Fermilab

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

Exascale and Exabytes: Computing at Fermilab for LHC and Intensity Frontier

Oliver Gutsche

Seminar at Universidad de los Andes - Bogota, Colombia

14. August 2015

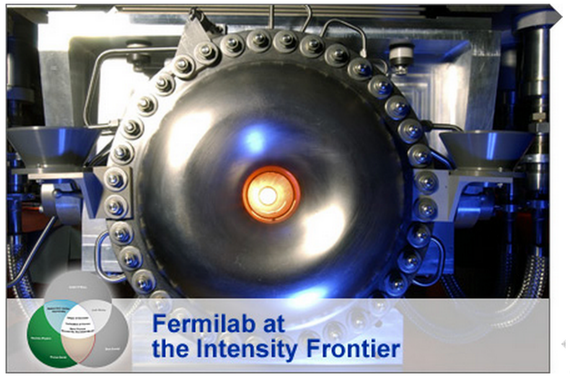

Fermilab at the Intensity Frontier

At the Intensity Frontier, scientists explore new physics in unprecedented breadth and detail using intense beams of neutrinos, muons, kaons and nuclei.

Fermilab at the Energy Frontier

At the Energy Frontier, scientists discover new particles and directly probe the architecture of the fundamental forces using high-energy colliders.

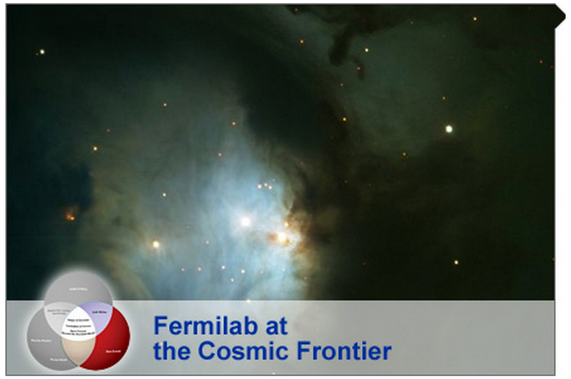

Fermilab at the Cosmic Frontier

At the Cosmic Frontier, scientists reveal the nature of dark matter, dark energy and other phenomena of new physics using particles and light from the cosmos.

Neutrino Experiments

- Detect Neutrinos and measure their masses
- Important Standard Model measurement, candidates for dark matter
- Produce neutrino beam and direct it to far detector, compare near and far detector measurements

- Current experiment: NOvA
  - 500 miles to Ash River, Minnesota
- Future experiment: DUNE/LBNF
  - 800 miles to Sanford Underground Research Facility, in Lead, South Dakota

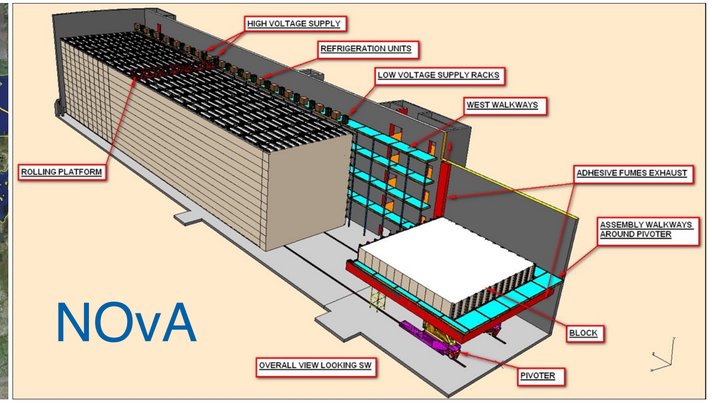

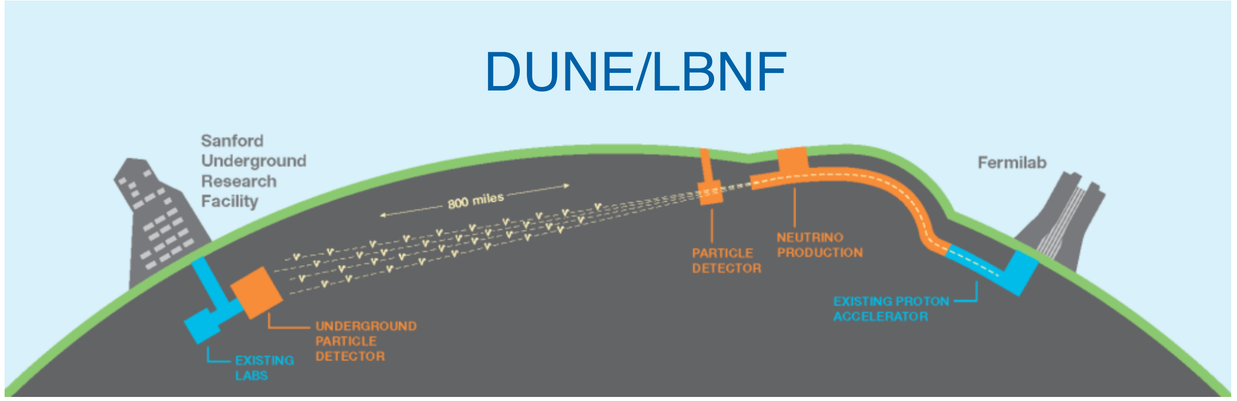

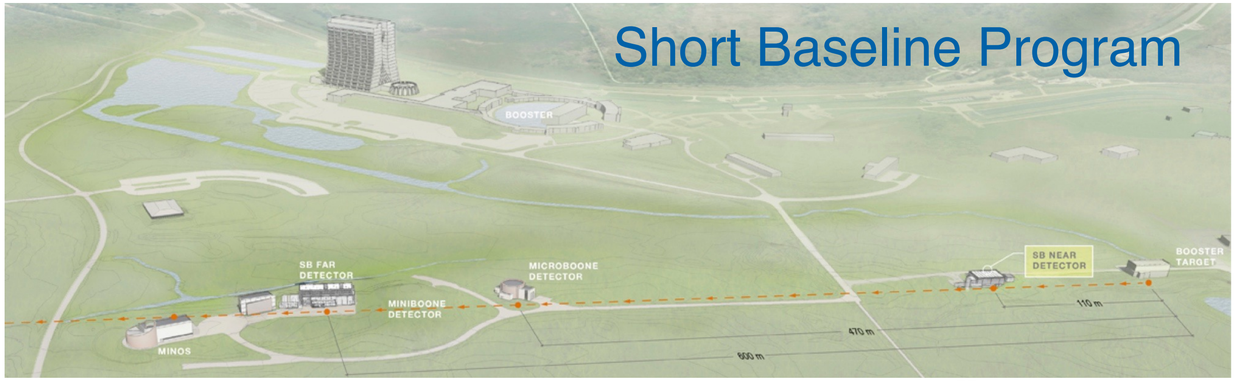

Energy Frontier

- Till September 2011: Tevatron at Fermilab
  - Circumference 6 km
  - Collision energy: ~2 TeV

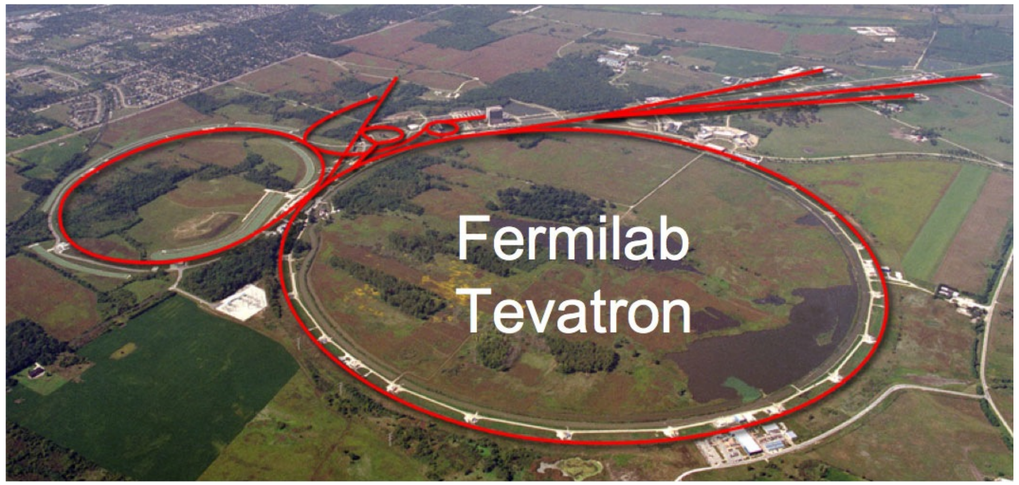

- Now: Large Hadron Collider (LHC) in Geneva, Switzerland
  - Circumference 27 km
  - Collision energy ~13 TeV
  - Fermilab is involved in the CMS experiment and contributes to the accelerator

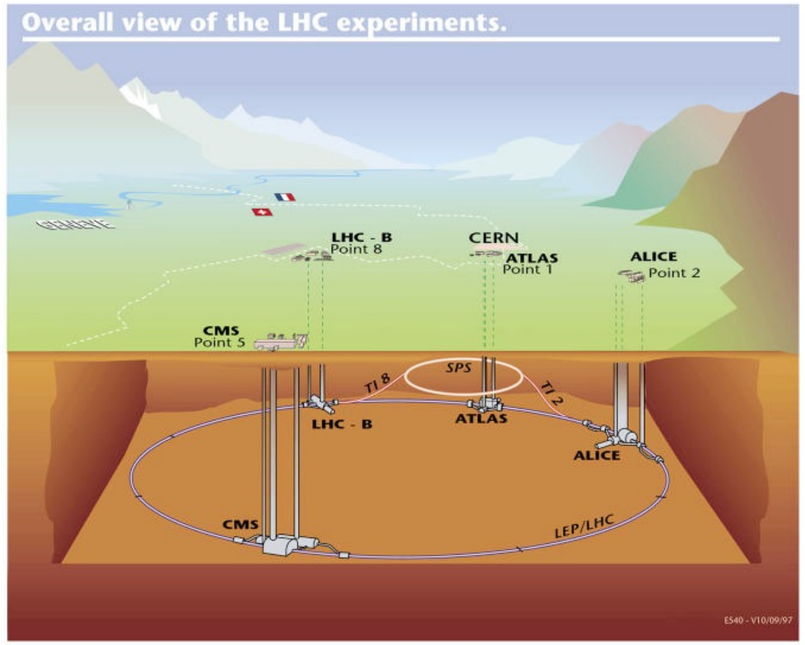

Cosmic Frontier

- Use the cosmos as a laboratory to investigate the fundamental laws of physics.
- Use detectors to study particles from space as they approach and enter our atmosphere
  - Cosmic rays, gamma rays and neutrinos emitted by the sun

COUPP

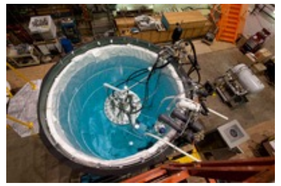

SDSS

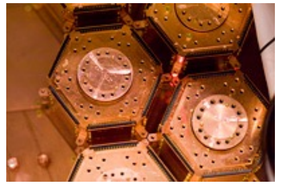

CDMS

Pierre Auger

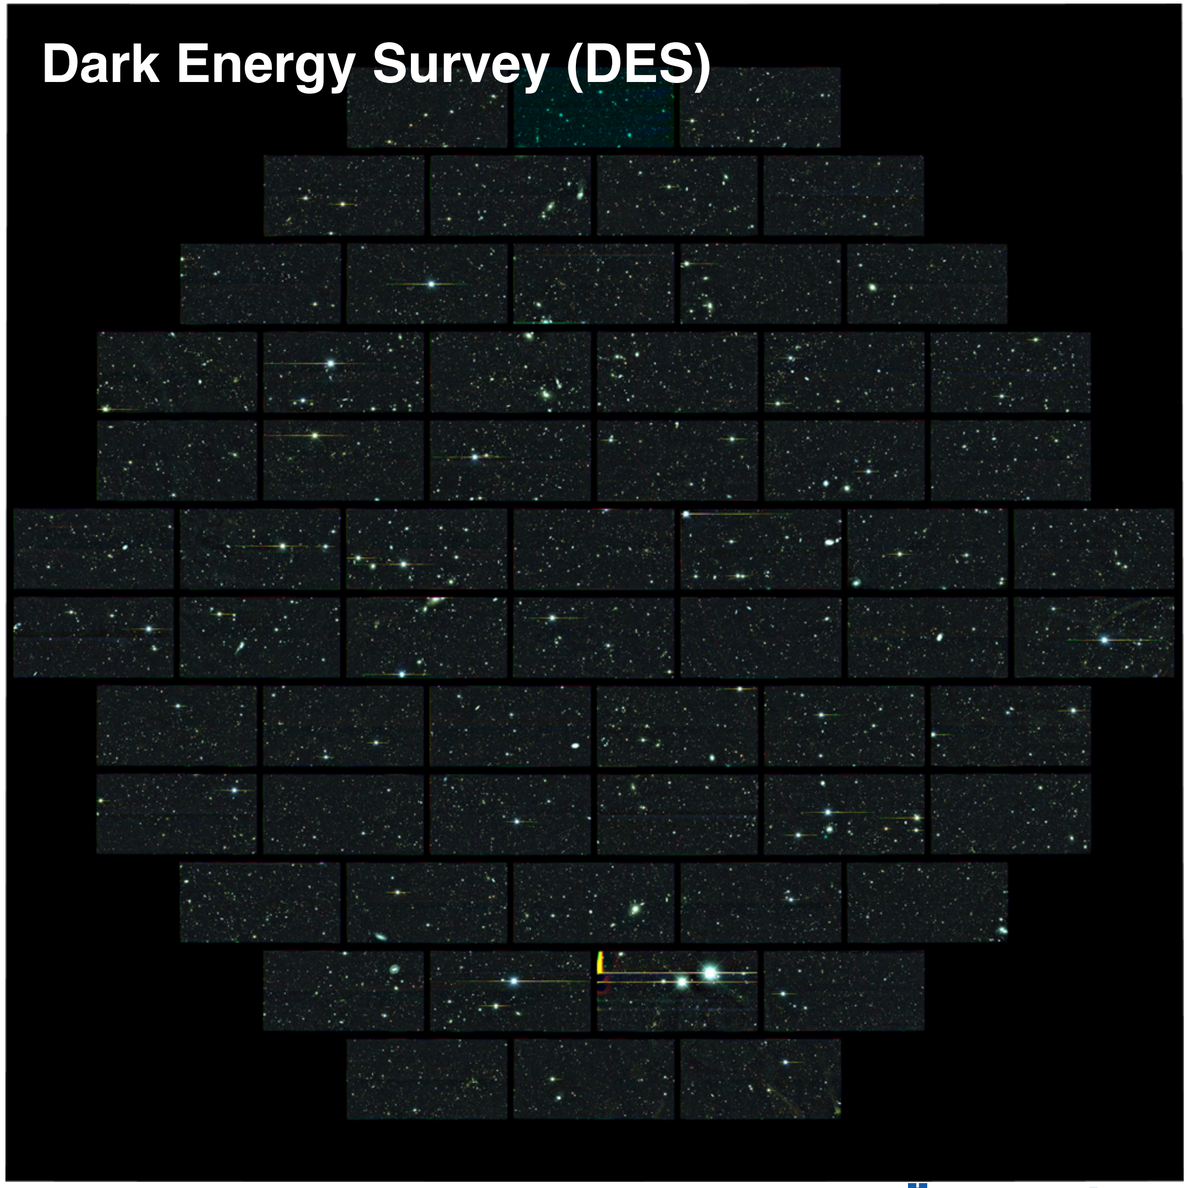

Fermilab Computing

- Provide and manage computing services and resources
- Data recording, storage, access
- Bulk processing, analysis
- Functionality analogous to LHC Tier-0 and Tier-1
- CPU Cores, Online (Disk) and Offline (Tape) Storage, Networking

