Prove-of-concept:
Cycling Data Assimilation Workflows in the Cloud

Stephen Lawrence, Henry R. Winterbottom, Sergey Frolov, and Jeff Whitaker

Cooperative Institute for Research in Environmental Sciences (CIRES)
National Oceanic and Atmospheric Administration (NOAA)
Physical Sciences Laboratory (PSL)

Motivation:
Cloud-Based Reanalyses and Reforecasts (R&R)

Offload Reanalysis and Reforecasts from NOAA's on-prem HPC:

- Insufficient capacity of NOAA on-prem to transition research to operations;
- 30 yrs of ensemble R&R required for major GEFS upgrades;
- GEFS v12 RNR took 2 years at lower resolution of atmosphere-only;
- GEFS v13 RNR will be higher-res and will be atmosphere/ocean/land/sea-ice.

Enable Community Collaboration:

- Make NOAA datasets available for community to enable post-processing and tuning efforts.
- Provide R&R software stack to the community outside NOAA's restrictive secured network.

Workflow footprint

Target:
High-resolution (1/4°) with ensemble

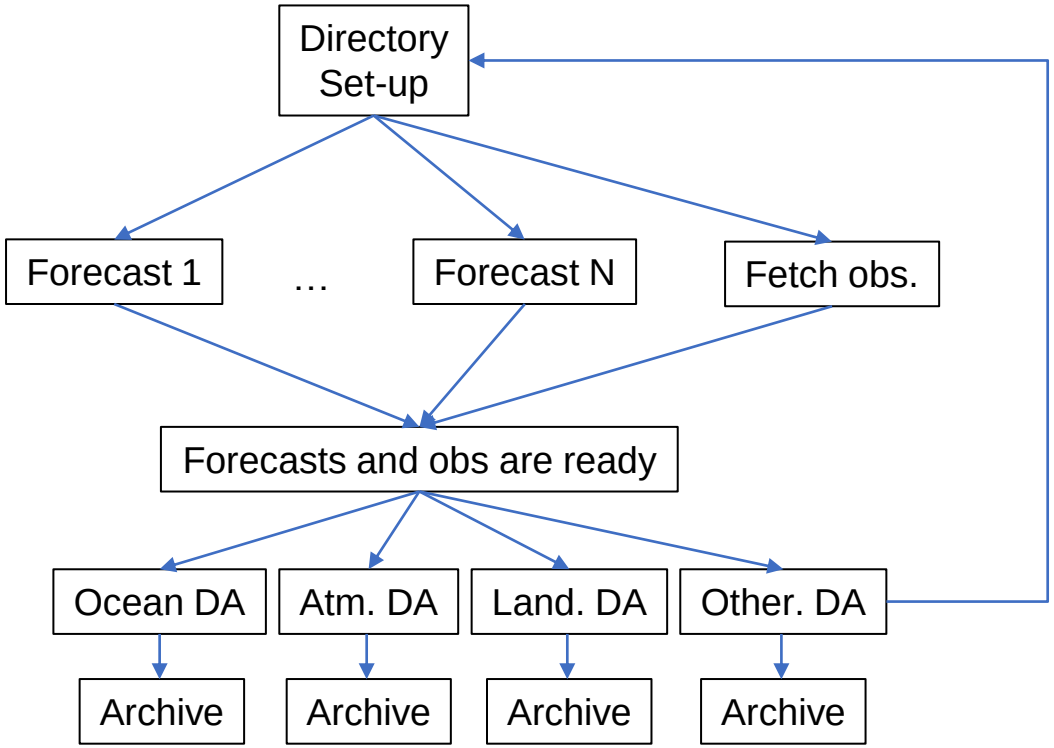

5K-30K CPU cores

Prototype (this presentation):
Low-resolution (1°) deterministic

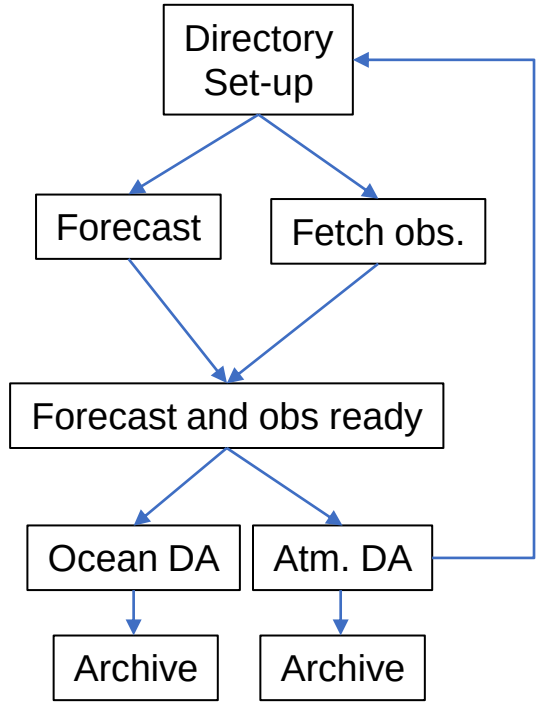

~200 CPU cores

Prototype 3DVAR Workflow

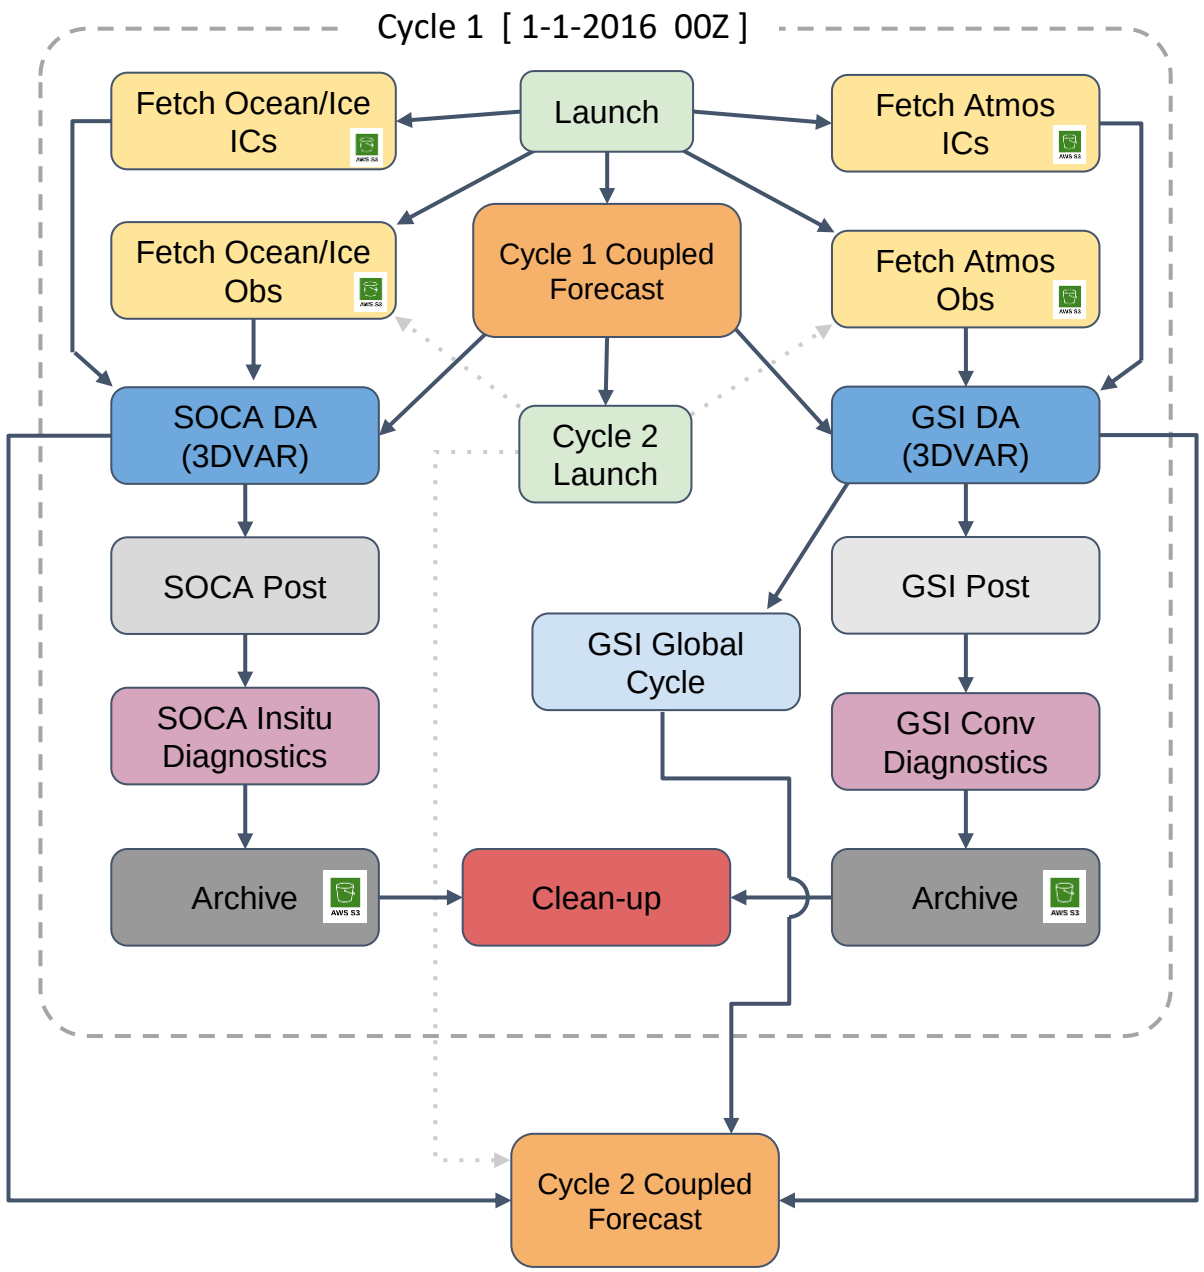

Task	CPU	MEM	IO	MPI Threads
Critical Tasks (must successfully run each cycle)				
Launch	Low	Low	Low	1
Fetch Atmos ICs Fetch Ocean/Ice ICs Fetch Atmos Obs Fetch Ocean/Ice Obs	Low	Low	**High**	1
Coupled Forecast	**High**	**High**	**High**	**208**
GSI DA (3DVAR) SOCA DA (3DVAR)	**High**	**Very High**	**High**	**80**
GSI Global Cycle	**High**	**Med**	**Med**	**6**
Non-Critical Tasks (can fail without affecting results)				
