Ad-hoc Filesystems

For Dynamic Science Workloads

Stig Telfer, CTO
John Garbutt, Principal Engineer

StackHPC

StackHPC Overview

StackHPC

Founded in 2016, based in Bristol, UK

- Convergence of HPC and Cloud
- Build, Maintain and Scale both Infrastructure and Applications
- Assist with transition to more efficient and sustainable approaches
- Sustained via collaboration upstream Open Source Communities

openstack. ceph openHPC ANSIBLE

Agenda

StackHPC

EUCLID

- Assumes a good Parallel File System
- Many IRIS OpenStack sites not providing (a good enough) one

Cambridge Data Accelerator

- #1 in io500 June 2019
- Initial SKA related work

Part 1:
EUCLID with Ceph inside Slurm

StackHPC

The Euclid Project

- **An ESA-led mission to study the nature of the dark universe.**
- **International body of more than 1,000 members who will oversee development of the instruments, manage science operations and analyse data.**
- **A 1.2m space telescope working in optical and infra-red wavelengths**
- **Mapping hundreds of millions of galaxies across the universe, including faint objects down to magnitude 24.5**
- **Launch scheduled for 2021**

https://www.euclid-ec.org/

Weak Gravitational Lensing

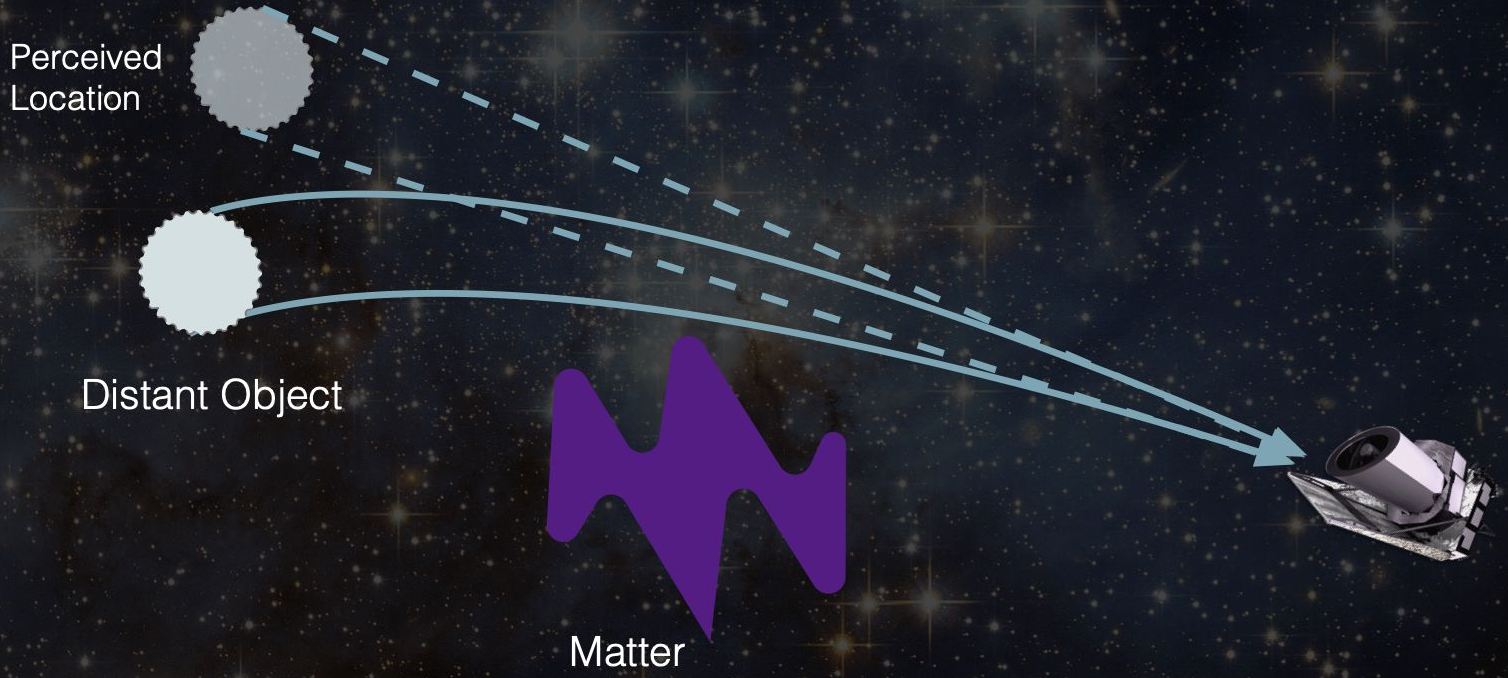

EUCLID's SHE Software Pipeline

