. This includes the ability to seamlessly integrate computational resources into its site-wide network, connecting to the comprehensive ecosystem of services essential for analyzing experimental outcomes. Additionally, PSI seeks granular control over the computational nodes for tasks such as monitoring node health, performing reboots, and managing base image updates and installations, ensuring optimal availability and security integration with their own policies.

For the HPC provider, the challenge lies in assimilating these specific requirements into a broader infrastructure that also serves the computational needs of various other institutions. This integration must maintain the delicate balance between customization for PSI and the operational efficiency of a shared HPC environment.

Adopting an IaaS model presents PSI with significant long-term benefits, including the delegation of hardware lifecycle management and the flexibility to scale resources both in capacity and through the incorporation of cutting-edge hardware. Moreover, it affords a greater degree of cost control, aligning computational expenditure with actual usage and future requirements.

2.2 Technology partnership

For over a decade, CSCS and Cray, which was later acquired by HPE, have fostered a collaborative relationship aimed at advancing HPC capabilities for the scientific community. This partnership has been at the forefront of technological innovation, notably in the large-scale integration of GPUs, establishing new capabilities in the field.

The advent of Cloud technology, introducing layers of abstraction in system management, combined with valuable customer

feedback, has driven Cray and subsequently HPE, to significantly overhaul their system management software and release the Cray System Management (CSM) [10] software product. This strategic shift aims to align more closely with modern computing paradigms and user expectations for flexibility and ease of use. An essential aspect of this evolution has been the decomposition of management workflows into independent microservices, adhering to a Service-Oriented Architecture (SOA) and enhancing security measures. Additionally, there has been a strategic emphasis on programmable APIs, enabling automation in accessing management services.

Simultaneously, HPE unveiled Slingshot [21], its next-generation high-speed networking technology. Beyond offering performance enhancements and advanced Quality of Service (QoS) management, Slingshot provides compatibility with Ethernet protocols without compromising performance, including support for virtual LANs and virtual network interfaces, marking a significant advancement in network flexibility, integration and isolation capabilities.

The integration of CSM and Slingshot technologies catalyzed CSCS to explore the provision of multitenancy and service isolation for its users. In the two years following the initial deployment of their large-scale system, CSCS leveraged both CSM and Slingshot to assess the viability of implementing strong (at the hardware-level) and soft (at the software-level) multitenancy on the HPE Cray Shasta EX system [1] [2]. A pivotal moment in this journey was the formalization of an agreement between HPE and CSCS, encapsulated in a statement of work that outlined binding conditions, including a dedicated work item to enable strong multitenancy capabilities.

To align HPE's development efforts with customer expectations and facilitate direct communication, HPE established a series of Special Interest Groups (SIGs) on key topics, including multitenancy. CSCS played a leading role in the multitenancy SIG, ensuring a focused and collaborative approach to overcoming the technical challenges associated with multitenancy on HPE supercomputers. As a direct outcome, inspired by CSCS's pioneering efforts, HPE initiated the development of Tenant and Partition Management System (TAPMS) [9], a multitenancy software designed to enable multitenancy capabilities on their EX series. The first version of TAPMS has been released, marking a significant milestone in HPE's software offerings.

The integration of multitenancy and network isolation capabilities, particularly through Slingshot, serves as a foundational pillar for strong tenant separation and introducing the IaaS model. This technology facilitates resource isolation at the network level, allowing for the management of independent services atop these segregated resources.

3 METHODOLOGY AND TECHNOLOGY SOLUTIONS

In this section, we will detail the specific requirements for implementing IaaS, focusing on those that present particular challenges for HPC systems. We will explore the microservices of the CSM and the Slingshot network, highlighting their critical features that facilitate multitenancy. In addition, we will discuss the developments undertaken by CSCS to deploy an IaaS service, thereby enabling PSI to leverage this advanced computing infrastructure.

3.1 Objectives and requirements

PSI's objective for their IaaS cluster is to manage and utilize the resources as though the cluster were an integral part of their own data center. This encompasses the ability to configure low-level services and node images, monitor performance, troubleshoot issues, and ensure a level of security that meets PSI's standards. Additionally, this encompasses the ability for system administrators to perform tasks on the nodes, such as rebooting. In essence, PSI seeks to seamlessly integrate the cluster into their existing ecosystem and services, necessitating the connection of compute nodes to their internal network for cohesive operation and security.

The workloads PSI plans to execute on the IaaS cluster span a diverse range of applications, including simulation workflows for cryogenic electron microscopy in biological research, offline analysis of particle physics experiments, and modeling and simulation for accelerator operation and development. These varied workloads introduce a spectrum of requirements, from traditional HPC batch scheduling with Slurm for compute-intensive tasks to support for interactive jobs facilitated through remote clients and Jupyter notebooks, and including high-throughput computing needs.

