



---

Managed by Fermi Research Alliance, LLC for the U.S. Department of Energy Office of Science

---


## **The next 5 years: Fermilab**

Panagiotis Spentzouris

OSG Council Meeting


Oct 8, 2015

# OSG a great success for the Fermilab program




Enabling access to large numbers of cycles in burst mode

Fermilab staff working with experiments to help get their workflows OSG “enabled”



# Fermilab resources contributing to the ecosystem



# Fermilab contributing to the OSG program

---

- Providing leadership to the program
- Providing excellent Program Management support and leadership
- Security leadership and technical contributions
- Monitoring and accounting
- Operations

# The next five years

## OSG evolution (Fermilab)

# Evolution of HEP experimental program

2015-2020 (large, mid-size programs)

## **FNAL MI-LB neutrinos (IF)**

- MINOS+, MINERvA, NOvA

## **FNAL Booster-SB neutrinos**

- uBooNE, SBND, ICARUS (IF)

## **FNAL Recycler-muons (IF)**

- g-2

## **KEK-heavy flavors (IF)**

- Belle II

## **LHC beams: Run 2**

- ATLAS, CMS (EF)
- LHCb (IF)

2020-2025

...

## **Recycler-muons (IF)**

- Mu2e

## **LHC Run 3 (phase 1 upgrade)**

...

2025-...

## **Long Baseline Neutrino Facility (LBNF)**


- DUNE (IF)

## **LHC Run4 (HL-LHC)**

...

# Future computing requirements: Compute and Data Needs


- Two new programs coming online (**DUNE**, **High-Luminosity LHC**), while new physics search programs (**Mu2e**, **Belle2**) will be operating
- Increased precision & event complexity, higher luminosity, will push computing needs to **~10X-100X of current HEP capabilities**
  - Lower value assumes optimized algorithms and new approaches*



Year	2013	2014	2015	2016	2017	2018	2019	2020	2021	2022
Tape [PB]	2.8	2.8	2.8	2.8	19.24	54.43	103.55	153.89	204.64	255.39
Disk [PB]	4.00	4.00	5.00	8.00	27.98	79.17	115.68	153.10	190.82	228.55
CPU [kHepSPEC1]	45.00	45.00	50.00	55.00	328.31	568.98	567.54	609.45	643.14	672.60

