

Federating OpenStack Powered Supercomputers

John Garbutt
@johnthetubaguy

StackHPC



Why Federate a Supercomputer?

IRIS: e-Infrastructure for STFC Science



STFC: UK Science and Technology Facilities Council

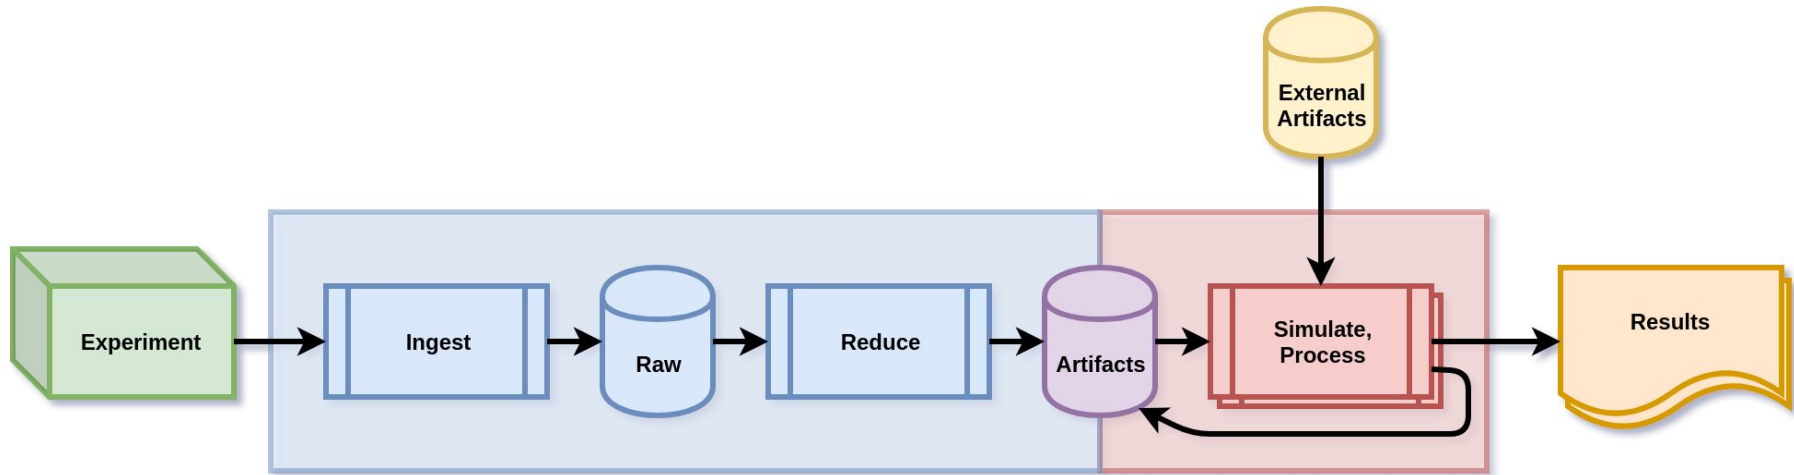
Understand requirements for a UK-wide e-Infrastructure for Science

Quicker, easier and more efficient access to infrastructure

Encourage projects to share:

- Infrastructure
- Expertise
- Software

Scientific Computing



Compute Requirements

What

- Scale
 - Part of a host
 - Many small jobs
 - Multiple hosts together
- Large memory, Shared Scratch, GPU
- Receive live data feed