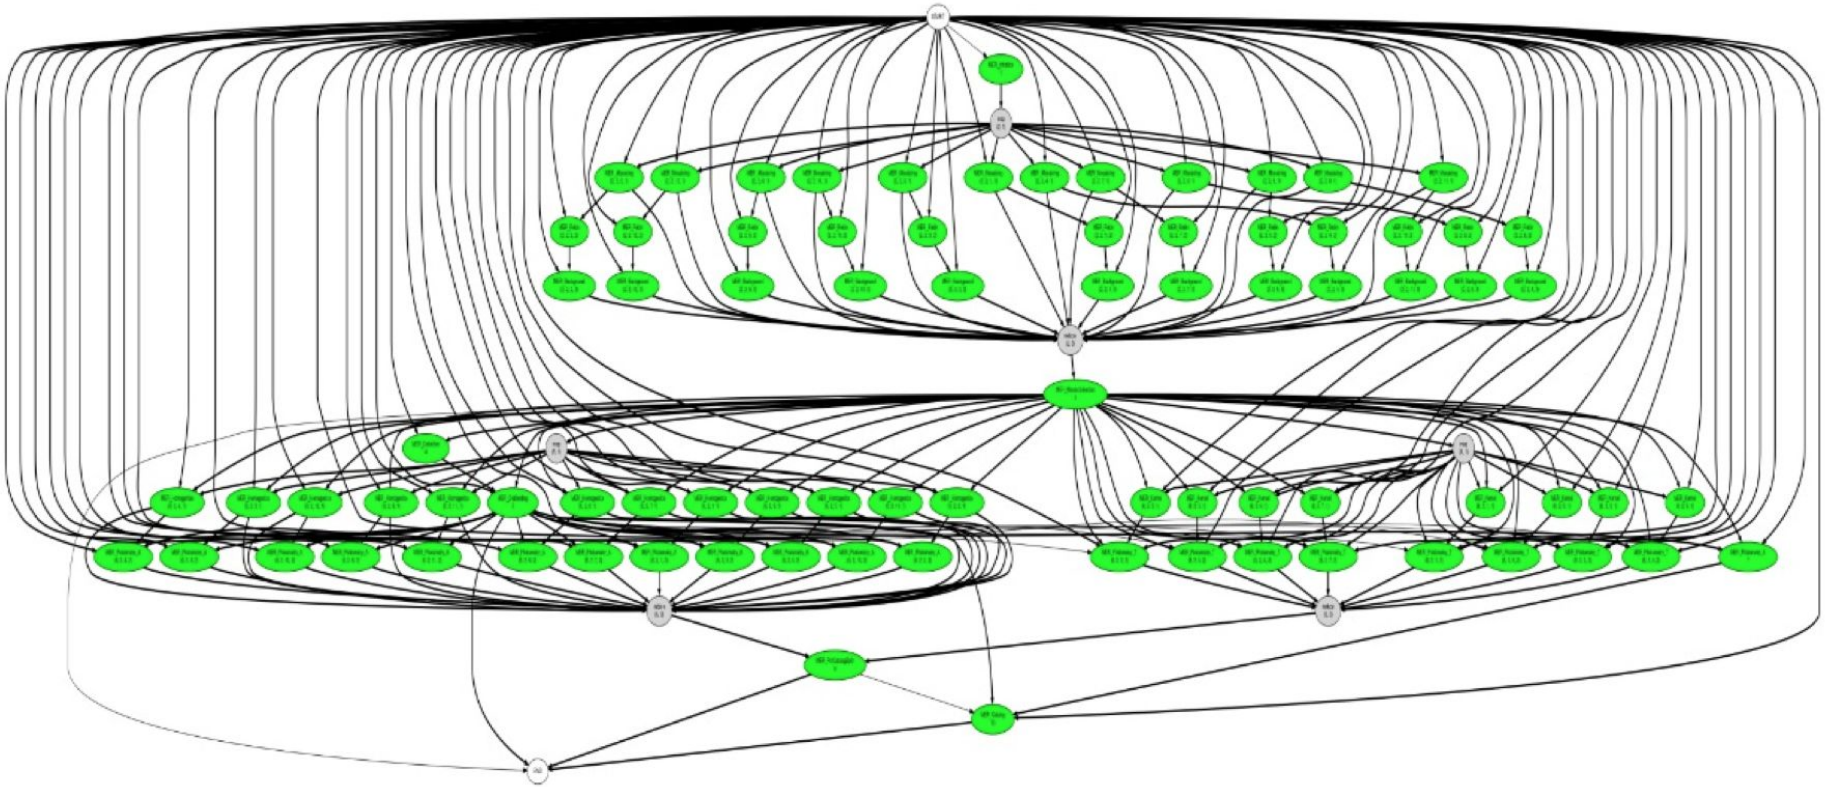

EUCLID: Federated Compute Platform

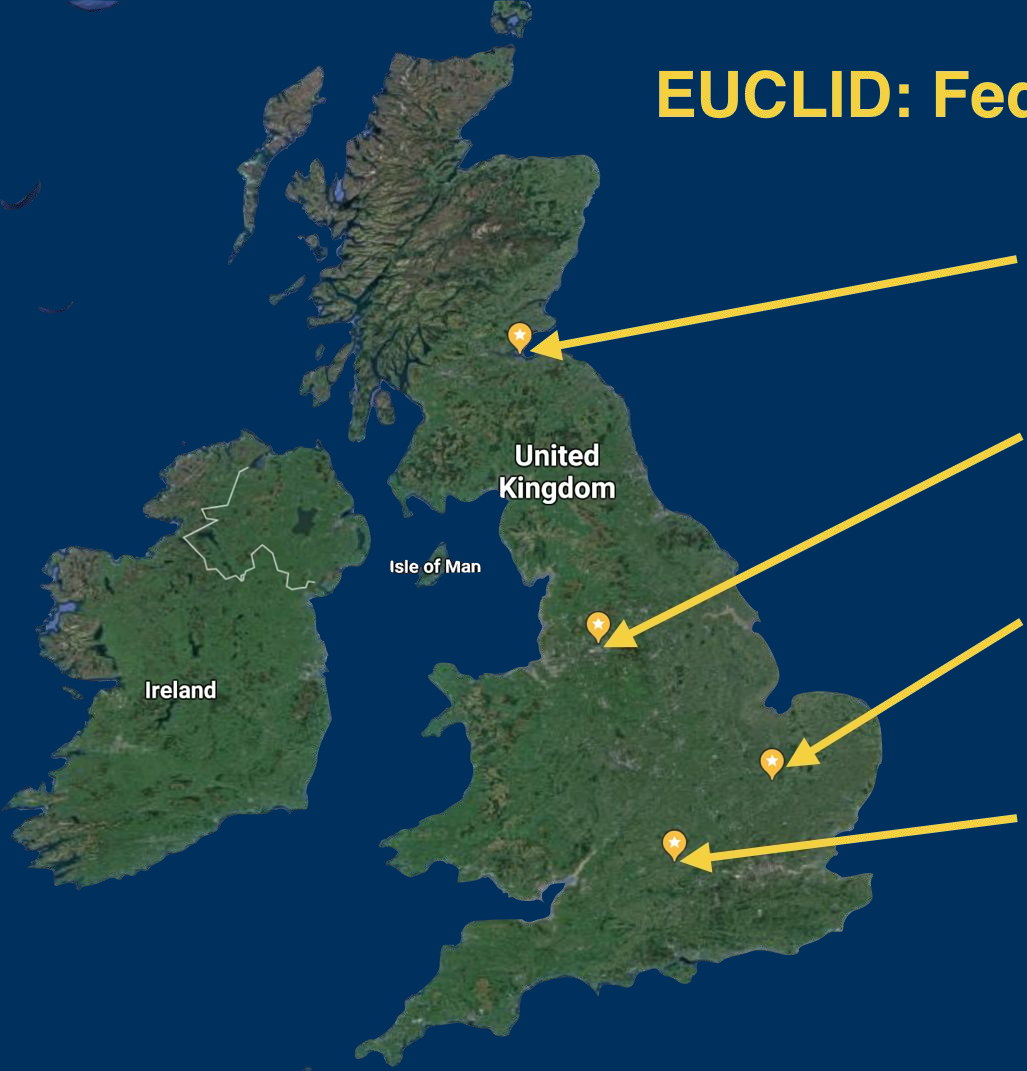

Royal Observatory, Edinburgh

University of Manchester

University of Cambridge

Rutherford Appleton Laboratory

EUCLID: Federated Compute Platform

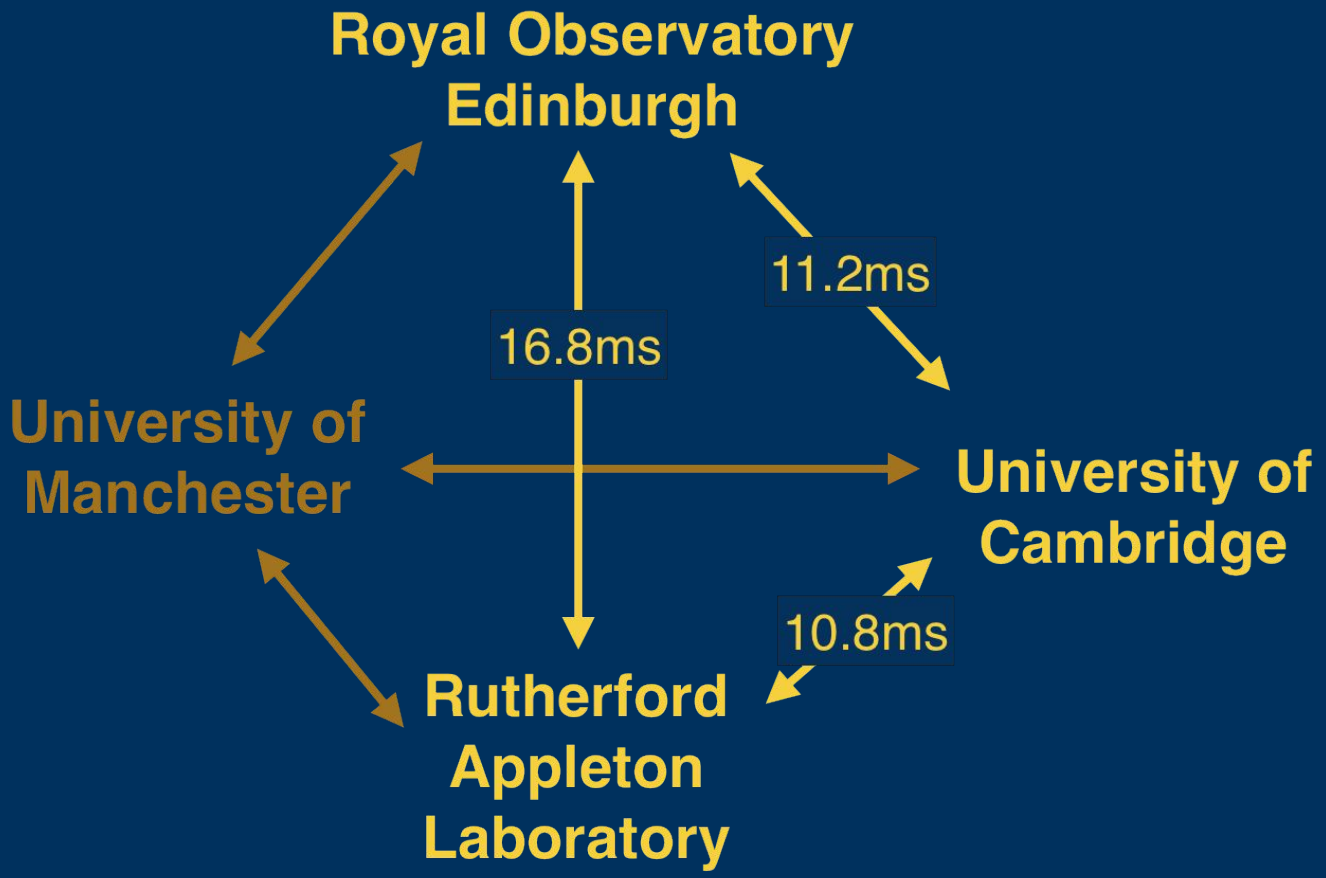

IRIS Compute Federation

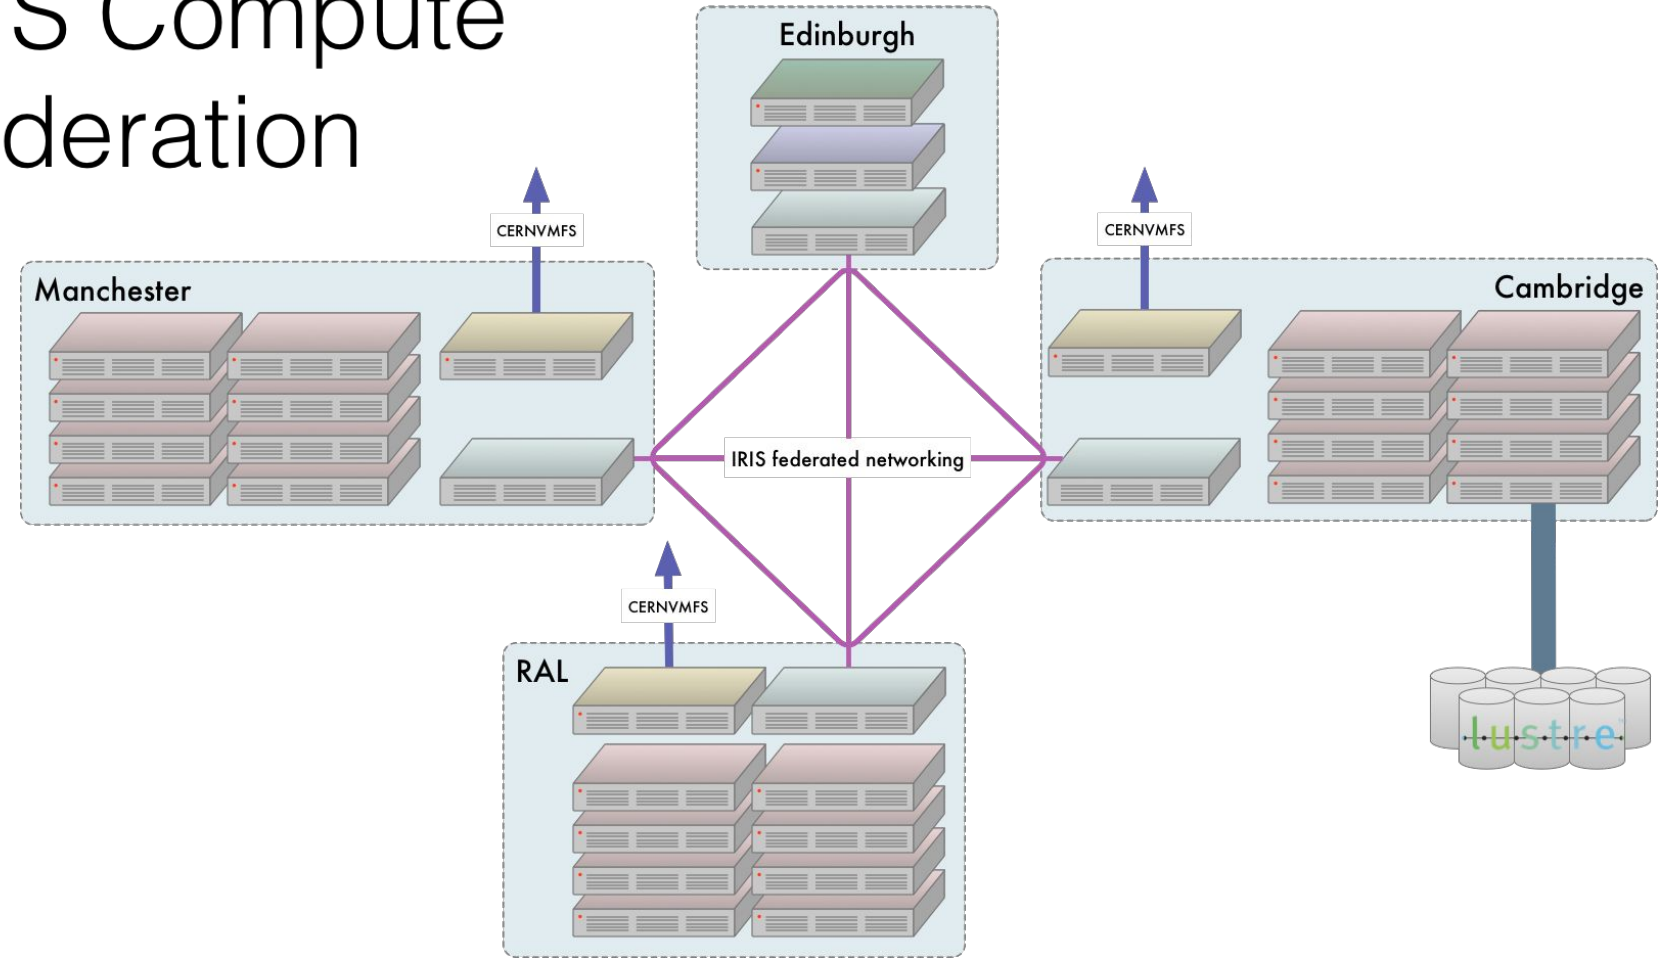

The Illusion of Proximity

- A functional solution using Ansible playbooks
- Create OpenVPN mesh on gateway nodes
- Define routes between subnets
- Open firewall rules to other sites

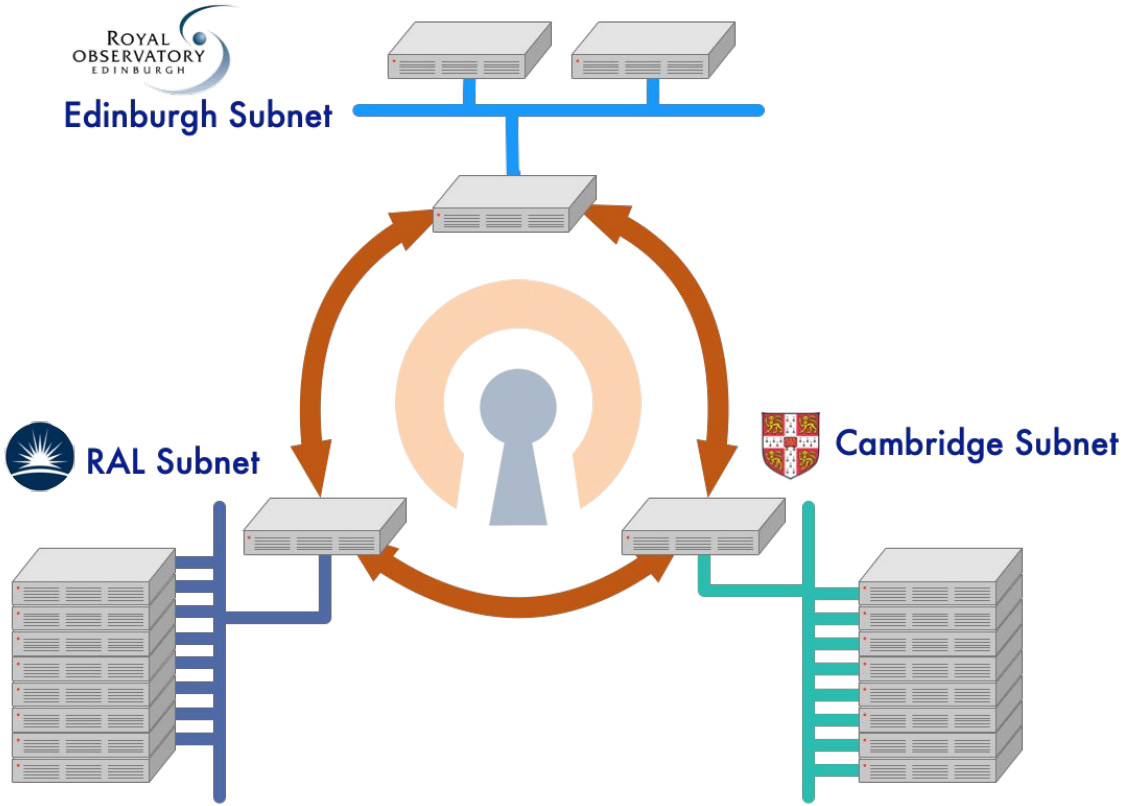

Multi-Site Ceph?