Belle2 computing requirements

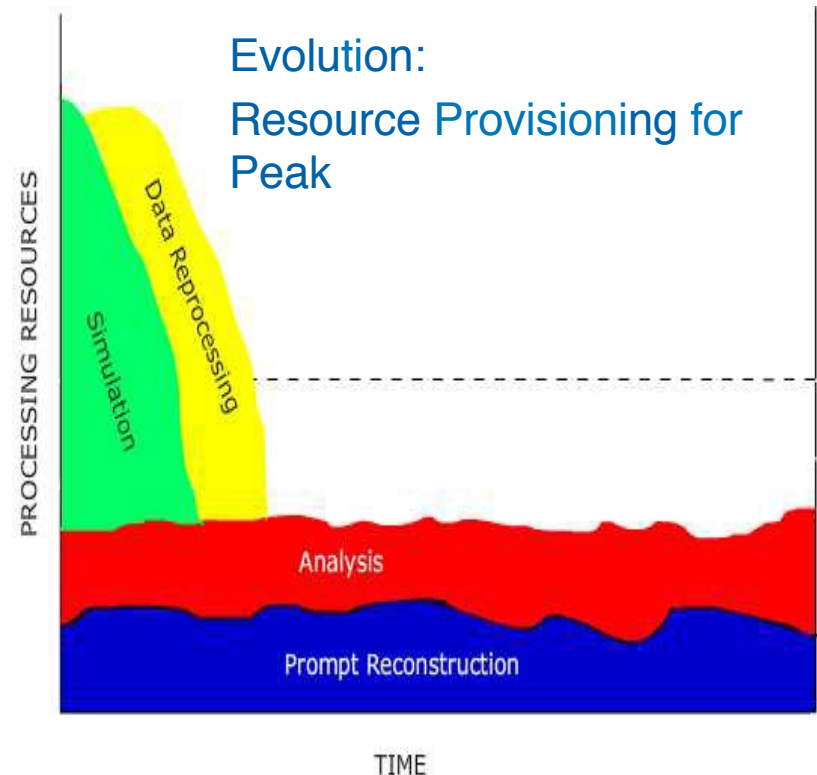
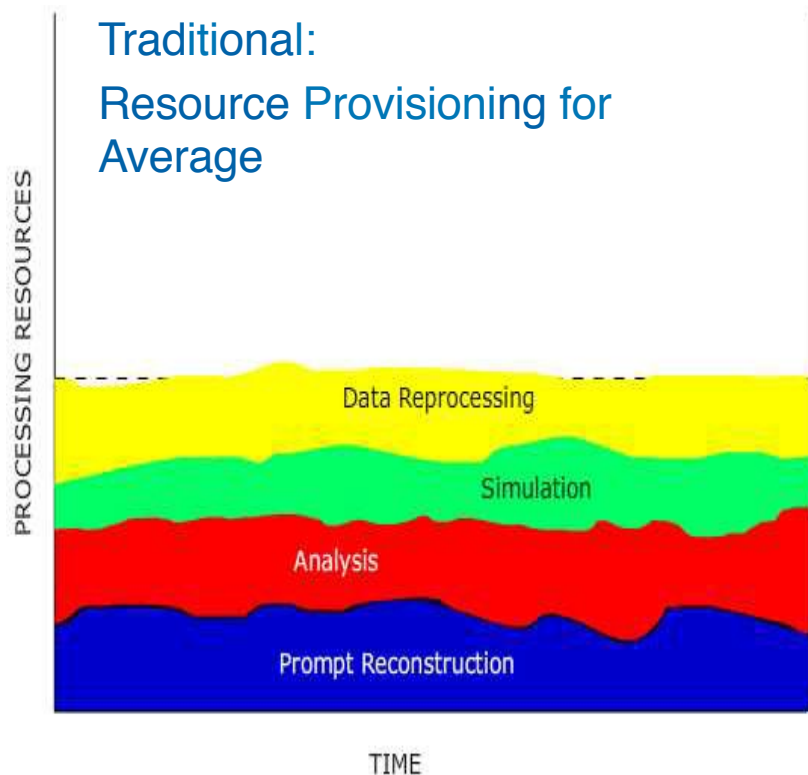
# LHC expected data volumes



- LHC Run 4 will start the exabyte era for HEP!
  - Derived data (reco, simulation) x8 of RAW...
- How do we process and analyze all these data?



# Evolving provisioning model (for efficiency)



- Provisioning needs to be adaptable, providing facility “elasticity”
- ➔ go beyond our current “elasticity” model (opportunistic resources):  
***new paradigm, incorporate and manage*** “rental” resources
  - including specific tape storage resources and capabilities such as archiving

## Fermilab: the next five years

---

- Would like to lead the effort in developing the concepts and deploying the “next generation” computing facility for HEP
  - “HEPCloud” concept and “HEPCloud” project as OSG components
- In an environment of many (tightly) linked “grants” we should also consider how to evolve cross cutting activities
  - Security, project management, monitoring, ...
    - Only listing Fermilab traditional contributions
- Fermilab will lead security activity efforts in the context of HEPCloud and contribute effort and expertise in program-wide security activities
  - Will need to define a cross-cutting entity to facilitate such contributions, since individual grants most likely won’t be able to cover the full spectrum