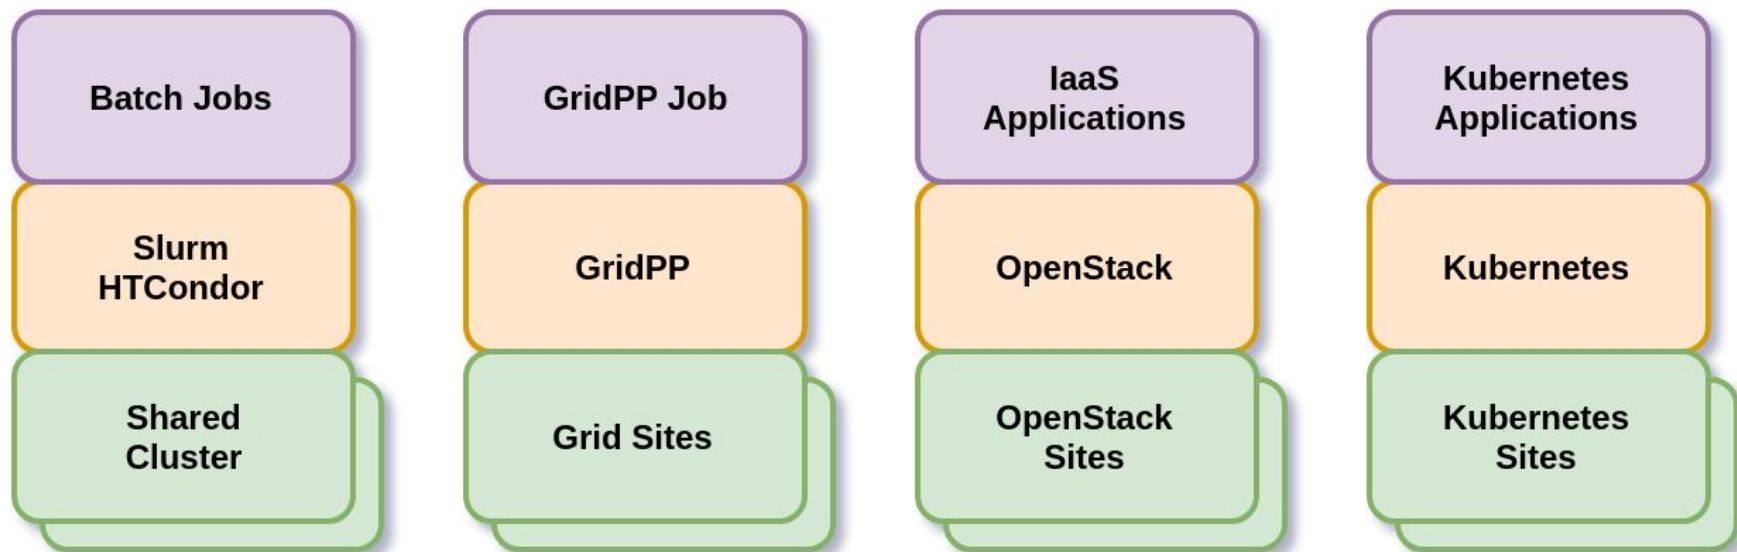
When

- Submit a Job
- Interactive
 - Scheduled
 - On-demand
- Web Service

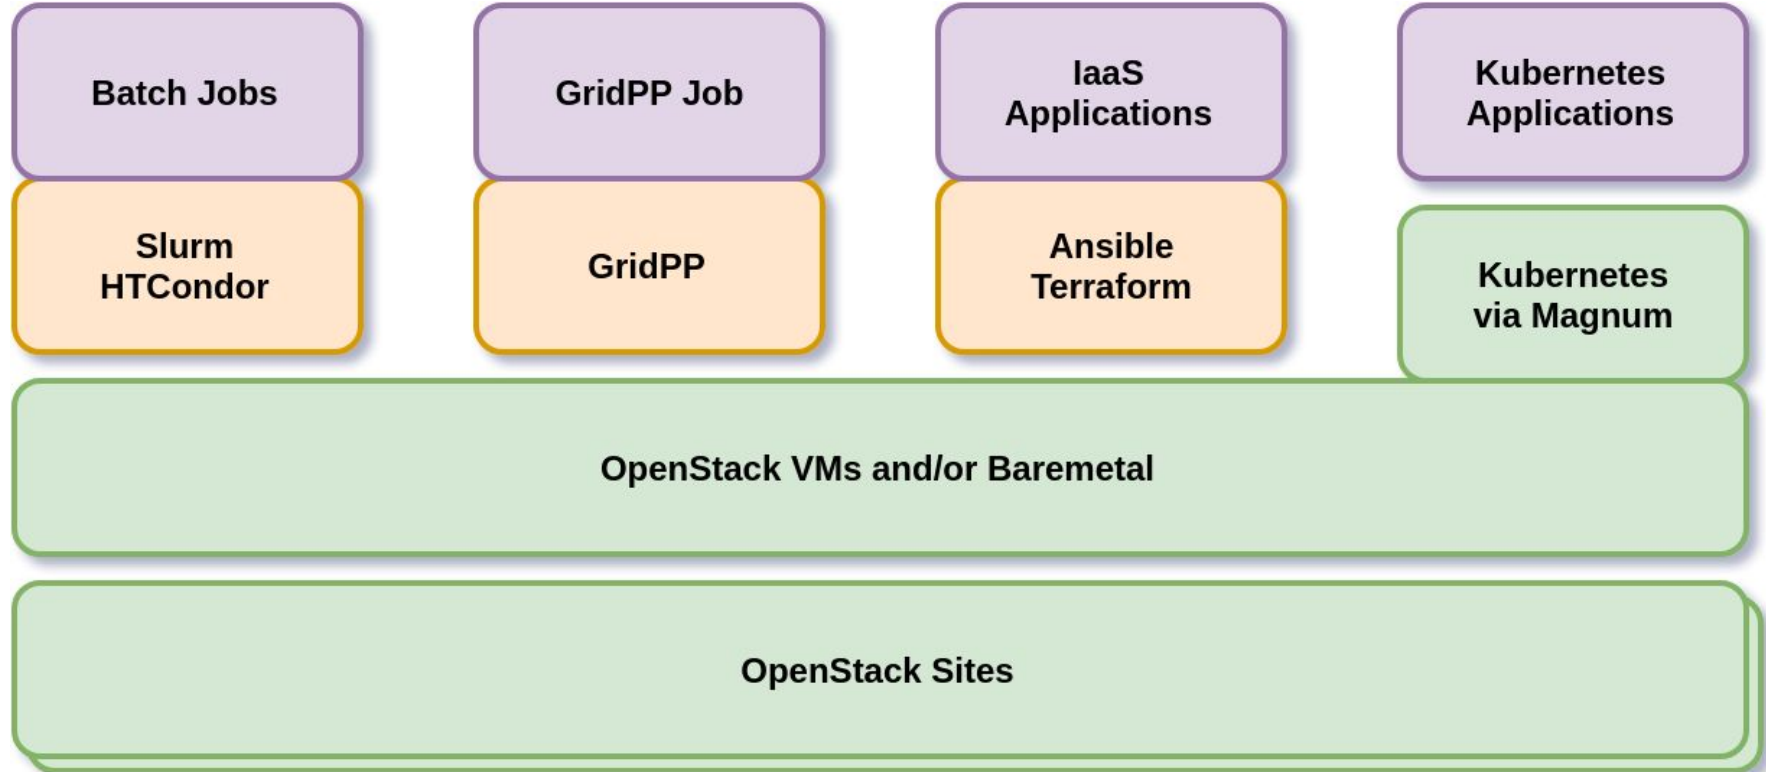
Resource Sharing Opportunity

- Facility
 - Large demand spikes for interactive processing
- Shared
 - Demand grows beyond availability