On the services front, PSI necessitates the integration of its own user identity management system. This integration helps maintain control over identity lifecycle management and ensures fine grained access control. Storage solutions are tailored to the specific needs of different workloads, incorporating HPC-enabled technologies such as Lustre for high-performance simulation runs and the Auristor File System (AFS) [4] for consistent distribution of software stacks on the compute nodes. The programming environment is made flexible and accessible through the use of modules, ensuring that users can easily load and manage software and libraries required for their research activities.

Regarding infrastructure, PSI's initial requirements encompass approximately 100 CPU-only AMD EPYC Milan compute nodes, alongside an additional 10 nodes, each equipped with 4 Grace Hopper GH200 superchips. For storage, there is a need for local SSDs on compute nodes, totaling around 200 TB. This need is complemented by a 10 PB of HDD storage with a Lustre file system accommodating extensive data storage and access demands. Given the distance separating the two institutions (approximately 230 Km), robust network connectivity is critical for supporting seamless service management. These specifications serve as a baseline, with expectations for scaling up as the new synchrotron device comes into full operation.

3.2 Cray System Management

The HPE Cray Shasta EX supercomputer is managed with the Cray System Management (CSM) [10] software. CSM incorporates traditional HPC system management functionalities but is designed with cloud computing principles in mind, facilitating the automation of management tasks alongside improved authentication mechanisms. As an open-source solution, CSM utilizes a microservices architecture and provides comprehensive RESTful API support, enabling detailed management and monitoring of all system components.

CSM is built on a microservices architecture, orchestrated via a Kubernetes cluster, which enhances its scalability and reliability. The architecture comprises several key microservices, each




Figure 1: CSM architecture: A microservices framework hosted on Kubernetes with comprehensive API accessibility. Image from the CSM [10] documentation.

responsible for distinct aspects of system management as shown in Figure 1.

The Version Control Service (VCS) oversees the internal Git repository and integrates GitOps solutions for streamlined operations and version control.

The Hardware State Manager (HSM) manages the organization of compute nodes into groups by adding labels to resources (HSM groups), oversees hardware inventory, and configures node Ethernet settings.

The Configuration Framework Service (CFS) is a core microservice designed to manage configurations, sessions, and components within a system. The following is a breakdown of its key components and functionalities:

- **Configuration Management:** CFS orchestrates the management of configurations, which are delineated as sequential layers of execution. Each layer encapsulates data sourced from a Git repository, as cataloged in the VCS, including essential details like the repository's URL, branch name, and commit ID.
- **Sessions:** CFS's operation are sessions, which are defined tasks that utilize a configuration to either generate an image or configure a node accordingly. This process involves cloning the specified Git repositories within the configuration and executing the scripts they contain. Currently, CFS

operates Ansible playbooks stored in Git repositories, leveraging their automation capabilities for session execution.

When constructing an image, CFS session tasks are orchestrated to execute within the CSM Kubernetes environment. Each session initiates two distinct pods: the initial pod is tasked with running Ansible playbooks, while the secondary pod functions as the target environment, equipped with a base image upon which the Ansible actions are executed. Following the sequential execution of Ansible scripts on the base image within the second pod, CFS diligently captures the modified image, subsequently storing it within an S3 bucket for future use.

In scenarios involving node configuration, CFS introduces the concept of “Components” which serve to associate a CFS configuration with a specific compute node. This setup enables Ansible scripts to directly target compute nodes, facilitating the deployment of configurations to designated nodes within the system.

The interplay among these components is graphically depicted in Figure 2. In essence, CFS streamlines the node configuration and image generation processes by managing the execution of Ansible playbooks sourced from Git repositories. Leveraging Kubernetes for efficient containerized job scheduling, CFS oversees the entire workflow from initial configuration through to the final deployment phase.




Figure 2: CFS sessions are set of tasks to build an image or configure a node. A session pulls a Git repository that contains Ansible scripts and then executes those scripts remotely with SSH to generate the image or configure the node.

The Boot Script Service (BSS) coordinates the management of kernel parameters, as well as kernel and initramfs files, essential for the booting process.

The Cray Advanced Platform Monitoring or Control (CAPMC) handles power management functions for nodes, including powering on, off, and resetting.

The Boot Orchestration Service (BOS) acts as an abstraction layer over HSM, CFS, BSS, and CAPMC, allocating configurations, boot images, and kernel parameters to node groups.

The Image Management Service (IMS) provides tools for creating system images using recipes, facilitating custom and efficient system deployments.