GSI Post SOCA Post	Low	**High**	**High**	1
GSI Conv Diagnostics SOCA Insitu Diagnostics	**Med**	**Med**	**Med**	1
Archive	Low	Low	**High**	1
Clean-up	Low	Low	Low	1

General Cloud Implementation Considerations

High Level Architecture

- Replicate on-premises setup (workflow runs as close to on-prem HPC as possible)
- Stage as much input data on blob storage as possible (S3 or Azure Block Blob)
  - Push initial conditions (IC) and observation data (Obs) to NOAA Big Data Program (BDP) with free egress
  - Use BDP for both On-Prem and Cloud cycling systems (much faster than On-Prem tape storage)
- Hybrid setup (some apps containerized - some apps natively compiled)
- Use Cylc workflow manager
- Use Slurm manager to acquire compute resources
- Except for resource configuration, On-Prem HPC workflow is identical to Cloud Workflow

Component Choices

- Compute instance types:
  - Network optimized: AWS - efa, Azure - Infiniband
  - AWS: Trade-off between compute and memory - C5n, R5n, M5n
  - AZURE: HPC optimized - hc44rs
  - AMD™ or Intel™ : On-Prem compile is optimized for Intel
- Scratch Storage:
  - Lustre
  - Network storage: AWS EBS, AZURE NetAPP
  - Local disk: Azure BeeOND
- Data and software stack location:
  - AWS: EBS or EFS
  - Azure: Block Blob or nfs 3.0

Access to Cloud Platforms

- Direct access: AWS vs Azure
  - More service options but also more setup work

- ParallelWorks : multi-cloud platform for AWS, Azure, GCP
  - Rapidly evolving and relatively young platform
  - Easy cluster deployment (caveat - only after fair amount of initial setup)
  - No custom instance images allowed
  - Restricted subset of configurations for node and service management