## Fermilab: the next five years

---

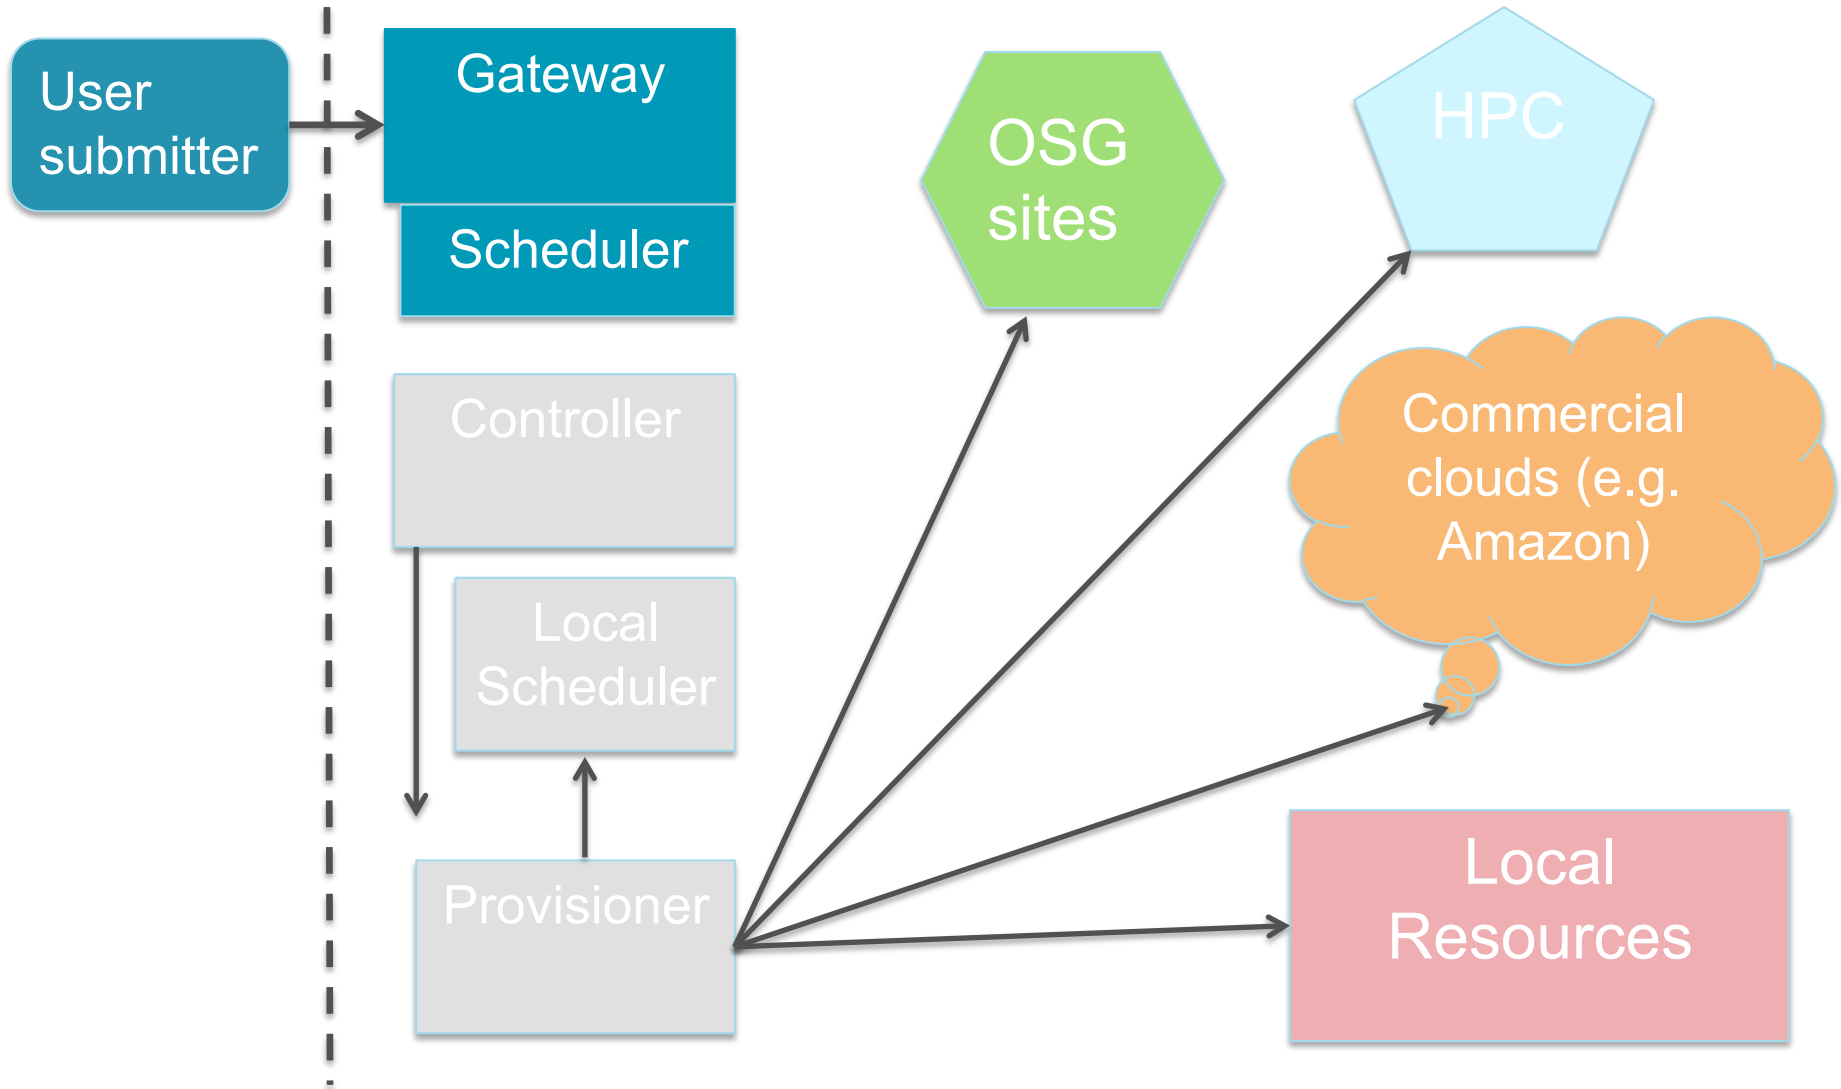
- In a similar fashion, we will contribute to program management (through the HEPCloud activities)
  - Federated program management sounds like a bad idea, obviously program wide coordination has to be present, but could be light-weight if “grants” coordinate and contribute.
- Would like to continue evolving monitoring capabilities, in collaboration with all other “grants” for program wide solutions

