

# The HPC + AI Cloud

Flexible and performant infrastructure  
for HPC and AI workloads

Matt Pryor, Senior Technical Lead, StackHPC



StackHPC


# StackHPC Company Overview



StackHPC

- Formed 2016, based in Bristol, UK
  - Based in Bristol with presence in Oxford, Cambridge, France and Poland
  - Currently around 30 people
- Founded on HPC expertise
  - Software Defined Networking
  - Systems Integration
  - Open Research Infrastructure Development and Operations
- Motivation to transfer this expertise into Cloud to address HPC & HPDA (AI)
- “Open” Modus Operandi
  - Upstream development of OpenStack capability
  - Consultancy/Support to end-user organizations in managing HPC service transition
  - Scientific-SIG engagement for the Open Infrastructure Foundation
- Hybrid Cloud Enablement

StackHPC





Getting Maximum Value from your Investment?

# Getting Maximum Value from your Investment




# StackHPC Three Pillars



Reconfigurable and isolated infrastructure



Performance to extract maximum value



Azimuth self-service platforms

# Open Source Co-Development



StackHPC



JASMIN




iris



GRAPHCORE



Science and  
Technology  
Facilities Council











DIRAC


# Diverse Use Cases



StackHPC




# Infrastructure



**Reconfigurable and isolated infrastructure**




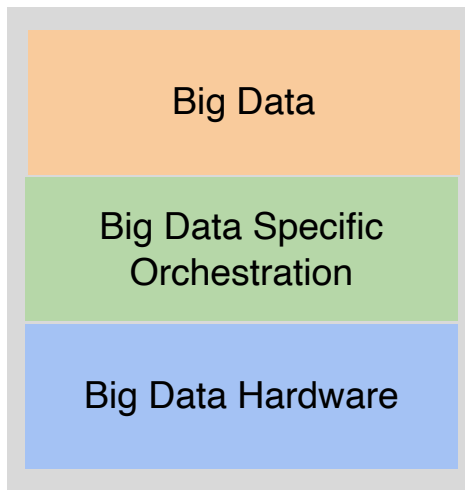
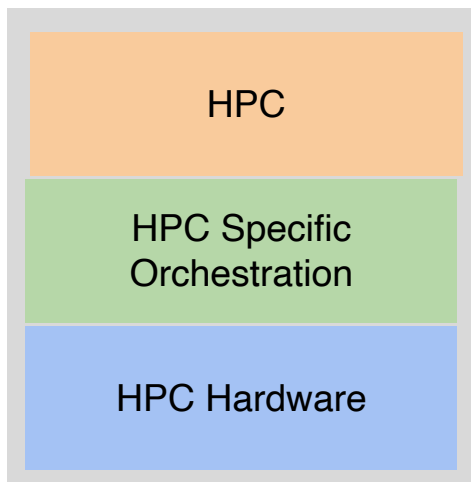
Performance to extract maximum value



Azimuth self-service platforms



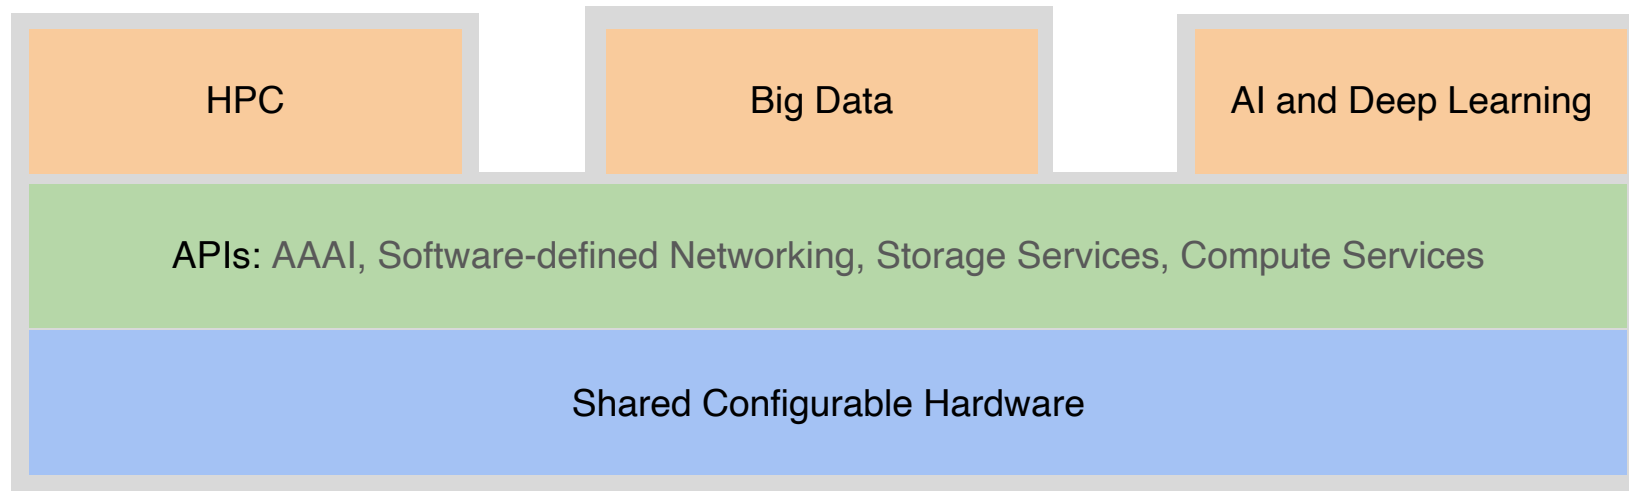
# Hardware silos



# Cloud-native HPC



StackHPC



# OpenStack to the rescue

- OpenStack is an open-source cloud platform
- Strong multi-tenancy guarantees
- APIs for compute, network and storage
  - Dashboard and command-line interface
- DevOps tools for OpenStack and platforms
  - Code reviewed changes
  - Continuous integration and delivery