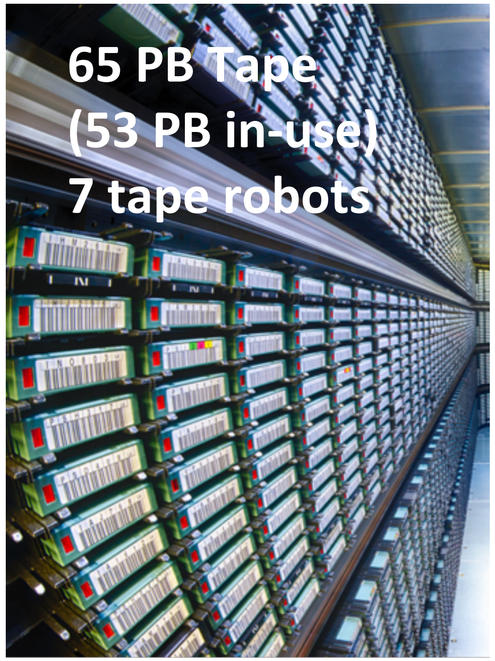

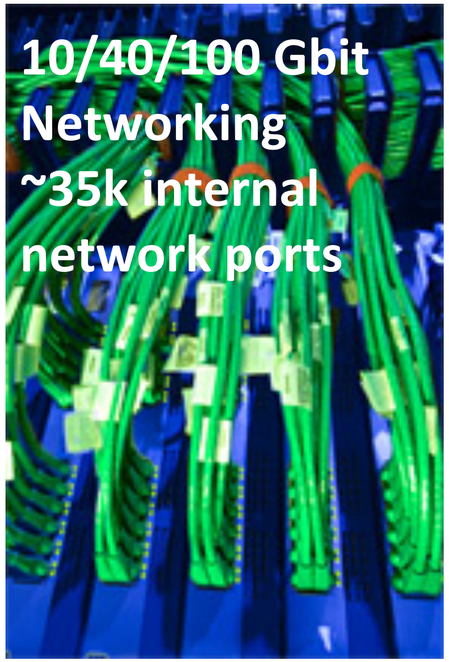

Computing in HEP

The Scientific Process

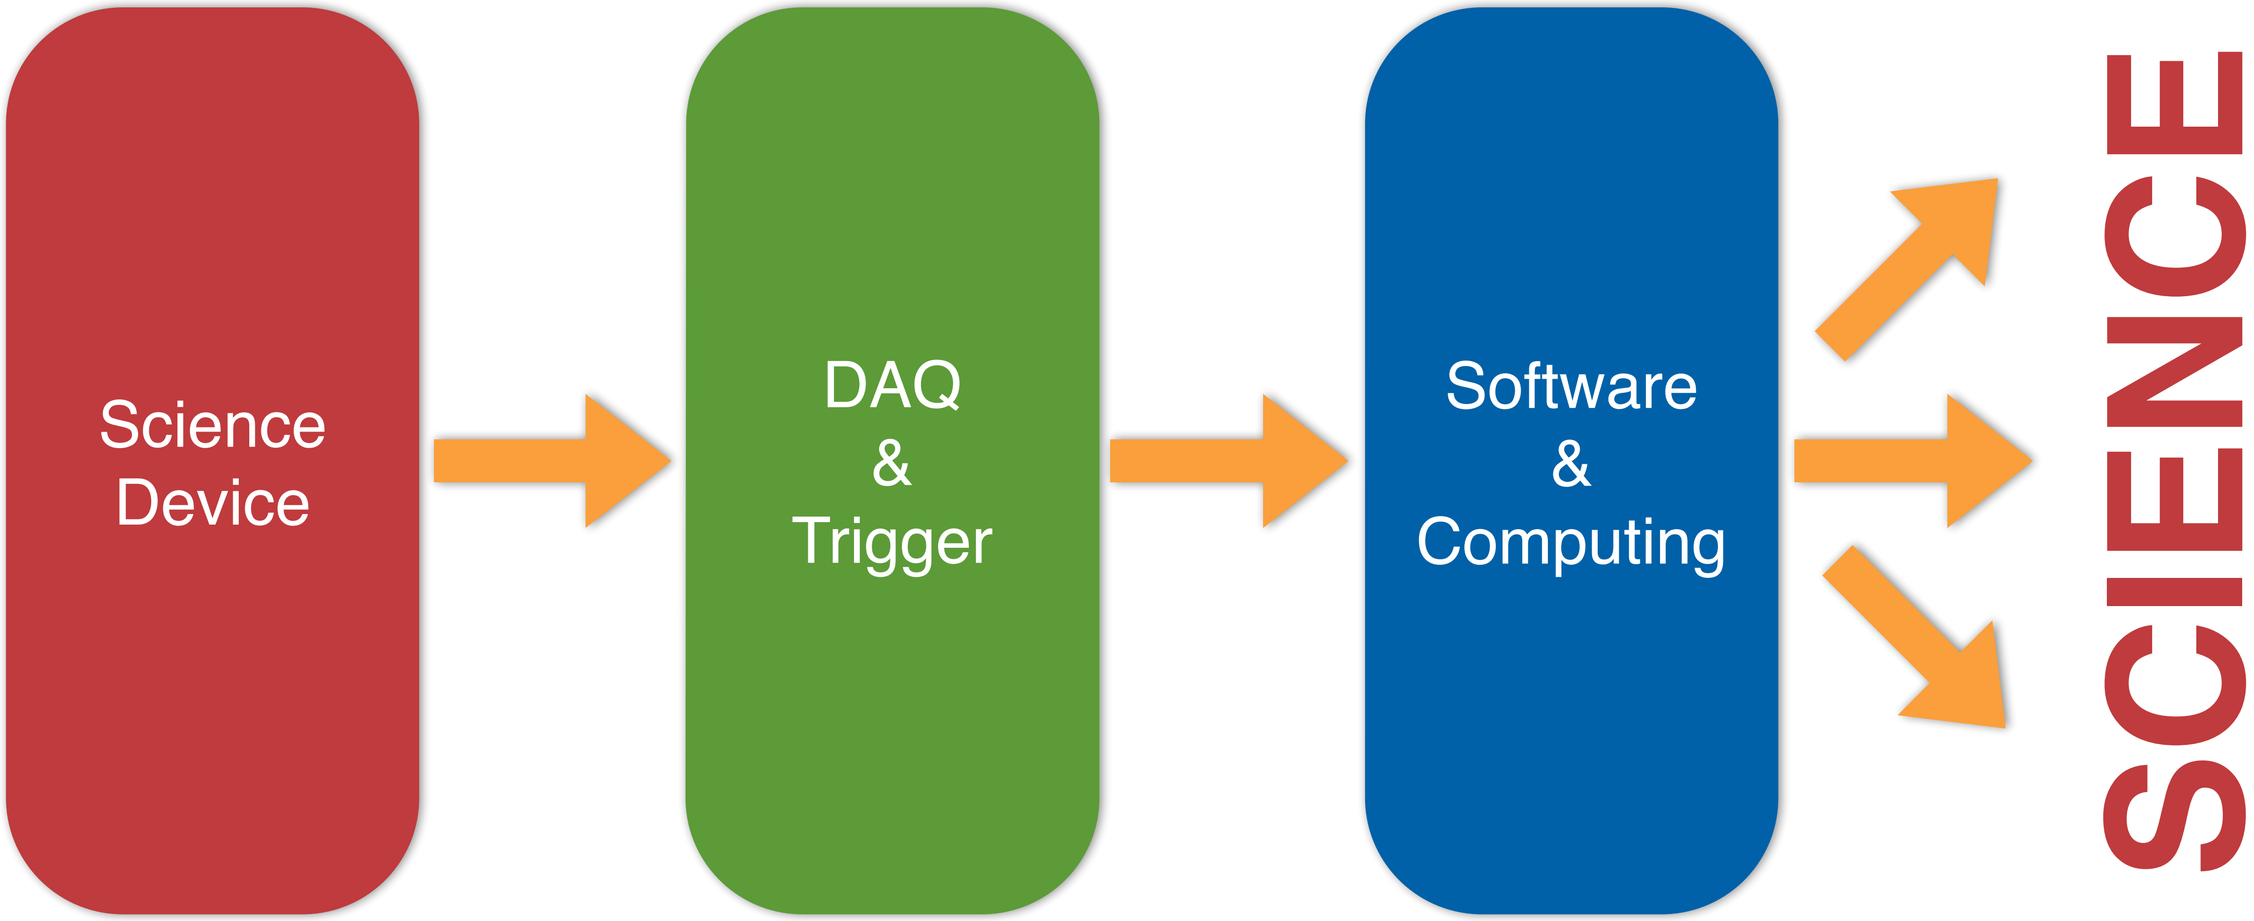

- Software & Computing is an integral part of the scientific process

Software & Computing

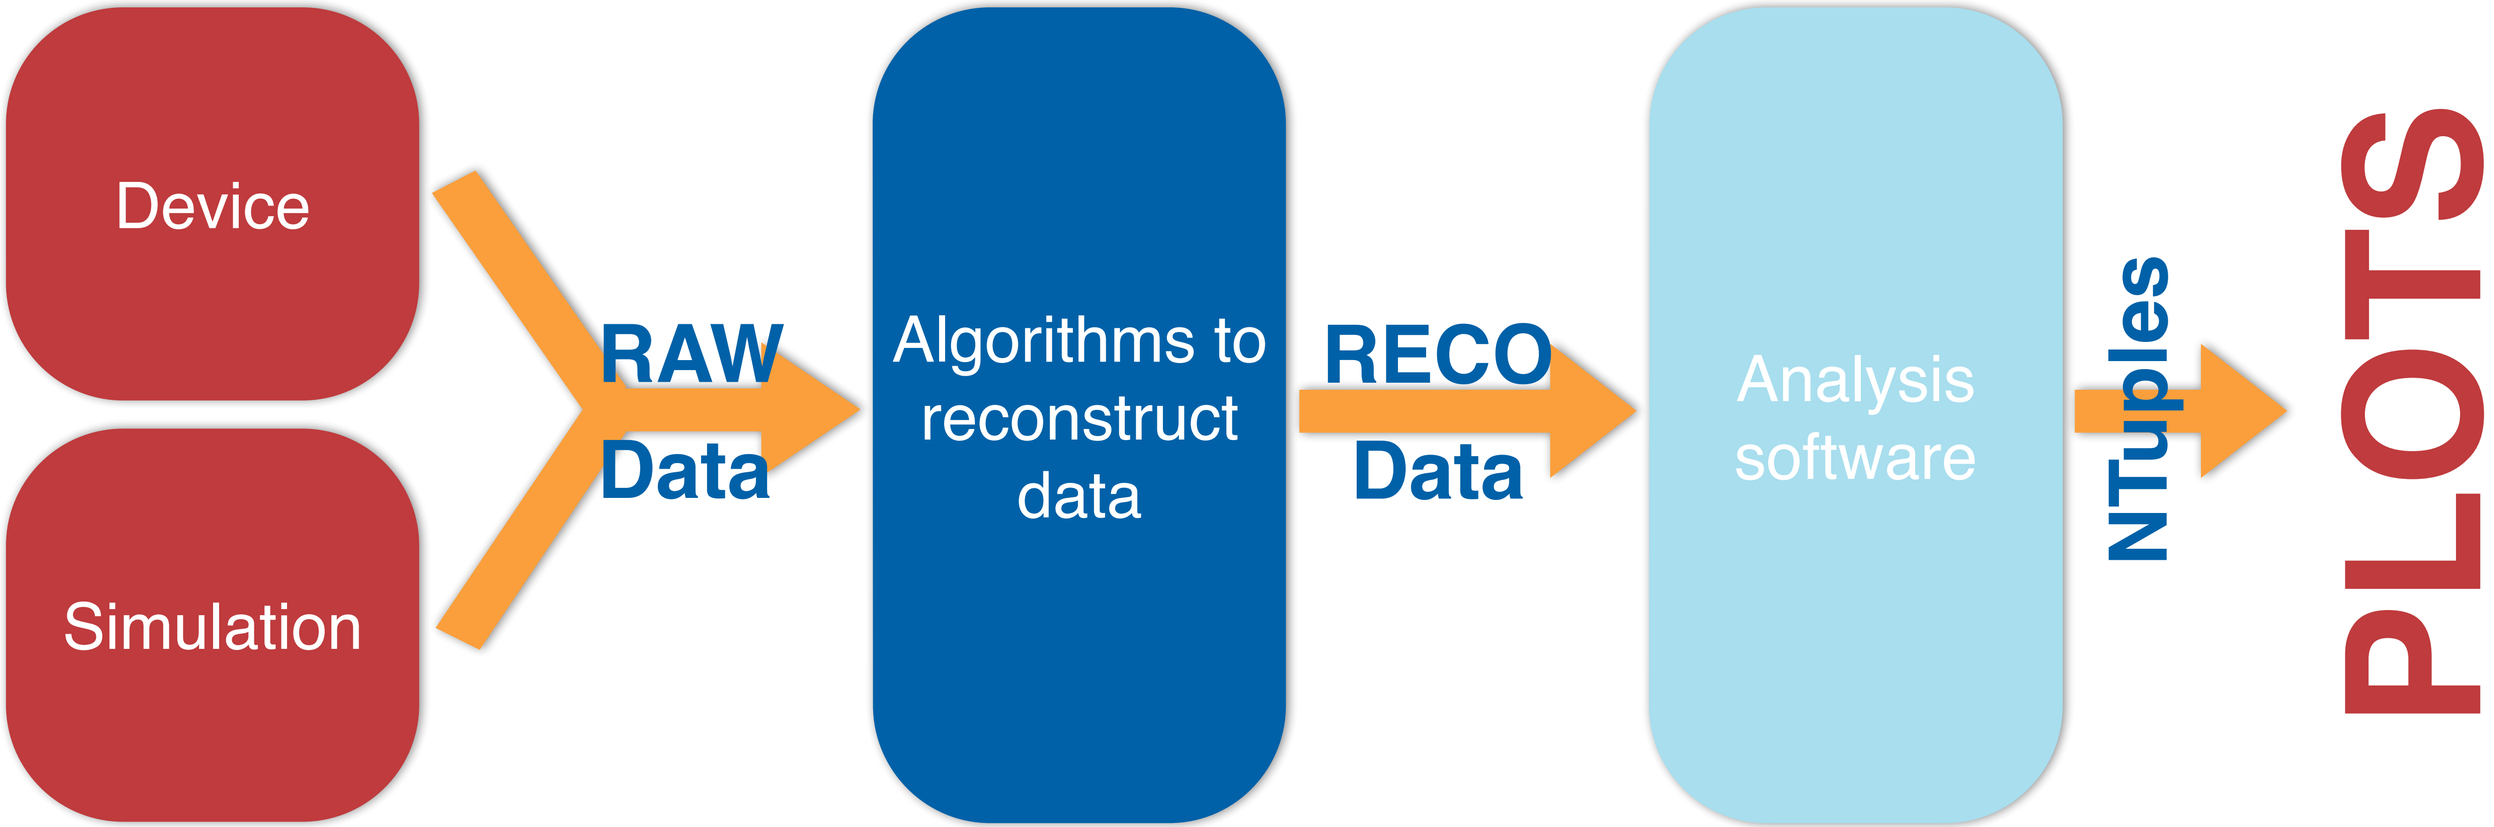

- **Software** is important for every step on the way to scientific results

Software & Computing

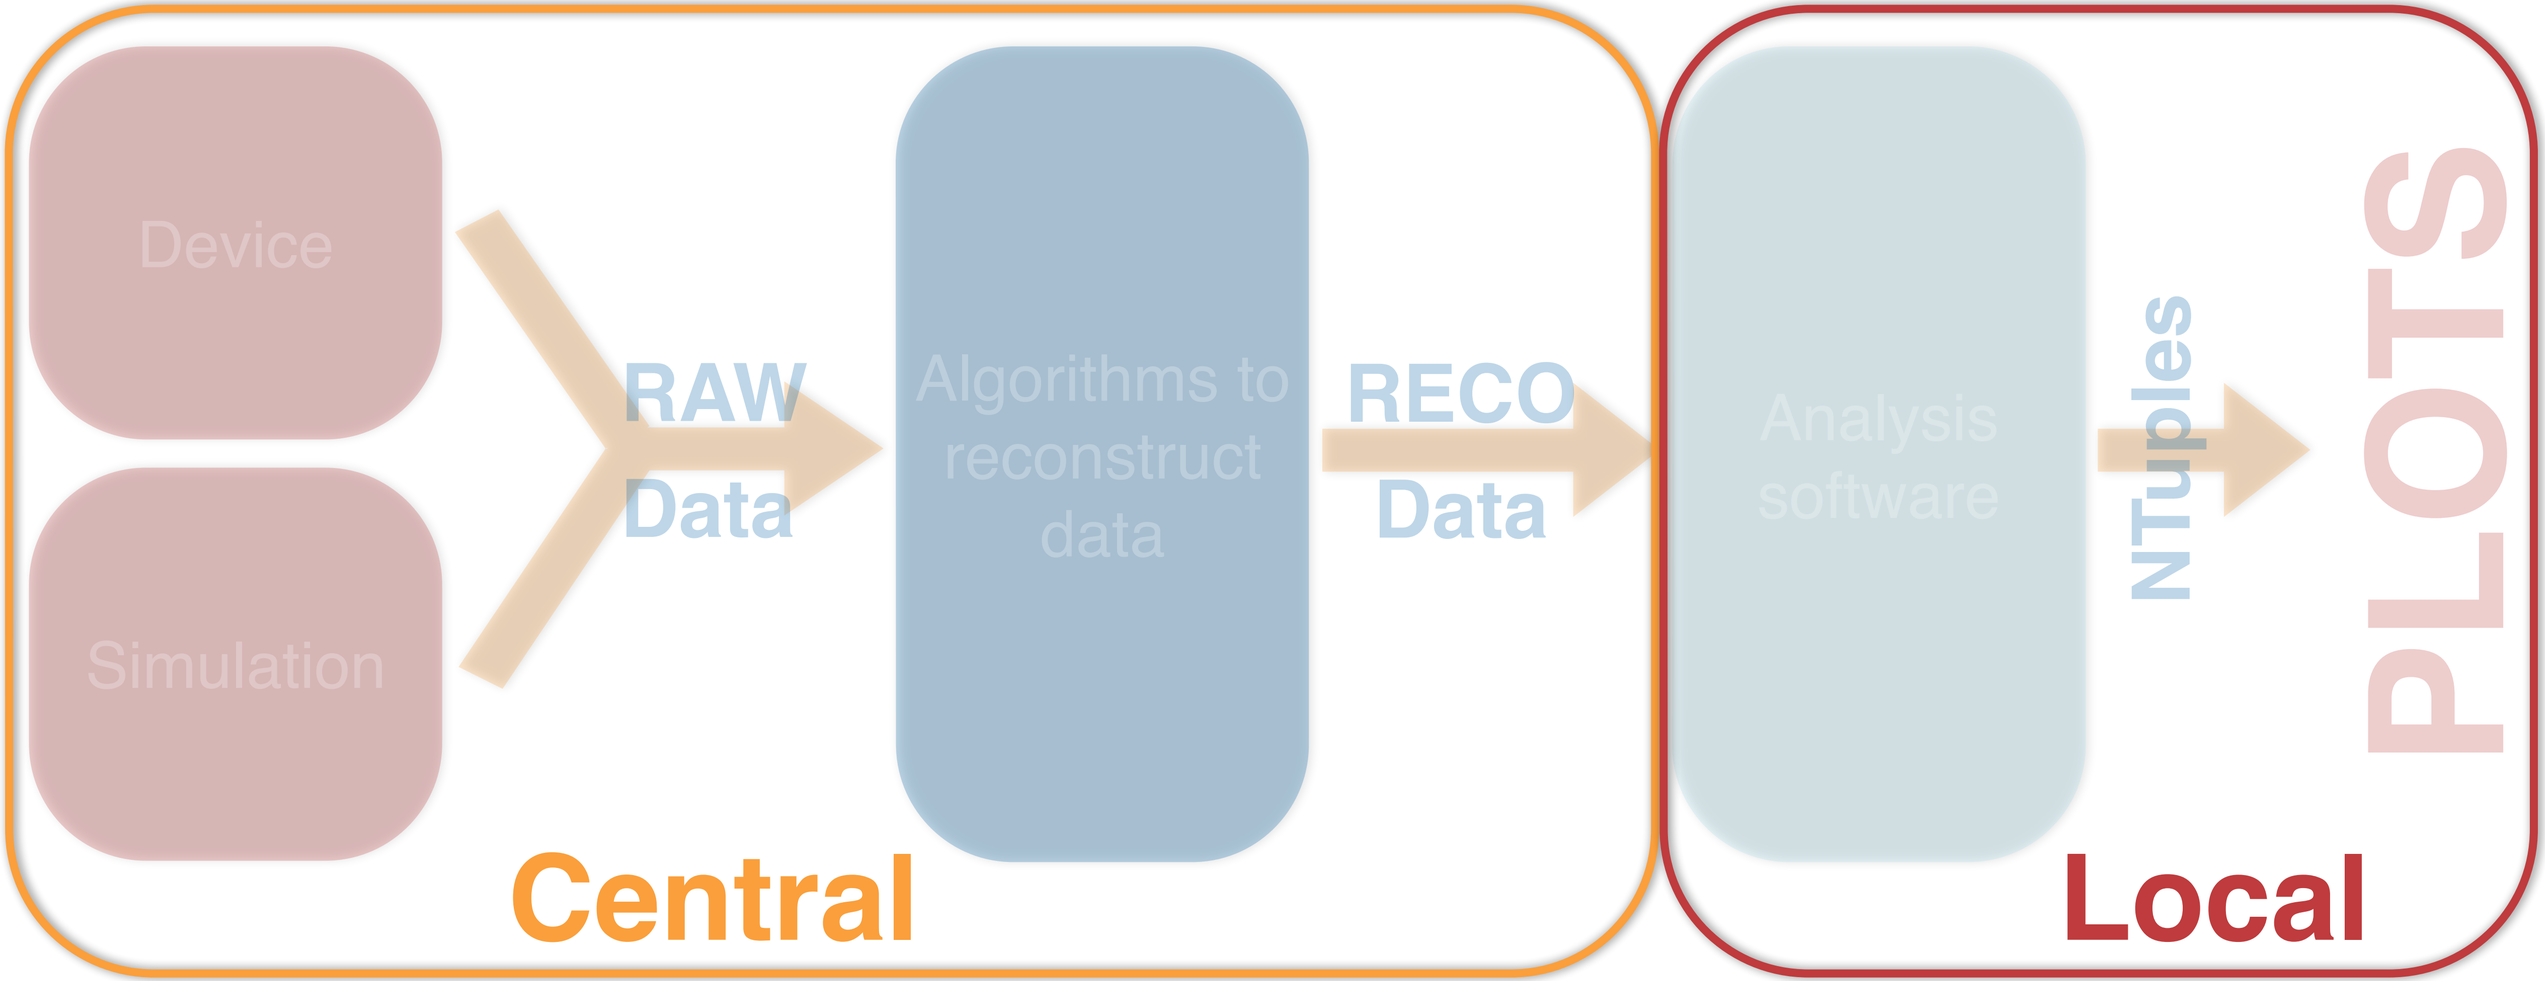

- Computing resources (**Storage and Network**, **Processing**, …) are needed for all steps

Computing at Fermilab

- Software and computing are needed at all steps from data collection to physics results of any experiment
- Performant and scalable solutions are required to extract physics results quickly and efficiently

- Fermilab is a world leader providing solutions for experiments, we concentrate on:
  - CMS, NoVA, MicroBooNE, MINOS, Mu2e, Muon g-2, DES, and other Fermilab experiments and projects and HEP community at large