- Compute nodes provide converged cluster storage
- Each site runs monitors and metadata servers
- Storage pool across entire deployment - 2-way replication for global data
- Storage pool per site - no replication for local “hot buffer”

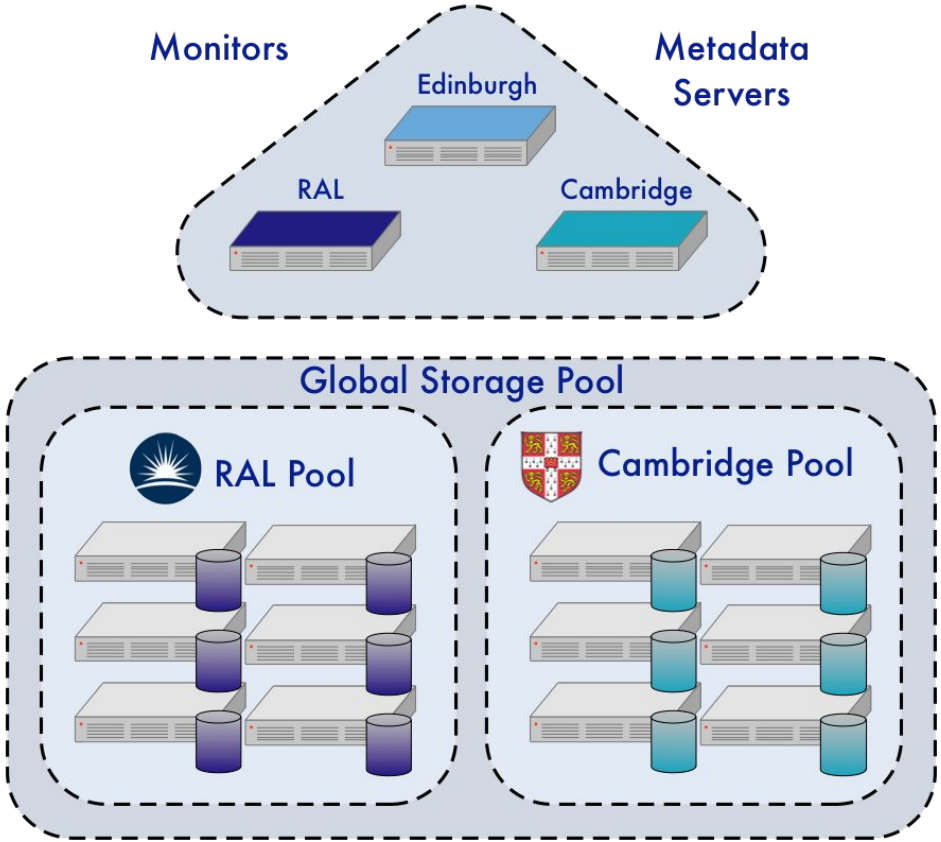

CephFS: A Federated Filesystem?

- Metadata pool across entire deployment
- Data pool across entire deployment
- Data pool per site
- Custom CRUSH rules for each data pool per site
- POSIX extended attributes set the back-end storage pool per file or directory

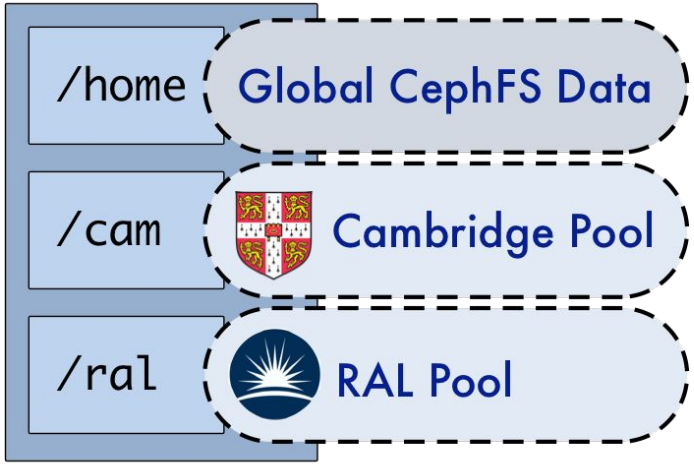

ceph

Federated Filesystem Experiences

StackHPC

- Single file namespace across the software platform provides user convenience.
- Has better data movement properties than object store.
- Pooled storage resource and IOPS of all nodes - scales with the deployment.
- SW-defined OpenVPN links can be slow and painful bottlenecks.
- Network latency can be erratic due to overloaded VPN tunnels.

It Works … kind of …

StackHPC

- Over 2.5 million jobs run through Slurm (and counting).
- 36 compute nodes at Cambridge, 90 compute nodes at RAL
- OOM-killer keeps stepping in when workload and Ceph OSD both get resource-intensive.
- Software VPN tunnels work but are sometimes a bottleneck.
- Ceph metadata ops is another current bottleneck.