These microservices work in tandem to facilitate cluster management tasks for users. The workflow within CSM typically follows a structured pipeline:

- (1) Resource Grouping with HSM: Users start by creating an HSM group to label resources, such as compute nodes, effectively organizing them for management purposes.
- (2) Version Control with VCS: Next, Ansible scripts are developed and uploaded to VCS Git, ensuring that configurations are version-controlled, and centrally managed. The Git repository is organized by the HPC provider owning the system.
- (3) Configuration with CFS: Utilizing the scripts stored in VCS, users initiate CFS sessions. These sessions are responsible for generating the compute node images based on the specified configurations.
- (4) Boot Orchestration with BOS: "BOS session templates" are then created and used to associate specific images, kernel parameters, and configurations with a set of compute nodes defined by the HSM group.
- (5) Node instantiation with CAPMC: Finally, CAPMC service is used to reboot nodes, applying the new configurations and effectively instantiating the nodes as per the defined setup.

The integration of API interaction and access control within the microservices architecture, coupled with the ability to group resources and manage configurations independently for each group, lays the foundation for an initial version of multitenancy at the management level. CSM has been at the forefront of leveraging this capability to implement multitenancy, driving the evolution of system requirements. This pioneering work has prompted HPE to incorporate enhanced multitenancy features directly into CSM, culminating in the development of TAPMS, a testament of a co-design development for advancing HPC system management. Together with the network isolation and connectivity, this capability is a main pillar in enabling IaaS on HPE Cray Shasta EX system.

3.3 Technical Challenges and Solutions

Implementing IaaS presents a spectrum of technical challenges, from ensuring robust network connectivity and system management access to integrating PSI’s specialized services through custom base images on compute nodes. Equally critical is the provisioning of storage resources, including the requirements of node-local SSD. This subsection delves into these challenges, outlining the innovative solutions and strategies deployed to overcome them.

3.3.1 Network connectivity and isolation. Network connectivity encompasses two types: firstly, establishing a link between the two institutions, separated by approximately 100 km; and secondly, integrating compute nodes via a Slingshot network into an external Ethernet network.

Over the years, PSI and CSM have embarked on numerous collaborative projects to enhance the integration between the two facilities. A notable project aimed at establishing a geo-redundant tape archive resulted in the creation of two dedicated links, each with a capacity of 100 Gb/s, significantly increasing the connectivity between the institutions. These two dedicated optical circuits

with predefined paths are provided by the national research and education network, SWITCH, utilizing a dense wavelength division multiplexing system along divergent routes. These connections achieve a latency ranging between 4 and 5 milliseconds, ensuring efficient and reliable data transfer across the significant geographical distance and remote connection for services.

The Slingshot network is seamlessly integrated with CSCS data center network through eight 100 Gb/s Ethernet connections. Plans are underway to augment this bandwidth by either expanding the number of links or transitioning to 200 Gb/s Ethernet interfaces for enhanced throughput.

The Slingshot network facilitates the definition of Virtual Network Identifiers (VNIs) and Virtual Local Area Networks (VLANs), enabling sophisticated network segmentation and isolation. The Slingshot network has been configured to support VLAN-based network isolation, with inter-VLAN routing managed within CSCS site-wide network. This arrangement allows for the implementation of tailored security protocols for each VLAN, enhancing network integrity and data protection.

CSCS network employs multiple Virtual Routing and Forwarding (VRF) instances to ensure precise traffic segregation. Specifically, PSI logical network has its own VRF, encompassing at least two distinct VLANs. One VLAN is dedicated to facilitating Ethernet connectivity within the PSI logical network, encompassing management interfaces and access to external services. The second VLAN has been established on Slingshot to segment the compute resource network.

Integrating nodes into PSI network requires a two-step process involving VLAN configuration and IP address assignment:

- (1) **VLAN Configuration:** Initially, for nodes allocated to PSI, it's necessary to identify and modify the VLAN settings for the Slingshot switch ports connected to each node. Depending on the node type, up to four ports per node may require VLAN reconfiguration to align with PSI's network segmentation. Then policies for each VLAN are created and applied to all ports that are part of the VLAN.
- (2) **IP Address Assignment:** Subsequently, new IP addresses and gateway settings for each node are established. This is achieved by creating a file for each node, containing the new network configuration. Although nodes initially boot using CSCS default IP range, the presence of a network configuration file triggers the application of PSI-specific network settings during the initial configuration phase, integrating the node into PSI's network.

These steps have been automated by CSCS, facilitating the easy addition or removal of nodes to the IaaS cluster.