Computing Lifecycle: CMS

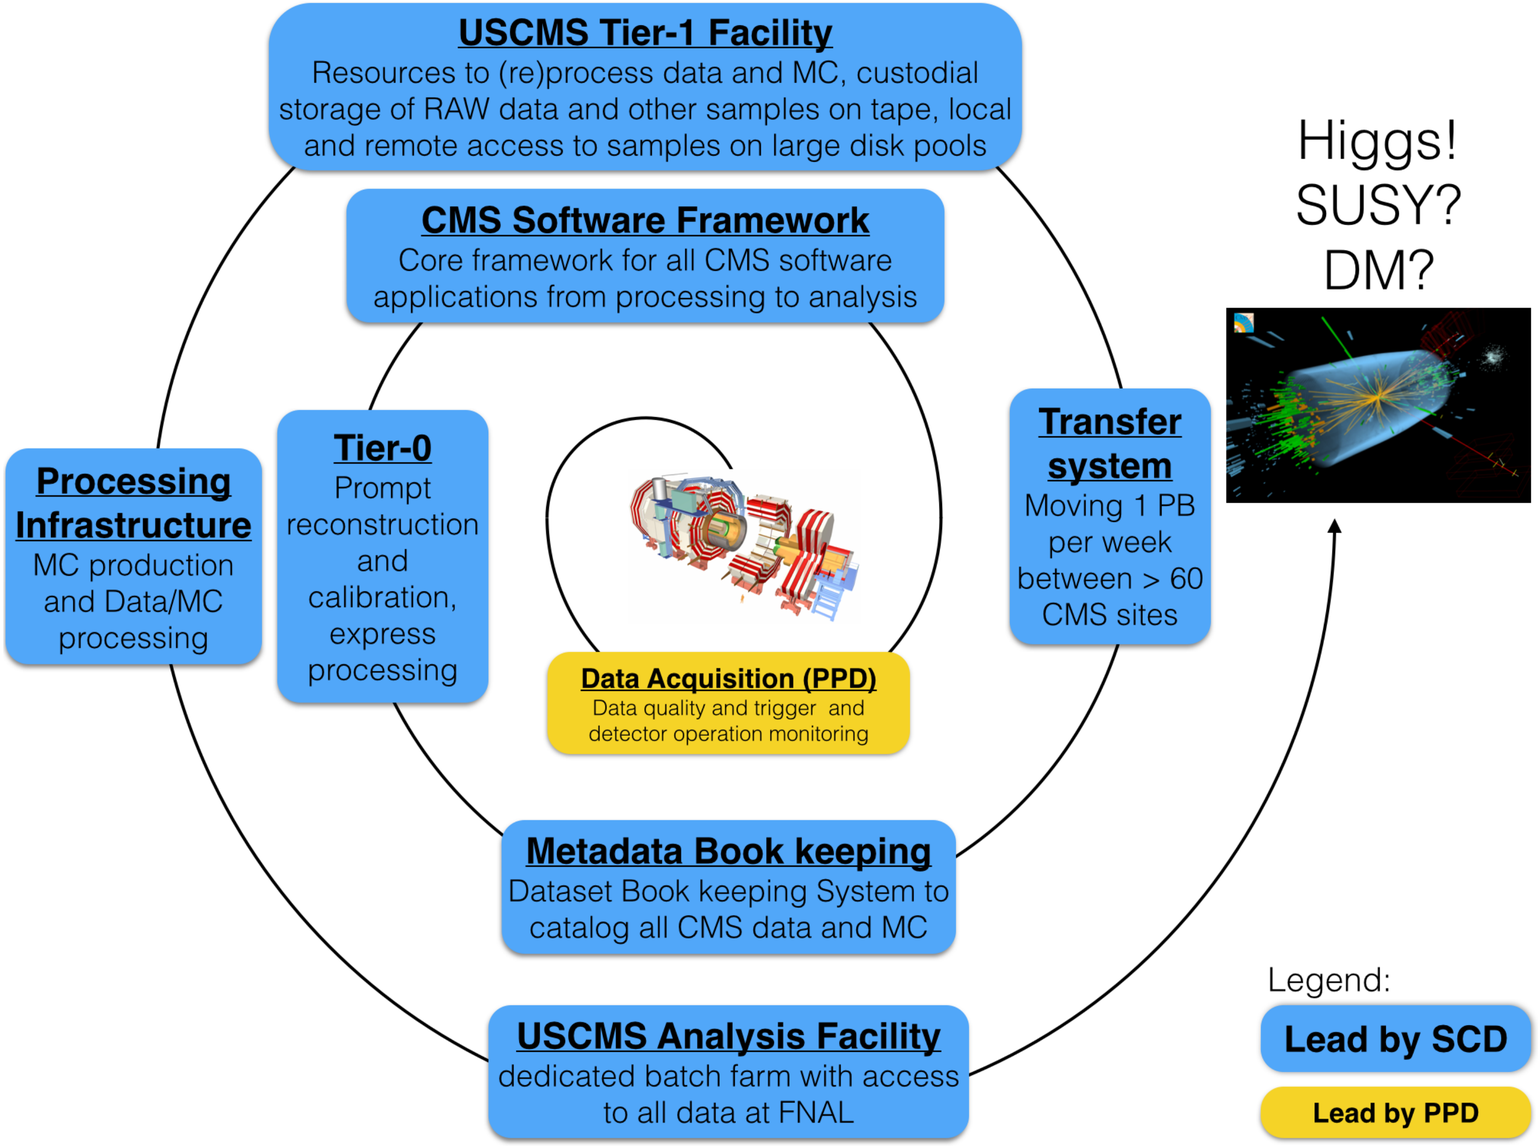

Computing Lifecycle: NOvA

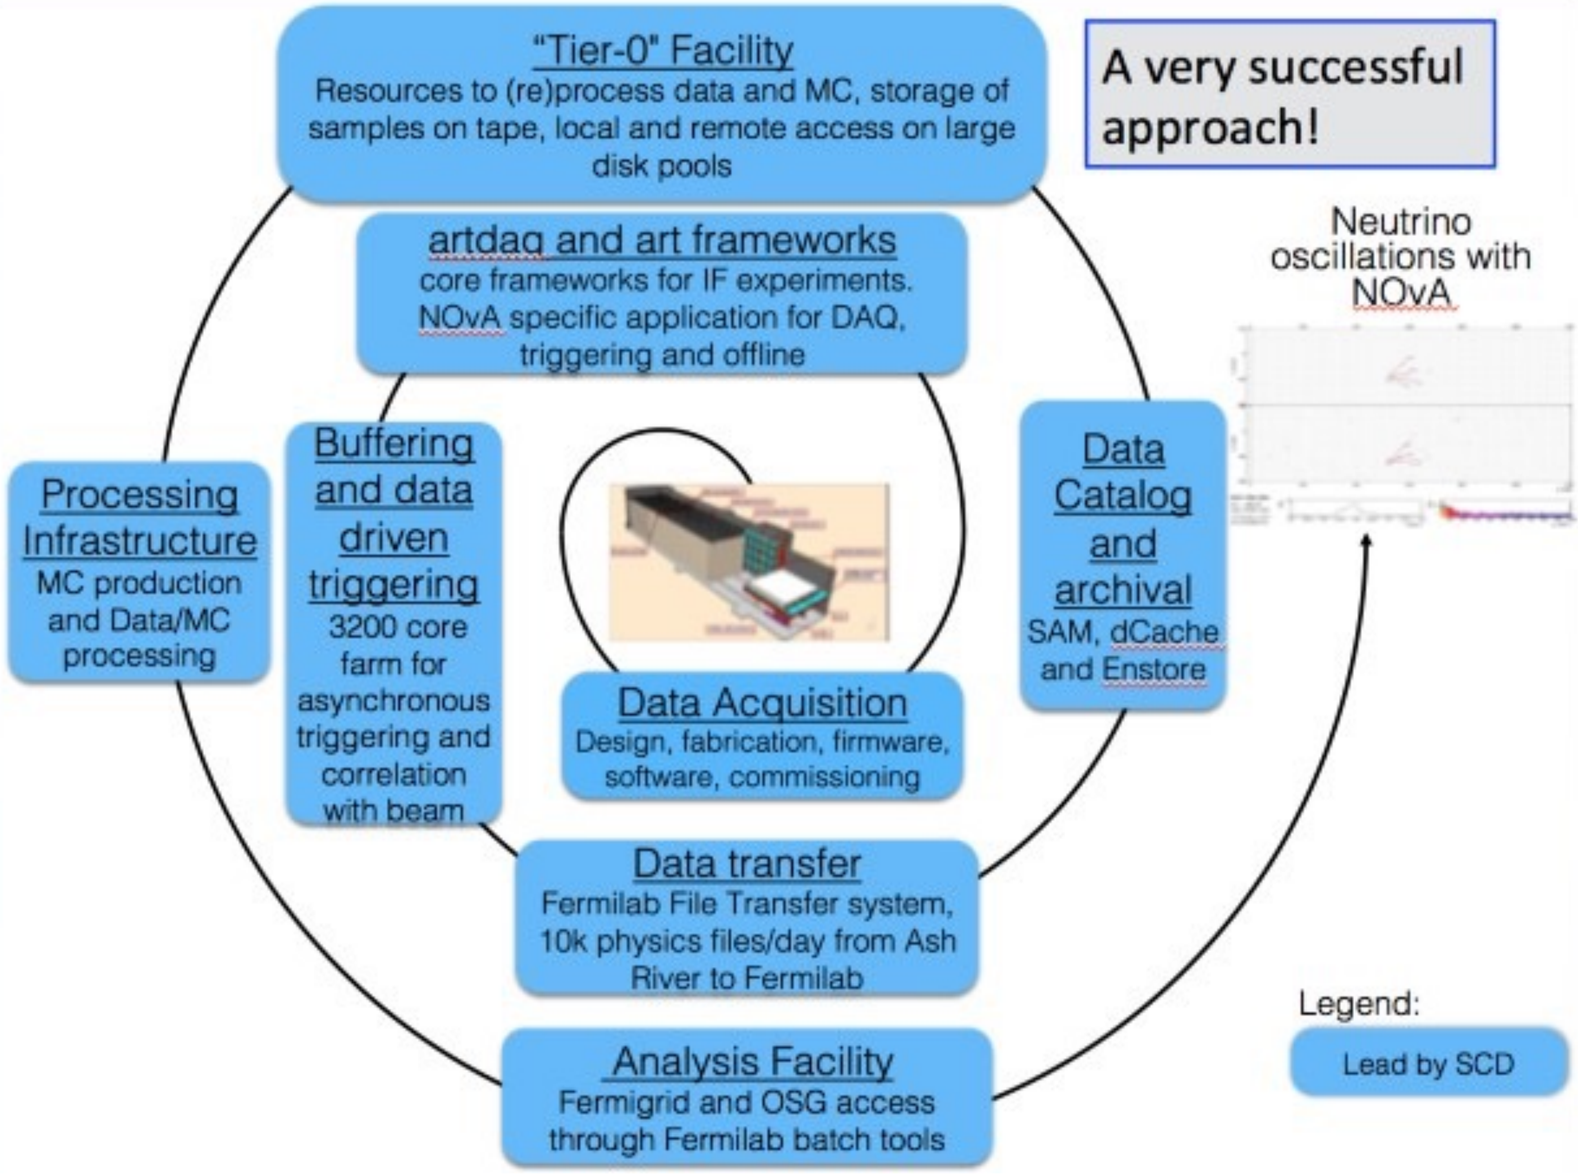

FIFE

Fabrlc for Frontier Experiments

NoVA, MicroBooNE, MINOS, Mu2e, Muon g-2, DES, …

- **Approach:**
  - Support offline & computing needs of experiments
    - We provide standard interfaces into Fermilab computing resources (CPU, disk, tape)
    - We enable science through a modular toolkit of services and applications

- **Goal:**
  - Help experiments focus on their science
    - Provide infrastructure to experiments that get them to computing resources
    - Provide software framework solution and support community-based reconstruction software
  - Utilize previous solutions and integrate everything into seamless model

- **FIFE project - Strategy:**
  - Address all of the computing needs for experiments
    - Modular enough so that experiments can take what they need
    - Spans all cross cuts, in particular the Neutrino, Muon and Cosmic cross cut
  - Provide mechanism for feedback from experiments to incorporate their tools and solutions
  - Help experiments utilize computing beyond the Fermilab campus
  - Integrate new tools and resources from outside Fermilab and other communities as they develop
    - **Synergy with CMS especially important**

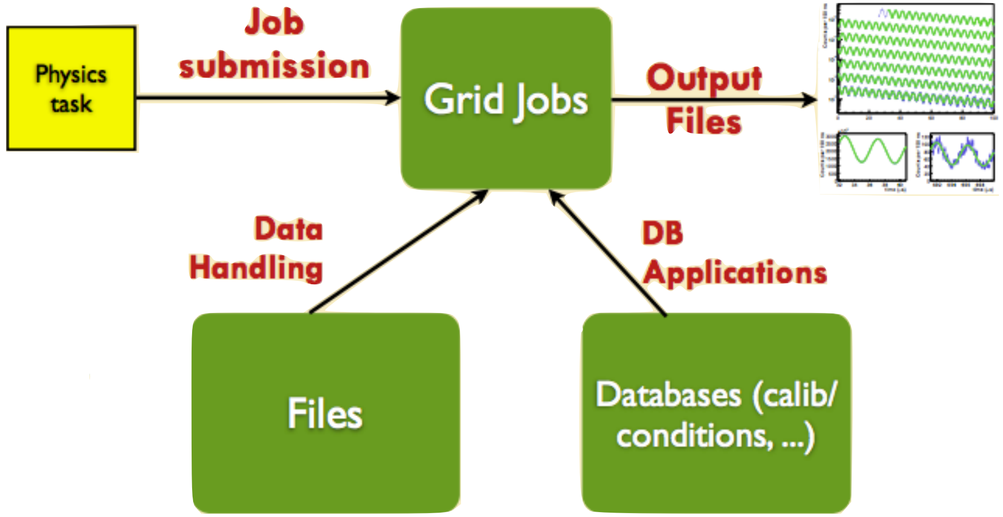

Common Services and Projects toolkit

- FIFE provides access and support for common tools in:

  - **DAQ and Controls**
    - DAQ Status, DAQ Framework
  - **Grid and Cloud**
    - Fermigrid, Fermicloud, Gratia, Job Sub, FIFEMON, OSG, Amazon Web Services (AWS)
  - **Scientific Data Storage and Access**
    - dcache/ enstore, Gridftp
  - **Scientific Data Management**
    - IFDH, SAM Web, F-FTS
  - **Scientific Frameworks and Software**
    - Art, Larsoft
  - **Physics and detector simulation**
    - Genie, Geant4
  - **Databases**
    - IF Beam, Custom Databases, mysql, postgres
  - **Scientific Computing Systems**
    - CVMFS, Interactive machine in GPCF, Experiment Control Rooms
  - **Scientific Collaboration Tools**
    - Redmine, CVS/ Subversion/ Git, ECL

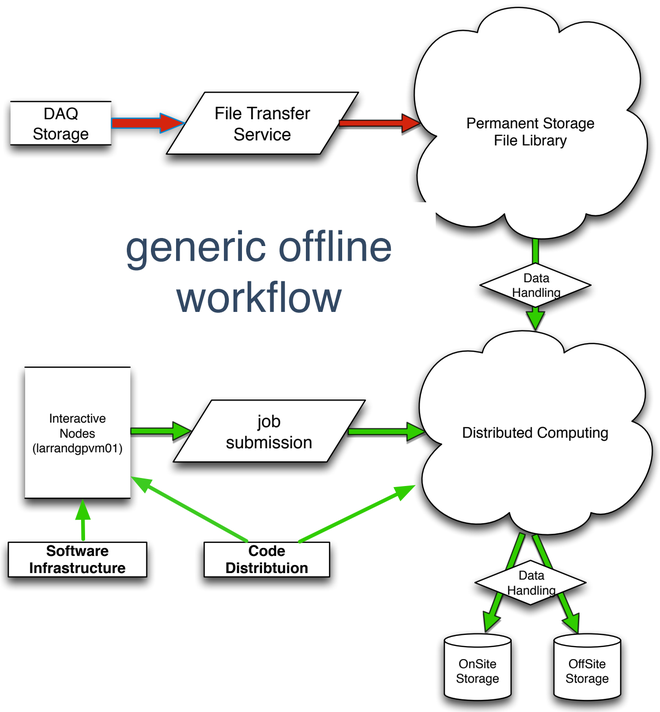

art & Larsoft

- ***art* is a modular event-processing framework for experiments**
  - Fork of the CMS framework: smaller experiments can base their software on a framework from one of the big LHC experiments
- **Many experiments are using *art* as their software framework**
  - This allows shared development and support among the experiments and the developers
    - Weekly stakeholder meetings, mailing list and issue trackers are used to coordinate development
  - Integration into Fermilab's data-handling system, I/O, etc. is an important part of the *art* project
- ***art-daq* is variant of the framework used as part of the daq system**

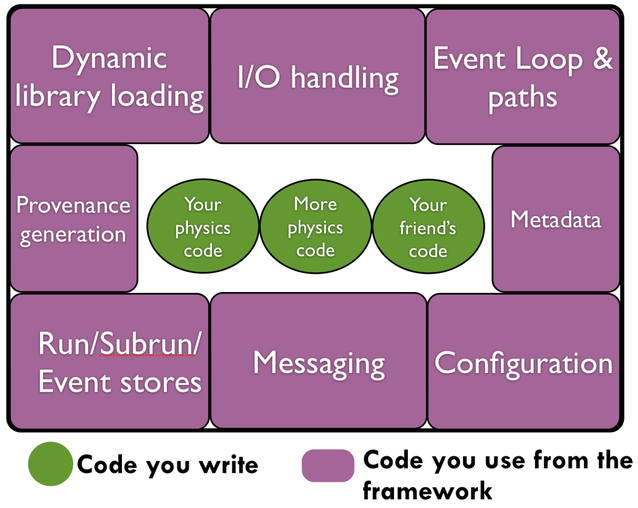

- **Larsoft**
  - **Technical goal:** Provide an integrated, *art*-based, experiment-agnostic set of software tools to be used by multiple LArTPC neutrino experiments to perform:
    - Simulation, Data reconstruction and Analysis
  - **Broader goal:** By developing common algorithms, services, data structures, architecture, dramatically reduce the cost of developing, maintaining and supporting the reconstruction and simulation software for collaboration members
  - **Community is collectively contributing to the larsoft software**
    - Resolve conflicts through process to integrate development from different stakeholders

Supported experiments and software

- **Collider experiments**
  - CDF
  - D0
- **Neutrino experiments**
  - ArgoNeut
  - LAr1-ND
  - LArIAT
  - LBNE
  - MicroBooNE
  - MINERvA
  - MiniBoone
  - MINOS
  - NOvA
  - NUMI
  - SNO+
- **Astroparticle/Cosmology Experiments**
  - CDMS
  - COUPP
  - DAMIC
  - Dark Energy Survey
  - DarkSide-50
  - DESI
  - Holometer
  - LSST
  - SDSS
- **Flavor experiments**
  - MIPP
  - Mu2e
  - Muon g-2
  - SeaQuest
- **Software**
  - Geant4
  - GENIE
  - MARS
- **R&D**
  - CHIPS
  - GENDEtRD
  - MCDRD
  - Next
  - SCENE

FIFE Experiments - Activity in 2014

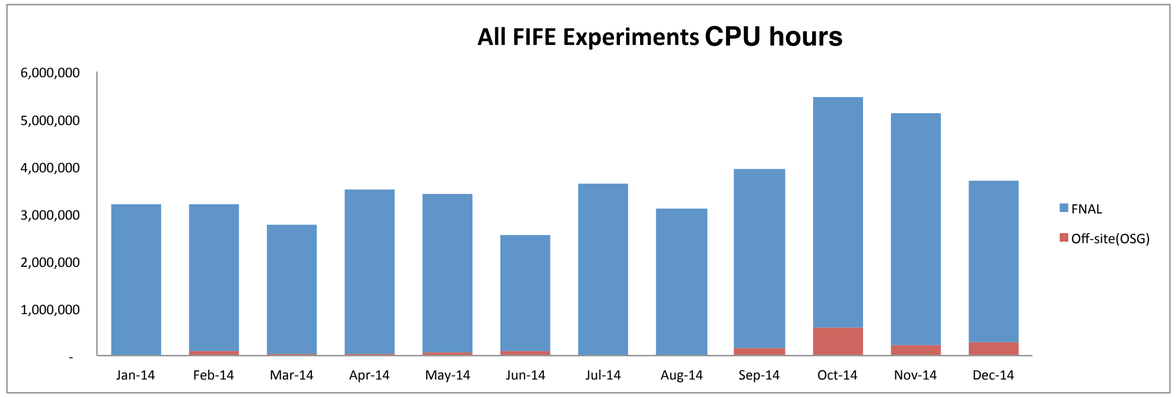

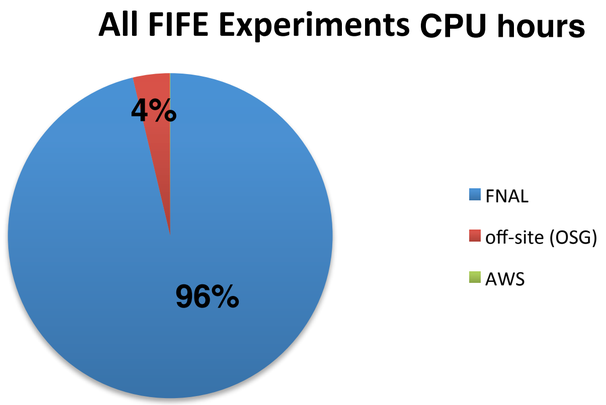

Off-site utilization is ramping up

In 2014 of 45M CPU hours, almost 2M hours were run off-site including first Amazon Web Services (AWS) usage of 20k CPU hours worth $3.3k as a pilot project

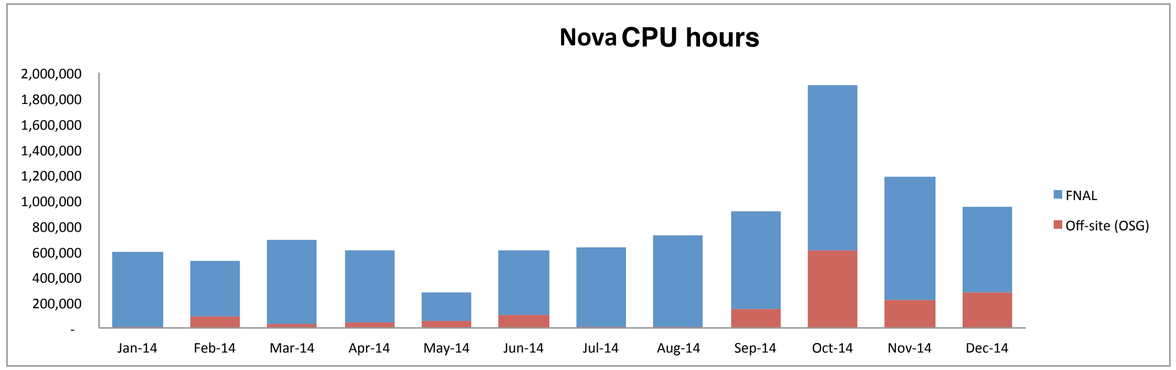

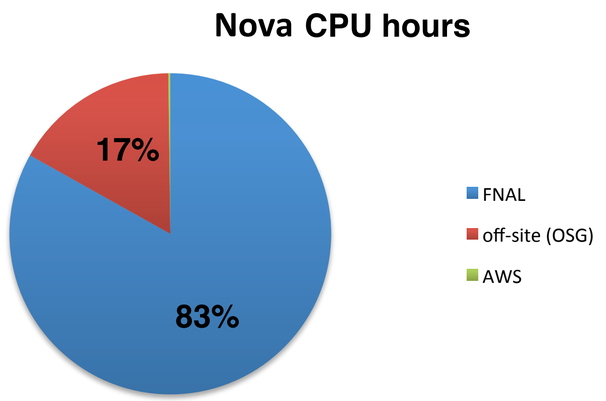

Operations group

- Operating the complex computing services requires significant knowledge and experience
  - Especially valid for job execution at high scales on various resources

- Fermilab founded the operations group to support NOvA, MINOS and Minerva (experiments that are currently taking data), helping in running crucial workflows like data keep-up processing and MC production

CMS

See Talk by Maria Girone

CMS - International and Distributed

- CMS resources worldwide
  - 7 Tier-1 sites, 54 Tier-2 sites
  - ~120,000 cores
  - ~75 PB disk
  - ~100 PB tape