The Effect of I/O on Federated Networking

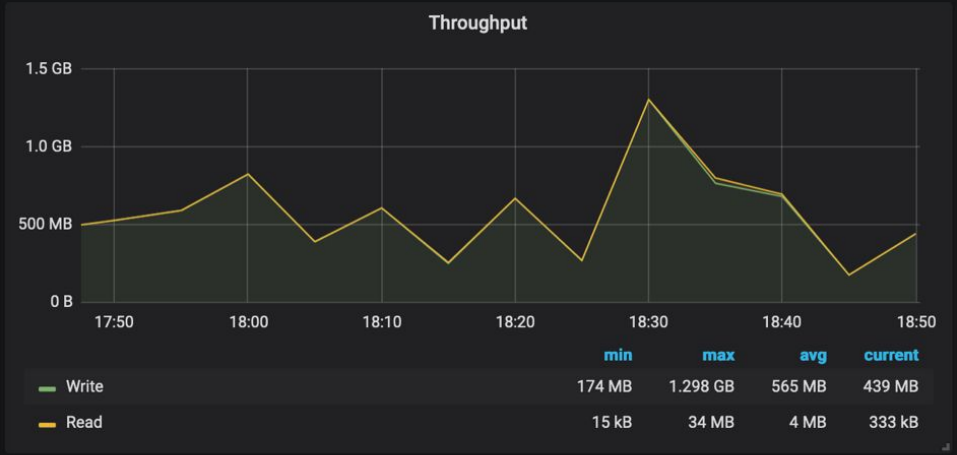

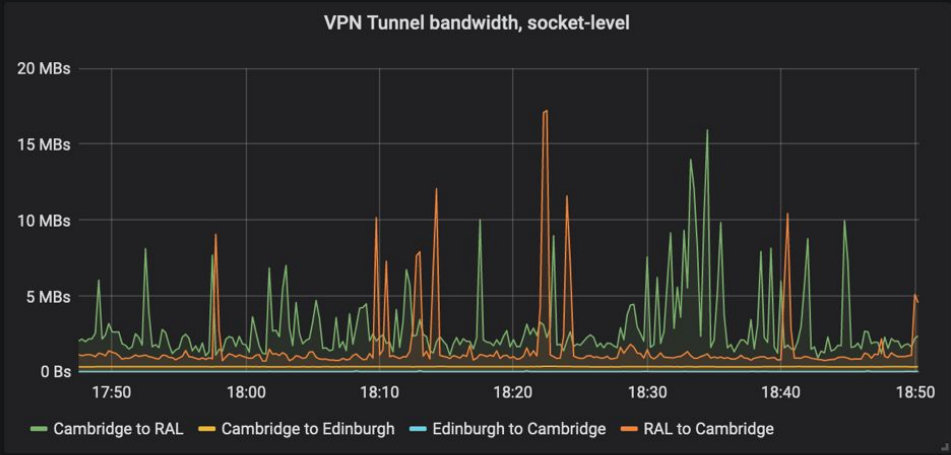

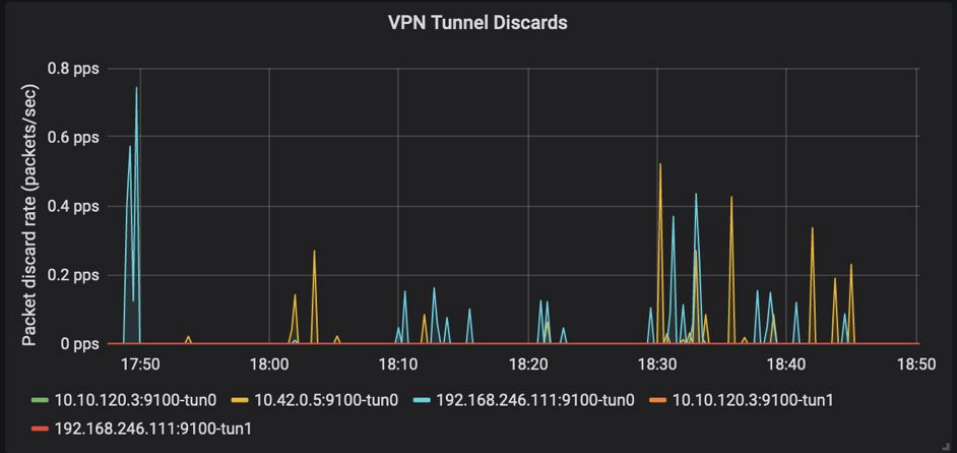

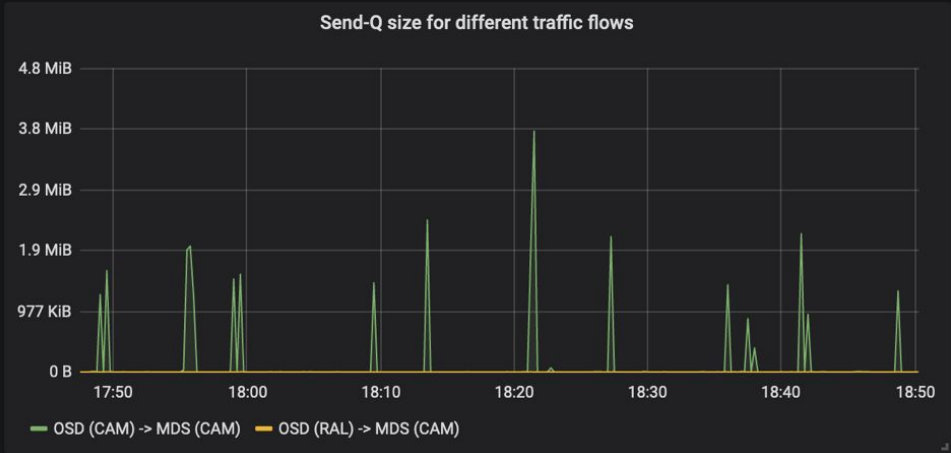

The Effect of Federation on Ceph I/O latency

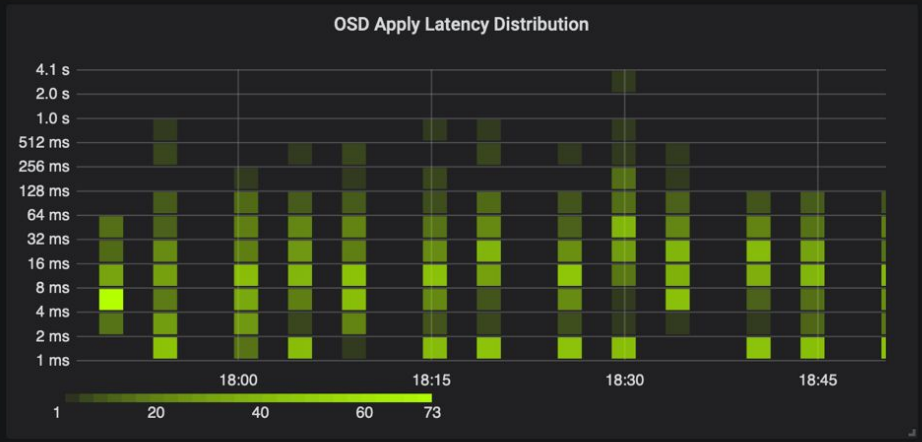

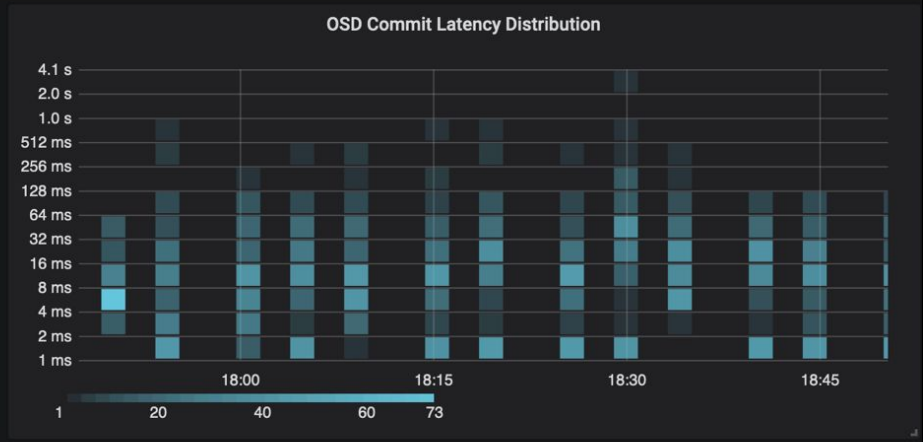

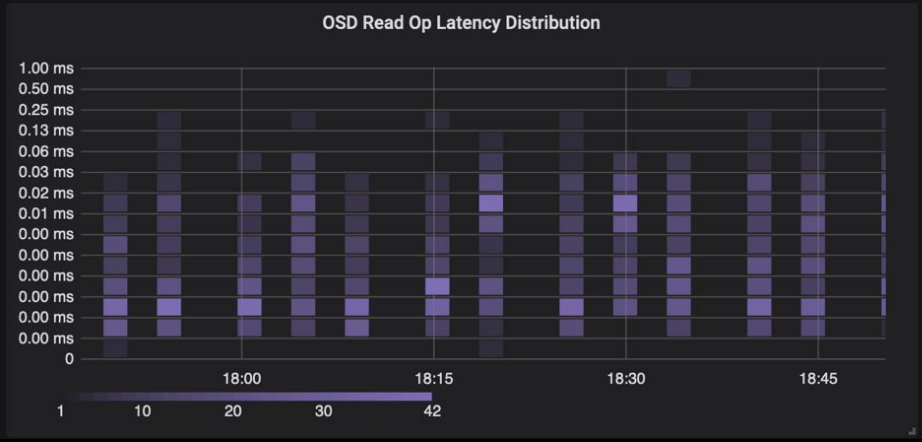

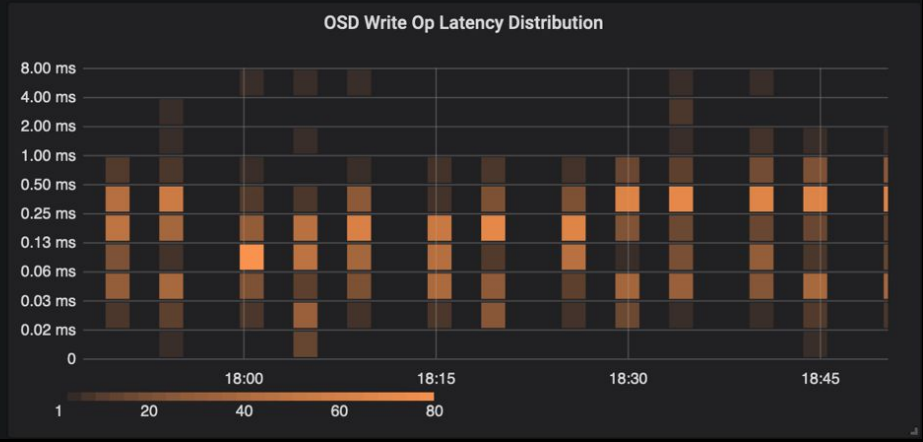

Summary

StackHPC

- A functional, cross-cloud interoperable design
- Provisions dynamically
- Scales with the deployment
- Provides the illusion of proximity between sites
- A performant solution would require infrastructure configuration changes or hardware support

Next Steps

StackHPC

- Speed up Ceph-Ansible for incremental addition/removal?
- Continue working on instrumentation for detecting back pressure
- Tighter memory management for OSDs
- TC traffic shaping on VPNs
- CephFS MDS subtree pinning (largely undocumented feature): https://ceph.com/community/new-luminous-cephfs-subtree-pinning/
- Hardware assisted federated networking à la Internet2 - via JISC?