AWS ParallelWorks Account Cloud Architecture

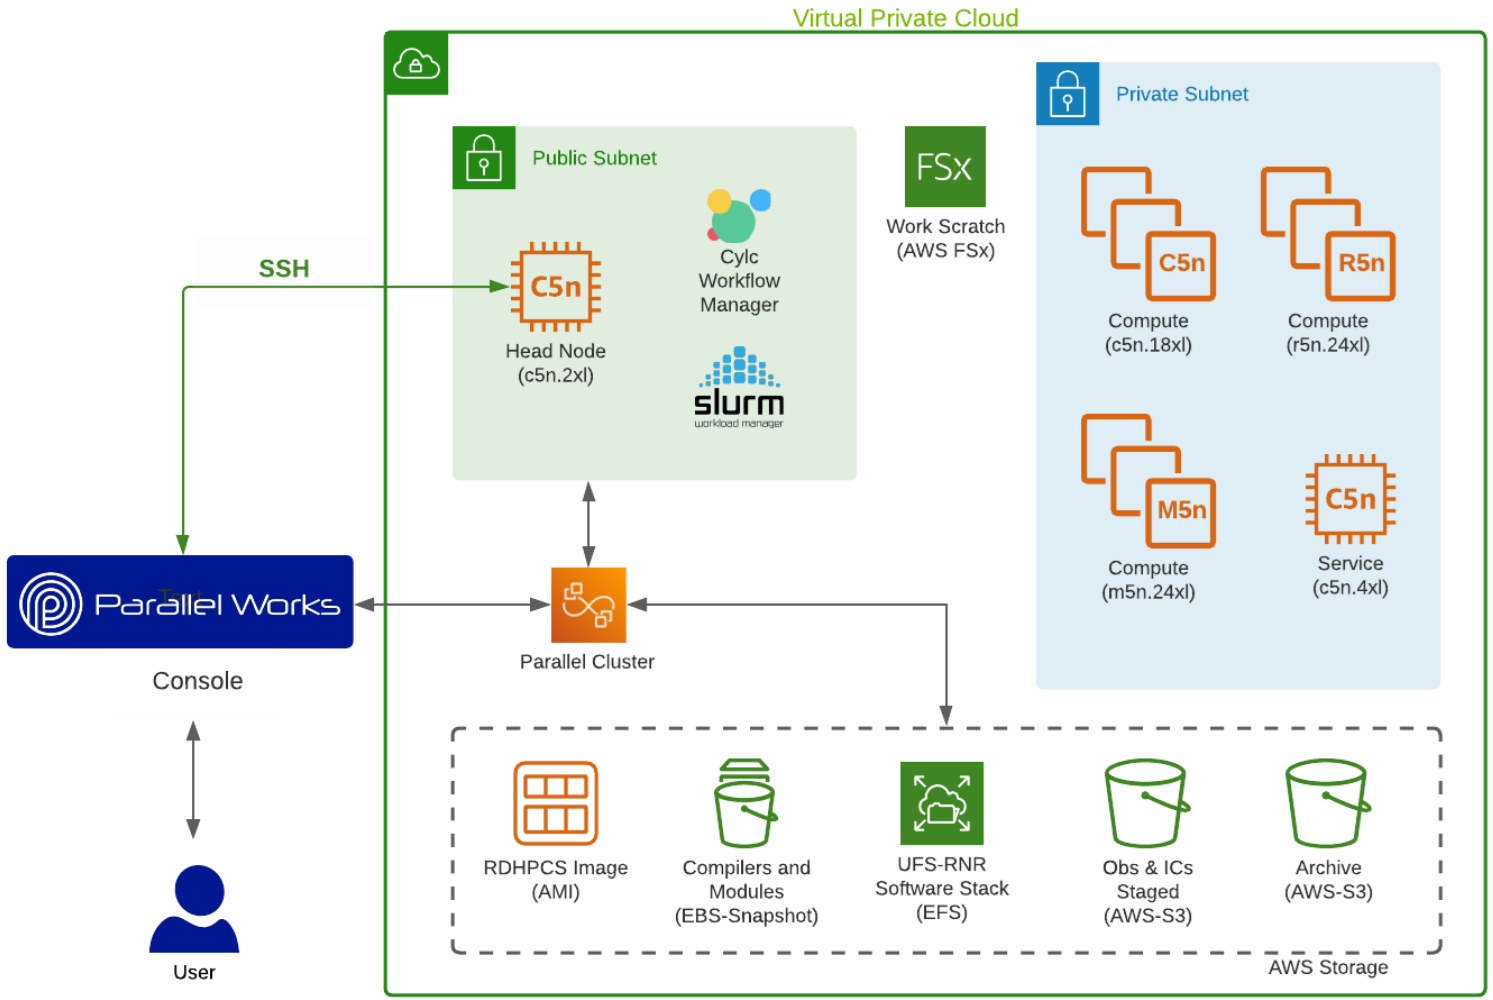

AWS Parallel Cluster™:	create and manage cluster
Master node (c5n.2xl):	compiling, cylc, slurm, NFS mounts
Compute queues:	compute - c5n.18xl big-mem - r5n.24xl generic - m5n.24xl service - c5n.4xl
Luster service (FSx):	classic work scratch
Persistent storage:	s3 for staged input data - ICs and Obs
Software stack:	OS - Alinux2 (CentOS compatible) PW stack - compiler, MPI Application - RnR