- CMS@FNAL
  - ~15,000 cores
  - ~15 PB disk
  - ~30 PB tape

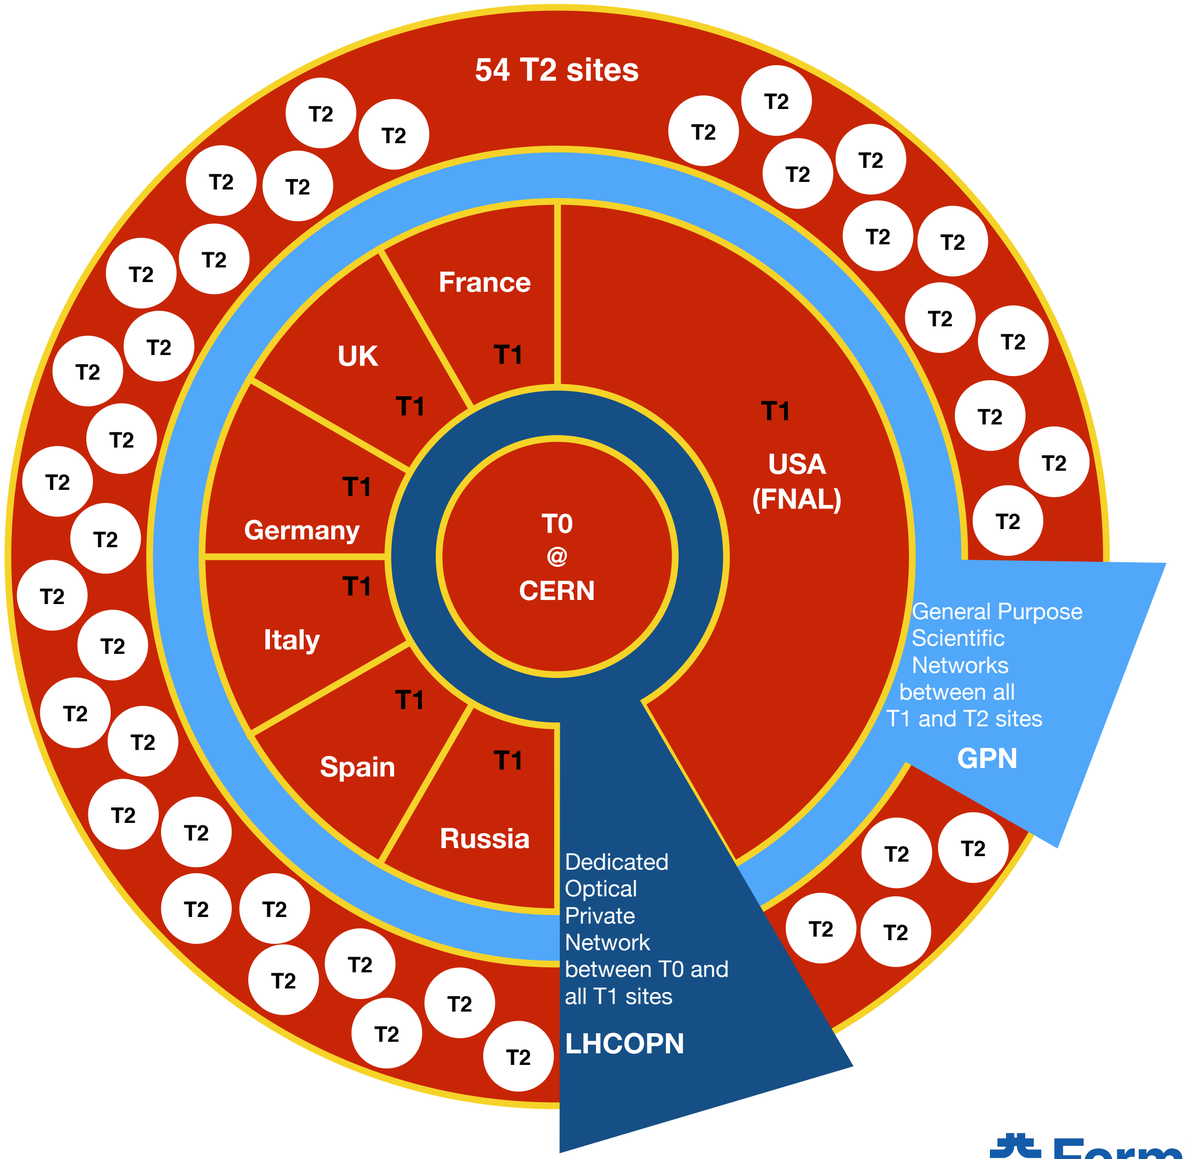

Services Overview

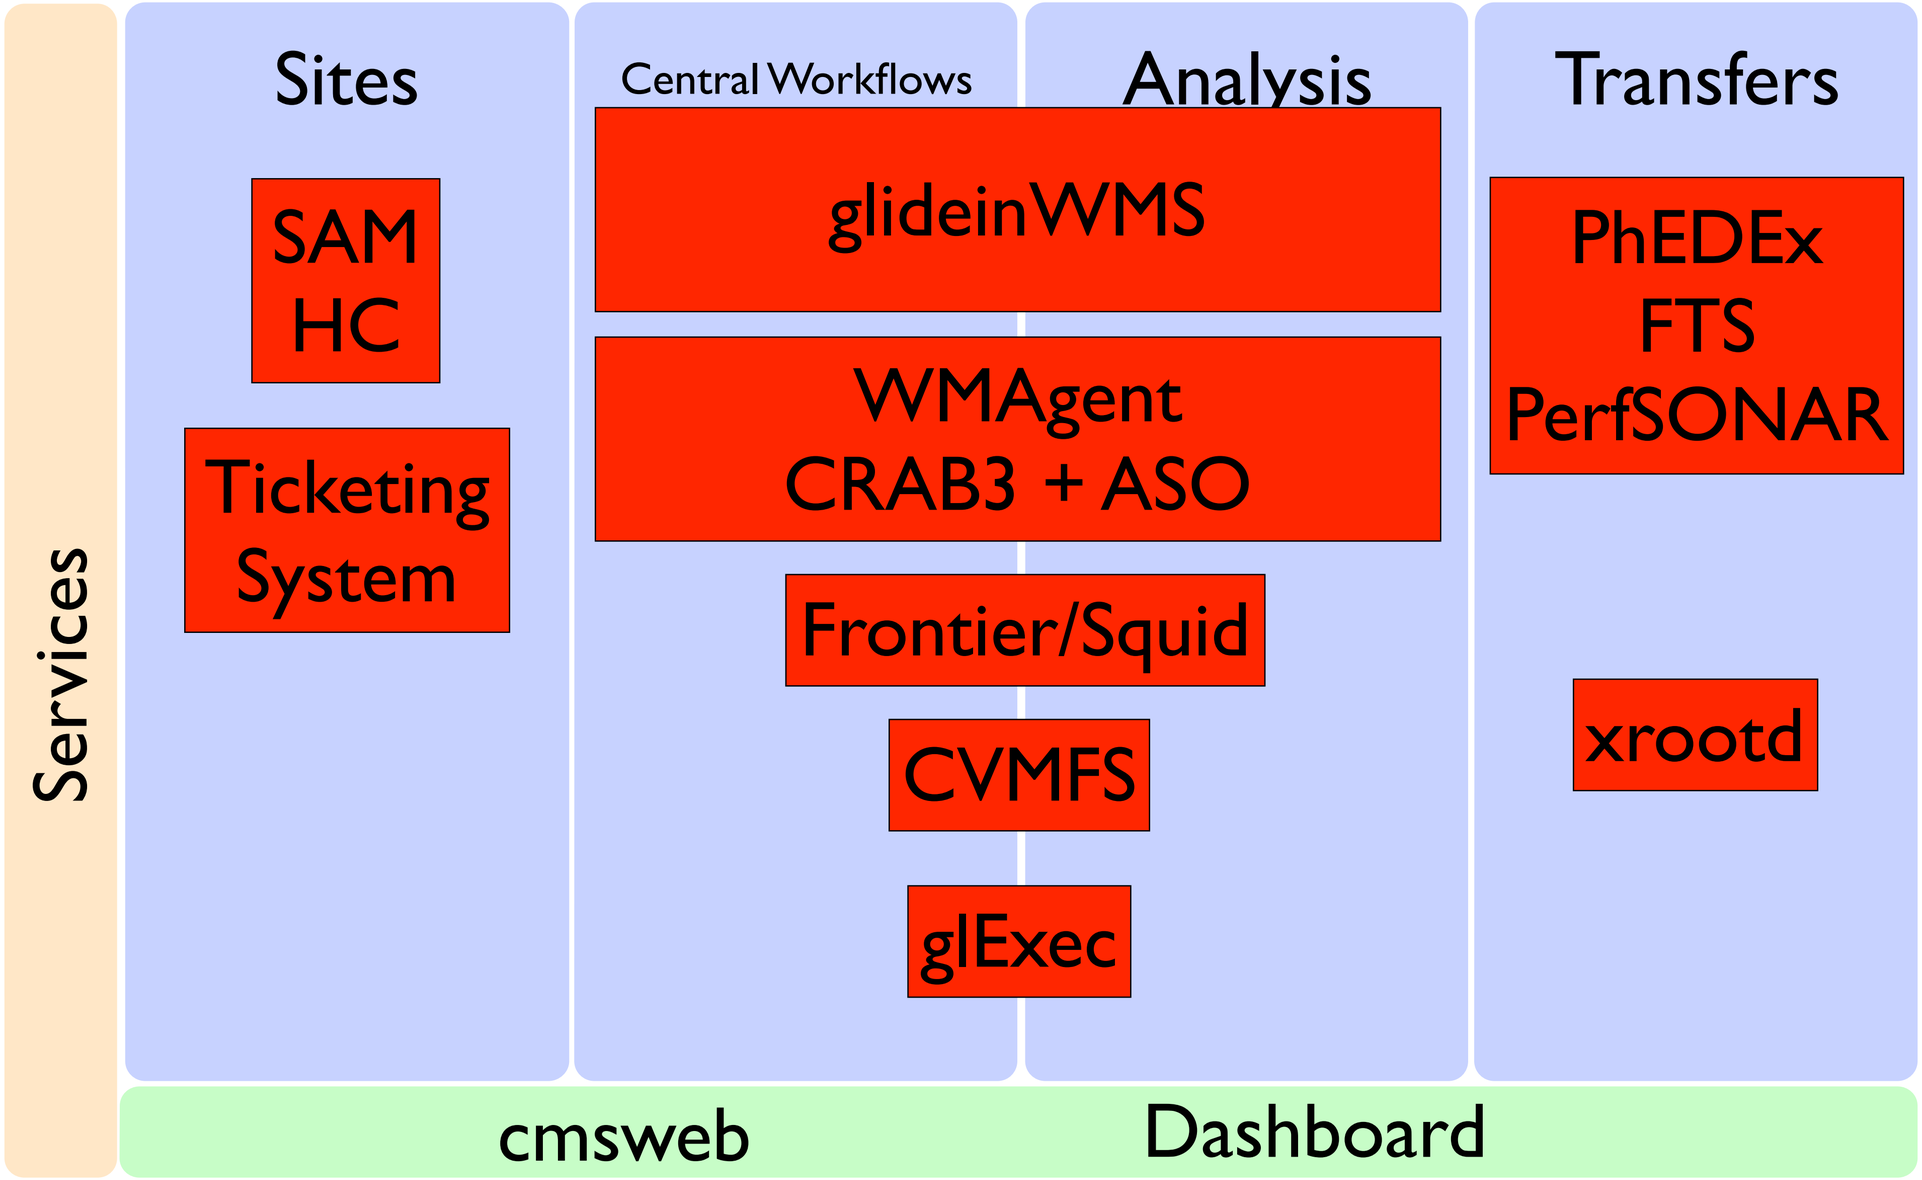

Global Pool

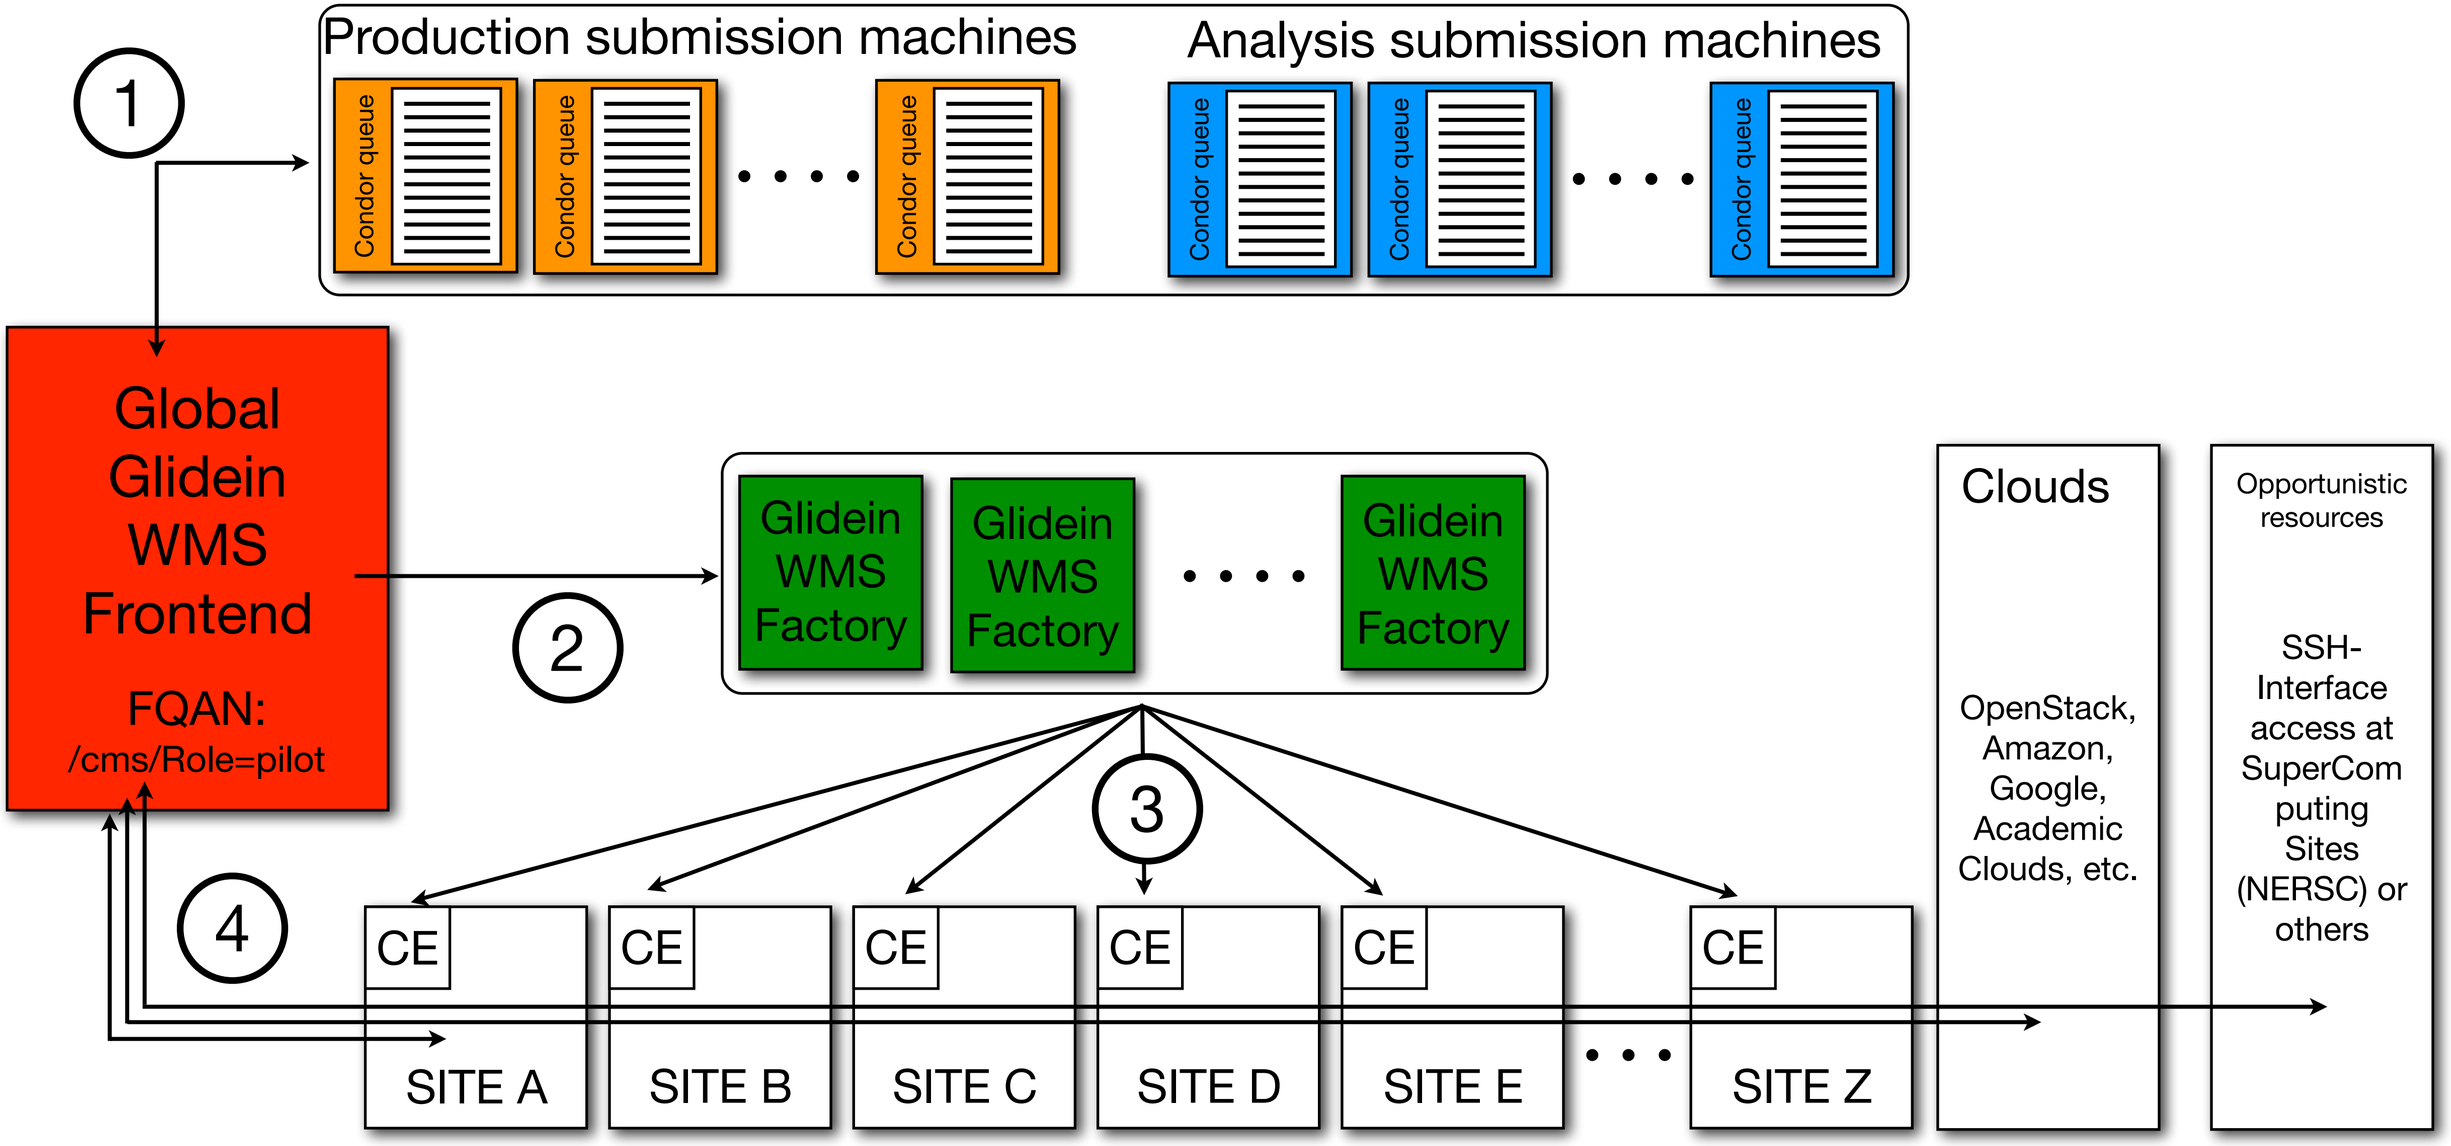
Production submission machines
Analysis submission machines
Condor queue
Condor queue
Condor queue
Condor queue
Condor queue
Condor queue
1
Global Glidein WMS Frontend
FQAN: /cms/Role=pilot
2
Glidein WMS Factory
Glidein WMS Factory
Glidein WMS Factory
3
4
CE
CE
CE
CE
CE
CE
SITE A
SITE B
SITE C
SITE D
SITE E
SITE Z
Clouds
OpenStack, Amazon, Google, Academic Clouds, etc.
Opportunistic resources
SSH-Interface access at SuperComputing Sites (NERSC) or others

Operations Group at Fermilab

- Taking care of
  - Running the Tier-0, Maintaining global transfers, Maintaining ≥ 60 sites worldwide, Operating the production infrastructure, …

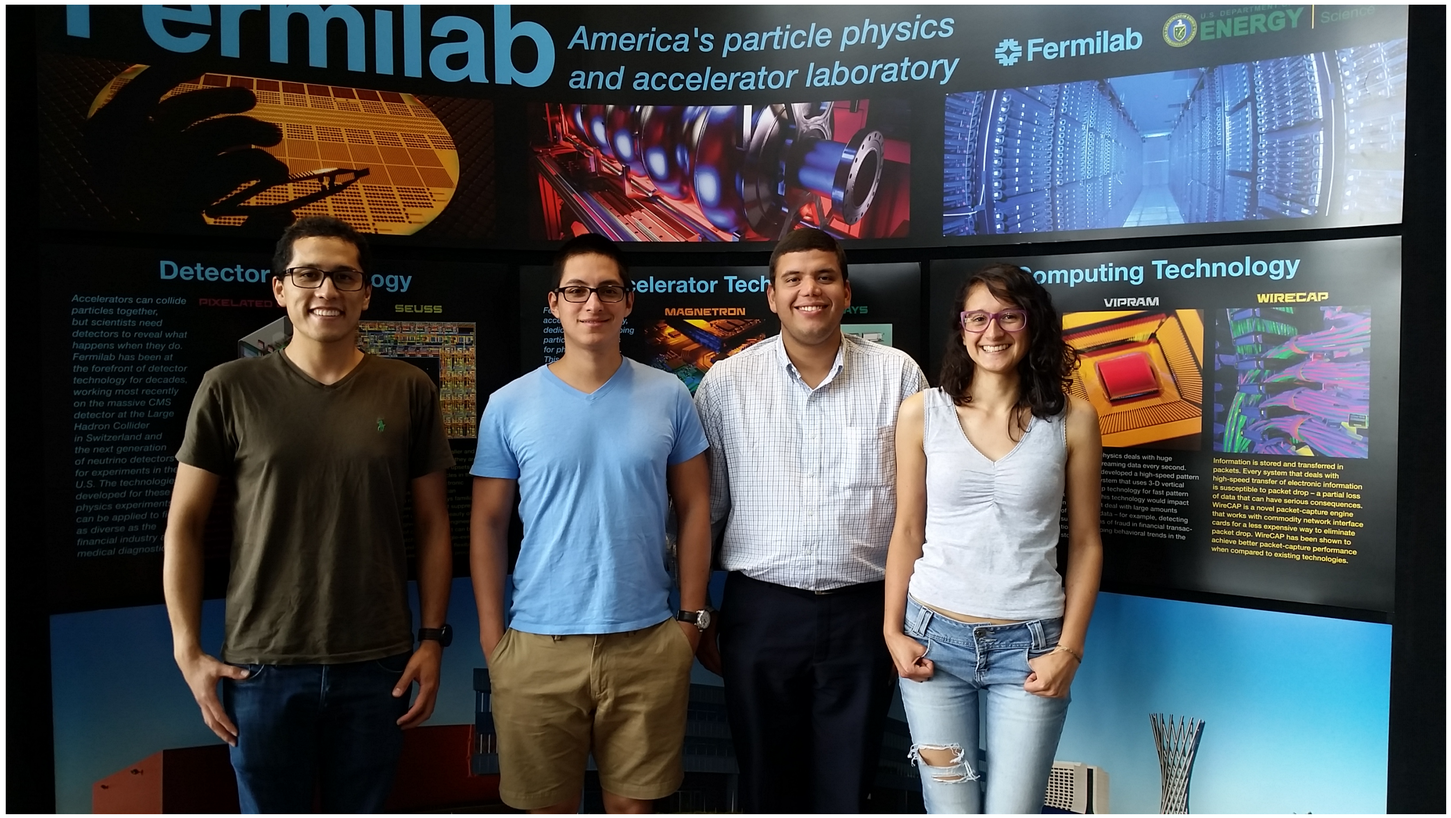

The Future

Software

Frameworks

- Underlying infrastructure, core of the software
  - Large experiments have their own Frameworks
  - Trend: community frameworks serving several experiments or detector technologies
    - Art: common framework for neutrino and muon experiments
    - LArSoft: common framework for liquid argon TPC (LArTPC) reconstruction software
    - Gaudi: common underlying framework for ATLAS and LHCb software
    - ALFA: the new ALICE-FAIR software framework
    - ILCroot: ILC, CLIC, MuC, ORKA
    - …

Moore's Law

- Traditionally, HEP software is optimized for a "simple" architecture
  - x86 based Linux
  - Machines:
    - ≥1 CPUs with ≥1 Cores
  - Shared memory