These actions and configuration ensure the complete segregation of both networks, with PSI's VRF and VLANs exclusively utilizing PSI IP addresses. This virtual network is fully integrated into PSI network infrastructure, maintaining operational independence and security. For instance, PSI benefits from the capability to access remote service such as mounting not only a local high-performance file system but also remote storage solutions present on their network, such as those used in experiment-related data acquisition end stations. This configuration approach also offers flexibility to CSCS, allowing for the entire system to boot in the default network mode

when necessary, such as during acceptance testing or benchmarking phases, while ensuring a smooth transition to PSI's dedicated network environment.

3.3.2 System management. CSM, with its microservices architecture with a multitude of APIs and tooling involved on single tenant operations, was not inherently designed to offer IaaS as a primary functionality. Granting PSI direct API access to CSM was deemed impractical for two primary reasons: security concerns, as actions could potentially impact the entire infrastructure, and the significant learning curve associated with mastering CSM's complexities. As a result, initial requests from PSI for configuration changes were managed by CSCS via a ticketing process. However, this approach was recognized as neither scalable nor efficient.

Given the absence of a built-in multi-tenancy feature and straightforward APIs for executing common IaaS actions within CSM, CSCS embarked on developing a dedicated tool to fulfill multi-tenancy needs, named Manta [22]. This API and associated Command-Line Interface (CLI) significantly simplifies user interaction with CSM APIs and tooling. Manta introduces an intuitive interface to efficiently manage essential tasks such as resource allocation, image generation, cluster configuration, power management, console access to compute nodes, and cluster definitions tailored to tenant's needs, all without external intervention in the infrastructure's management domain. Manta requires a straightforward configuration within CSM's Keycloak. This involves creating specific roles for each HSM group within the CSM realm and assigning these roles to users, thereby authorizing access to those HSM group resources. This configuration ensures that users can only access their permitted tenant resources, leveraging Keycloak's Role-Based Access Control (RBAC) for secure and targeted interaction.




Figure 3: Manta simplifies the aggregation of CFS session logs into a straightforward CLI command, efficiently retrieving information from CSM while adhering to robust RBAC protocols to manage multitenant clusters.

Developed through a co-design effort between CSCS and PSI, Manta is optimized for system engineers, and it is used by CSCS to manage its infrastructure more effectively.

Illustrating the complexity of managing a cluster within CSM, consider the process of retrieving logs from a CFS session. An operator must first utilize the CSM CLI to obtain CFS session details

within CSM, then extract the Kubernetes job name. This name is used with the Kubernetes CLI (kubectl) to fetch job information and identify the pod running the Ansible and Git containers. Only then can another kubectl command be issued to access the logs. Both Kubernetes and the CSM CLI lack multi-tenancy awareness, granting operators system-wide access. Consequently in a IaaS mode, granting external users access to these interfaces poses significant security risks from CSCS's standpoint. Manta simplifies this type of complex operations and safeguard tenant operations from system-wide access. For instance, acquiring CFS logs via Manta is condensed into a single command. This command not only validates the authorization of the request, but also aggregates the necessary log information, as illustrated in Figure 3. Manta facilitates many more multitenant operations than just retrieving logs. Here is a summary of high-level key features:

- **Power management:** Offers comprehensive control over power states, including the ability to power on, off, and reset both individual nodes and HSM groups of nodes.
- **Cluster configuration management:** Simplifies the process of node selection and clustering, enabling users to define and manage cluster configurations and boot images dynamically.
- **Data querying:** Allows for the retrieval of detailed historical data on clusters and nodes, including CFS configurations for image building or node configuration, and access to CFS session histories for tracking purposes.
- **CFS session logs:** Facilitates the troubleshooting of Ansible logs, providing insights into errors encountered during the image building or node configuration processes.
- **Node's console access:** Enables remote access to a node's interactive console, allowing users to diagnose and resolve boot issues or hardware malfunctions.
- **Multi-site management:** Empowers a single Manta instance to oversee multiple CSM instances, supporting functionalities like cluster migration, as well as backup and restoration of clusters across various CSM instances.

In summary, Manta facilitates interaction with CSM and introduces IaaS multitenant management with RBAC.

3.3.3 Storage. Addressing storage access within an IaaS model posed significant challenges on an HPE Cray Shasta EX system. On the one hand, the technology available at the time for storage solution on a HPE Cray Shasta EX system lacked multitenant capabilities; on the other, the compute nodes within the infrastructure did not support local disks as required by PSI.

Given the absence of multitenant support in storage, a decision was made to allocate dedicated storage resources. A Lustre file system, utilizing HPE ClusterStor technology, was provisioned and dedicated to the IaaS cluster.