openstack®


{ REST }

OpenTofu<sup>📦</sup>

# Cloud-native HPC site




StackHPC



# Cloud-native HPC site



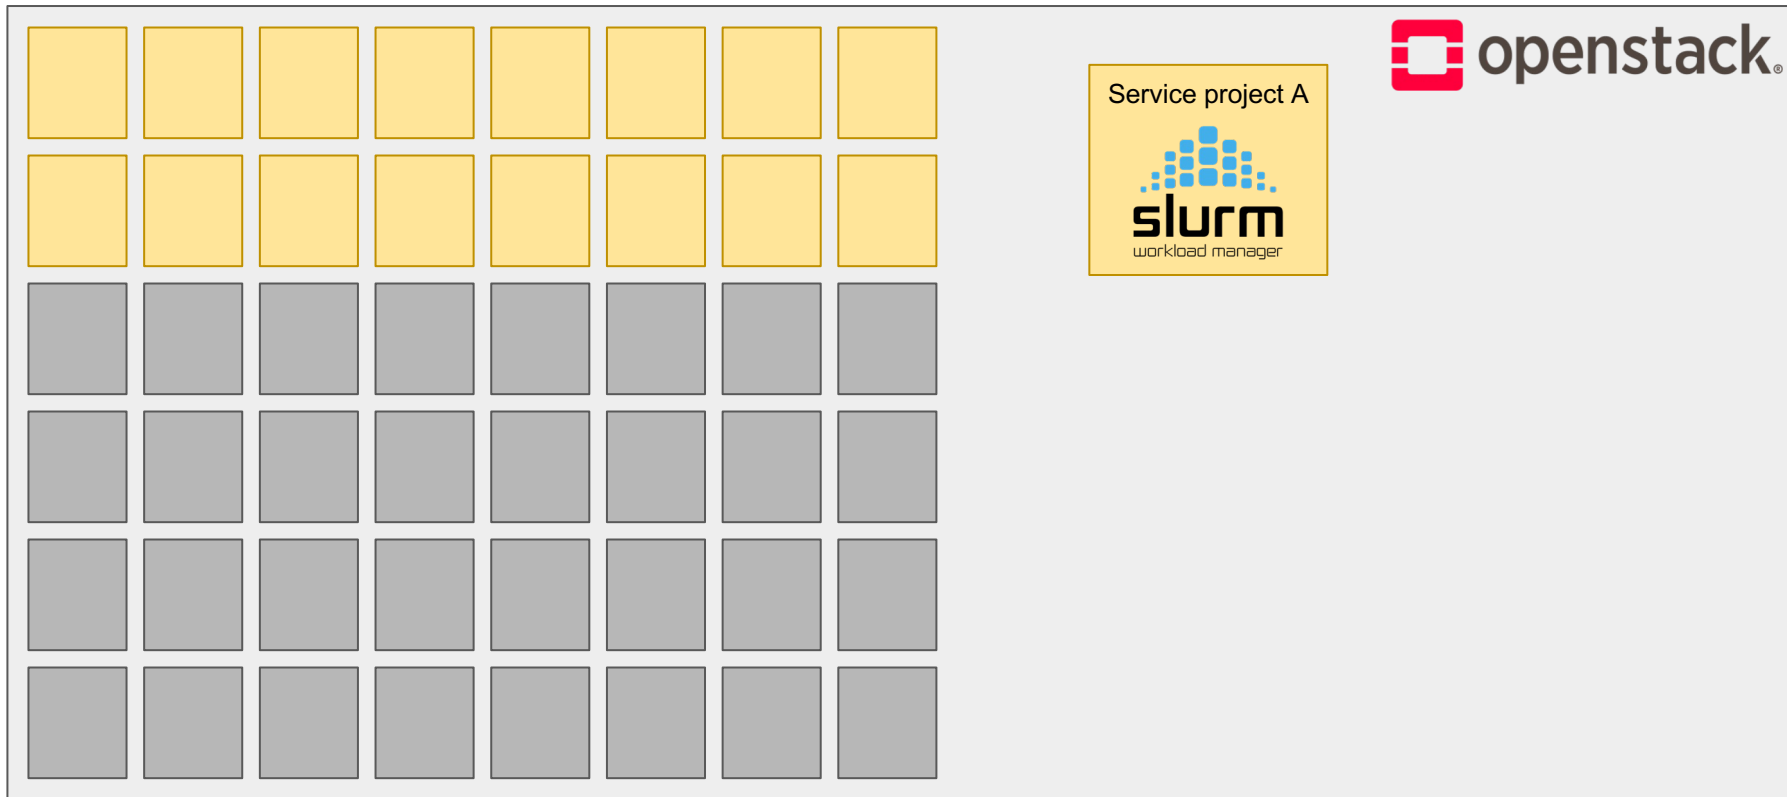
StackHPC



# Cloud-native HPC site



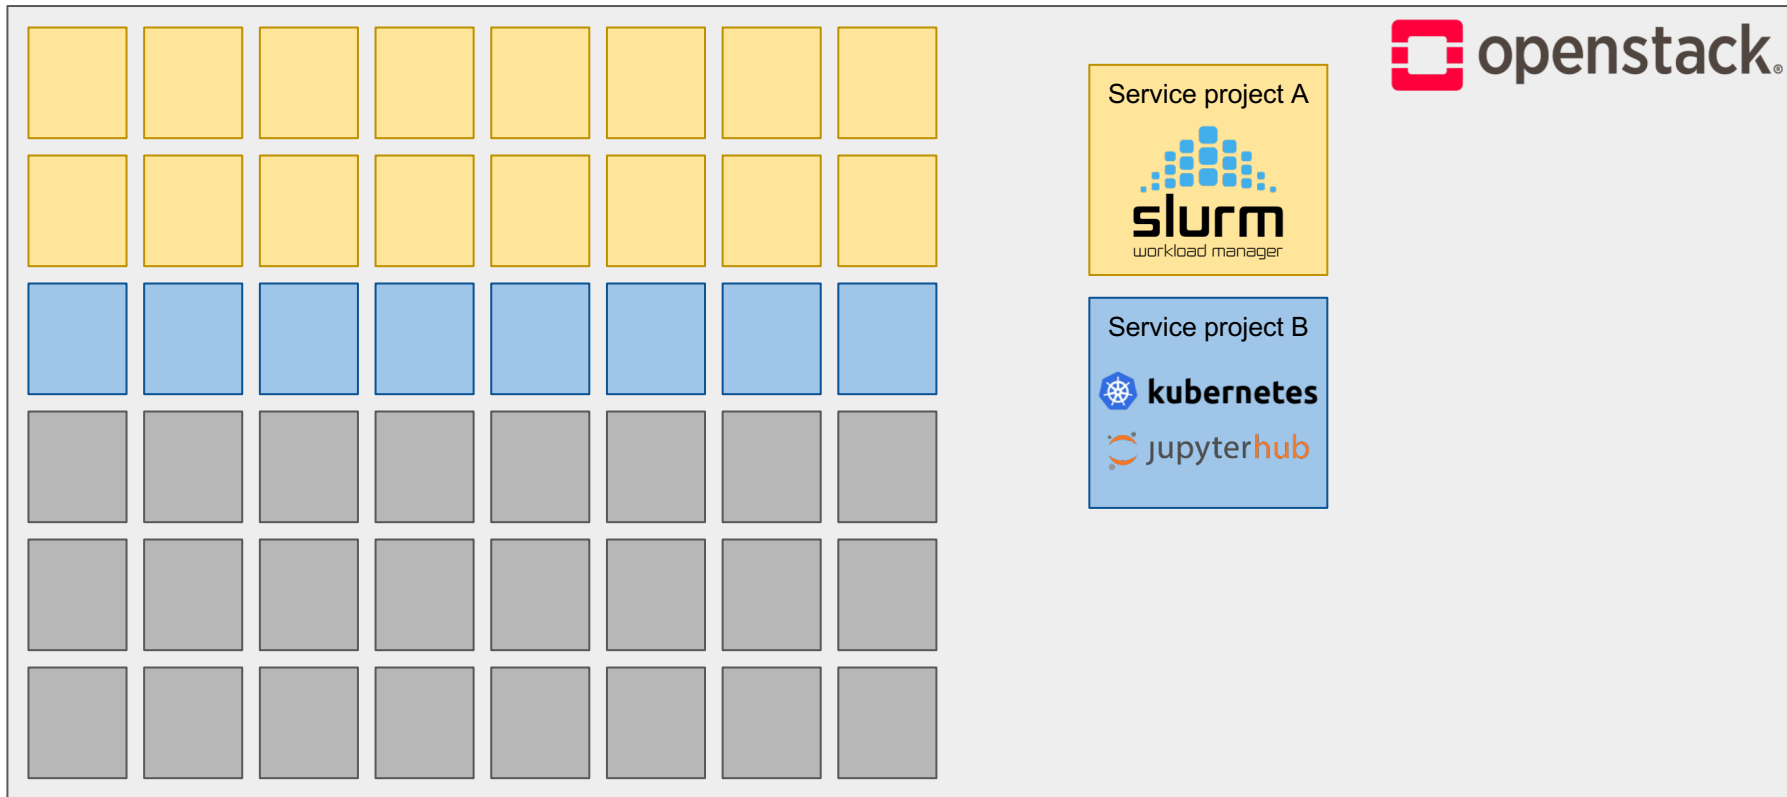
StackHPC



# Cloud-native HPC site



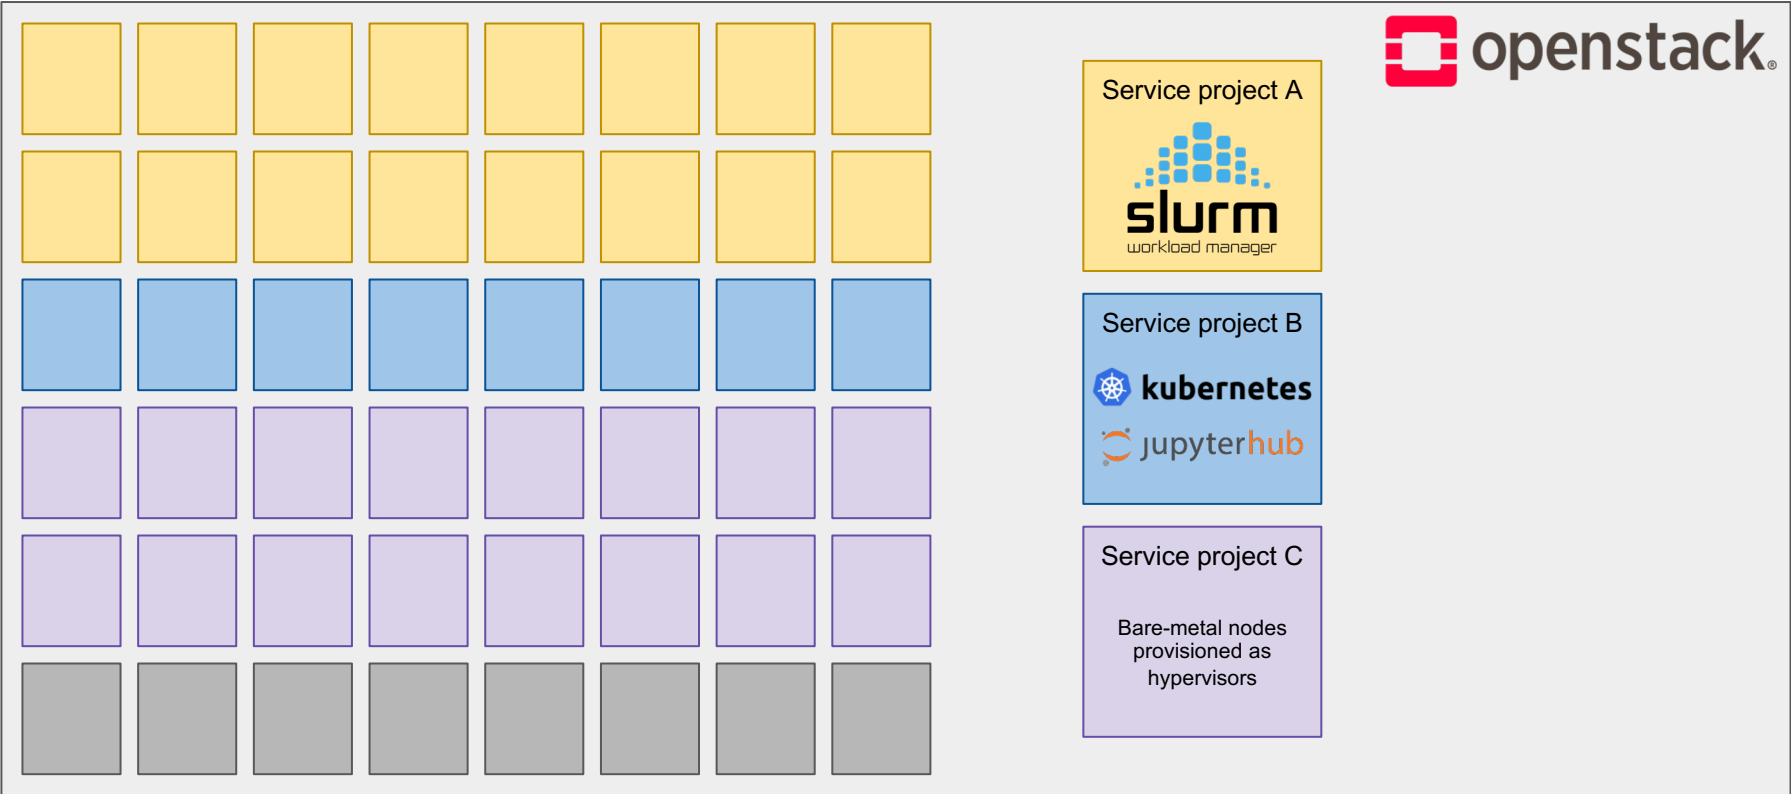