  - Shared local disk space
  - An application uses one core and memory and local disk space

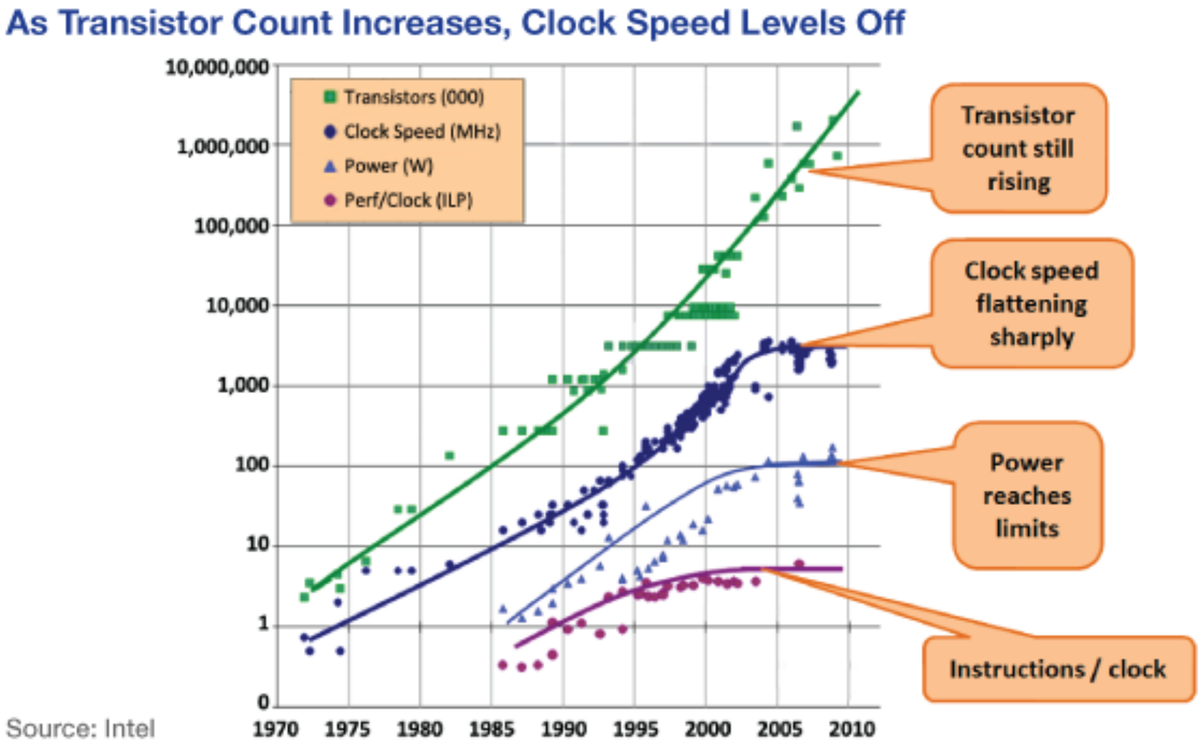

What we see: more and more cores, but less powerful individually.

New technologies: more and more cores!

- x86-based machines: running into limitations
  - Each application needs
    - “A lot” of memory (~2GB for LHC experiments) and corresponding bandwidth from memory to a core
    - The more cores in a single machine ➜ the more memory and bandwidth is needed
- New technology: GPGPU: General-purpose computing on graphics processing units
  - Use of a graphics processing units (GPUs) optimized for parallel processing ➜ using many cores per application
  - To perform computation traditionally handled by the central processing unit (CPU)
- New technology: Co-Processor architectures
  - Keyword: Intel MIC (Many Integrated Core) Architecture
- Consequence: **We need to use more cores in parallel for our applications!**

Multi-threading: frameworks

- Advantage: save memory by sharing between threads
- current state: run each event in own thread

- future: run parts of events in different threads ➜ higher optimization results with even less memory usage

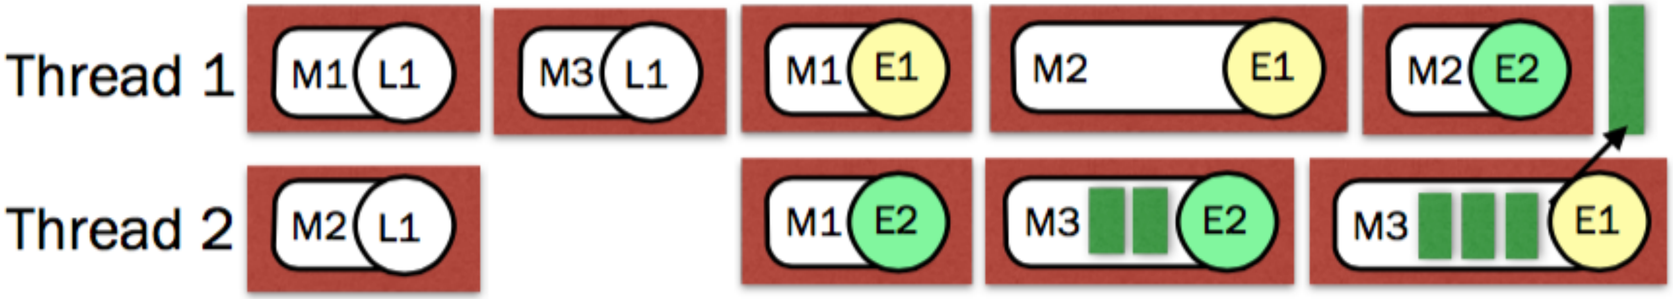

Thread-safe programming

- New technologies: multi-threading, GPGPU, Co-Processors
  - Require new programming skills!
  - My opinion: comparable to Fortran ➜ C++ switch
- Multi-threaded programming needs to be done right
  - Small amounts of non-thread-safe code reduces the efficiency significantly ➜ Amdahl's law
- **Go and learn thread-safe programming!**