# HEPCloud Facility concept

---

- The HEPCloud Facility is a portal to an ecosystem of computing resources, commercial or academic.
- Provides “complete solutions” to all users, with agreed upon levels of service
  - The Facility routes to local or “rental” resources based on efficiency, cost, workflow requirements and target compute engine policies.
  - Manages allocations of users to target compute engines
- Provides storage services appropriate to the system that the workflow is routed

# HEPCloud Facility concept



## Example: the Fermilab HEPCloud Facility project

---

- The goal is to integrate “rental” resources into the current Fermilab computing facility in a manner transparent to the user. Objectives include
  - A seamless user environment for all resource types, including necessary tools and infrastructure
  - The architecture, including network, needed to support required data rates.
  - The policies and middleware for efficiently using and prioritizing the use of different resources
  - The information security policies, procedures and monitoring.
- Partnership with resource providers necessary to understand requirements, architecture and software development needs

# Evolving HEP computing facilities: the Fermilab HEPCloud Project

---

- Goals: (from Charter)
  - The goal of the Fermilab HEP Cloud Facility Project is to extend the current Fermilab Computing Facility to transparently run on disparate resources.
- Scope: (from Charter)
  - The project scope for the Facility includes at least the abilities to:
    - Manage access to the facility
    - Execute production workflows
    - Monitor operation of the facility
    - Extend the resources of the facility
    - Financially account and plan for use of the facility
    - Manage contractual and similar agreements related to the facility
  - Includes demonstrating that this works for production workflows for at least:
    - CMS experiment
    - NOvA experiment
    - DES experiment
    - NOvA-Big Data workflow