StackHPC



# Cloud-native HPC site



StackHPC

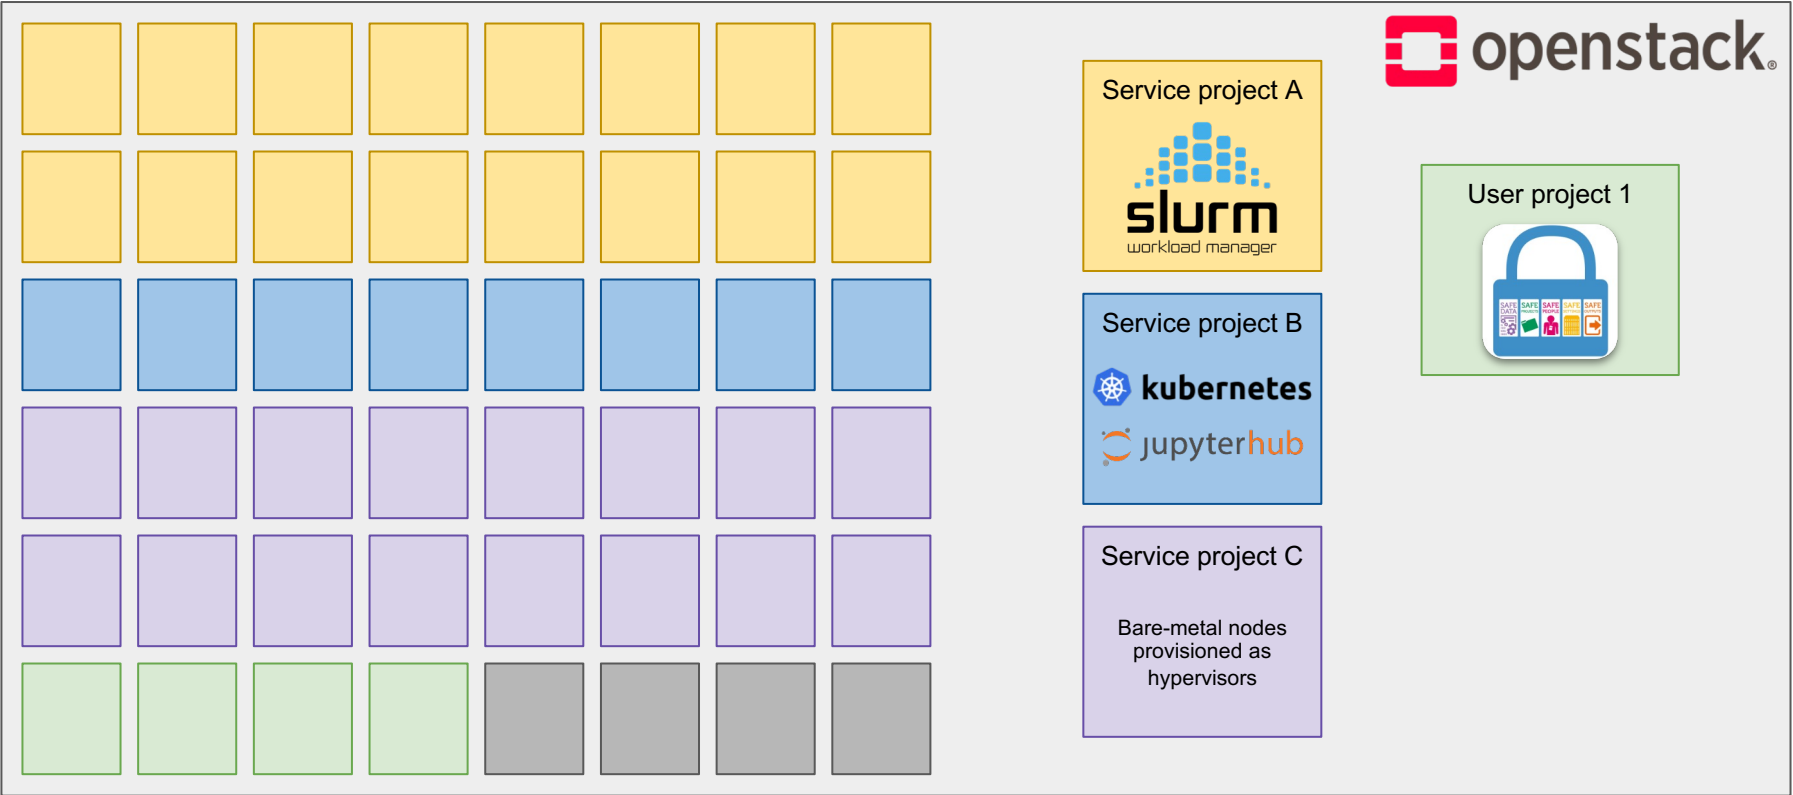




# Cloud-native HPC site



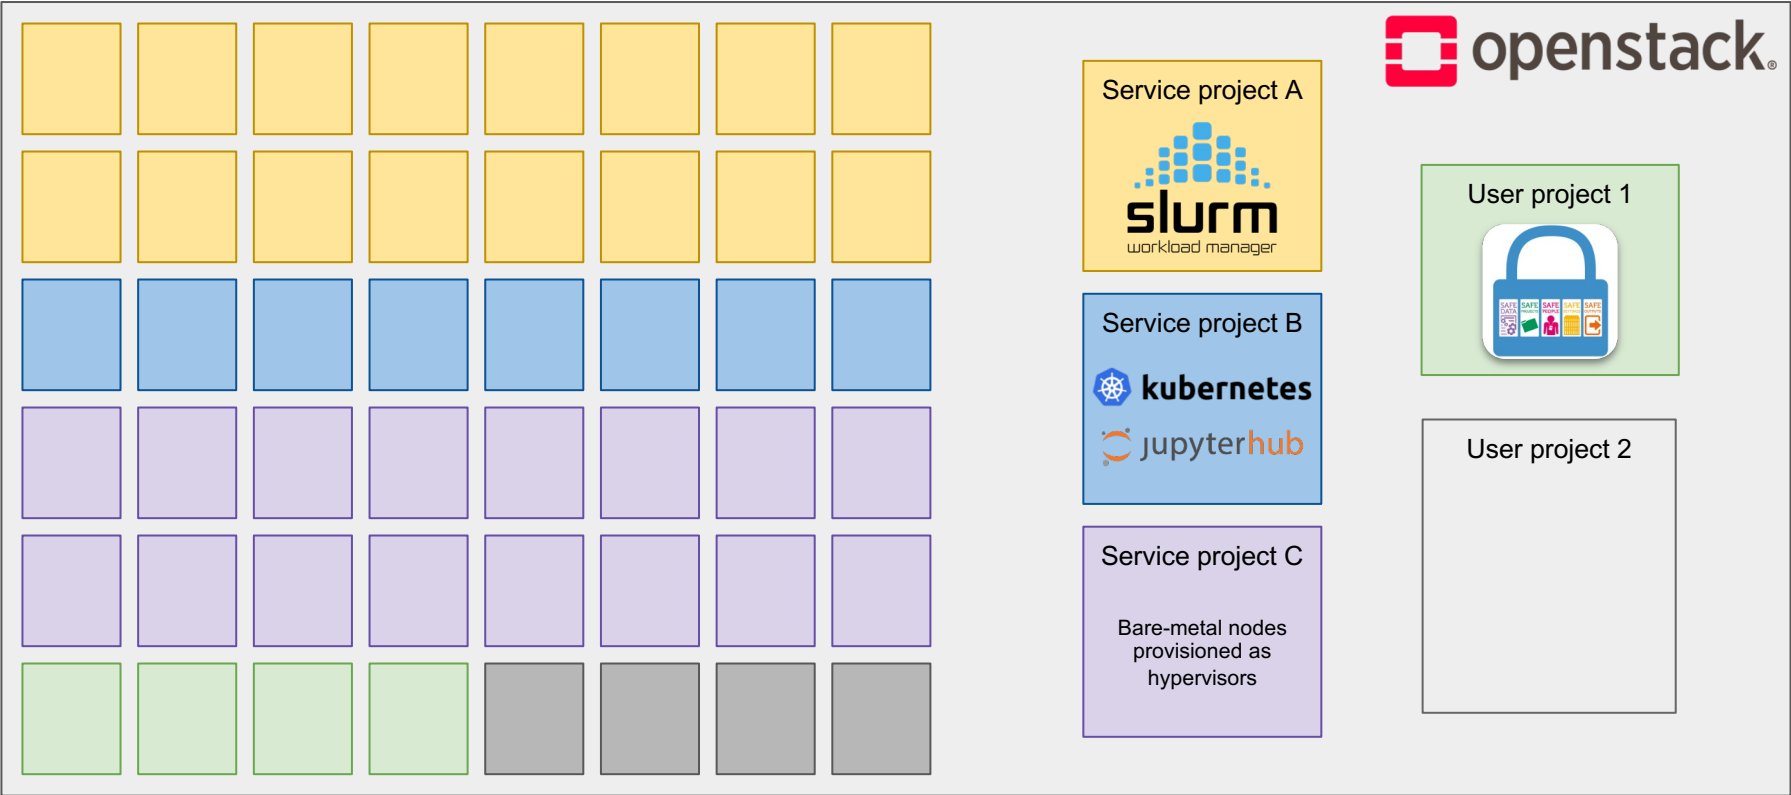
StackHPC



# Cloud-native HPC site



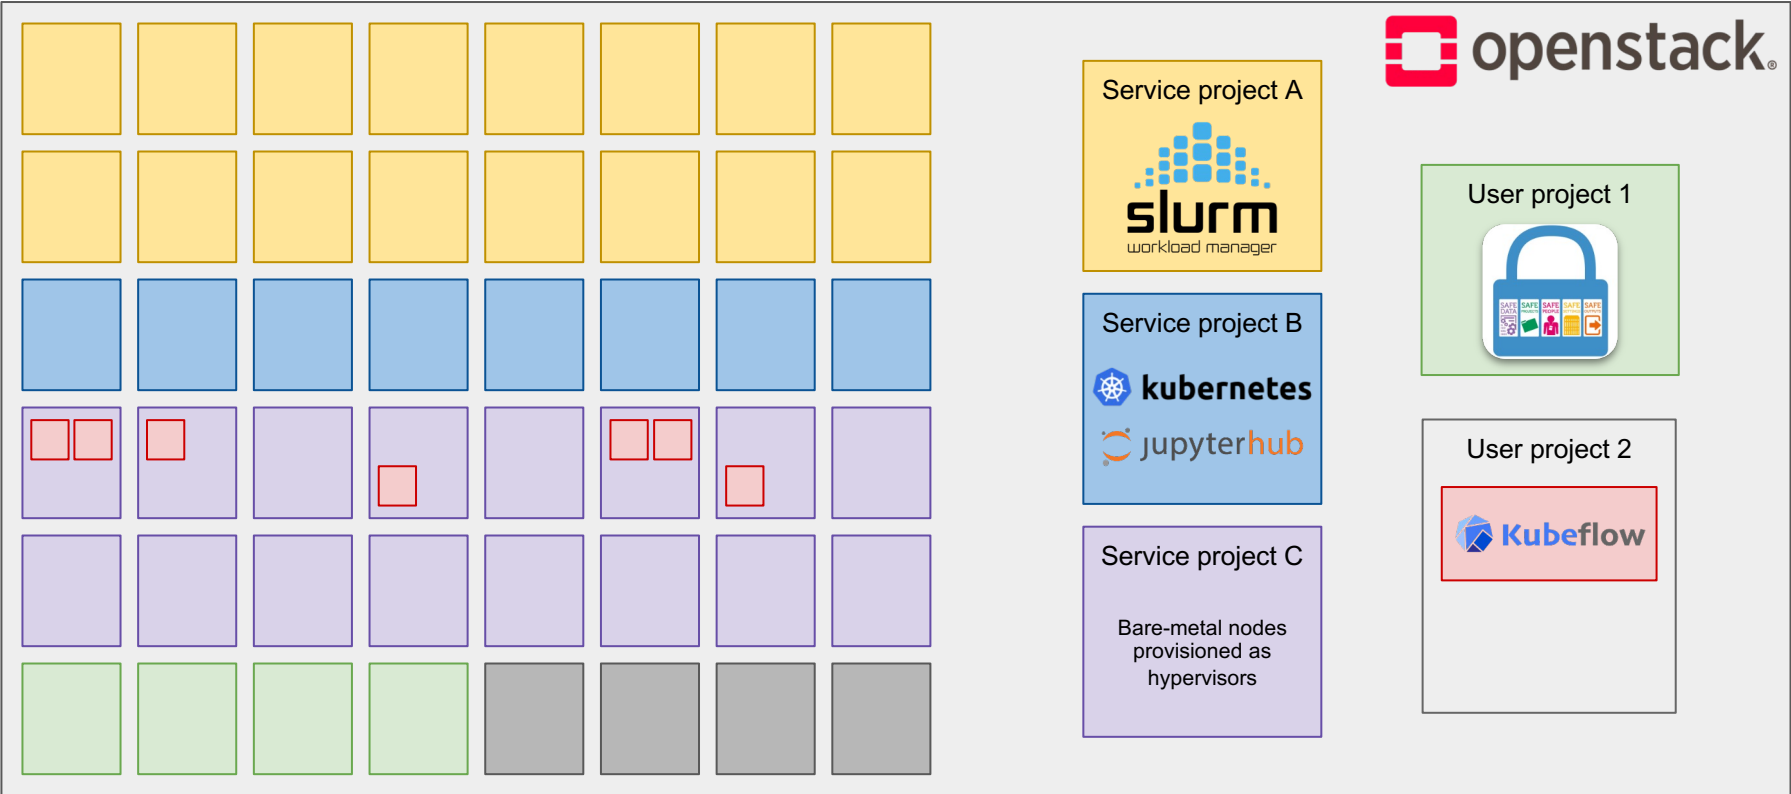
StackHPC