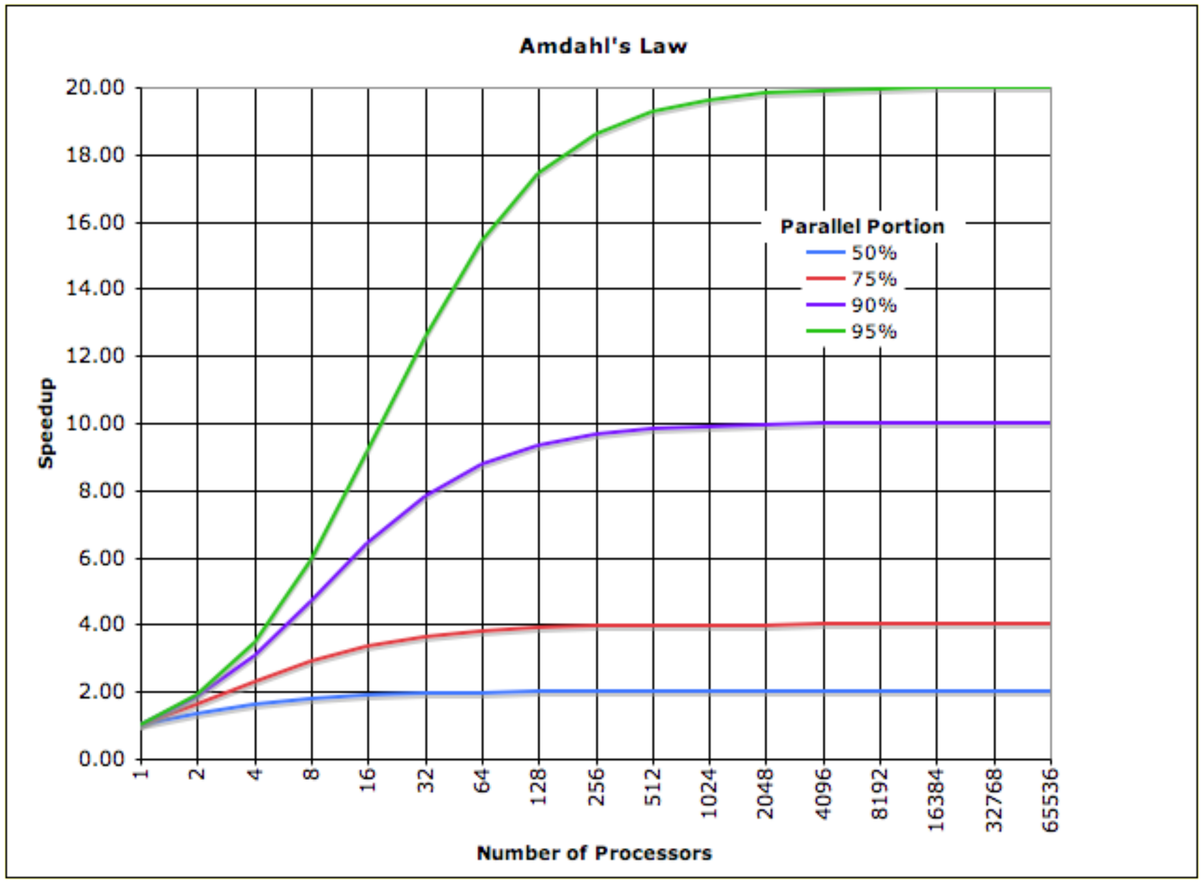

Storage

What is a Petabyte?