## NoVA Processing

Processing the 2014/2015 dataset

16 4-day “campaigns” over one year

**Demonstrates stability, availability, cost-effectiveness**

Received \$30,000 academic grant from Amazon Web services

## Dark Energy Survey - Gravitational Waves

Search for optical counterpart of events detected by LIGO/VIRGO gravitational wave detectors (FNAL LDRD)

Modest CPU needs, but want 5-10 hour turnaround  
Burst activity driven entirely by physical phenomena (gravitational wave events are transient)

**Demonstrates provisioning to peak**

## CMS Monte Carlo Simulation

Generation (and detector simulation, digitization, reconstruction) of simulated events for Run 2 after beam spot determination

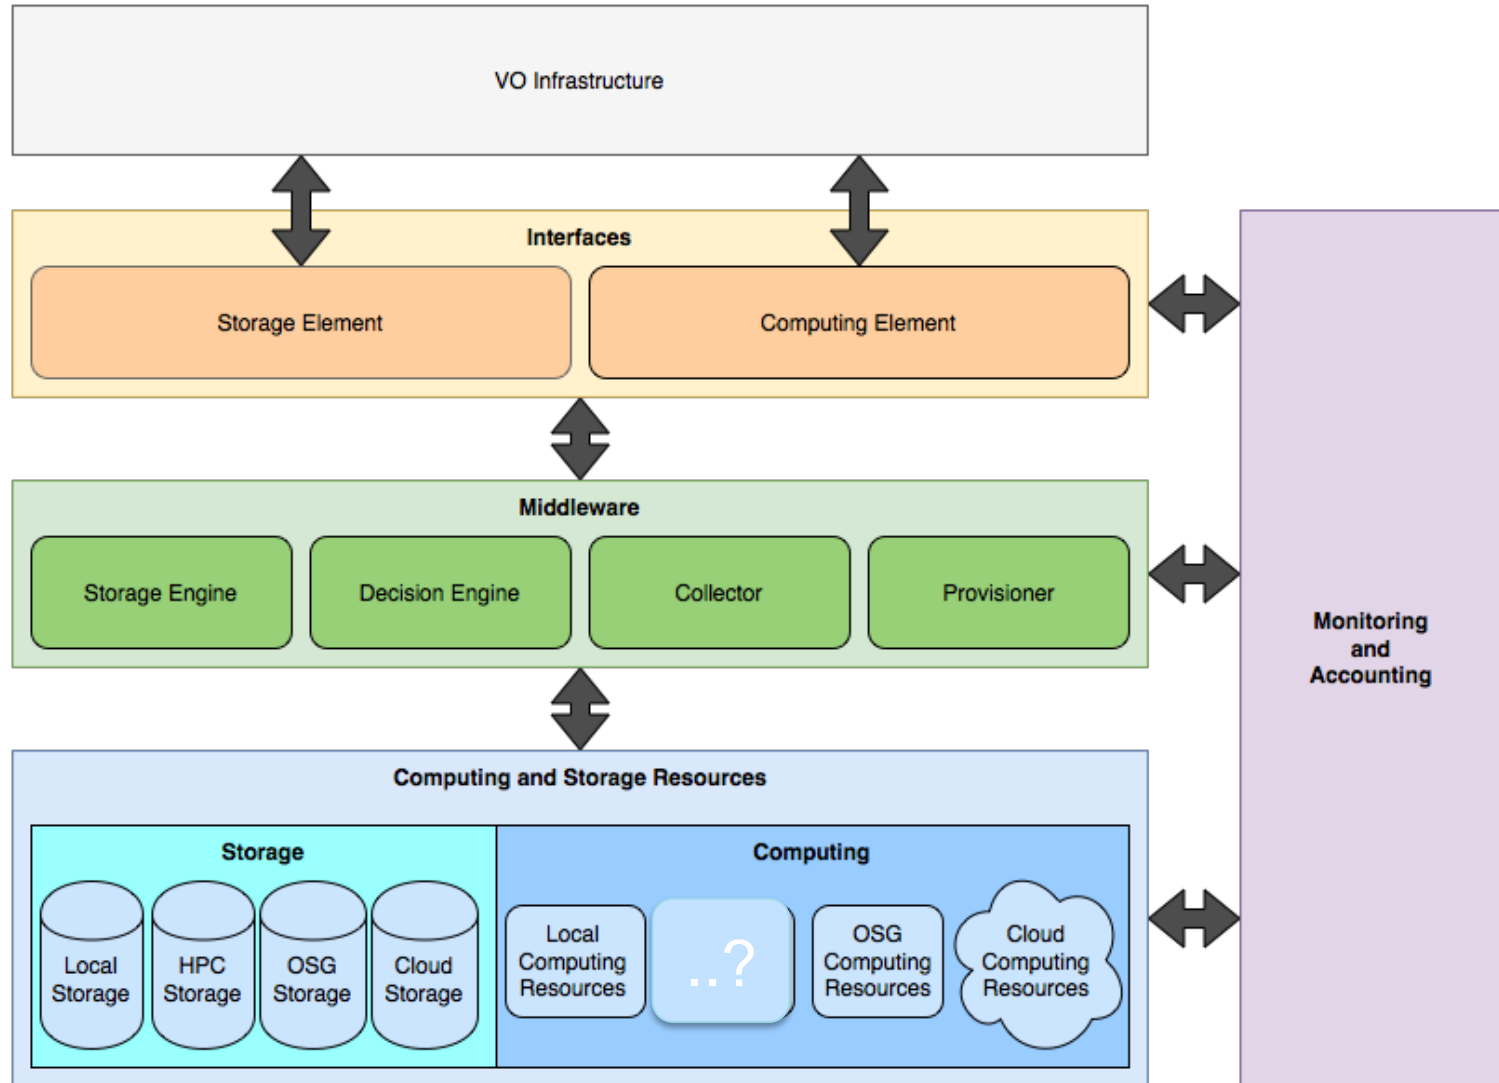
56000 compute cores for 1 month, steady-state

