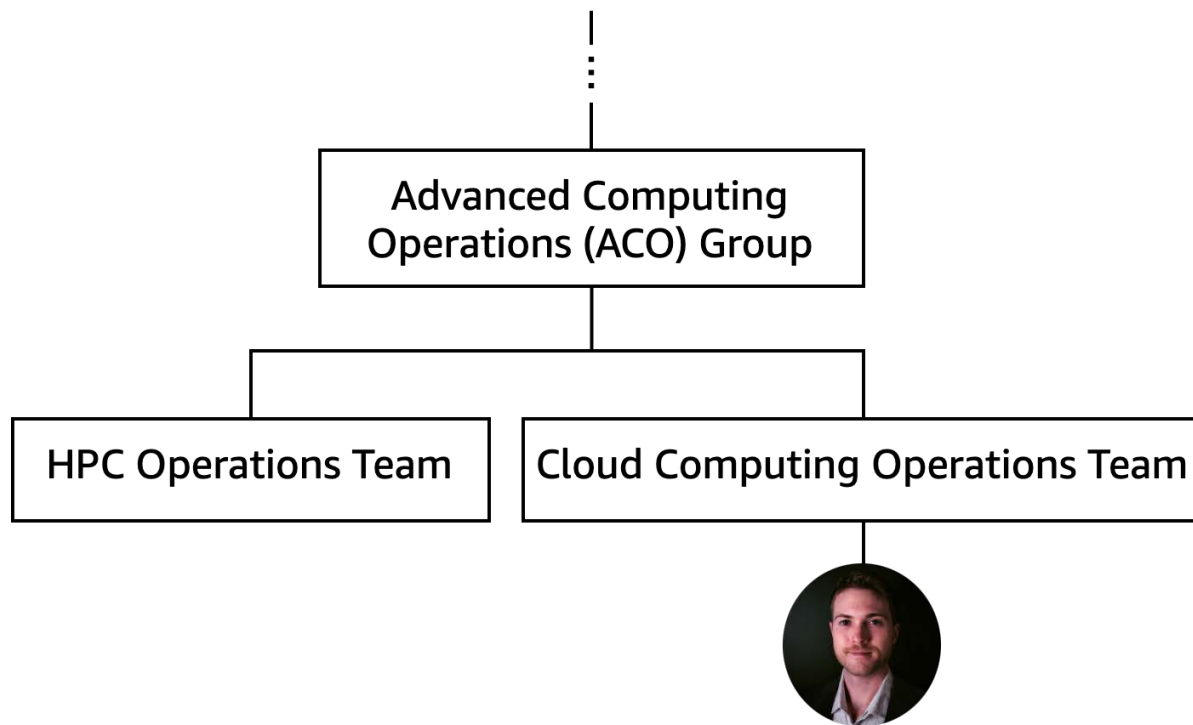




NREL's Journey with HPC in the Cloud and Hybrid Computing

Michael Bartlett
NLIT Summit 2025
May 7th, 2025

My Role at NREL



Why Hybrid HPC?

System availability

While on-prem system is offline during maintenance periods, cloud system could continue running jobs in the interim.