 - Inflexible
- Dedicated
 - Hard to predict required size


Today: Siloed Infrastructure Sites



Remove Silos




OpenStack Powered Supercomputer



What does “**Federated**” mean?

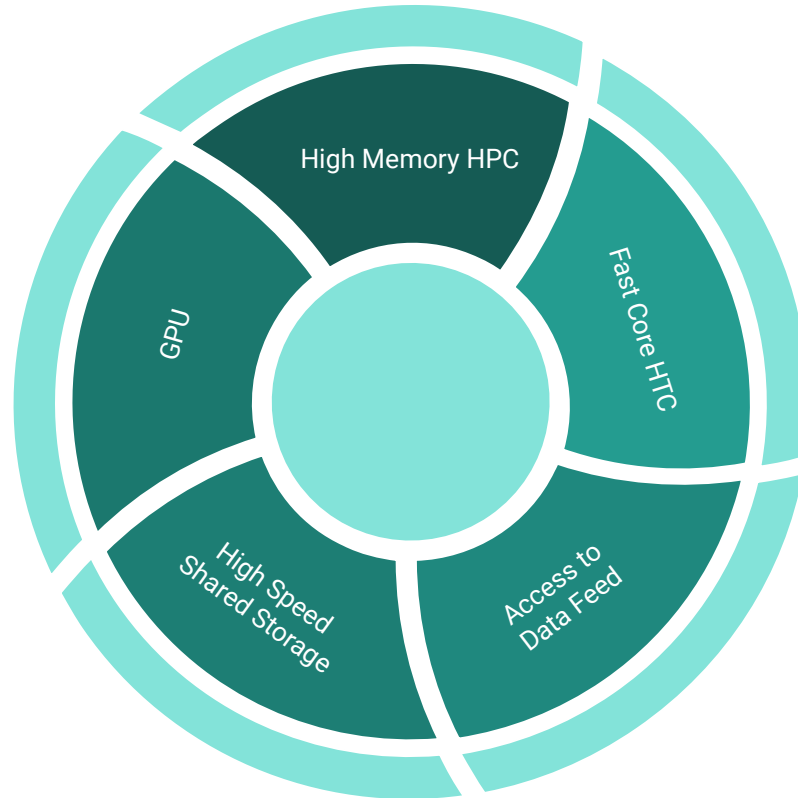
Federated OpenStack Powered Supercomputer