Cost Comparisons

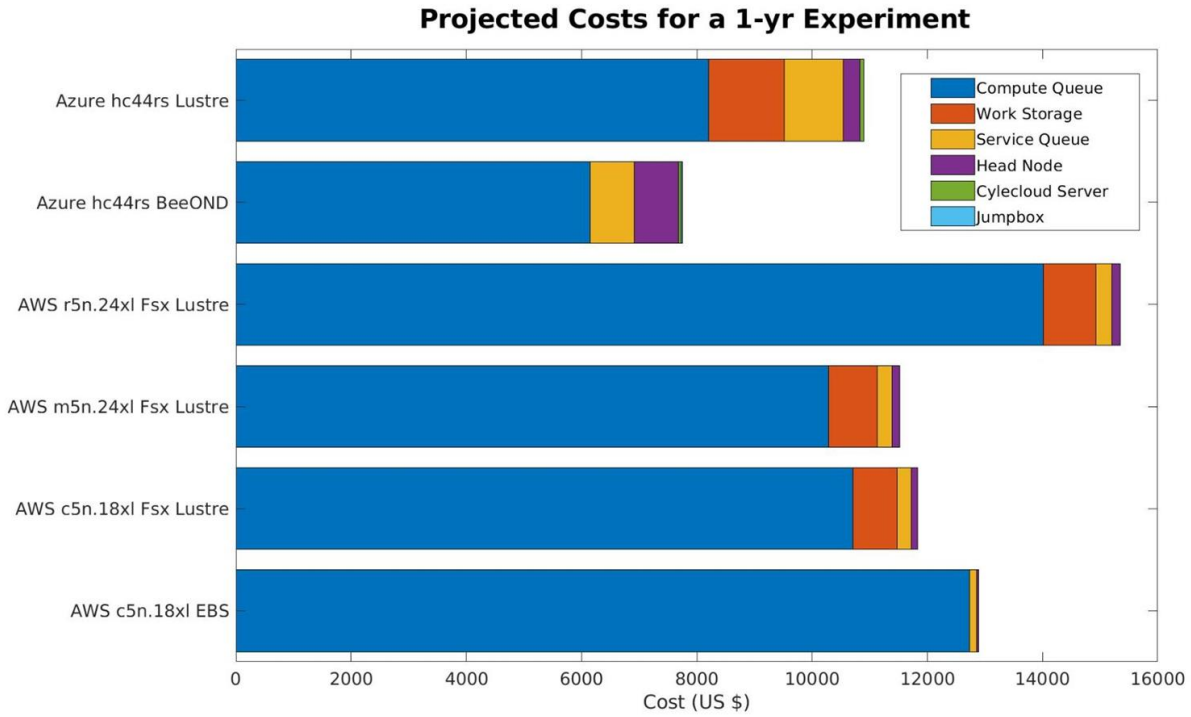

- Compute costs dominate on all platforms
- Lustre: costly but can be configured to be cheaper
- Azure requires more vm overhead (cyclecloud) but instances are slightly cheaper per hour
- BeeOND: faster and free but not fault tolerant
- Using generic mode with EFA (m5n) was more cost effective than compute-optimized or memory-optimized nodes.

Throughput Comparisons

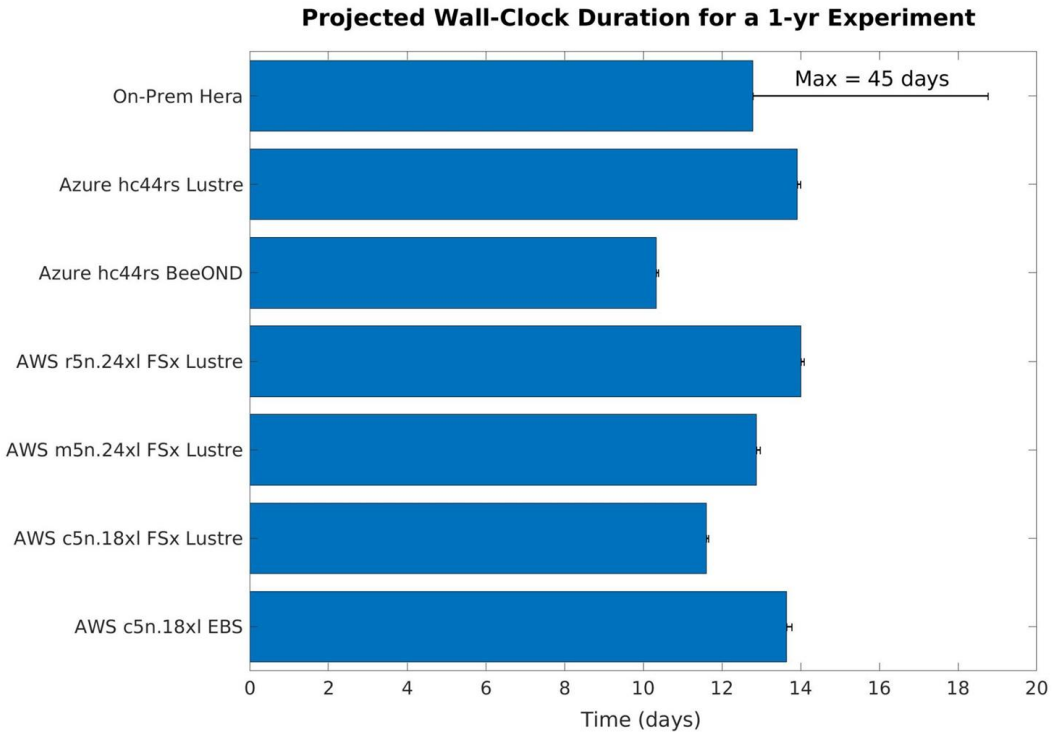

- On-Prem can be fast and is already budgeted, but in practice resource competition between researchers can degrade throughput substantially
- All Cloud platforms show consistent and predictable throughput that is comparable or faster than On-Prem

Lessons learned

- Fast interconnect is essential for DA workflow (e.g. DA solvers);
- Scaling compute footprint up and down is impractical in cycling workflow;
- Memory-optimized nodes are expensive;
- Storage choices affect performance and cost;
- Containers are tricky for multi-node MPI jobs;
- Spot availability is spotty (at least on AWS)
- Compiling natively can be time-consuming (but we were successful at using binaries compiled on Azure to run on AWS)

Conclusions from Cloud Testing

- Scientific results indistinguishable from On-Prem.
- Cost of the scaled-down workflow is affordable.
- Throughput matches or beats practical On-Prem performance.
- Hybrid (CSP/on-prem) deployments are possible:
  - E.g. can run reanalyses streams using combination of on-prem and CSP.
- Next steps:
  - Run extensive simulations (multiple decades);
  - Explore using spot computing on Azure;
  - Develop cloud-enabled diagnostic capabilities.