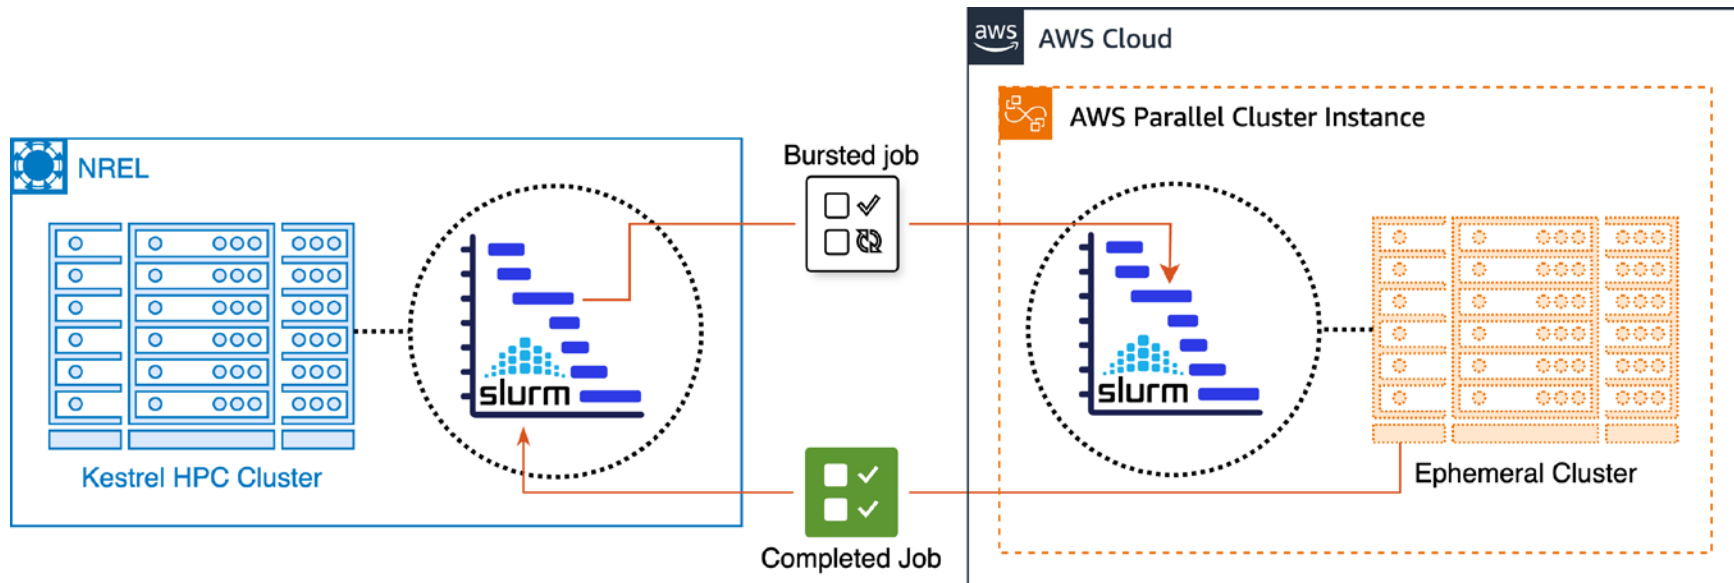
Reduce on-prem scheduler congestion

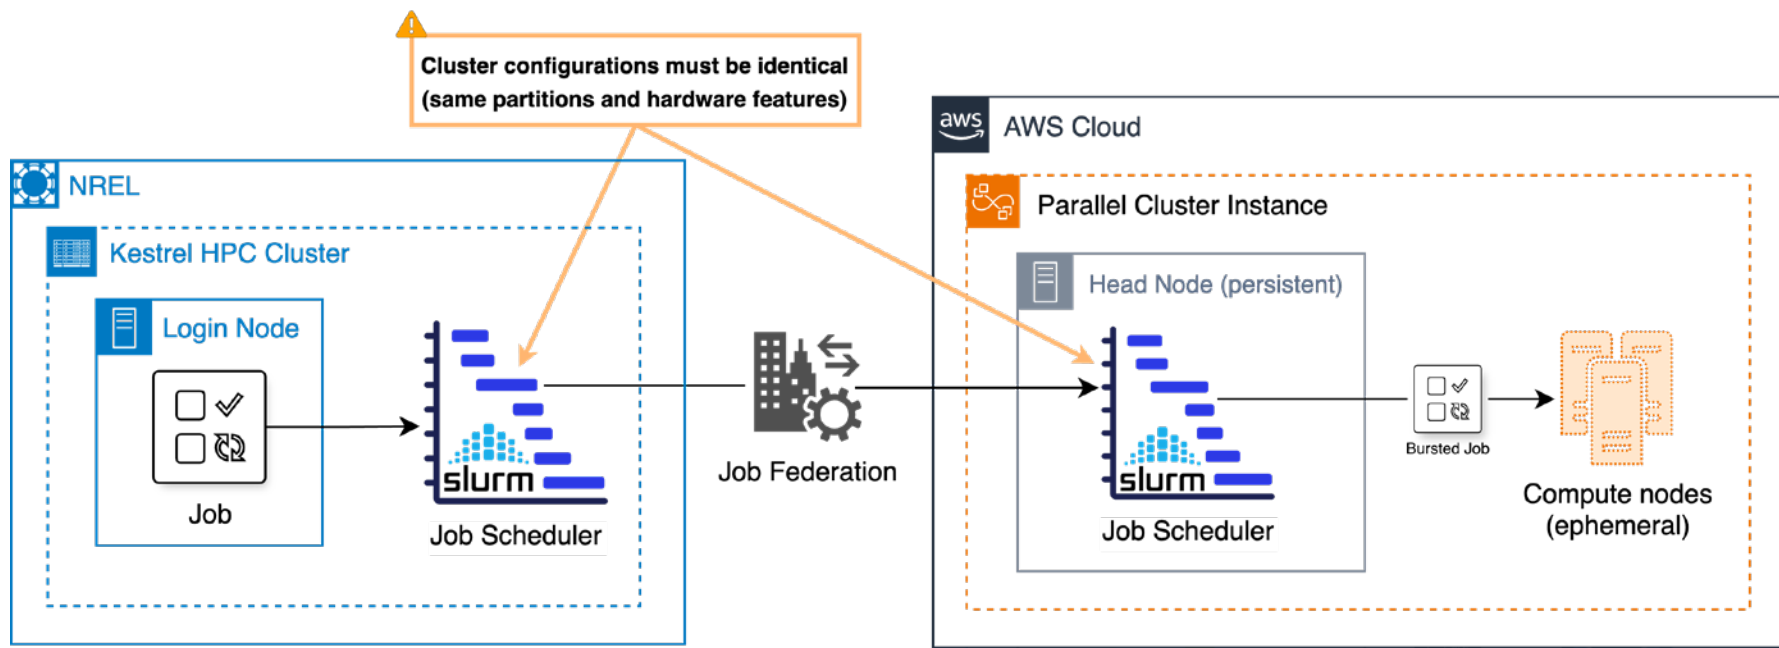
Not all jobs are well suited for large-scale HPC and create noise for the scheduler.

Hardware flexibility

On-prem system hardware is fixed, while the cloud system can dynamically offer different hardware features.