# Cloud-native HPC site



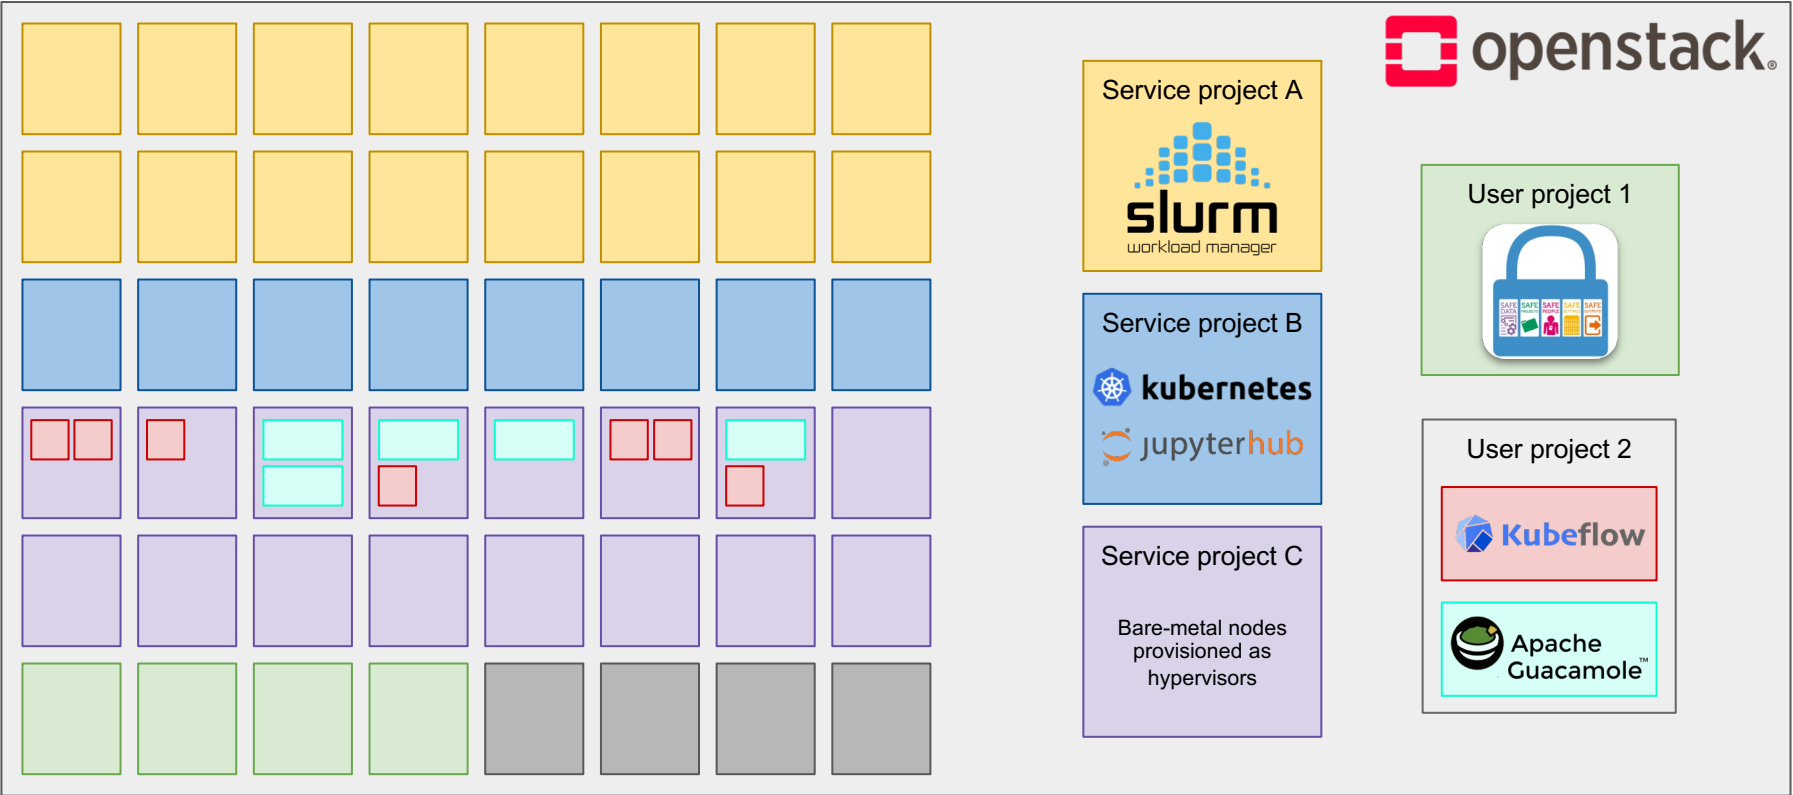
StackHPC



# Cloud-native HPC site



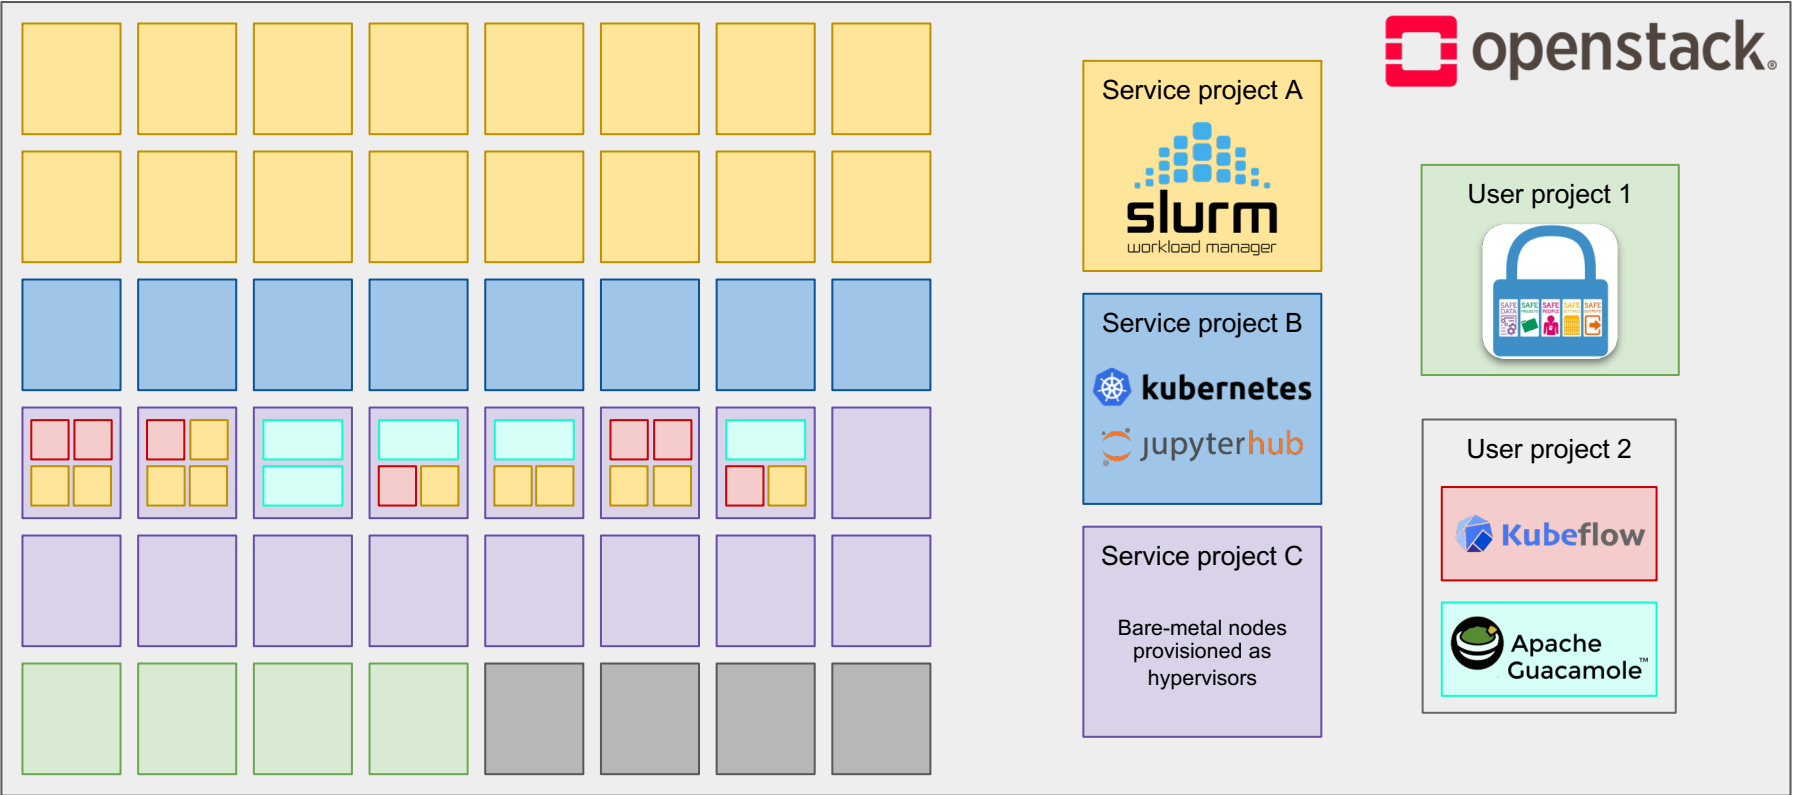
StackHPC



# Cloud-native HPC site



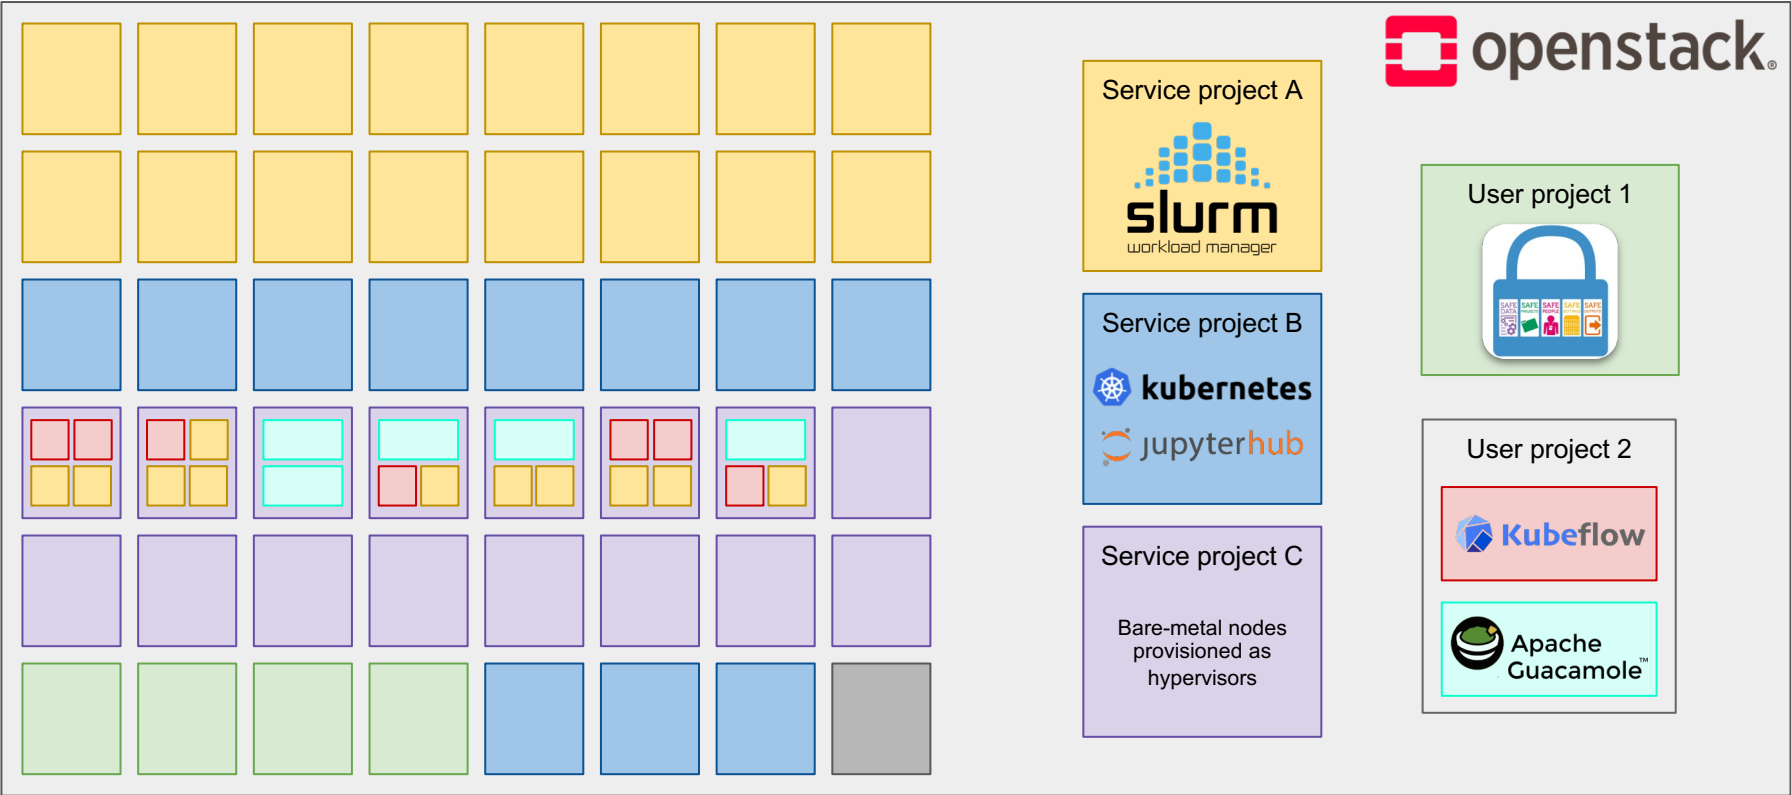
StackHPC



# Cloud-native HPC site



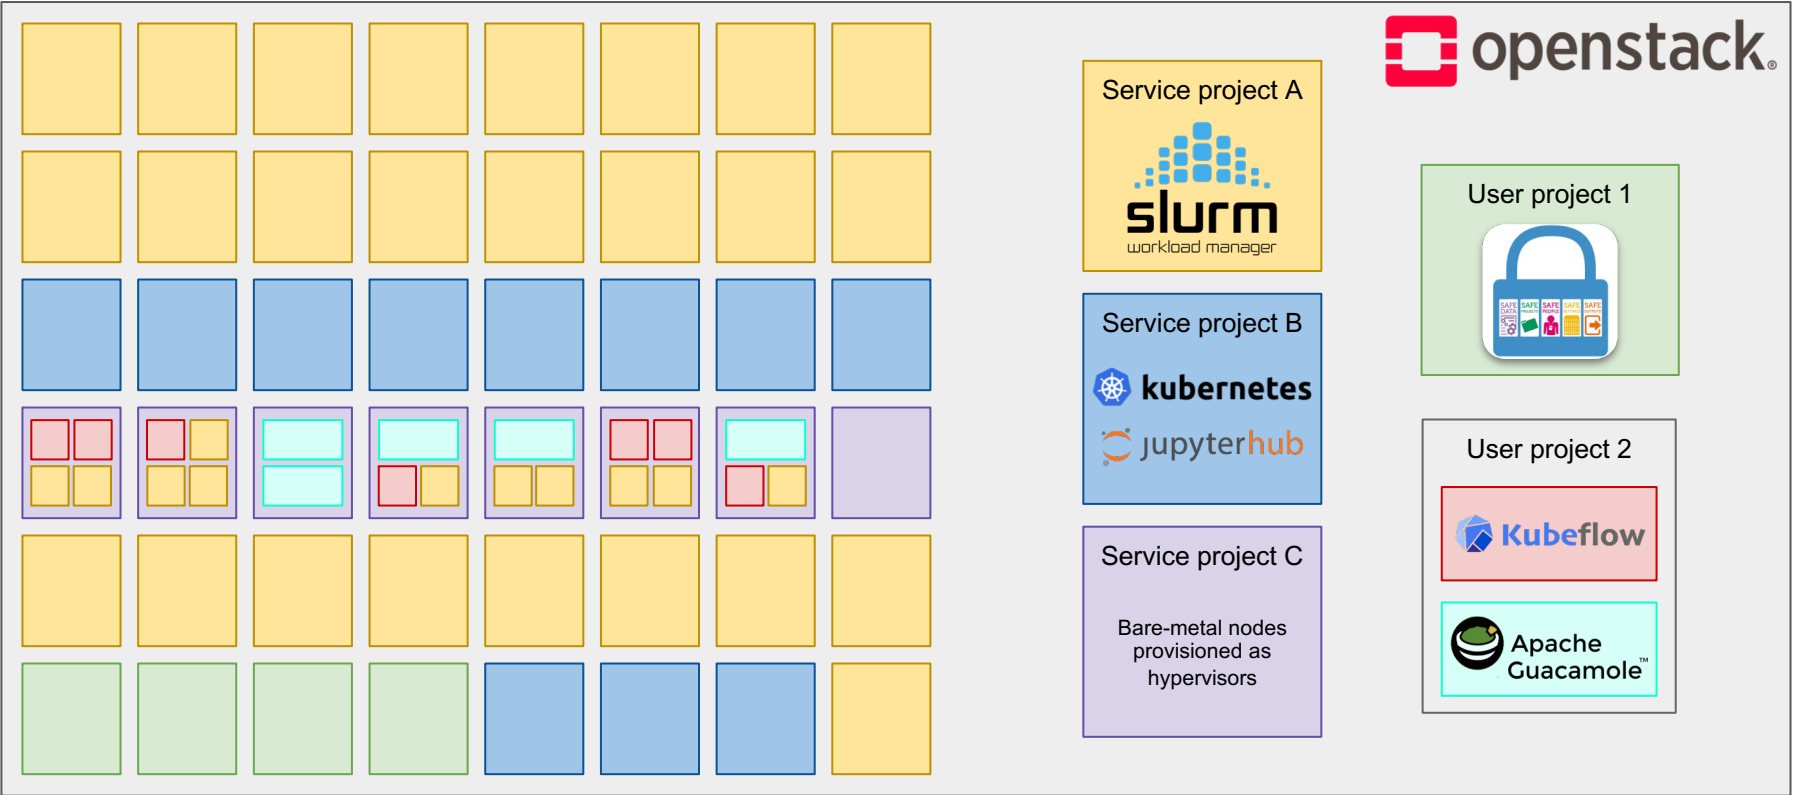