Share and Enjoy

- Euclid's "Euclid-as-a-Service" playbook: https://github.com/astrodb/euclid-saas
- Ansible role for deploying software clusters with Heat: https://galaxy.ansible.com/stackhpc/cluster-infra
- Ansible role for OpenVPN federated mesh: https://galaxy.ansible.com/stackhpc/openvpn
- Ansible role for OpenHPC/Slurm deployment: https://galaxy.ansible.com/stackhpc/openhpc
- Prometheus exporter for TCP stream back-pressure: https://github.com/stackhpc/sockpuppet (WIP)

StackHPC

Part 2:
Cambridge Data Accelerator

StackHPC

CUMULUS
UK SCIENCE CLOUD

UNIVERSITY OF CAMBRIDGE
Research Computing Services

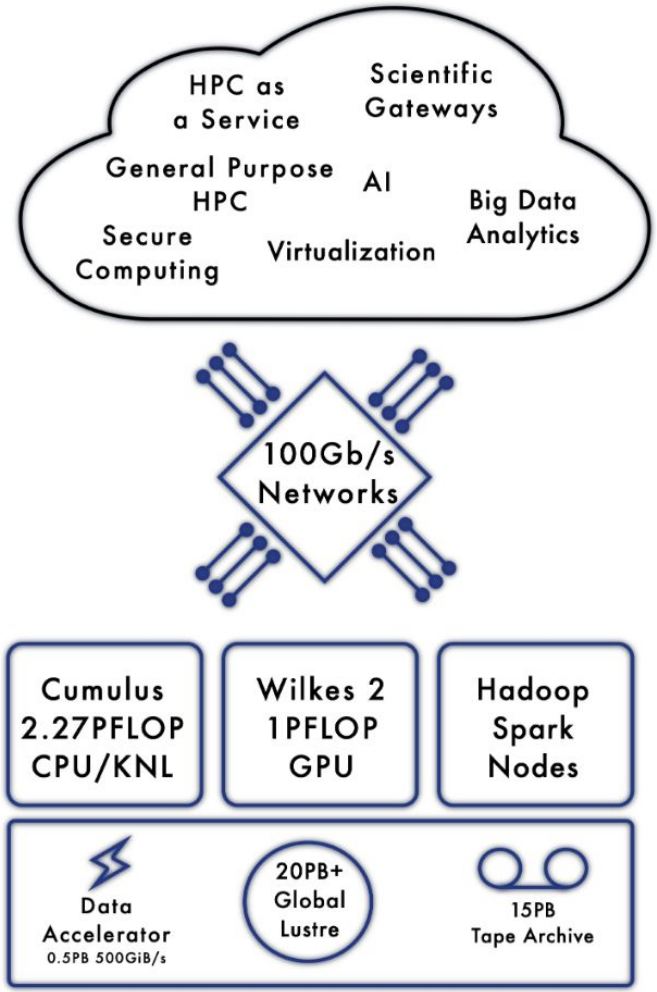

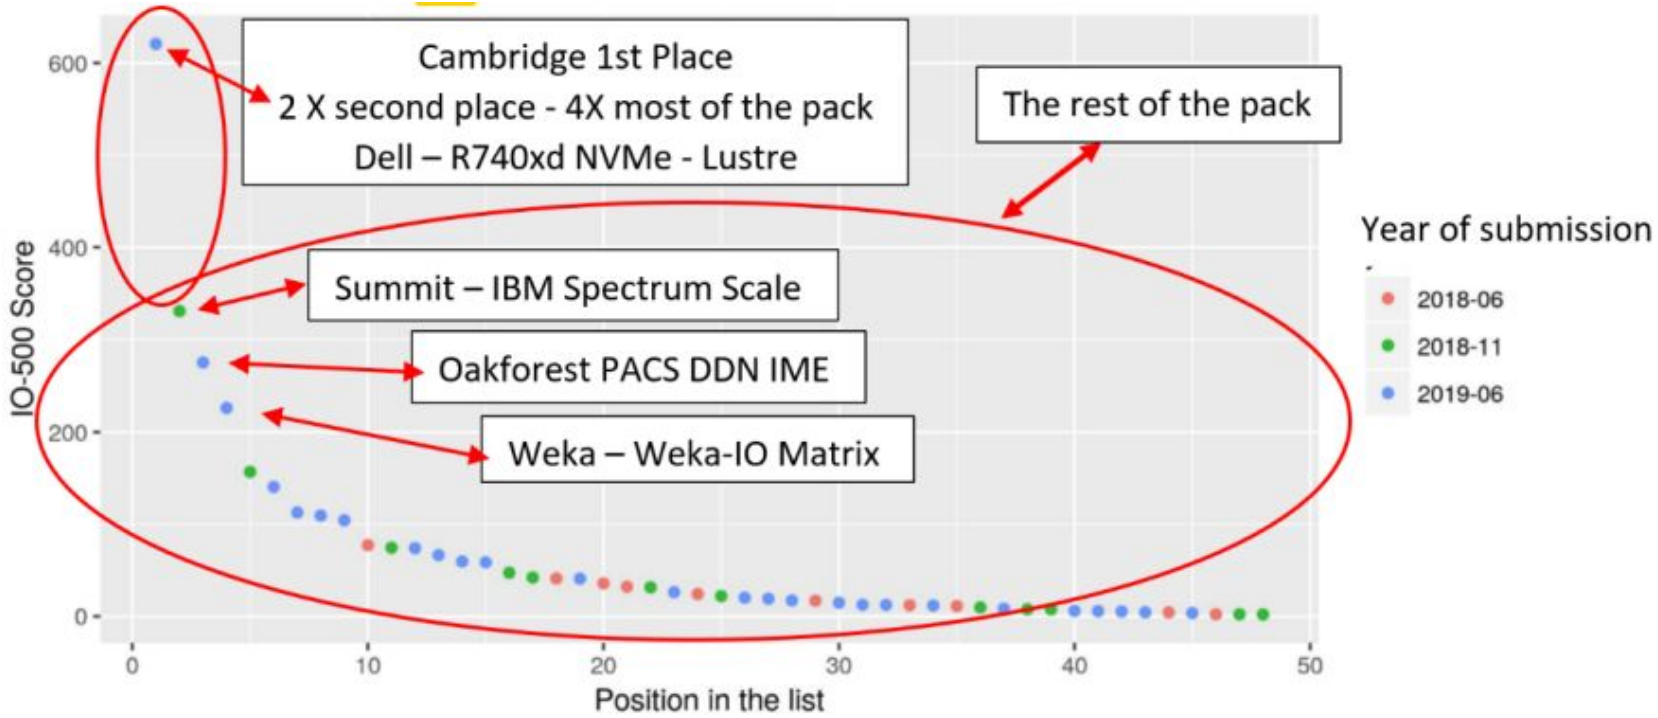
Cambridge 1st Place
2 X second place - 4X most of the pack
Dell – R740xd NVMe - Lustre
The rest of the pack
Summit – IBM Spectrum Scale
Oakforest PACS DDN IME
Weka – Weka-IO Matrix
IO-500 Score
600
400
200
0
0
10
20
30
40
50
Position in the list
Year of submission
2018-06
2018-11
2019-06

How?

StackHPC

Cambridge Data Accelerator Hardware

StackHPC

- 24 x Dell EMC PowerEdge R740xd
- 12 Intel SSD P4600 1.4 TB
- 2 Intel Omni-Path Adaptors

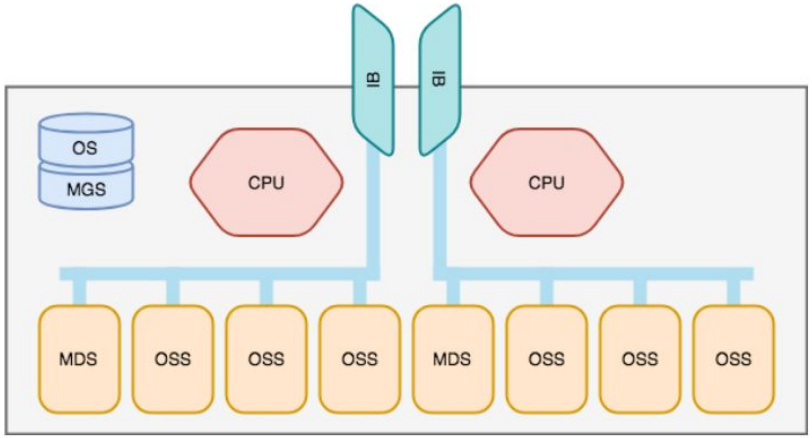

Lustre®

- User requests number of SSDs
- Partitioned into Storage & Metadata
- Lustre: multi-rail
- BeeGFS: Pick NUMA local IB card

- IO-500 Bandwidth: 509 GB/s read, 328 GB/s write
- IO-500 IOPS: 2377 kIOPs

Journey to IO-500

StackHPC

- Balanced Hardware
- Network fabric optimizations
- Iteration: Ansible to rebuild Lustre, Slurm job for benchmark component
- Understanding efficiency of each layer (LNET Self test)
- Experimental Patches from Lustre community