End

Parallel Works Managed Azure Cloud Architecture

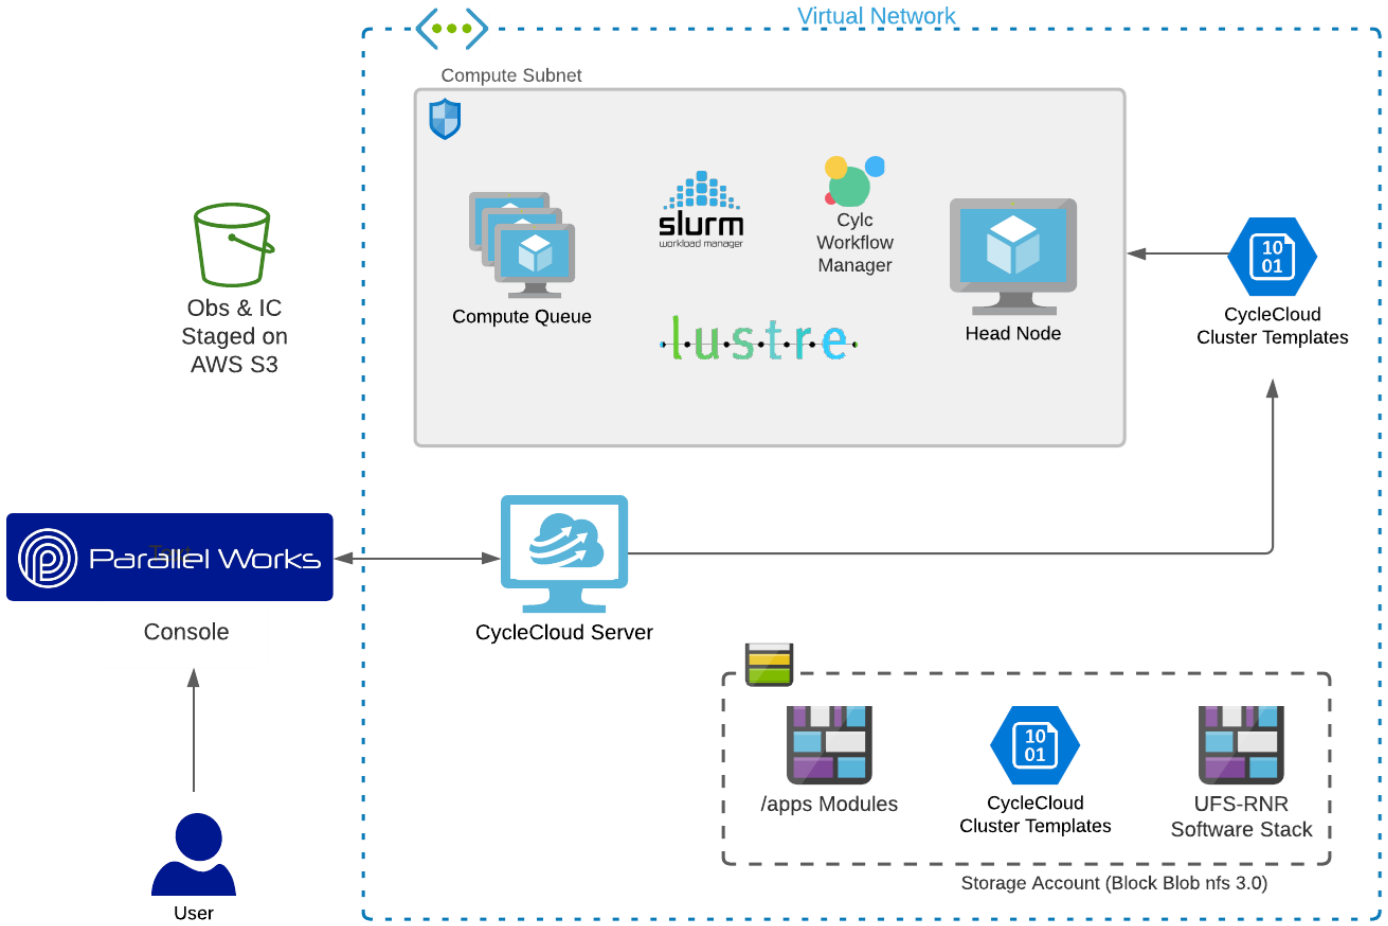

Azure CycleCloud™:	create and manage cluster
Master node (h8):	compiling, cylc, slurm, NFS mounts
Compute queue (hc44rs):	balanced compute and memory
Luster cluster (D8ds_v4):	1 mds node and 2 oss nodes - classic work scratch
Persistent storage (aws s3):	staged input data (ICs and Obs)
Software stack (block blob):	OS - CentOS 7, PW stack (compiler, MPI), application (RnR) on block blob nfs 3.0 storage.