The need for local disk provisioning stemmed from the requirement to handle high metadata operations and high IOPS. Since integrating local disks into nodes was not feasible, a solution was found in provisioning the dedicated Lustre file system with two different pools. One pool used traditional HDDs for scratch space, providing high bandwidth and a large capacity. A second pool utilized SSDs to better meet the requirements for a high number of IOPS. In order to shield the Lustre services from excessive metadata accesses, the compute nodes will use scratch file systems based

on loopback mounted files on the SSD pools. Although other approaches like projecting disk views on nodes using NVMe over Fabrics (NVMeoF) were considered, they did not align with the project's schedule due to integration with the Slingshot roadmap.

The decision was made to connect the management of the HPE ClusterStor system to the PSI network, enabling PSI to fully oversee the appliance and monitor it using their own tools. This approach also mitigated potential security risks that could arise from bridging the networks between the two partners. Consequently, CSCS agreed to offer high-level support in the form of consultancy as needed, while HPE committed to providing both hardware support and additional high-level assistance to PSI via their ticketing system.

Additionally, PSI's storage solutions include securely and exclusively mounting their remote AFS file system. This setup enabled PSI users to access their private datasets for stage-in into the cluster's high performance Lustre, and facilitate collaboration with other services.

3.3.4 Service management. To facilitate the integration of PSI's internal services with the IaaS cluster, several prerequisites must be met. These services, including LDAP, DNS, NTP, Slurm, AFS mounting points, and others, necessitate PSI having three key capabilities: root privileges on compute nodes to create service deployment recipes; the ability to deploy these recipes effectively; and a dedicated network connection linking PSI's internal network to the IaaS cluster. It is crucial that the configuration of these services remains persistent across node reboots and are automatically applied to minimize manual intervention.

In the IaaS model, as the cluster owner or tenant, PSI bears the full responsibility for the service deployment and associated node configuration. By segregating the high-speed network of the compute nodes, which then facilitates their connection to PSI's internal network, these nodes can access a Kubernetes cluster and virtual machines managed by PSI, which hosts a suite of essential services. Post-boot, once the nodes are verified as operational via the CSM, additional specific configurations are applied. This includes executing custom Ansible scripts to configure services such as AFS mounting points.

CSCS grants PSI high-level access privileges to the compute nodes, thanks to the network segregation that prevents PSI engineers to access other segments of the infrastructure. Another requirement for PSI team is to use SUSE distribution, which prevents the team from leveraging PSI's existing deployment infrastructure tailored to the RHEL distribution. CSM currently supports only the SUSE distribution provided by HPE, along with a custom kernel. The option to use alternative distributions is anticipated to be available in the future.

3.3.5 Security. From the CSCS perspective, security is primarily focused on ensuring that no actions on the Alps management system can be executed by PSI outside of those explicitly permitted. This is enforced through Manta, which restricts the set of permissible actions, and by employing role-based access control to ensure PSI operators can only access resources labeled for their specific IaaS cluster.

For PSI, a key security requirement is the jointly defined security zone established through coordinated network configurations across both institutions. This zone is securely managed using PSI's

own policies and rules. Technologies such as VLANs, VRFs, and IP-based segmentation within the CSCS network allow PSI's security team to maintain full control over network traffic, enabling them to enforce their security standards effectively.

Additionally, since the storage infrastructure is fully dedicated to and managed by PSI, CSCS does not bear any responsibility for storage-related security concerns – this further reinforces the separation of responsibilities and simplifies both institution security postures.

3.4 Tenant operational view




Figure 4: Tenant operational view, the IaaS cluster belongs to a jointly security zone and is integrated into PSI network. Operators can configure their services with existing mechanisms whereas they use an identity provided by CSCS to execute system management actions through Manta. For PSI's users access to the IaaS cluster is transparent.

Figure 4 illustrates PSI operational perspective of the IaaS cluster. PSI operators, utilizing identities granted by CSCS, employ Manta to modify the configuration of compute resources providing a minimal integration layer needed for the service provisioning. Thanks to that layer, PSI configures specific services running on the compute nodes like Slurm or LDAP, using existing deployment and configuration mechanisms. While CSCS maintains a separate configuration layer on the compute nodes for network access control, PSI operators leverage DevOps pipelines running from PSI service infrastructure to manage both service settings at PSI operational Kubernetes cluster and at the IaaS cluster compute resources.

End-users accessing the IaaS cluster do so through the PSI network, authenticated with credentials provided by PSI. They interact with the resources via login nodes, experiencing seamless integration as if using a standard cluster hosted by PSI. This setup conceals the underlying complexity and the fact that the resources are physically located elsewhere, within the broader research infrastructure provided by CSCS.