Scale out Bandwidth

StackHPC

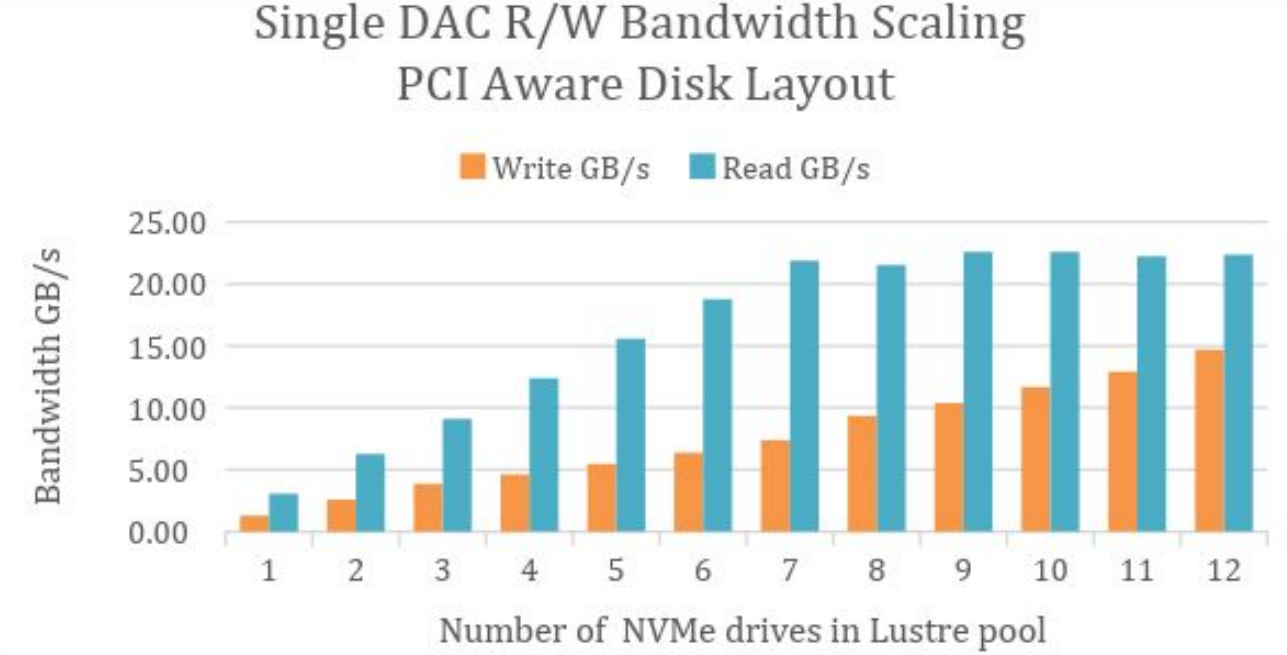

High IOPS

StackHPC

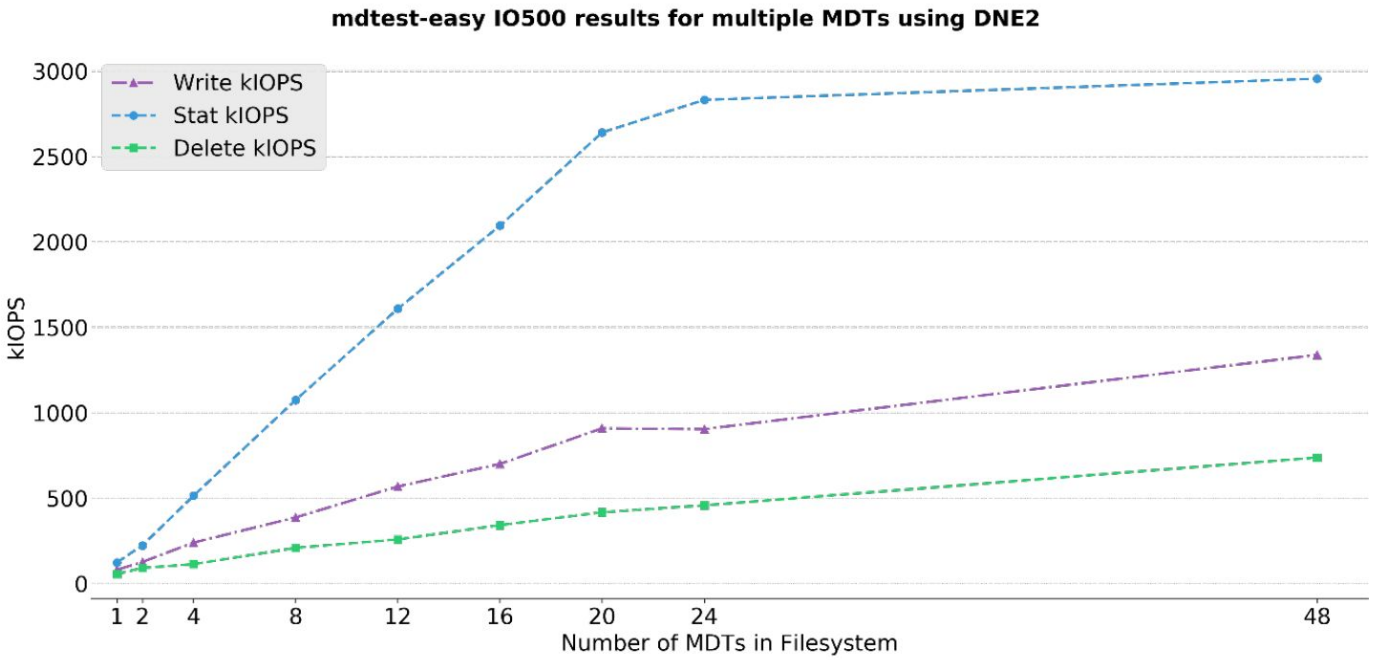

User Experience?