StackHPC



# Cloud-native HPC site



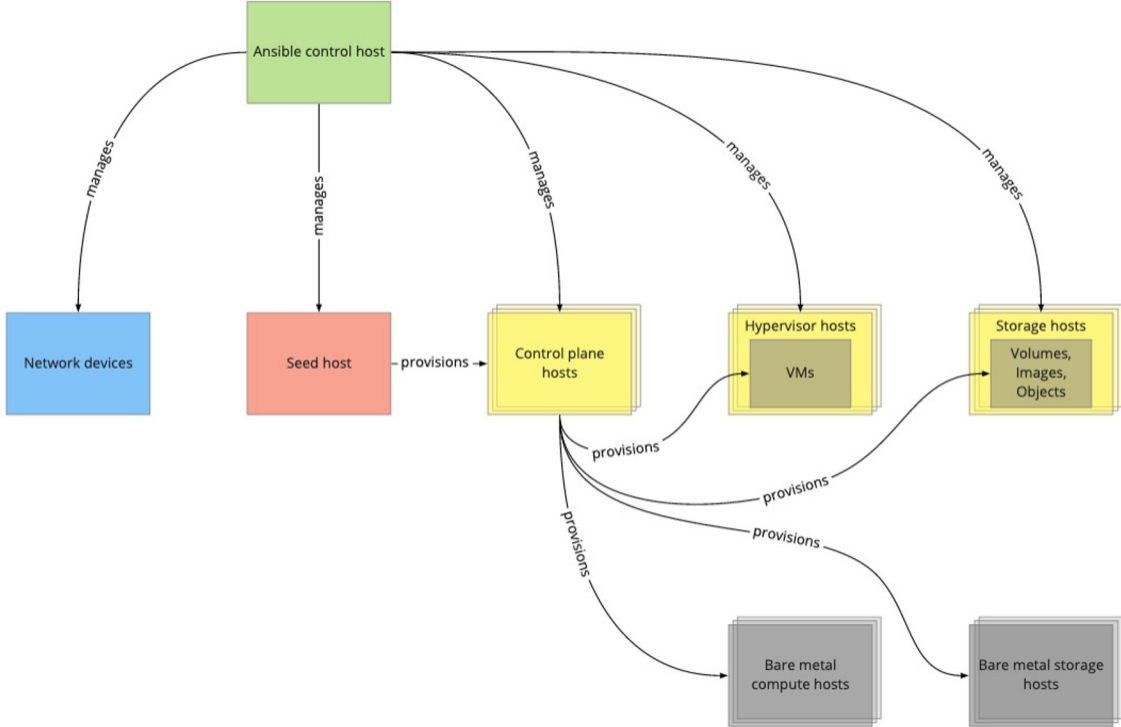
StackHPC



# OpenStack Kayobe Infrastructure as Code



StackHPC






# Performance



Reconfigurable and  
isolated infrastructure



**Performance to extract  
maximum value**



Azimuth self-service  
platforms

# Bare metal vs VMs

- Ironic bare metal adds zero overhead
  - Direct access to hardware
  - Optional dynamic VLAN and cleaning
  - But nodes are big
- Virtual machines
  - Deploy faster
  - GPU passthrough, SR-IOV and pinned CPUs
  - No access to firmware patching