IRIS Compliance Tests

- Built on OpenStack Interoperability Tests
- Add extra optional Manila and Magnum tests
- Make Cinder optional

Best Fit Resources



Location Transparency

Workflow

Describe required processing steps

Platform

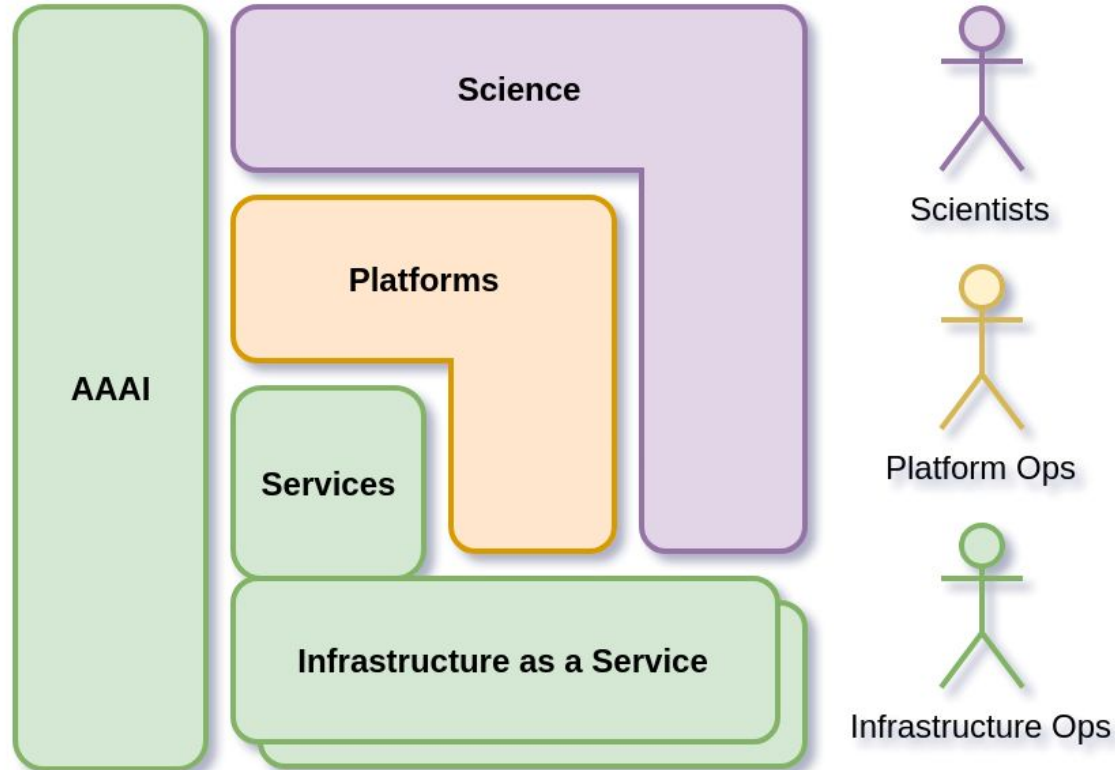
Choses Region, Optimises Data Flow, Orchestrates workflows

OpenStack

Infrastructure split between Regions

What is “AAAI”?