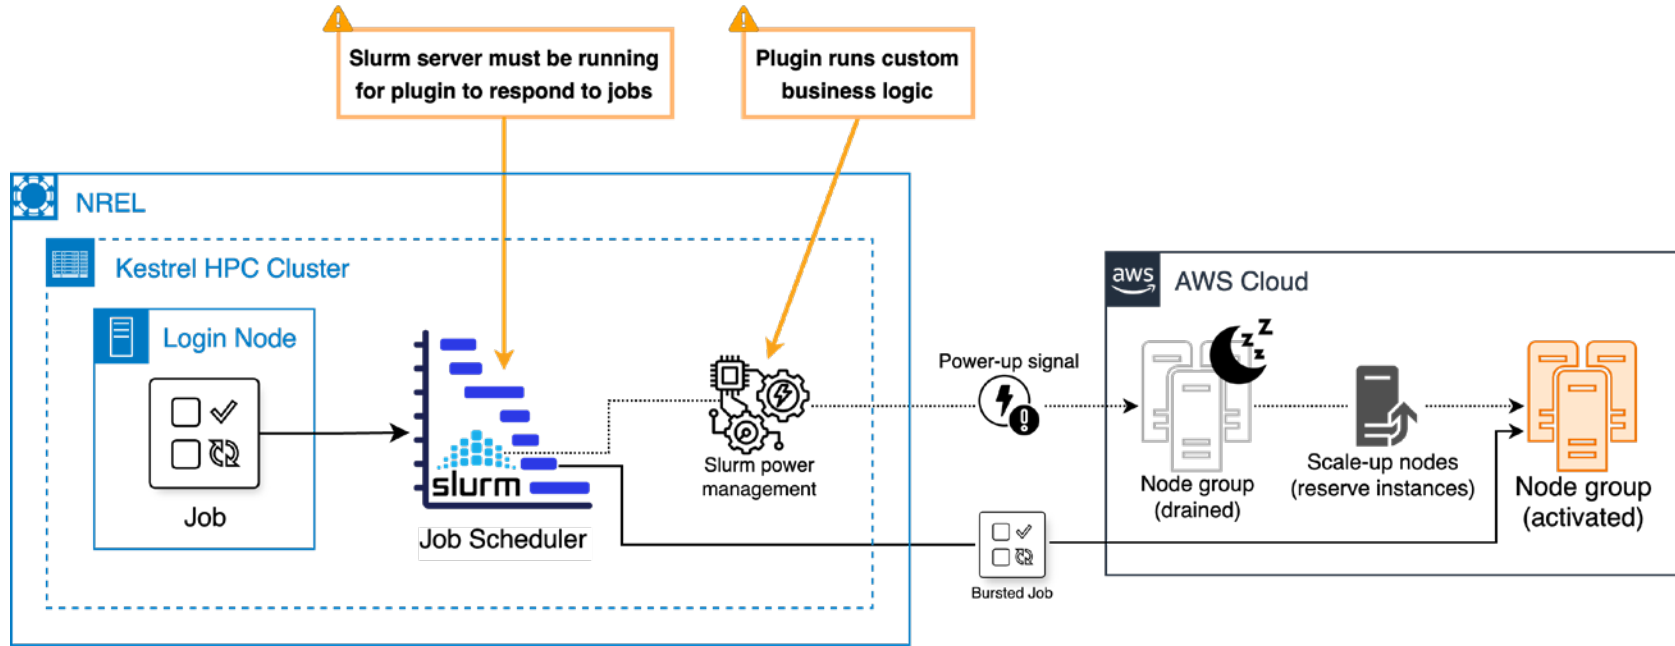
Design Overview





Implementation: Slurm Federated Scheduling

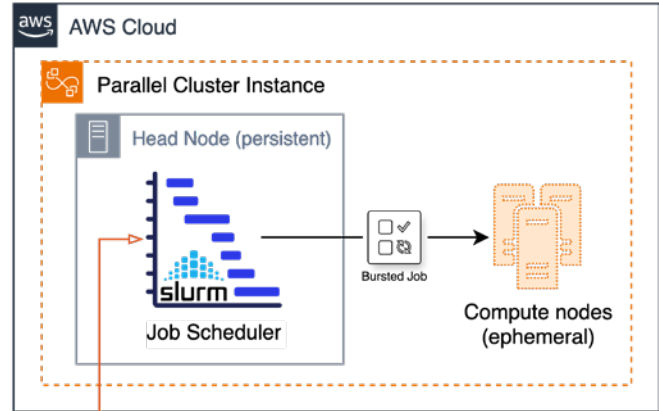
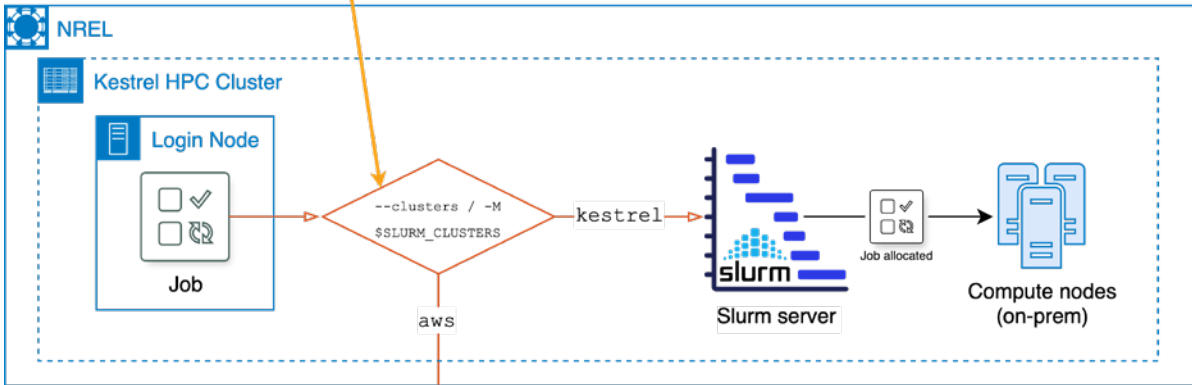
- Clusters must have identical configuration.
 - This was a non-starter as this restriction prevents right-sizing and hardware flexibility.
- Job plugin should add a penalty to cloud cluster, or else jobs prefer it since it's always "ready."