**IRONIC**

*an OpenStack Community Project*

# High-speed networking

- Utilise hardware offloads
- Single Root I/O Virtualisation (SR-IOV)
  - VMs can utilise hardware offloads
- Remote Direct Memory Access (RDMA)
  - Low latency, high bandwidth
  - Requires application support
  - MPI and storage transport
  - Typically RoCE (RDMA over Ethernet)
  - Infiniband supported by StackHPC



StackHPC

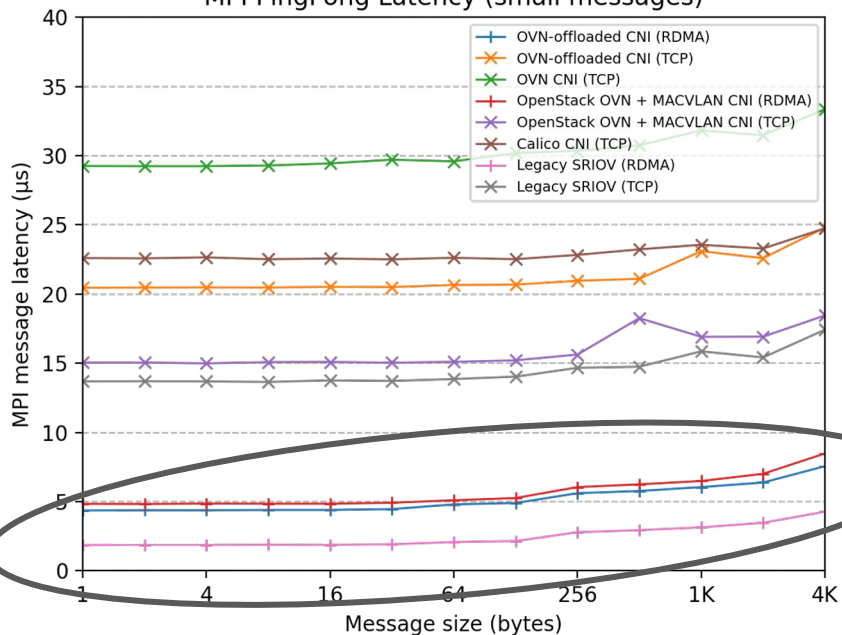


# MPI PingPong

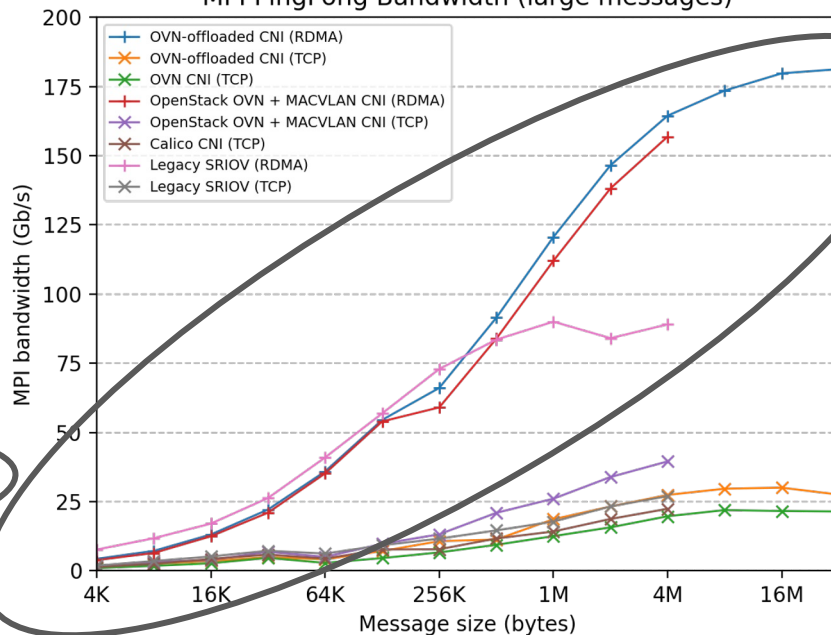


StackHPC

### MPI PingPong Latency (small messages)



### MPI PingPong Bandwidth (large messages)





<https://github.com/stackhpc/kube-perftest>

# OpenFOAM via kube-perftest



StackHPC



<https://github.com/stackhpc/kube-perftest>

# Self-service platforms



Reconfigurable and isolated infrastructure




Performance to extract maximum value



**Azimuth self-service platforms**

# Azimuth

- Web portal for self-service platforms
- Configurable catalogue of curated platforms
  - StackHPC reference platforms
  - Site-optimised platforms
  - Automation using standard tools
- Platform services exposed using Zenith
  - Tunneling application proxy
  - No public IP required
  - SSO and TLS
- Manage platform users with Keycloak

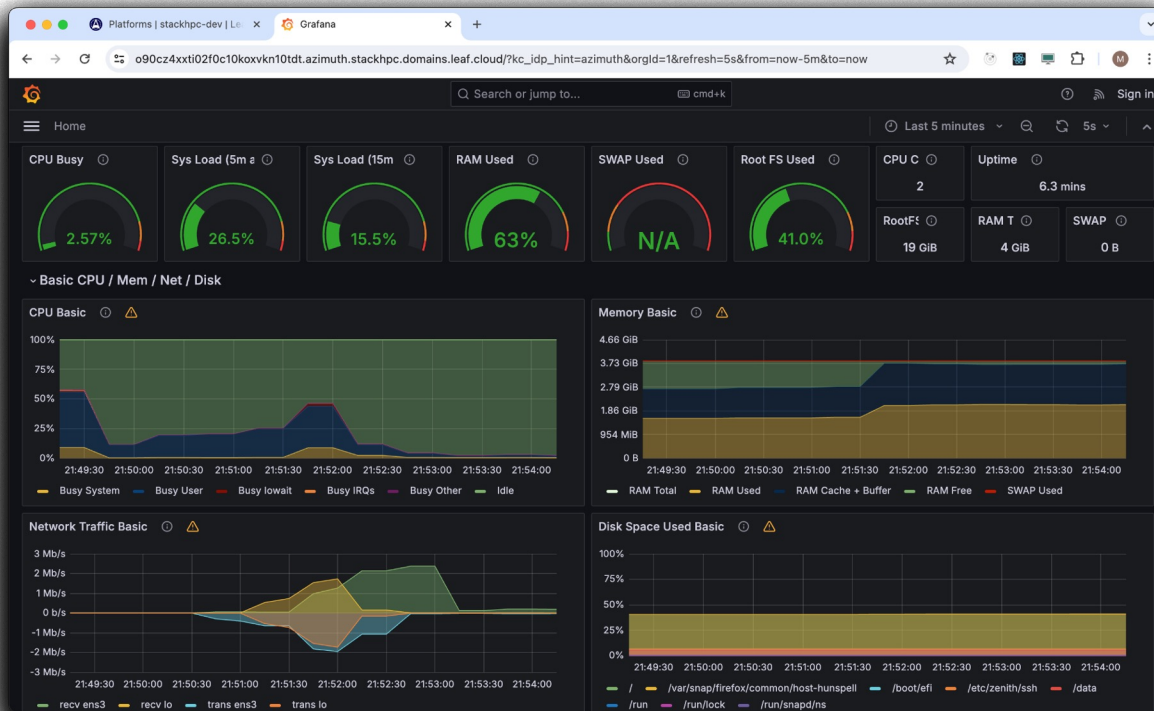


# Workstation



StackHPC

- Web-based shell and desktop
- Secure access via Zenith
- Monitoring stack
- Platform lifetime
- User gets sudo
- Apptainer and podman
- Optional SSH with public IP
- Access to project share



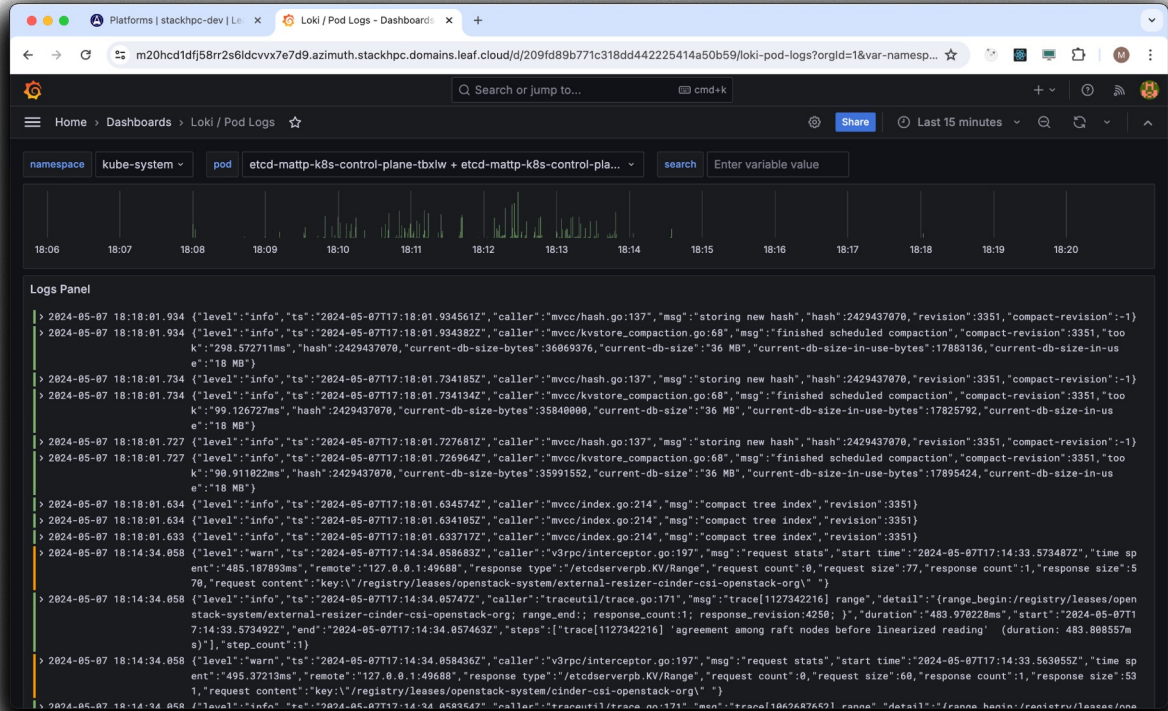


# Kubernetes



StackHPC

- Built on Cluster API
- HA control plane
- Multiple node groups
- Download kubeconfig
- Autoscaling, autohealing
- Rolling upgrades
- NVIDIA GPU + NIC support
- Kubernetes dashboard
- Monitoring and logging
- Secure access via Zenith