AWS OAR-TigerTeam Account Cloud Architecture

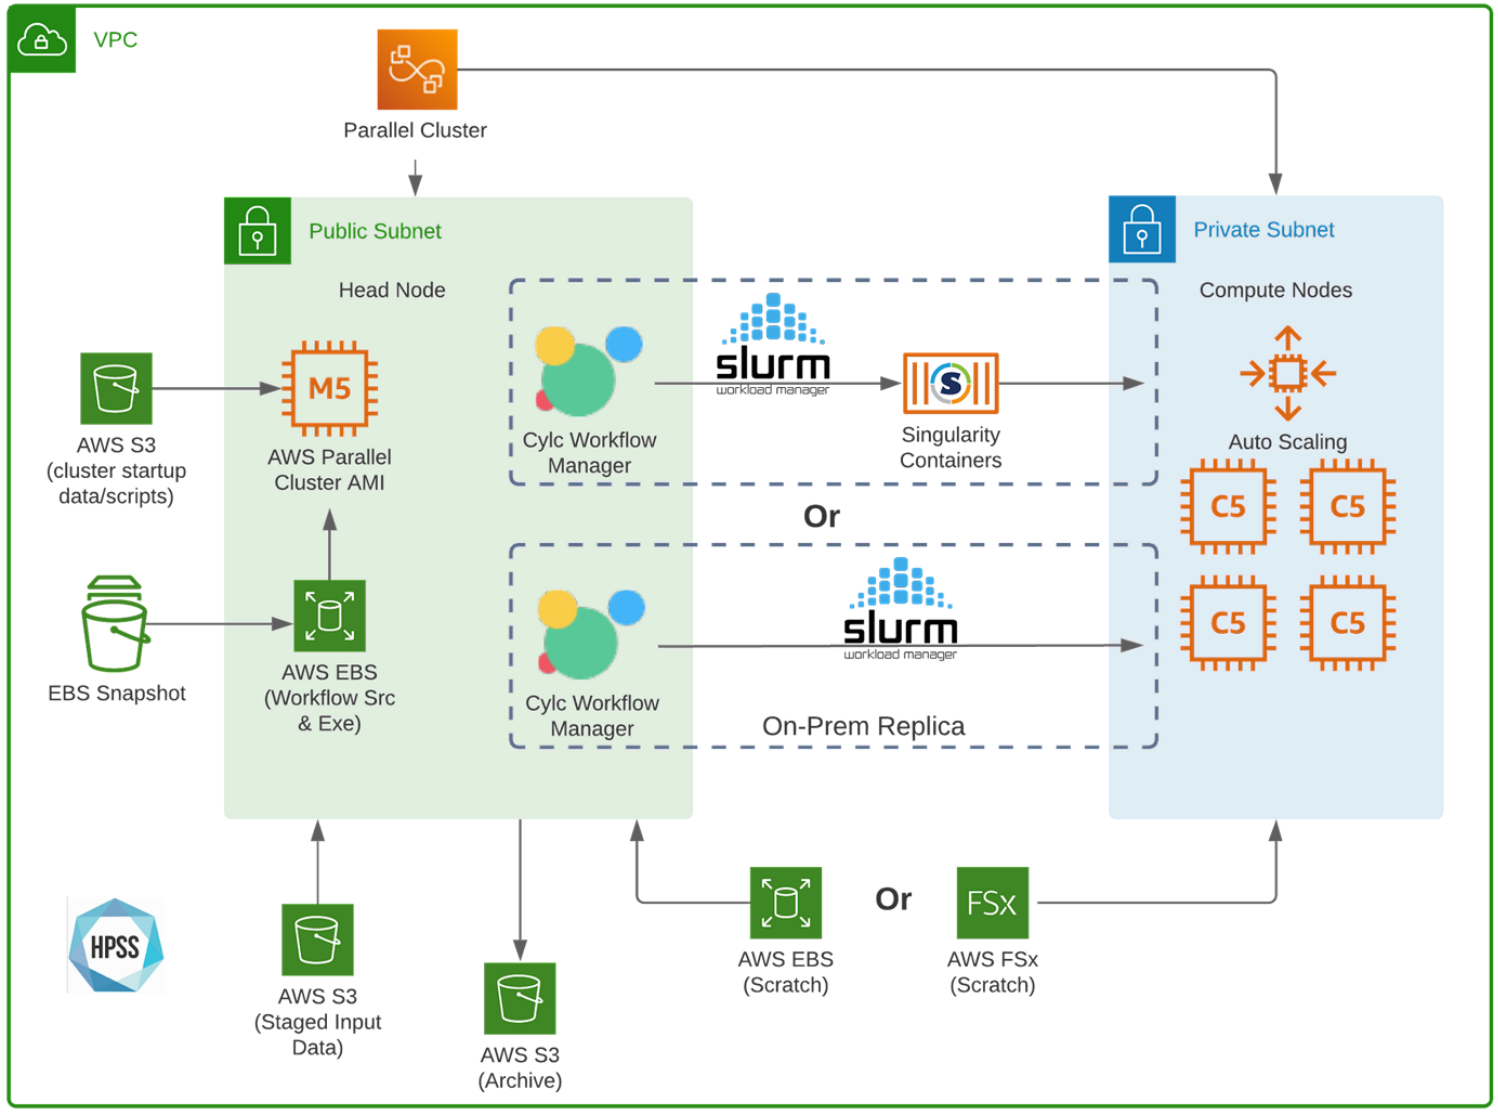

MS Azure AI for Earth Account Cloud Architecture

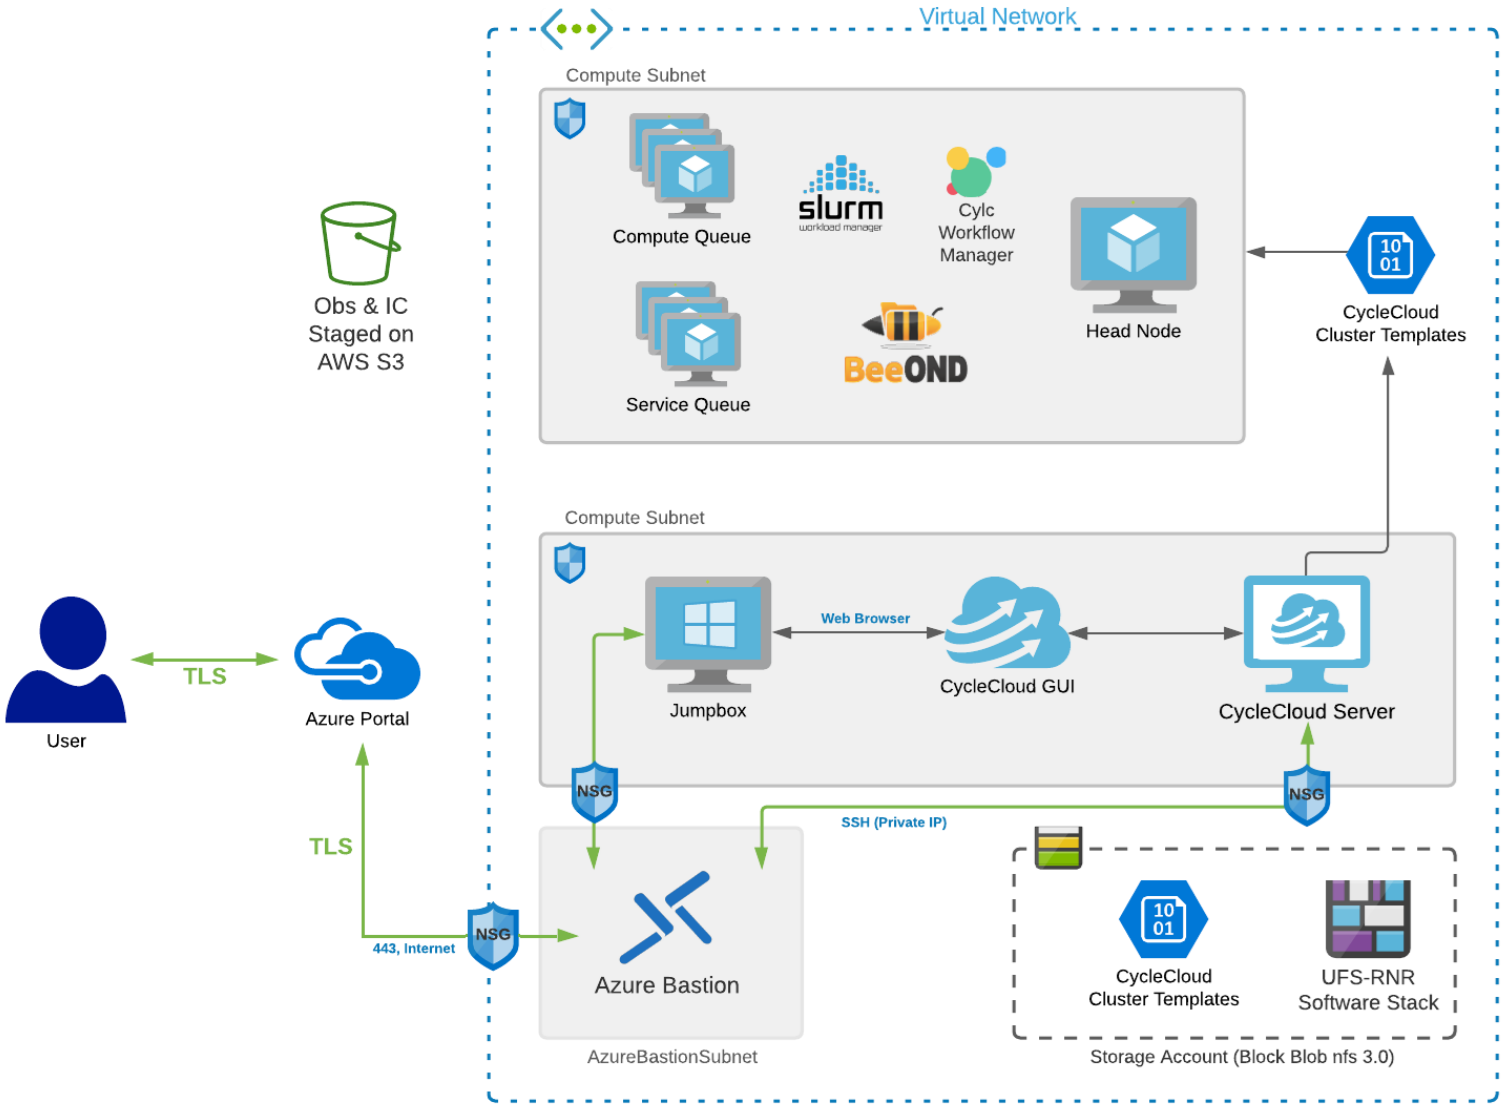

Cloud-Based Workflow Is Scientifically Accurate
15-day Experiment Starting on 1-1-2016 (Cloud Matches Hera)