Implementation: Custom power-management plugin

- Plugin requires Slurm server is running on-prem.
 - This prevents providing job scheduling during maintenance periods.
- The plugin would be a custom solution which would require on-going development and maintenance of the logic, which is likely to be time-consuming and error-prone.

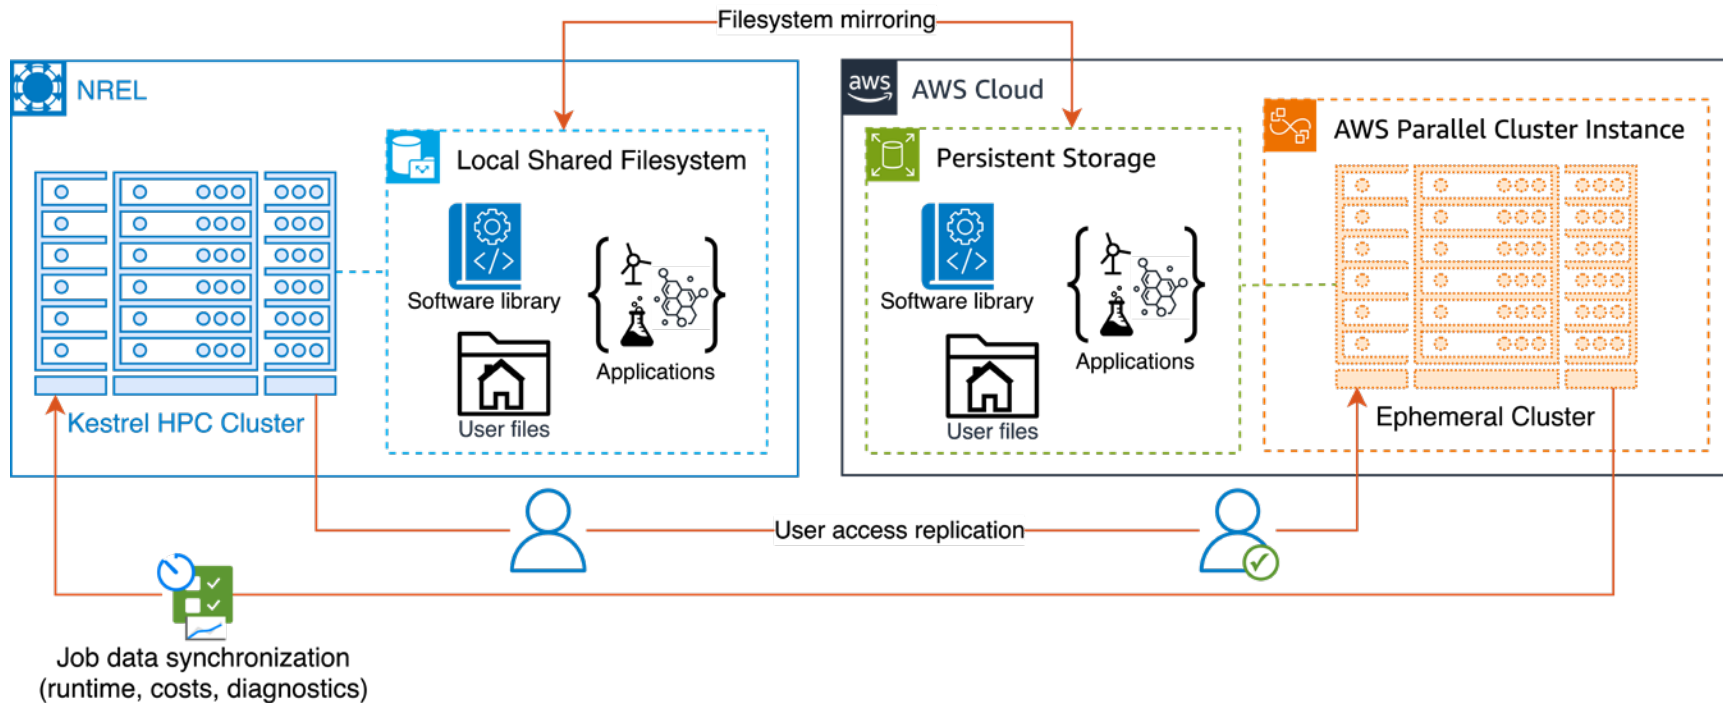
⚠ Users must specify cluster for all Slurm commands (sbatch/sinfo/queue,etc.)