**Demonstrates scalability**

Received academic grant covering 90% of incurred costs



# HEPCloud Architecture



## DOE HPC as a resource

---

- A very appealing option is to consider ASCR HPC facilities as potential compute resources in the ecosystem
- Identify use cases with workflows that allow such utilization within the constraints of allocation, security and access policy of HPC facilities
  - Some work already done (ATLAS event generation on ALCF)
- Develop partnerships with ASCR to fully understand constraints and requirements and develop the necessary process, policies and tools necessary for HPC access (through HEPCloud)
  - FNAL has an established relationship with ASCR and HPC
  - LQCD, Cosmology, Accelerator Modeling, Data Management (through SciDAC), Networking (ESnet, other projects)

# Data management and storage as a resource

---

- Fermilab provides scientific data management capabilities developed by the HEP community
  - The “active archive infrastructure” technologies utilize the wide-area transfer protocols and cached storage systems at Fermilab
- These services would be integrated in HEPCloud, for HEP experimental program
- The Fermilab Active Archival Facility currently provides these services to other science activities in the US, with cost recovery and through specific agreements
  - Simons Foundation, U. of Wisconsin
- Understand how (if) this model works in the evolving ecosystem

# Summary

---

- OSG has been a valuable resource to the Fermilab HEP program (and the HEP program overall)
  - and, I believe, Fermilab an (at least) useful contributor to the OSG program
- Fermilab will push the evolution of HEP facilities to respond to the significant computing challenges of the evolving HEP program
  - Would like to do that in the context of the OSG, and in collaboration with OSG participants