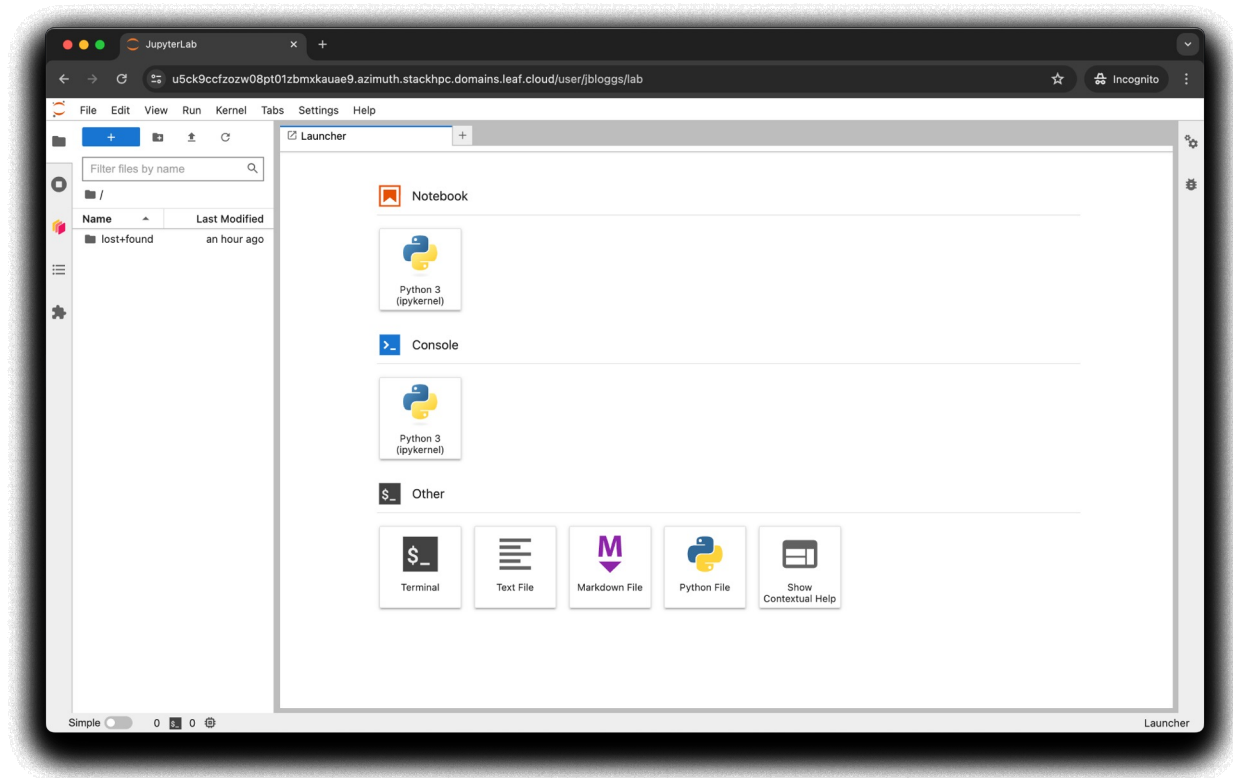


# DaskHub



StackHPC

- Runs on Kubernetes cluster
- Each user gets their own notebook server
- Secure access via Zenith
- Grant access to external users using tenancy Keycloak realm
- Dask clusters for parallel computing using Dask Gateway

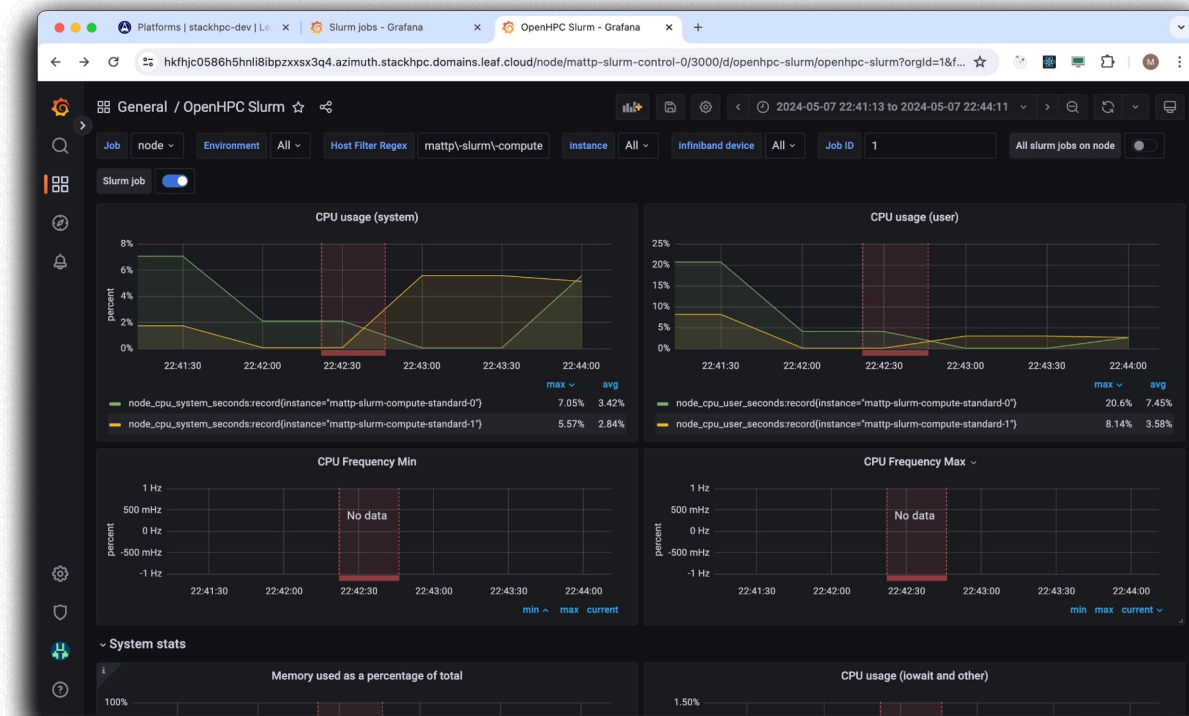


# On-demand Slurm



StackHPC

- Single-user Slurm cluster
- No waiting for queues
- Image-based updates
- OpenHPC, Apptainer, EESSI
- Open OnDemand UI
- Job aware monitoring
- Access to project share



How to get started?

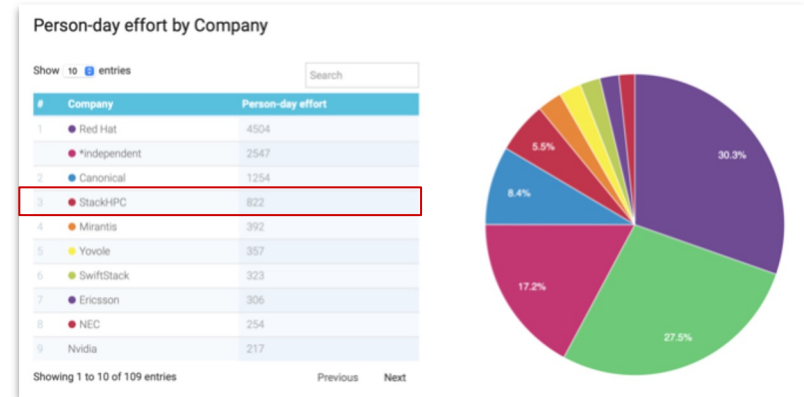


StackHPC

# Open Source



- Entire stack is open-source
- OpenStack governed by OpenInfra Foundation
  - StackHPC third-largest contributor
- Azimuth licenced under Apache 2.0
  - Putting together CNCF Sandbox application
- Collaborative development process





# Workshops and design



StackHPC

- Workshops available
  - OpenStack
  - Platforms and Azimuth
- Requirements capture
- Architecture design
- Implementation
- Training and knowledge transfer
- Support and maintenance





# StackHPC

The Rise of the HPC Cloud

# Thank You

Come and visit us at  
Booth 12!

<https://www.stackhpc.com>