3.5 Collaborative Framework

The collaboration between CSCS and PSI aimed to tailor the IaaS service to meet PSI's specific needs, while also partnering with HPE to enhance multitenant capabilities and associated technologies. A memorandum of understanding was jointly signed by PSI and CSCS, outlining the plan to integrate PSI's next research cluster as a tenant within CSCS's research infrastructure. Subsequently, dedicated teams from each institution were formed, adhering to project management principles through regular meetings, technology workshops, risk assessments, and steering committee discussions over a three-year period.

Throughout this time, CSCS's infrastructure underwent phased upgrades, incorporating updates in HPE technology such as new version of CSM, slingshot and new hardware. These updates required downtime of the infrastructure and occasionally disrupted PSI's progress, necessitating the creation of mitigation plans. To maintain continuity, PSI was transitioned between multiple CSM instances. This evolving technological update of CSM required ongoing adaptation, and knowledge transfer was facilitated by direct communication channels between the two teams.

A notable challenge involved granting PSI high-level privileges for node configuration before network segregation was fully implemented. This was resolved through a security agreement between the institutions, specifying permitted actions and monitoring protocols for allocated nodes. Additionally, a significant knowledge transfer occurred regarding Lustre file system management as the PSI team was not familiar with Lustre management. Later, an agreement was concluded that CSCS would procure storage capabilities through HPE, but PSI would have the right to directly address storage-related issues with HPE. This tripartite arrangement between HPE, CSCS, and PSI exemplifies the collaborative effort to develop a robust and responsive multitenancy.

4 LESSONS LEARNED

Reflecting on the project, three key lessons emerged from the collaboration between CSCS and PSI in developing IaaS and strong tenant separation capabilities:

Adaptability in project roadmap and features: The project's journey highlighted the necessity of adaptability driven by the evolution of the resource availability and feature set roadmaps. Challenges inevitably arose, requiring the formulation of mitigation strategies. However, an early start to the project proved advantageous, allowing work to progress despite obstacles. Such as the decision to open early access to Manta to PSI leading to the second insight.

Early user engagement drives development: Involving PSI as an early user of the Manta tool was a strategic decision that paid dividends. PSI's active use and feedback on Manta were instrumental in shaping its development, ensuring that the tool was aligned with user needs and capable of addressing real-world challenges.

The value of strong collaboration: The partnership between CSCS and PSI was marked by a high level of cooperation and close collaboration. This strong bond facilitated a rich exchange of ideas and fostered a re-enforced trust between the two institutions. Furthermore, the collaboration between CSCS and HPE in implementing multitenancy has yielded significant benefits for both parties.

For HPE, this partnership has opened new opportunities. For CSCS, it has secured a commitment from HPE to deliver a comprehensive set of features for facilitating tenant management.

These experiences demonstrated that beyond technical alignment, building relational bridges is crucial for the success of complex co-design projects.

5 FUTURE DEVELOPMENT

Regarding future work, CSCS and HPE are committed to contributing to the development of OpenCHAMI (Open Composable Heterogeneous Application Management Infrastructure) [19]. OpenCHAMI, an open-source HPC system management software, adopts a microservices architecture similar to that of CSM and is designed to be infrastructure-agnostic. Managed by a consortium that includes US national laboratories, European institutions, and HPE, OpenCHAMI leverages CSM as its foundational technology. CSCS aims to influence the feature set of OpenCHAMI, emphasizing multitenancy, and plans to integrate Manta as a user-friendly interface to enhance its accessibility.

Additionally, the development of the Manta API is set to continue. Efforts will focus on expanding its user interfaces, including the introduction of a web interface. This enhancement aims to empower users with self-service capabilities, simplifying the process of managing and accessing HPC resources. Through multiple technologies such as Manta, CSCS is paving the way for more versatile and software-defined HPC system management solutions [15].

For the IaaS cluster, future enhancements will address the specific needs of machine learning (ML) and artificial intelligence applications. Based on PSI's requirements for storage efficiency for ML workloads, CSCS is exploring software-defined multitenant global storage solutions based on SSD technology to improve efficiency of IO for ML tasks. This initiative demonstrates CSCS's commitment to evolving its IaaS platform to meet the growing demands of ML tasks from PSI.

CSCS is planning to offer tenant capability with IaaS clusters to other institutions with a second project already on-going. In particular, such tenant separation capability facilitate technology deployment and responsibility in managing sensitive data while enabling HPC capability, essentials for use cases such as processing medical data with ML models.

6 RELATED WORK

Exploring the alternative of deploying HPC cluster on IaaS-capable cloud platforms [3] reveals significant inefficiencies, particularly for applications requiring intense computational power and high-speed networking. Cloud environments often fall short in both performance and cost-effectiveness for HPC tasks [18]. This shortfall is attributed to the comparatively lower performance of cloud network hardware [8], and the overhead associated with virtualization [5].