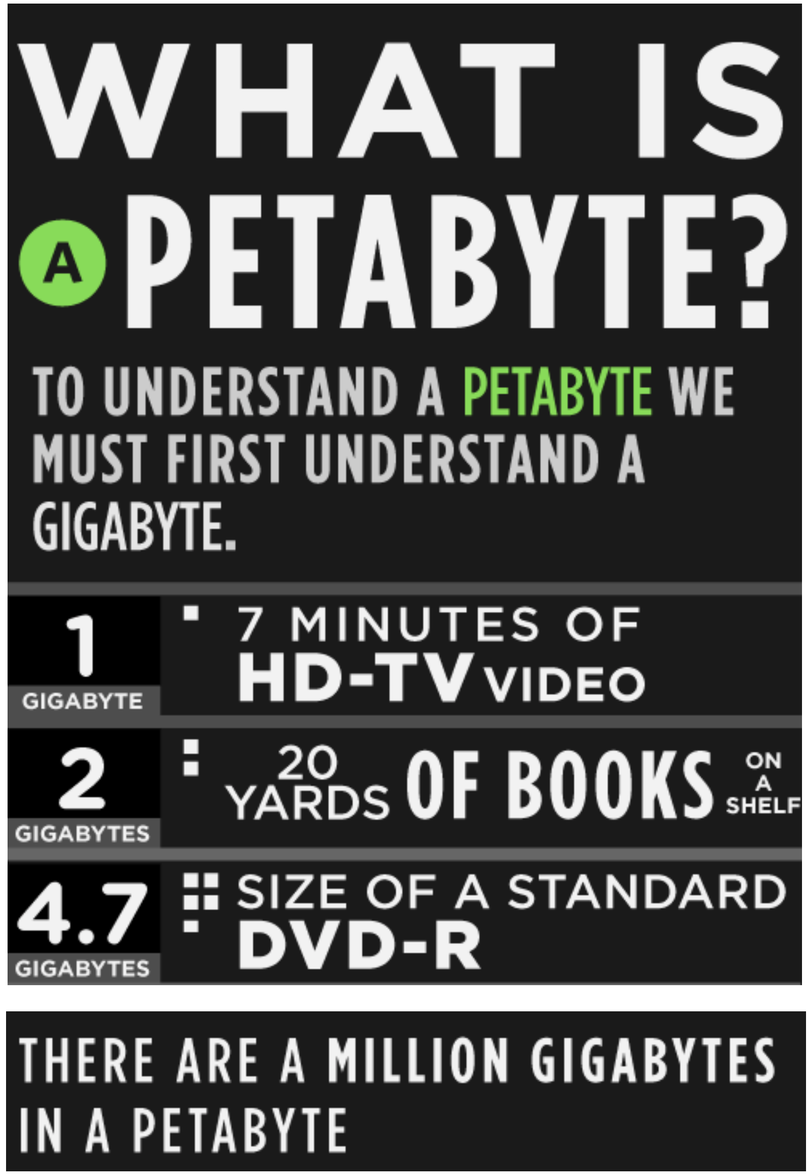

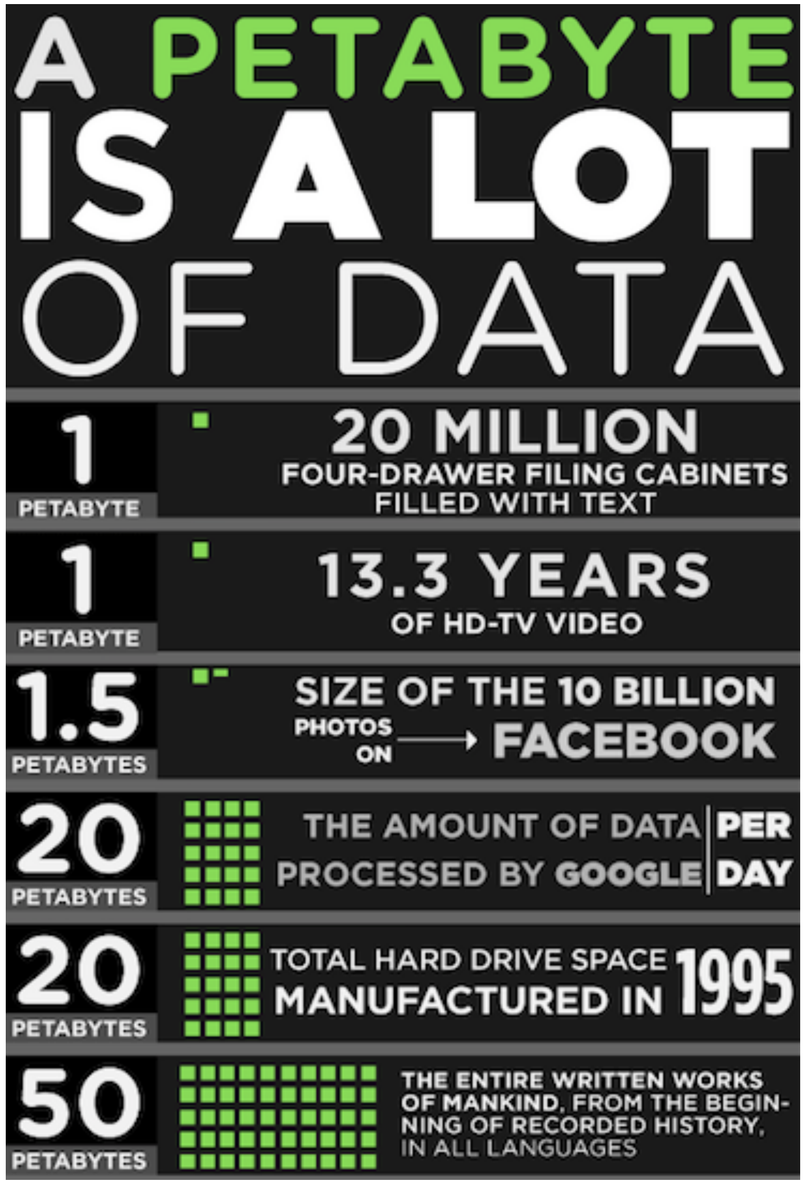

LHC schedule

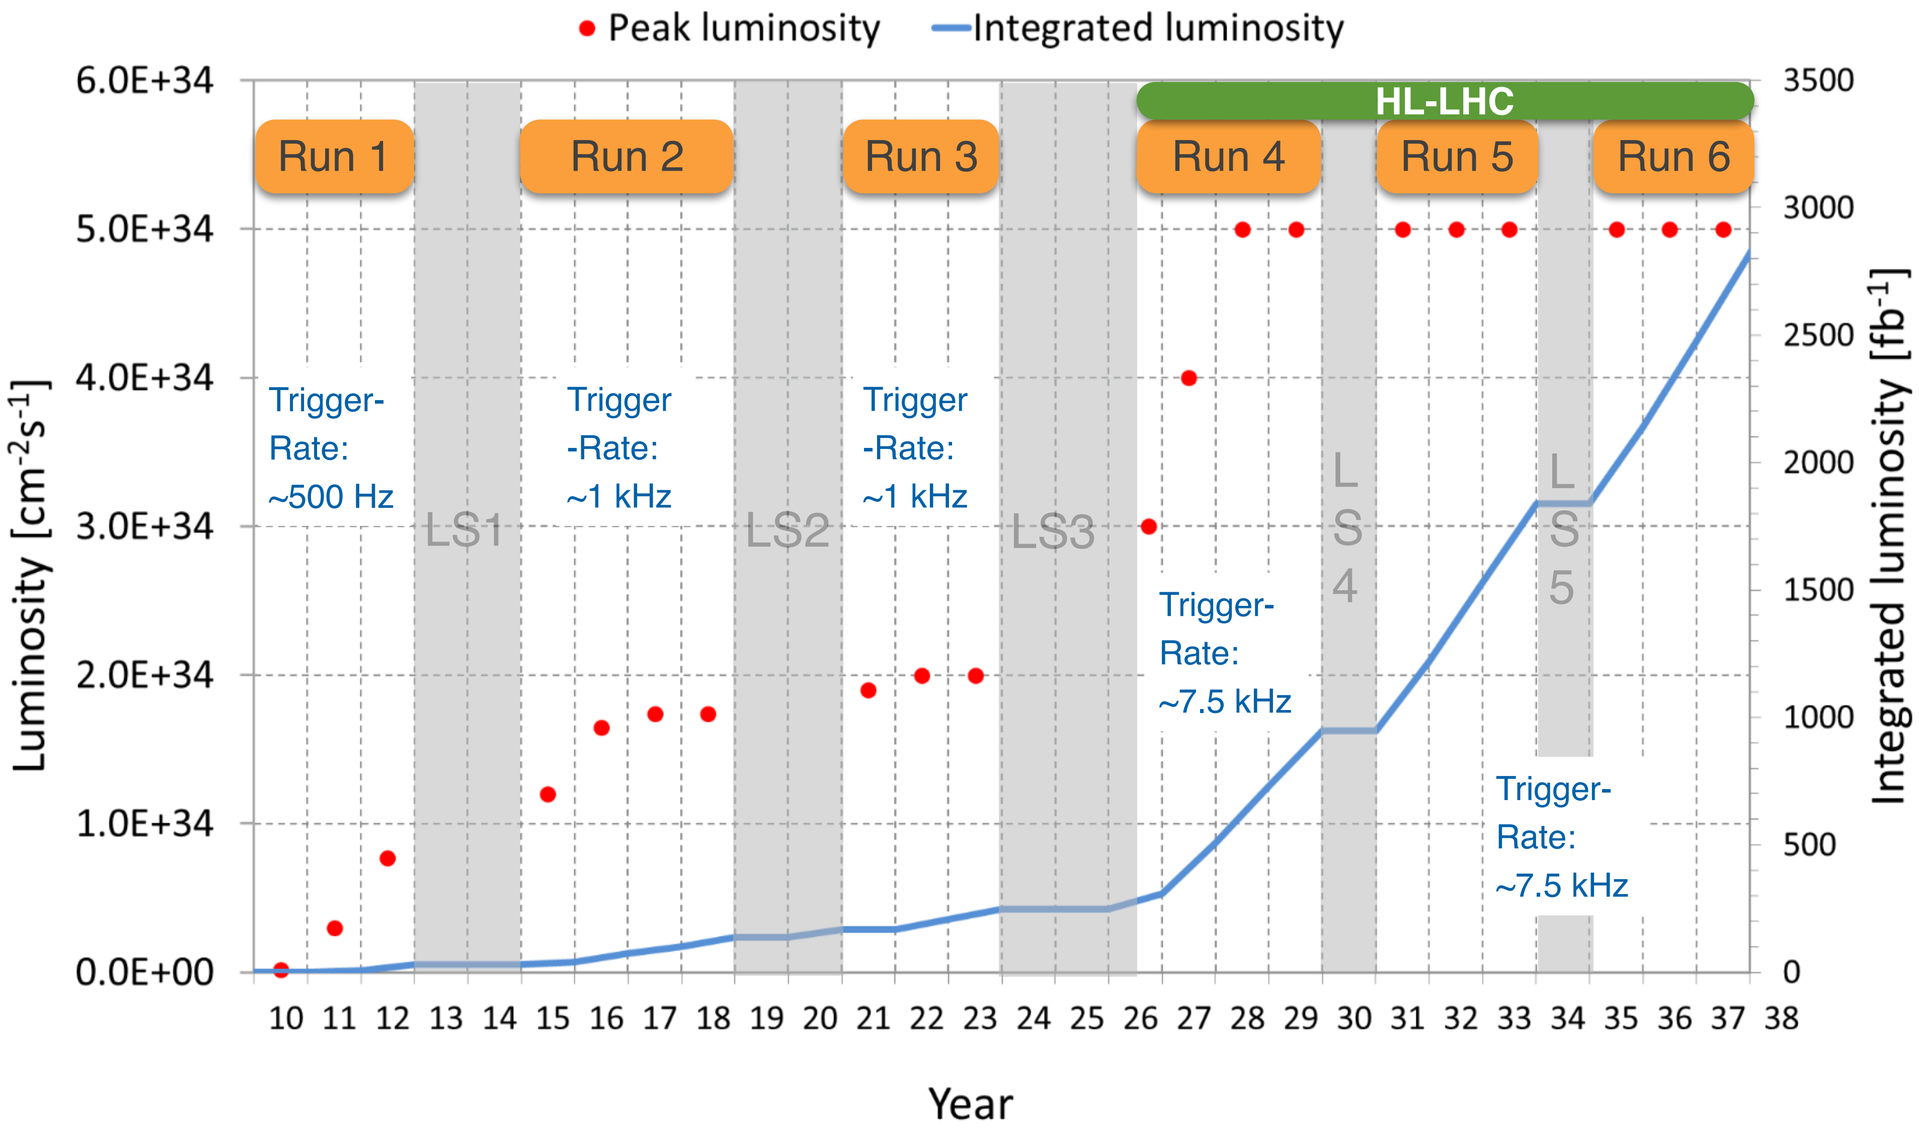

LHC expectation data volumes

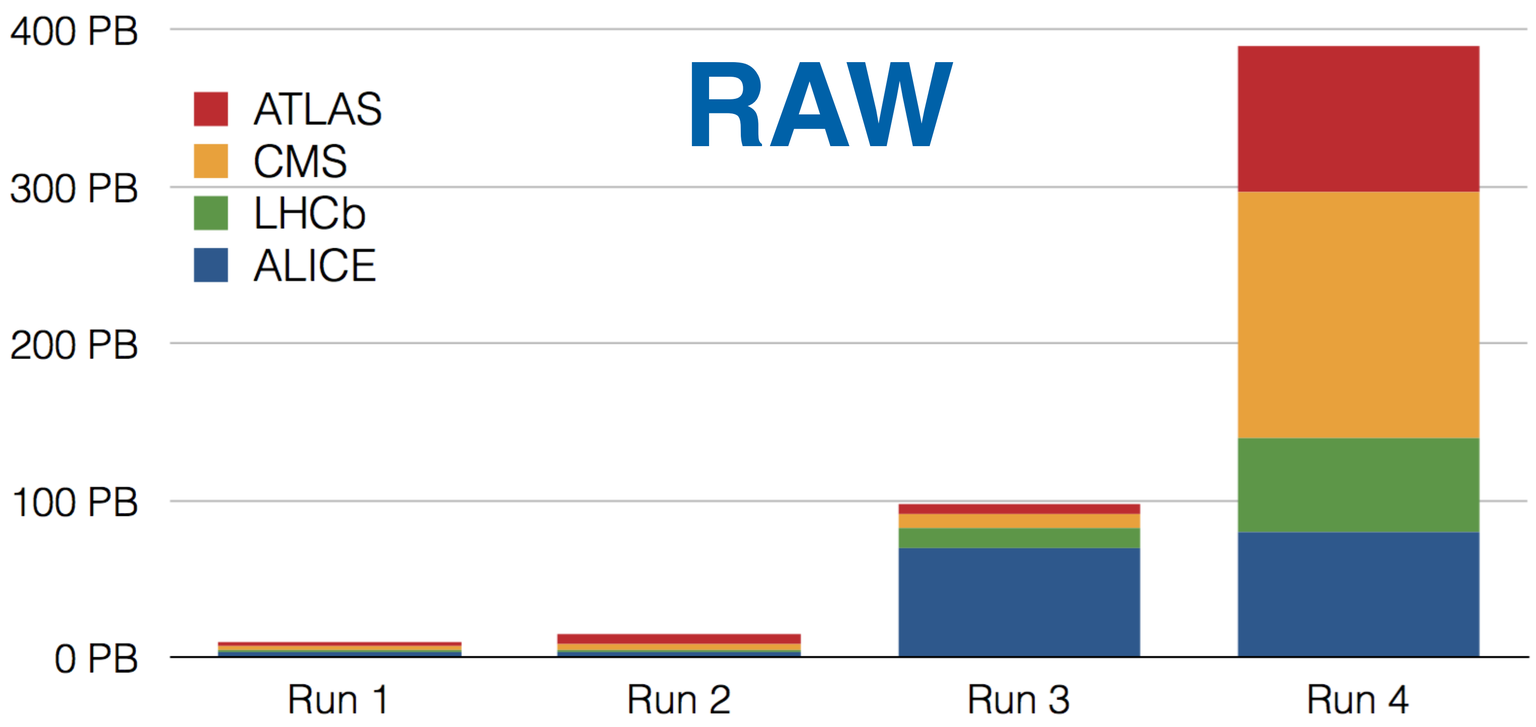

- Shown: RAW expectations
  - Derived data (RECO, Simulation): factor 8 of RAW
- LHC Run 4 is starting the exabyte era
- How do we analyze that much data in the future?

Strong networks: ESNet

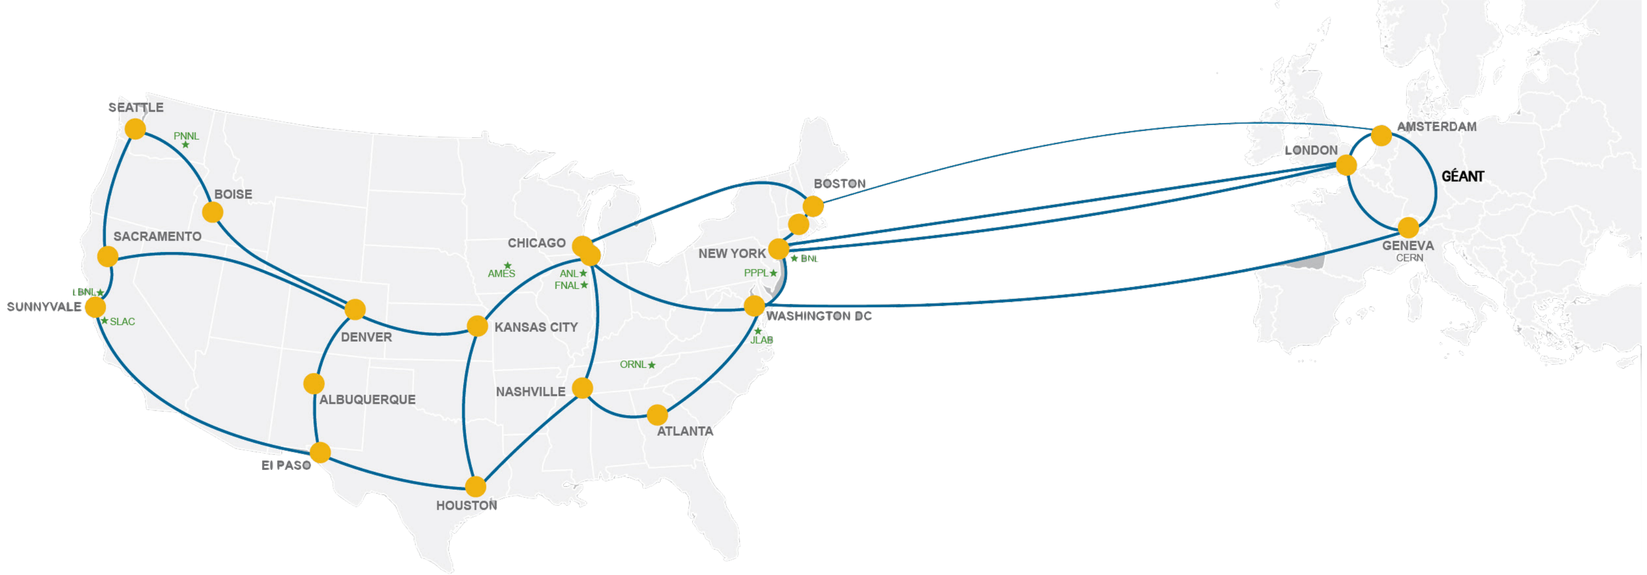

ESnet
ENERGY SCIENCES NETWORK

The Office of Science supports:

- 27,000 Ph.D.s, graduate students, undergraduates, engineers, and technicians
- 26,000 users of open-access facilities
- 300 leading academic institutions
- 17 DOE laboratories

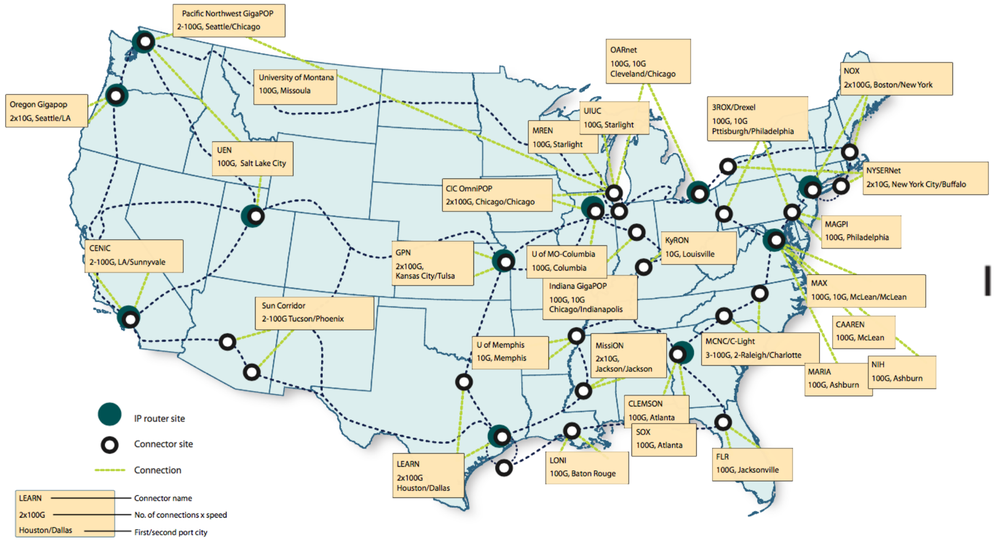

INTERNET2

Fermilab

Distributed infrastructures and transfer systems

Example: Worldwide LHC Grid (WLCG)

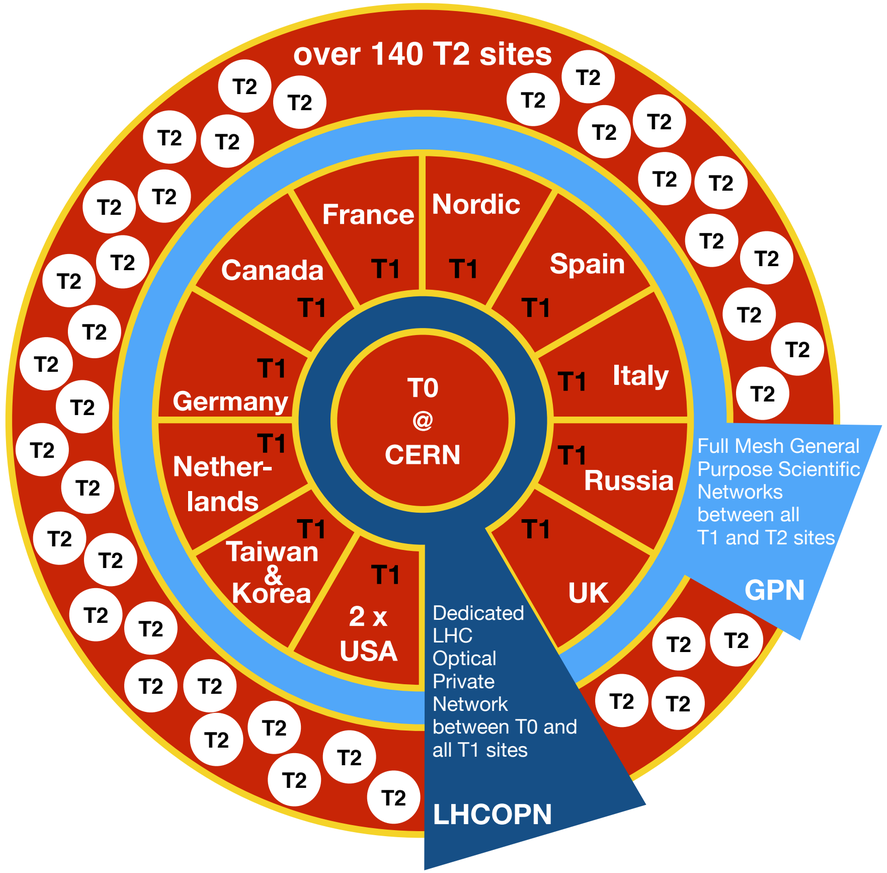

Community uses various solutions to provide distributed access to data:

Experiment specific: Atlas (Rucio), CMS (PhEDEx), …

Shared: SAM (Neutrino and Muon experiments)