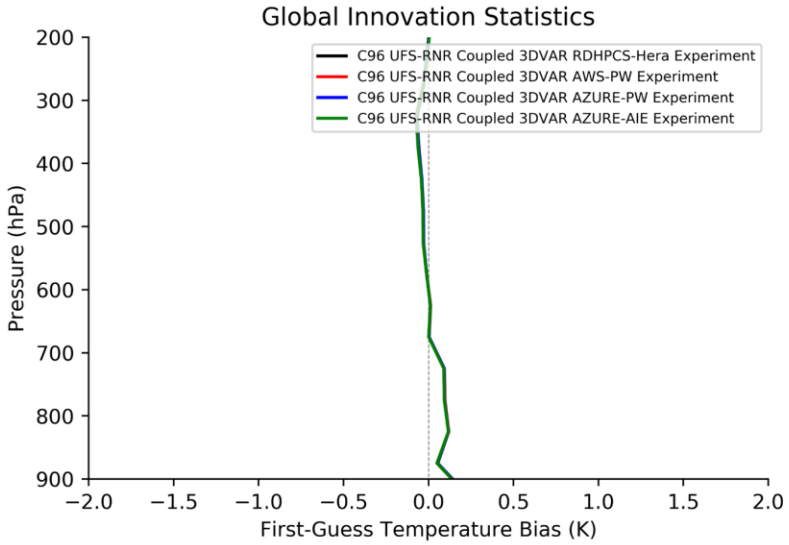

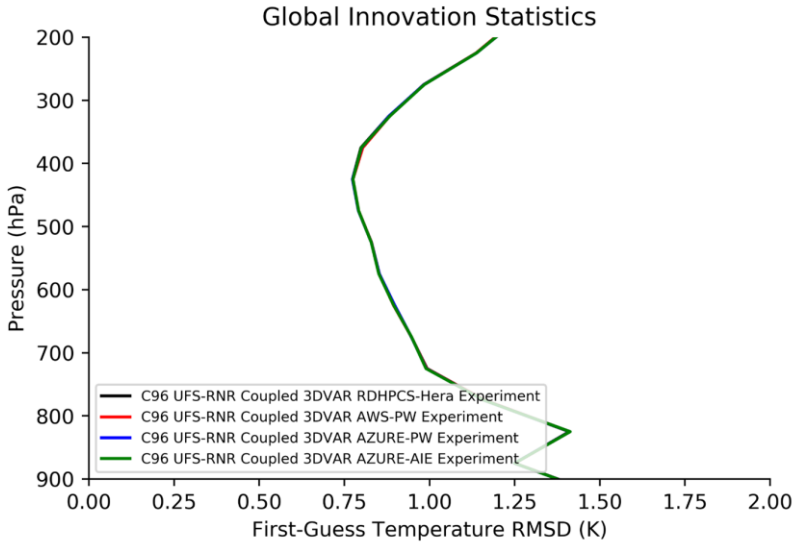

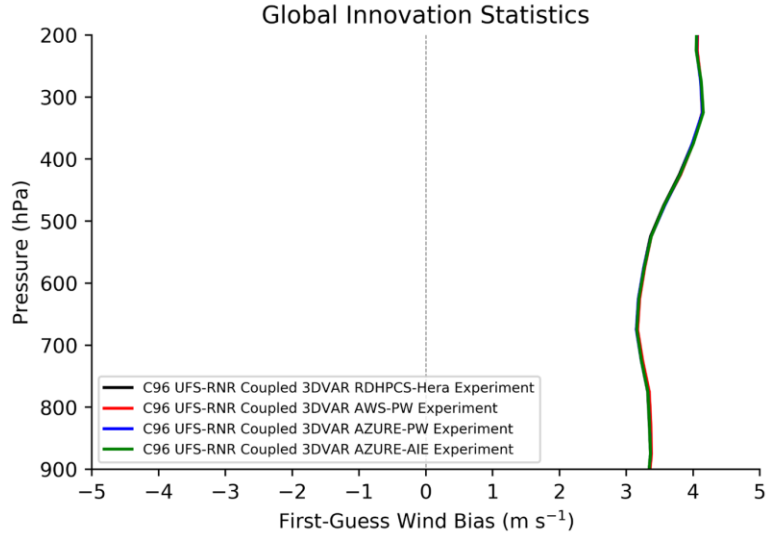

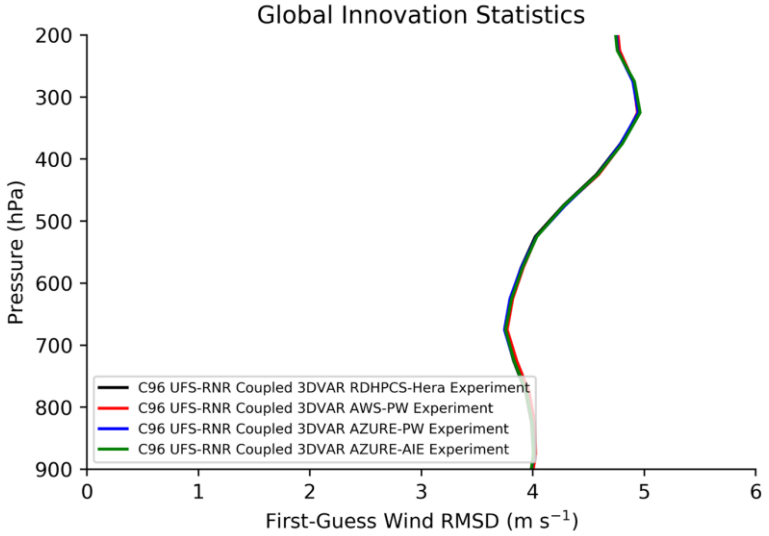

AWS Elastic Compute Cloud (EC2)
Instance Type / Capacity Considerations