Federated OpenStack Powered Supercomputer



Authentication and Authorization



Non-Interactive Authentication

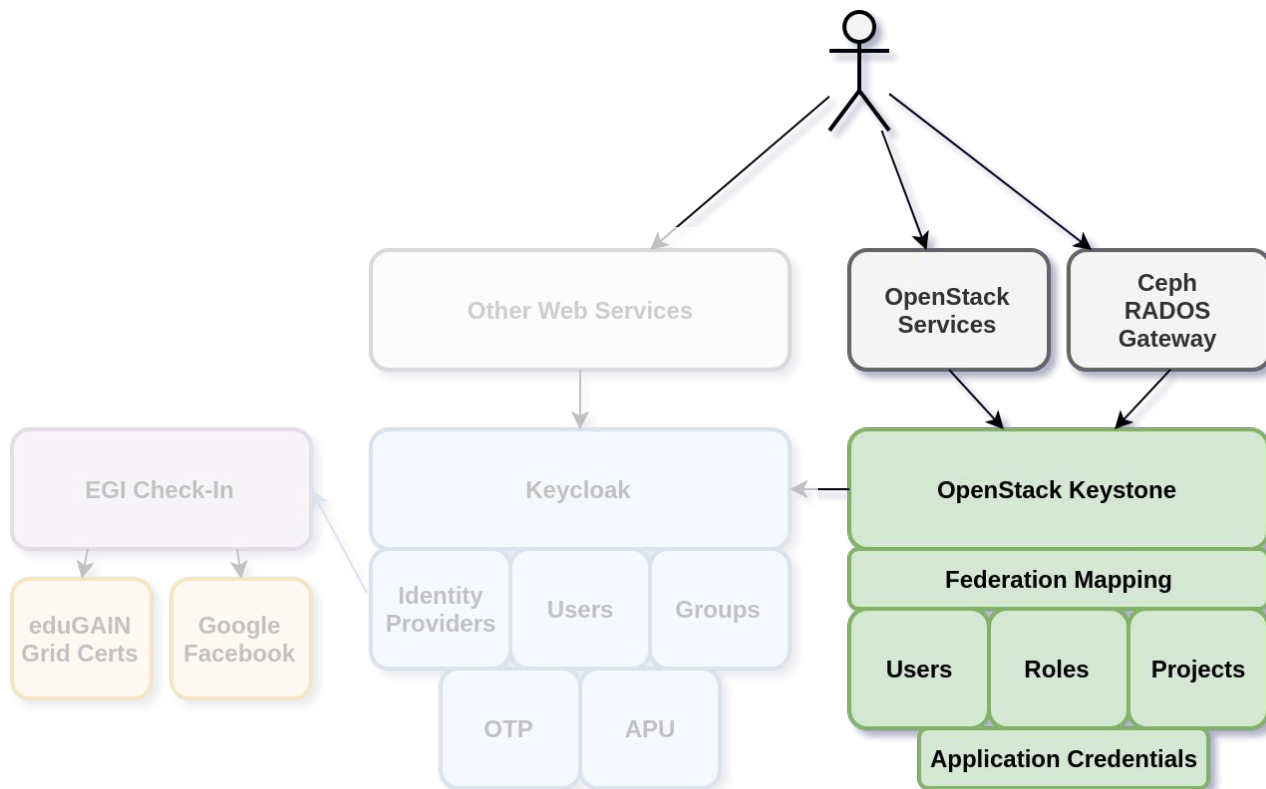


Building Blocks of Federated Identity

- Authenticate via OIDC
 - Keycloak OIDC to EGI Check-in
 - Indigo IAM
- Authorisation via Federation Mapping
 - Concrete users and roles
 - Avoid Groups
- Application Credentials
 - Non-interactive authentication for Keystone

