



Exceptional service in the national interest

# Take Me to the Clouds Above: Bridging On Site HPC with Clouds for Capacity Workloads

**Jay Lofstead** and Dmitry Duplyakin

CLOUD COMPUTING 2021

April 18, 2021



Sandia National Laboratories is a multission laboratory managed and operated by National Technology and Engineering Solutions of Sandia LLC, a wholly owned subsidiary of Honeywell International Inc. for the U.S. Department of Energy's National Nuclear Security Administration under contract number DE-NA0003525.



# A Tale of Two Compute Infrastructures

Traditional HPC	Cloud Computing
In house “large” scale compute	Off-site “large” scale compute
Optimized to local workloads	Optimized to general workloads <ul style="list-style-type: none"><li>• Small subsets dedicated to particular workloads</li></ul>
100% use dedicated to organization	Shared with other cloud customers
Fully isolated from external users	Only as isolated as security settings dictate <ul style="list-style-type: none"><li>• User errors can expose things to the world accidentally</li></ul>
If you can use nearly 100% of the compute, cheapest option	If demand is variable, probably the cheapest option
Must have in house expertise to manage <ul style="list-style-type: none"><li>• Can be expensive and difficult to find and keep</li></ul>	Management provided as part of the service



# Traditional HPC Additional Benefits

---

- Specialized architectures can be fully deployed to optimize performance
  - DE Shaw MD machine from several years ago
  - Dedicated ML hardware
  - Dedicated high DRAM nodes
  - Dedicated on-node storage
  - Specialized interconnects
- Optimized for performance
  - Bare metal hardware access, potentially containers for deployment ease
- Security possible
  - US DOD classified cloud computing only allows “secret” level on certified clouds (several levels beyond that exist)
  - Air gapping and other physical security



# Cloud Additional Benefits

---

- Evolving platform
  - Up to date hardware being deployed continuously
  - Specialized hardware, in small quantities deployed to enable testing
  - VM abstraction for consistent software interface on varying hardware
    - Typically containers on top of that
  - Evolving security and privacy standards addressed by dedicated experts and worked into cloud defaults
    - Just don't break it without knowing what you are doing
  - Large, public data sets may live here allowing quick access at low cost
    - Climate data sets, among others, are made available for others
  - Sharing between organizations without specialized VPNs and security possible
    - E.g., Snowflake and similar services
  - Third party arbiter for inter-group and intra-group conflicts
    - When bills come in, decisions about what is important can change quickly



# Combining Traditional HPC and Cloud

On Site	Cloud
High security sensitive workloads	Low security workloads
Isolation critical	Sharing with other organizations more easily
Scale up capability workloads	Scale out capacity workloads
Specialized hardware	No special hardware needs
Special machine reservations	Potentially wait for special hardware reservations

