



ESnet

ENERGY SCIENCES NETWORK

Report from the USATLAS-USCMS HPC/Cloud Blueprint workshop

Dale W. Carder

ESnet Network Engineering

Lawrence Berkeley National Laboratory

LHCOPN/LHCONE Meeting #49

2022-10-25



U.S. DEPARTMENT OF
ENERGY

Office of Science



Workshop for the USATLAS-USCMS HPC/Cloud Blueprint

- US LHC Operations Program review (Feb 2022) recommendation
- **Mandate to conduct a blueprint process about the usage of Cloud/HPC resources in U.S. ATLAS and U.S. CMS collaborations**
- Produce and submit a report by December 1, 2022 to US ATLAS & US CMS S&C ops program managers
 - S&C ops program managers will send a version to funding agencies and decide how findings will be published to a wider audience

Breakdown of the Charge

- Workflows
- Integration
- Total cost of operation
- Further R&D and development
- Recommendations to ATLAS and CMS collaborations
- International experience
- Desired facility features and policies

DOE HPC (and US Tier-1) Connectivity

- ESnet6 built physical network into each DOE national lab
 - ex: ANL, BNL, FNAL, NERSC (LBNL), ORNL
 - ESnet optical system is open, can source modems from any vendor
 - extremely cost-effective
 - ESnet6 routers on-site, can offer all services uniformly
 - site upgrades from $n \times 100\text{G}$ to $n \times 400\text{G}$ underway now
- personal opinion:
 - site infrastructure (border routers, security stuff, dtn's, storage) will be a more limiting factor than the wide-area network during ESnet6
 - wide disparity in HPC support for data-centric workflows
 - collaborate to remove technical barriers
 - increase pressure on administrative barriers



Cloud Connectivity Options Overview

Public Cloud

- Direct Peering with cloud provider's public facing network - \$
- Commodity service

Private Cloud

- why? BYO (Private) addressing, connect back to home institution, control of routing and/or tunneling, avoid public internet, QOS?
- Dedicated Interconnect (1Gbps or greater) -\$\$\$\$
- Cloud Exchange - \$\$\$
- Partner Interconnection (Hosted & Dedicated models) -\$\$



Public Cloud Interconnection Issues

- commercial interconnections require mutual agreements
- cloud providers expect commodity diurnal traffic, and "plan" accordingly
 - data science can look much more like a square wave
- physical infrastructure doesn't like surprises
 - more collaboration and multi-party communication is required
- use of regions
 - some providers accept all traffic in all locations
 - some think "we are smarter than thou" and dynamically move traffic w/o coordination
 - some providers only send/receive traffic per-region
- a soup of traffic
 - hosted applications, cat videos, HEP workflows, all mixed

ESnet Cloud connectivity today

ESnet6 built physical network into major commercial facilities

- via private fiber interconnects
 - 3x100G to Google
 - 6x100G to Oracle
- via fabrics
 - 5x100G to Microsoft
 - 5x100G to Amazon

Private Cloud

- Use case is typically administrative computing
- ESnet Partner Interconnection (Hosted & Dedicated) -\$\$
 - 10x100G (5 locations, each 2x100G) to PacketFabric
- Cloud Exchange: 2x10G, likely will be deprecated



Additional items

- API for dynamic Layer 2 circuits
- Internal automation for dynamic L3vpn instantiation
 - * new in ESnet6
- Integration with Rucio via SENSE
- DOE Integrated Research Infrastructure (IRI)
- FPGA acceleration
- Packet Marking for accounting
- DTN as a Service & Xcache nodes

Broader Ecosystem connectivity

- US University HPC sites
 - some do truly understand high-throughput jobs
 - may have more flexible security model vs DOE
 - very few are particularly well connected
- ESnet's Trans-atlantic scale
 - practical to consider using US resources from EU
- LHCONE and AUP issues
 - defacto usage of IP addresses as authorization tokens
 - R&D needed for further integrating cloud sites
 - FNAL's HEPCloud does show a working model