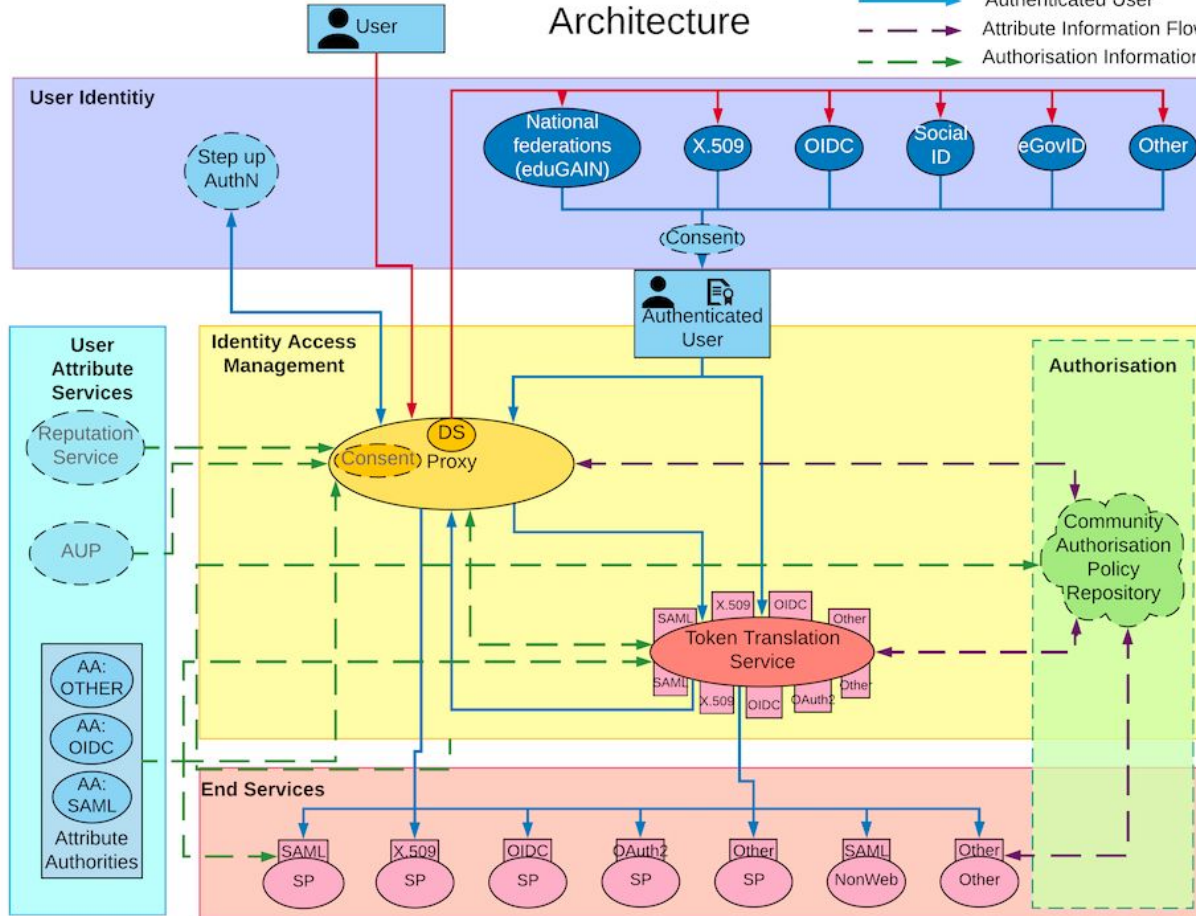


Keystone Federation



AARC Blueprint Architecture

- Unauthenticated User
- Authenticated User
- - - Attribute Information Flow
- - - Authorisation Information Flow




Accounting

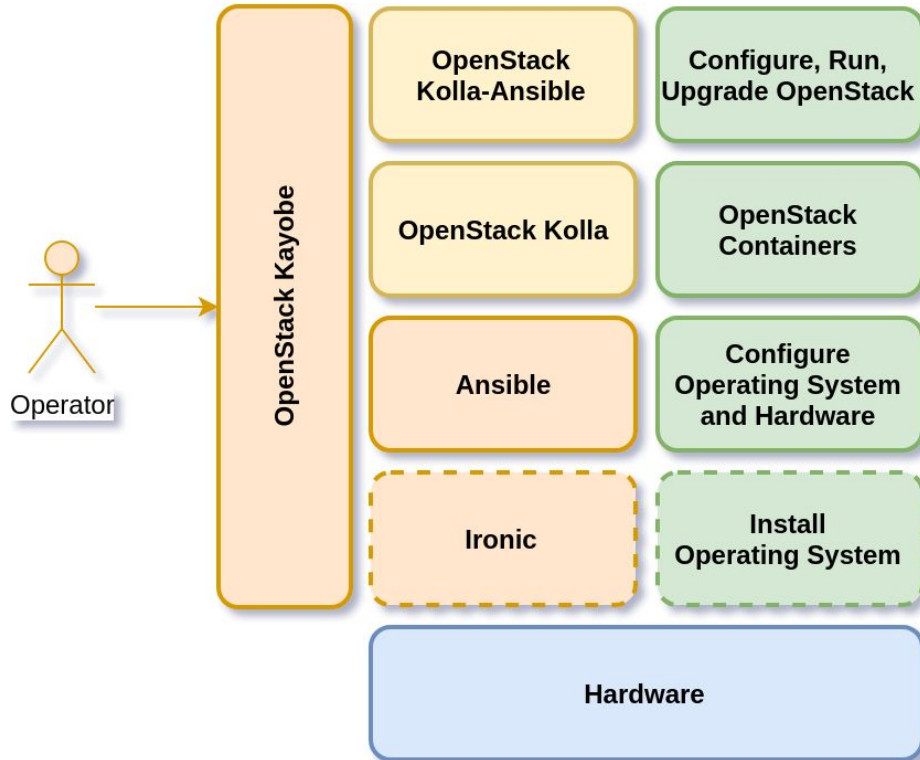
- Focus on Traceability
- Usage: cASO sends to Fluentd (and APEL)
- Quota: limit maximum concurrent usage
- Allocation: allowed usage over given duration

How to pass IRIS compliance tests?