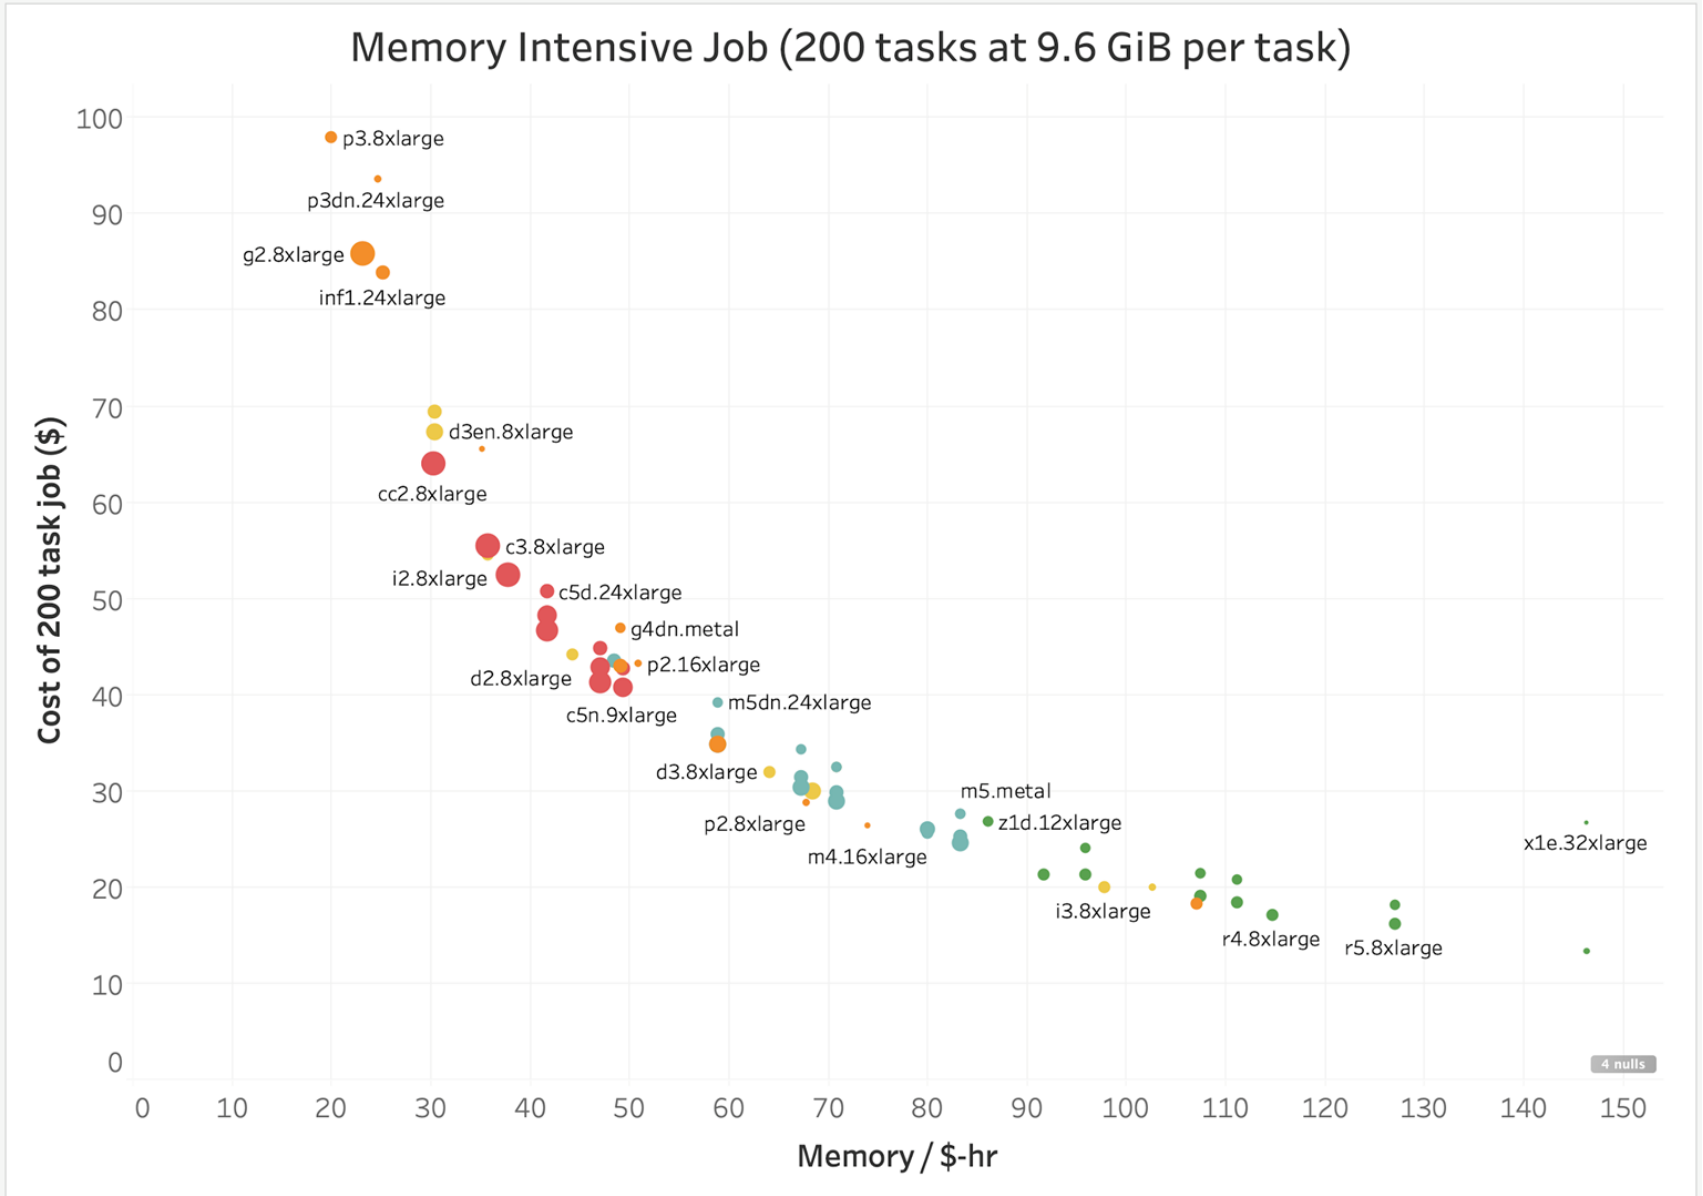

Balance

- Capability
- Availability
- Cost

Instance Family
- Accelerated Computing
- Compute Optimized
- General Purpose
- Memory Optimized
- Storage Optimized

Nodes Required
- 1.00
- 10.00
- 20.00
- 33.00