In contrast, Azure offers access to Cray XC systems, albeit through a dedicated, single-tenant model, which does not leverage the full potential of IaaS for flexible resource allocation and management [16]. Furthermore, for the HPE Cray EX series, a partnership has been established to provide a dedicated machine to the UK Met

Office, reinforcing the trend towards single-tenant supercomputing solutions [17].

Alternative system management solutions like OpenStack provide network flexibility and virtualization capabilities. However, their complexity, particularly in API layers and underlying technology, restricts access to bare-metal performance levels [25]. StackHPC represents an effort to integrate HPC functionalities within the OpenStack ecosystem [13]. Similarly, Metal-as-a-Service solutions offer another approach to managing hardware resources [6]. Despite these initiatives, neither solution fully addresses the unique requirements of high-end, large-scale HPC systems, such as the HPE Cray Shasta EX series.

7 CONCLUSION

In conclusion, this paper explores the pioneering implementation of strong tenant separation using an IaaS model on the HPE Cray Shasta EX supercomputer. Despite the inherent challenges posed by the lack of virtualization technologies in traditional HPC systems, our collaborative co-design approach with the partner institution has led to the development of a suite of technologies that exposes and simplifies resource isolation, granular control, and dynamic resource management from the CSM. Our work underscores the importance of collaboration, adaptability, and technological innovation in bridging the gap between cloud computing and supercomputing, offering valuable insights and directions for future research and development in making multitenancy a viable and effective model for HPC systems.

ACKNOWLEDGMENTS

An AI-generated tool based on ChatGPT built on GPT-4 architecture has been used to enhance the readability of all sections of this document. The authors have carefully integrated the AI-suggested edits to preserve the intended meaning. The AI tool was not used to generate ideas or data.

REFERENCES

- [1] Sadaf R. Alam, Miguel Gila, Mark Klein, and Maxime Martinasso. 2021. Multitenancy Management and Zero Downtime Upgrades using Cray-HPE Shasta Supercomputers. In *2021 SC Workshops Supplementary Proceedings (SCWS)*, 87–94. doi: 10.1109/SCWS55283.2021.00021.
- [2] Sadaf R. Alam, Miguel Gila, Mark Klein, Maxime Martinasso, and Thomas C. Schulthess. 2023. Versatile software-defined HPC and cloud clusters on Alps supercomputer for diverse workflows. *The International Journal of High Performance Computing Applications*, 37, 3–4. doi: 10.1177/10943420231167811.
- [3] Rawan Aljamal, Ali El-Mousa, and Fahed Jubair. 2018. A comparative review of high-performance computing major cloud service providers. In *2018 9th International Conference on Information and Communication Systems (ICICS)*, 181–186. doi: 10.1109/IACS.2018.8355463.
- [4] AuriStor, Inc. 2024. AuriStor The Global Namespace File System. <https://www.auristor.com/filesystem/>. Accessed: 2024-03-15. (2024).
- [5] Aditya Bhardwaj and C Rama Krishna. 2021. Virtualization in cloud computing: Moving from hypervisor to containerization—a survey. *Arabian Journal for Science and Engineering*, 46, 9, 8585–8601.
- [6] Canonical. 2015. Metal as a service (MAAS). (2015).
- [7] Felipe A Cruz, Alejandro J Dabin, Juan Pablo Dorsch, Eirini Koutsaniti, Nelson F Lezcano, Maxime Martinasso, and Dario Petrusic. 2020. FirecREST: a RESTful API to HPC systems. In *2020 IEEE/ACM International Workshop on Interoperability of Supercomputing and Cloud Technologies (SuperCompCloud)*. IEEE, 21–26.
- [8] Daniele De Sensi, Tiziano De Matteis, Konstantin Taranov, Salvatore Di Girolamo, Tobias Rahn, and Torsten Hoefler. 2022. Noise in the Clouds: Influence of Network Performance Variability on Application Scalability. *Proc. ACM Meas. Anal. Comput. Syst.*, 6, 3, Article 49, (Dec. 2022), 27 pages. doi: 10.1145/3570609.