Implementation: Multi-cluster

- Clusters are independent, one being down will not impact availability of the others.
- Users must specify `--clusters / -M` or use `$$SLURM_CLUSTERS` for all Slurm commands.
- Requires persistent Slurm daemon.

Next Steps




Hybrid Environment Research and Operations (HERO)




Auth

Provide authentication and authorization.




Data Repo

Store and organize public and private data.



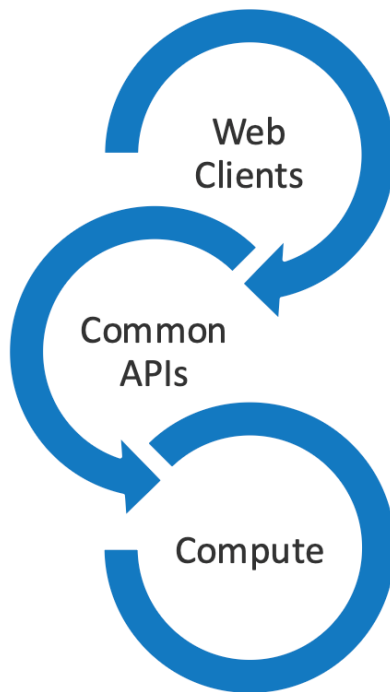
Task Engine

Define distributed and hybrid compute.



Search

Quickly find resources using metadata.



Simple Design—Multiple projects that have similar requirements don't need to individually solve the same problem.

Effective Tooling—Use NREL hardware when either necessary or desirable without leaving behind the benefits that cloud provides.

Sustainable Development—Following the Well Architected Framework helps lower the barriers to compute resources and leveraging infrastructure as code simplifies cumbersome and expensive processes.



<https://hero.nrel.gov/>

Thank you

www.nrel.gov

NREL/PR-2C00-93600

This work was authored by NREL for the U.S. Department of Energy (DOE) under Contract No. DE-AC36-08GO28308. Funding provided by the National Renewable Energy Laboratory (NREL). The views expressed in the article do not necessarily represent the views of the DOE or the U.S. Government. The U.S. Government retains and the publisher, by accepting the article for publication, acknowledges that the U.S. Government retains a nonexclusive, paid-up, irrevocable worldwide license to publish or reproduce the published form of this work, or allow others to do so, for U.S. Government purposes.

Photo from Getty-181828180