StackHPC

Slurm Burst Buffer Integration

StackHPC

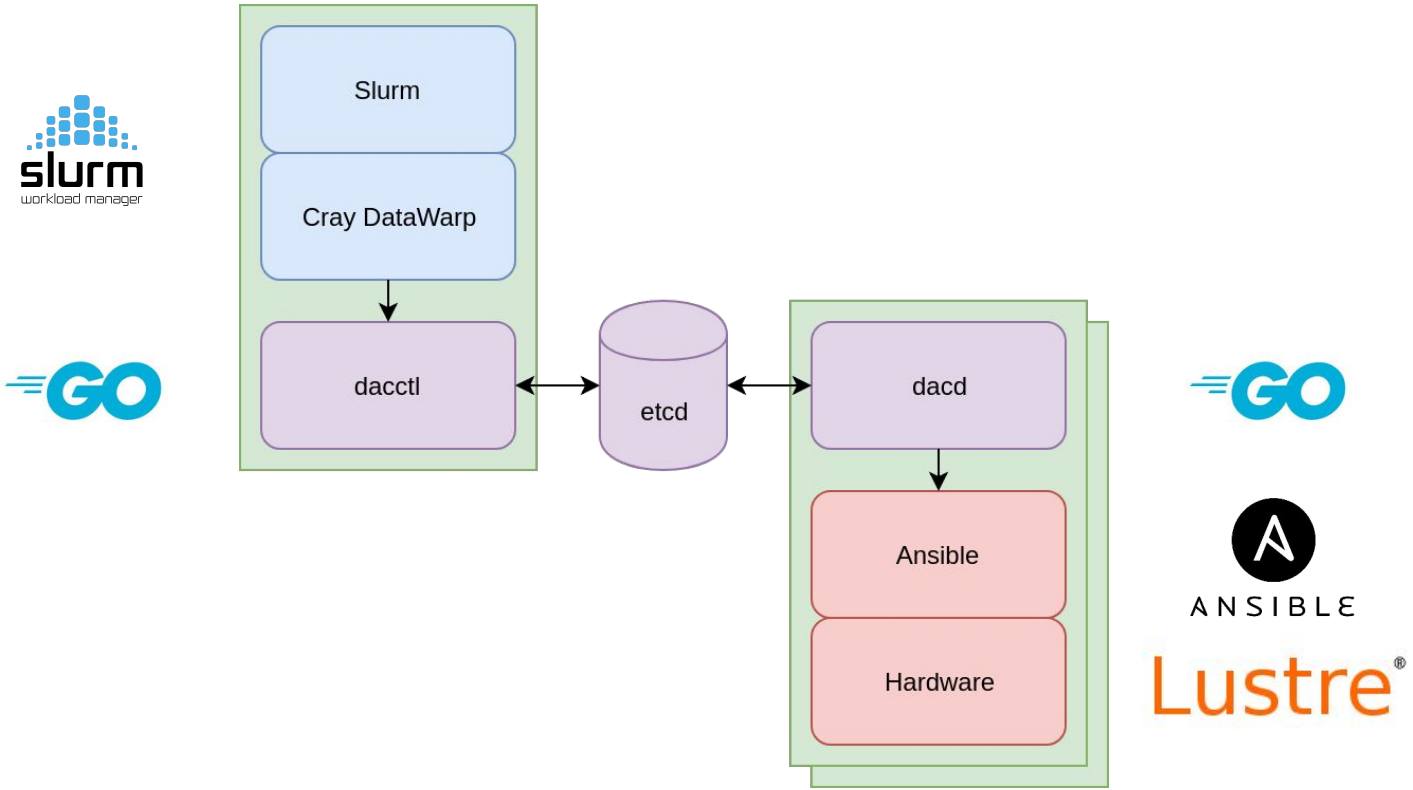

Slurm Burst Buffer Workflow

StackHPC

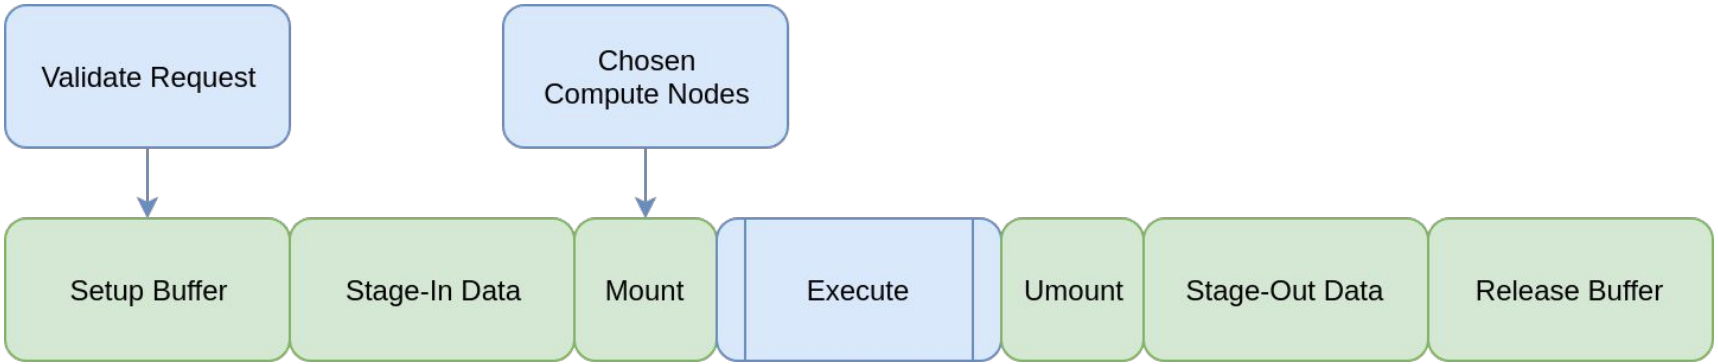

How useful has it been useful?

StackHPC

Square Kilometre Array (SKA)

StackHPC

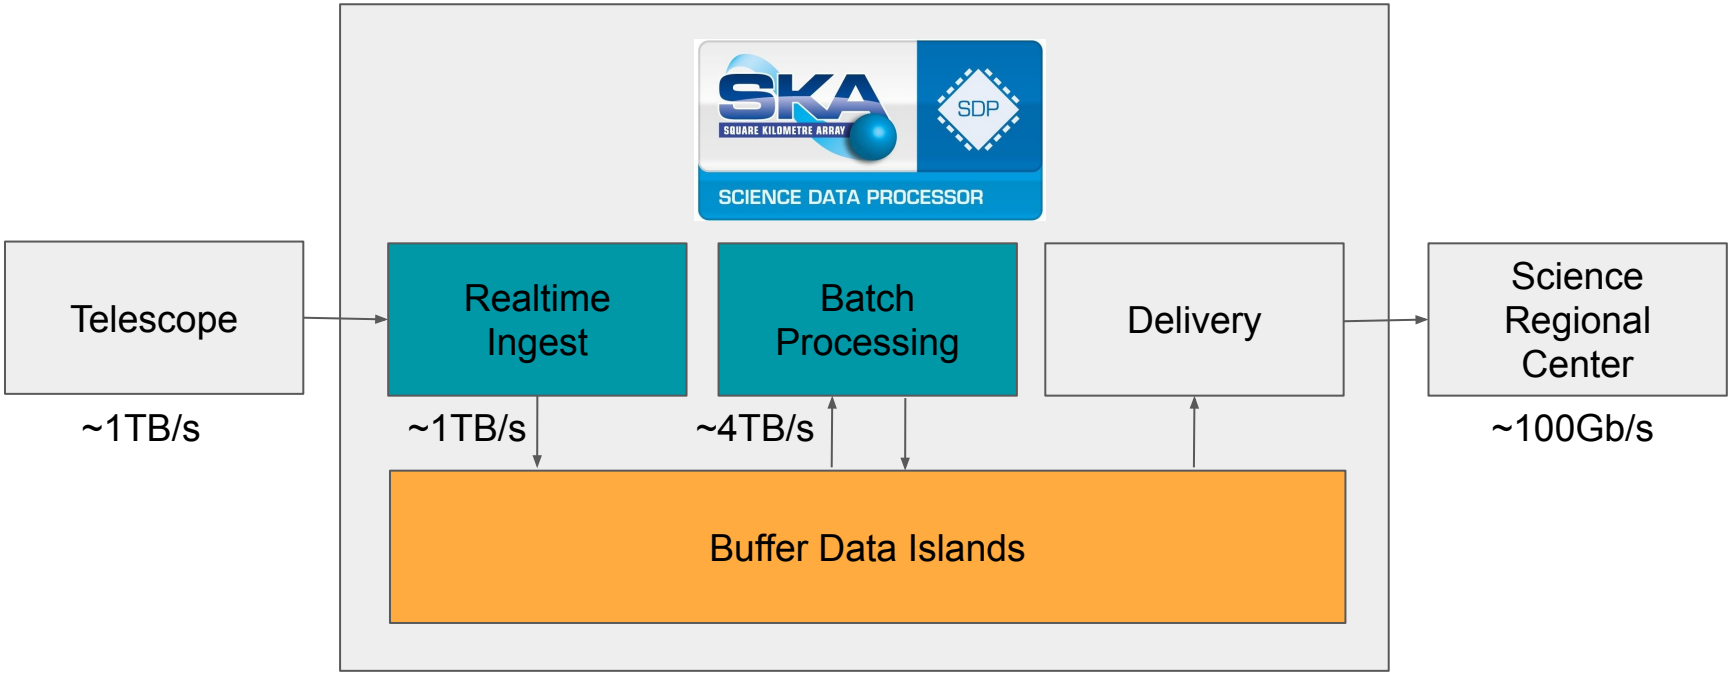

SKA SDP Prototyping

StackHPC

Understanding **Buffer** Management

Feasibility of SDP **Data Islands**

SDP Memo 102

Distributed Predict I/O Prototype

By Peter Wortmann

- Originally limited by Storage
- DAC found limit of current code
- Uses HDF5 and OpenMPI
- Not reaching ior_easy peak rate

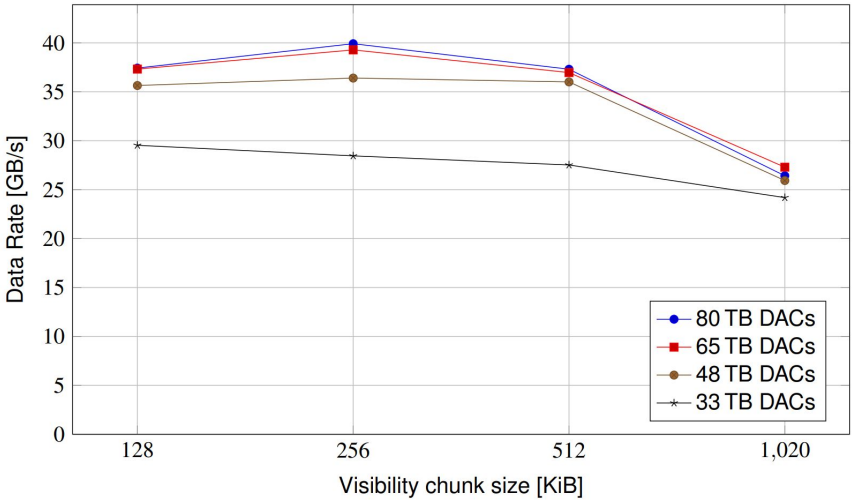

Figure 10: Effective visibility write rate per storage configuration (CSD3/DAC)

What next?

StackHPC

Data Accelerator

StackHPC

Get Involved: Open Source (Apache 2)

Current focus is general availability for CSD3 users

Previously: looked at GlusterFS, BeeGFS

Future: CephFS? OpenStack Manila and K8s PVC?

https://github.com/RSE-Cambridge/data-acc

Thank you!

StackHPC

Special thanks to:

Mark Holliman, *ROE*

Alex Dibbo, *STFC/RAL*

Paul Browne, *Cambridge University*

John Taylor, Mark Goddard, John Garbutt, Doug Szumski, Bharat Kunwar, *StackHPC Ltd*

Special thanks to:

Matt Rasò-Barnett, Paul Browne, Alasdair King, *Cambridge University*

@oneswig
@johnthetubaguy
@stackhpc

StackHPC