- [9] Jeremy Duckworth, Vinay Gavirangaswamy, David Gloe, and Brad Klein. 2023. Software-defined Multi-tenancy on HPE Cray EX Supercomputers. In *Proceedings of the Cray User Group (CUG) 2023*. Accessed: 2024-03-15. https://cug.org/proceedings/cug2023_proceedings/includes/files/pap132s2-file2.pdf.
- [10] Hewlett Packard Enterprise. 2021. Cray System Management Documentation. <https://cray-hpe.github.io/docs-csm/en-10/>. Accessed: 2024-03-15. (2021).
- [11] Hewlett Packard Enterprise. 2023. HPE Cray Supercomputing EX. <https://www.hpe.com/psnow/doc/a00094635enw>. Accessed: 2024-03-15. (2023).
- [12] Shamim Hossain. 2013. Infrastructure as a Service. (2013). <https://api.semanticscholar.org/CorpusID:63680470>.
- [13] John Garbutt. 2022. Using OpenStack to reduce HPC service complexity. https://archive.fosdem.org/2022/schedule/event/openstack_hpc/attachments/slides/5136/export/events/attachments/openstack_hpc/slides/5136/OpenStackHPCSlides.pdf. Accessed: 2024-03-15. (2022).
- [14] S. Leong, H. Stadler, M. Chang, J. Dorsch, T. Aliaga, and A. W. Ashton. 2020. SELVEDAS: A Data and Compute as a Service Workflow Demonstrator targeting Supercomputing Ecosystems. In *2020 IEEE/ACM International Workshop on Interoperability of Supercomputing and Cloud Technologies (SuperCompCloud)*. IEEE Computer Society, Los Alamitos, CA, USA, (Nov. 2020), 7–13. doi: 10.1109/SuperCompCloud51944.2020.00007.
- [15] Maxime Martinasso, Mark Klein, Benjamin Cumming, Miguel Gila, Felipe A. Cruz, Alberto Madonna, Manuel Sopena Ballesteros, Sadaf R. Alam, and Thomas C. Schulthess. 2024. Versatile software-defined cluster for HPC using cloud abstractions. *Comput. Sci. Eng.*, 26, 3, 20–29. doi: 10.1109/MCSE.2024.3394164.
- [16] Microsoft Azure. 2024. High Performance Computing (HPC) Solutions on Azure. <https://azure.microsoft.com/en-us/solutions/high-performance-computing/#cray-supercomputing>. Accessed: 2024-03-15. (2024).
- [17] Microsoft Azure. 2024. Microsoft Brings Azure Supercomputing to UK Met Office. <https://azure.microsoft.com/en-us/blog/microsoft-brings-azure-supercomputing-to-uk-met-office/>. Accessed: 2024-03-15. (2024).
- [18] Marco A. S. Netto, Rodrigo N. Calheiros, Eduardo R. Rodrigues, Renato L. F. Cunha, and Rajkumar Buyya. 2018. HPC Cloud for Scientific and Business Applications: Taxonomy, Vision, and Research Challenges. *ACM Comput. Surv.*, 51, 1, Article 8, (Jan. 2018), 29 pages. doi: 10.1145/3150224.
- [19] OpenCHAMI. 2023. Open Composable Heterogeneous Adaptable Management Infrastructure. <https://www.ochami.org/>. Accessed: 2024-03-15. (2023).
- [20] Paul Scherrer Institute. 2023. The Upgrade Project: SLS 2.0. <https://www.psi.ch/en/media/the-upgrade-project-sls-20>. Accessed: 2024-03-15. (2023).
- [21] Daniele De Sensi, Salvatore Di Girolamo, Kim H. McMahon, Duncan Roweth, and Torsten Hoefler. 2020. An In-Depth Analysis of the Slingshot Interconnect. *CoRR*, abs/2008.08886. <https://arxiv.org/abs/2008.08886> arXiv: 2008.08886.
- [22] Swiss National Supercomputing Centre (CSCS). 2024. Manta: Another CLI tool for Alps. <https://github.com/eth-cscs/manta>. Accessed: 2024-03-15. (2024).
- [23] Swiss National Supercomputing Centre (CSCS) and Hewlett Packard Enterprise (HPE) and NVIDIA. 2021. CSCS, Hewlett Packard Enterprise, and NVIDIA Announce World’s Most Powerful AI-Capable Supercomputer. <https://www.cscs.ch/science/computer-science-hpc/2021/cscs-hewlett-packard-enterprise-and-nvidia-announce-worlds-most-powerful-ai-capable-supercomputer>. Accessed: 2024-03-15. (2021).
- [24] Swiss National Supercomputing Centre (CSCS) and Paul Scherrer Institute (PSI). 2018. CSCS Will Store Petabyte Data for the Paul Scherrer Institute. <https://www.cscs.ch/publications/press-releases/2018/cscs-will-store-petabyte-data-for-the-paul-scherrer-institute>. Accessed: 2024-03-15. (2018).
- [25] S Telfer. 2016. The Crossroads of Cloud and HPC: OpenStack for Scientific Research: Exploring OpenStack cloud computing for scientific workloads. (2016).