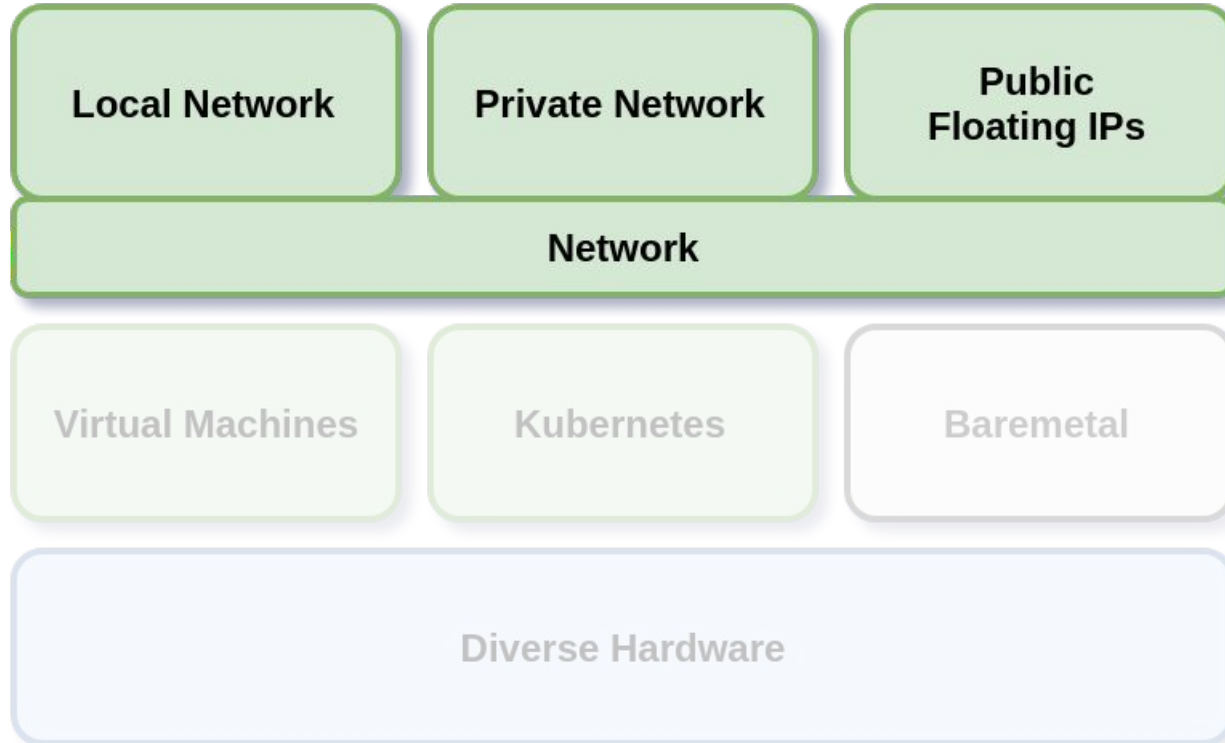
Shared Operational Tooling



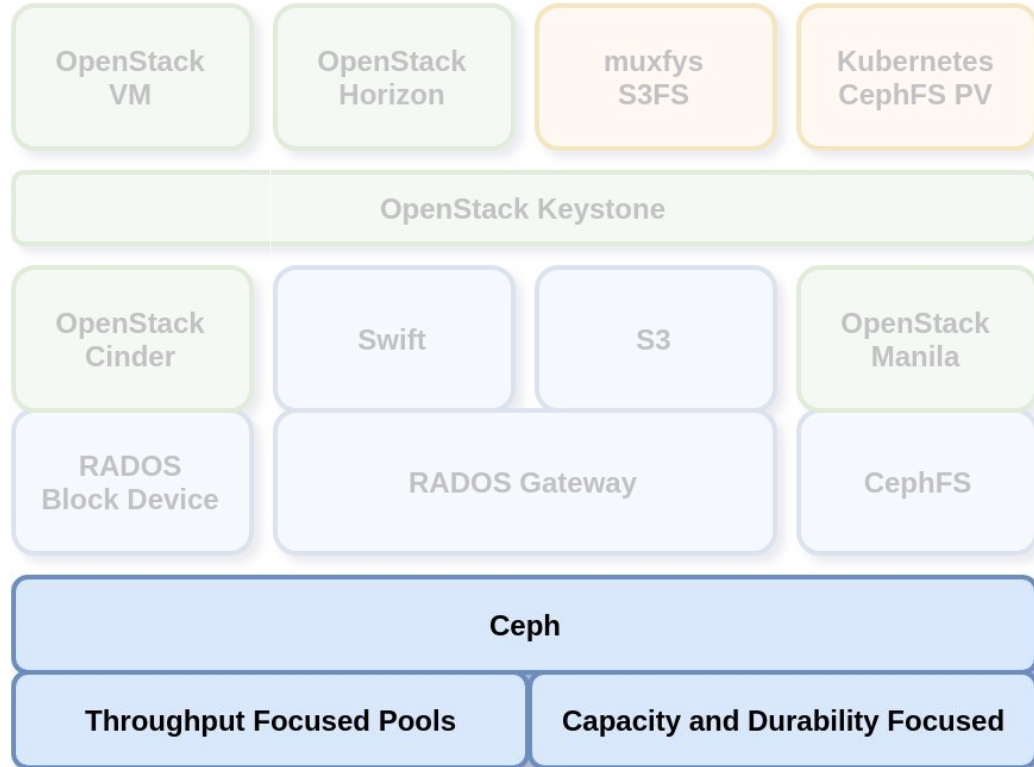
OpenStack Deployment



Scientific OpenStack Compute



Scientific OpenStack Storage



Any unsolved problems?

Authorization of Federated Identity

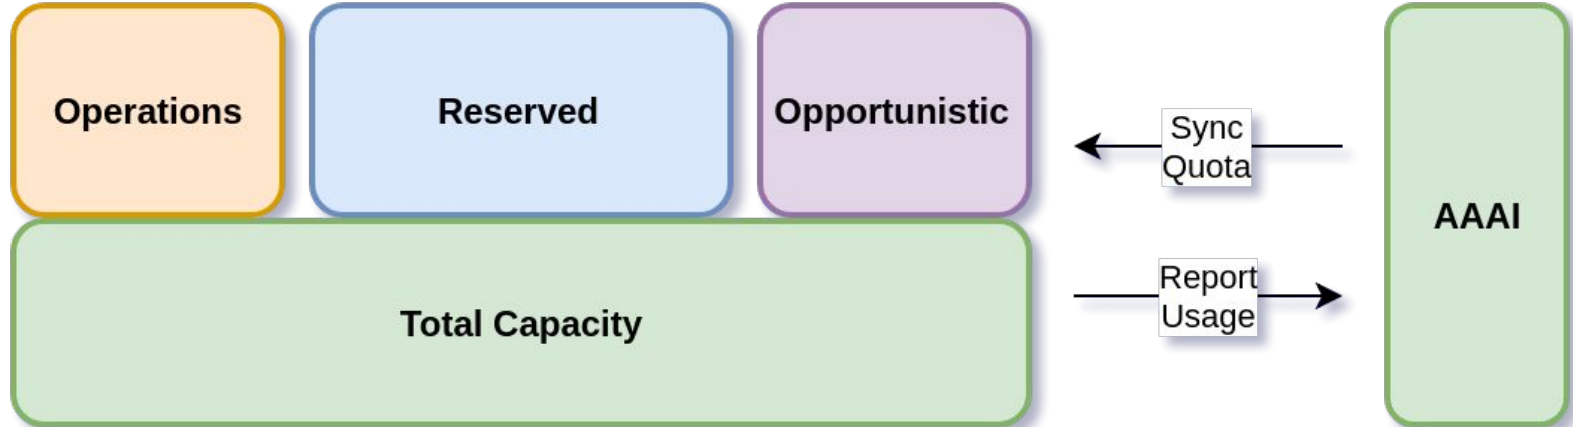


FIM4R

“Every researcher is entitled to focus on their work and not be impeded by needless obstacles nor required to understand anything about the FIM infrastructure enabling their access to research services.”

FIM4R version 2: <https://fim4r.org>

Improve Resource Sharing



Lessons Learned?

- (1) Building a Community Matters
- (2) Federated OpenStack works
- (3) Application Credentials can work

StackHPC

@johnthetubaguy
johng@stackhpc.com