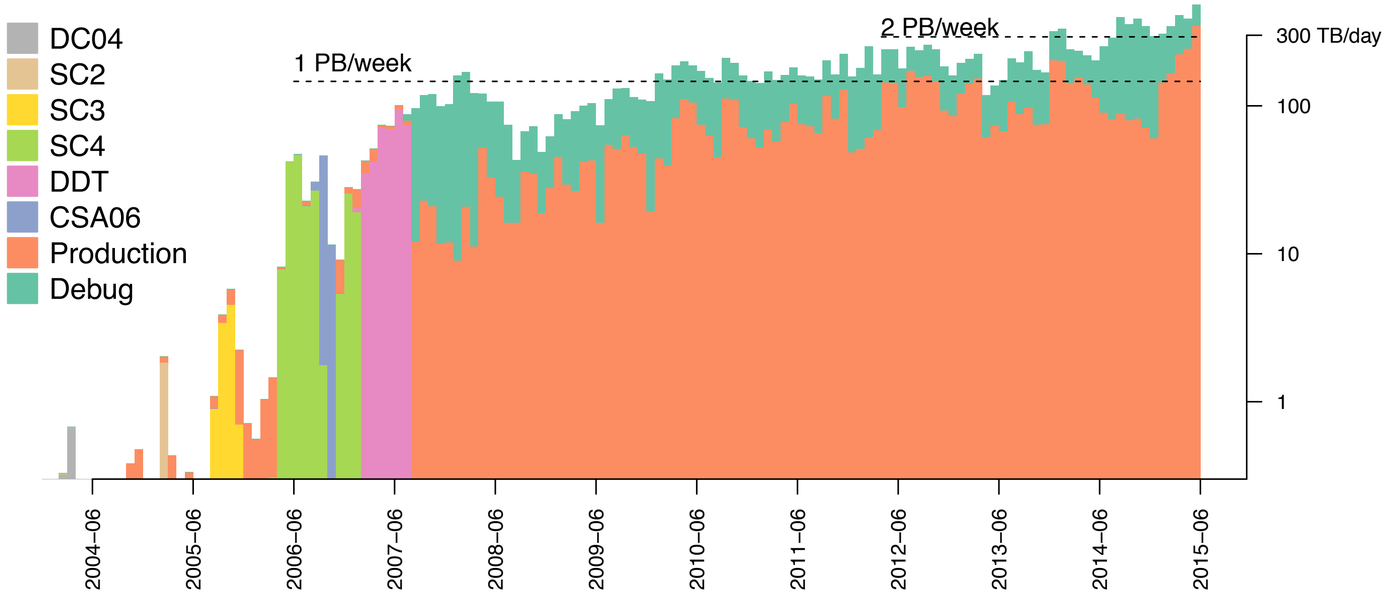

CMS transfers: more than 2 PB per week

Dynamic Data Management

- Subscription based transfer systems
  - PhEDEx (CMS) and Rucio (Atlas)
  - LHC Run 1: mostly manual operations
  - LHC Run 2: **dynamic data management**
    - Popularity is tracked per dataset
    - Replica count across sites is increased or decreased according to popularity
- Fully integrated distribution system
  - SAM (shared amongst Neutrino and Muon experiments)
  - All movement is based on requests for datasets from jobs.
  - Interfaces to storage at sites, performs cache-to-cache copies if necessary
- Data is distributed automatically for the community

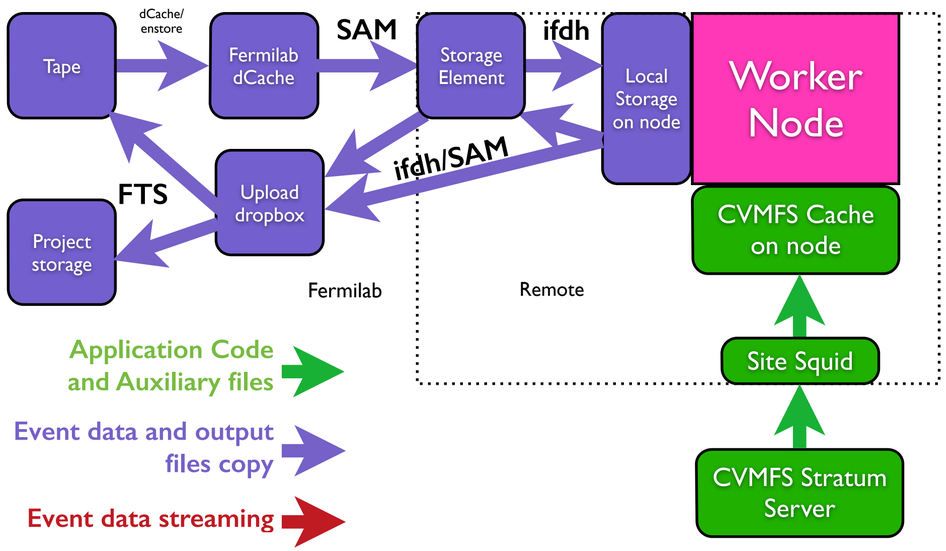

Data Federations

- xrootd: remote access to files
- ALICE based on xrootd from the beginning
- CMS and Atlas deployed xrootd federations
  - AAA for CMS, FAX for Atlas
  - Allows for remote access to all files on disk at all sites
  - Use cases:
    - Fall back
    - Overflow for ~10% of all jobs

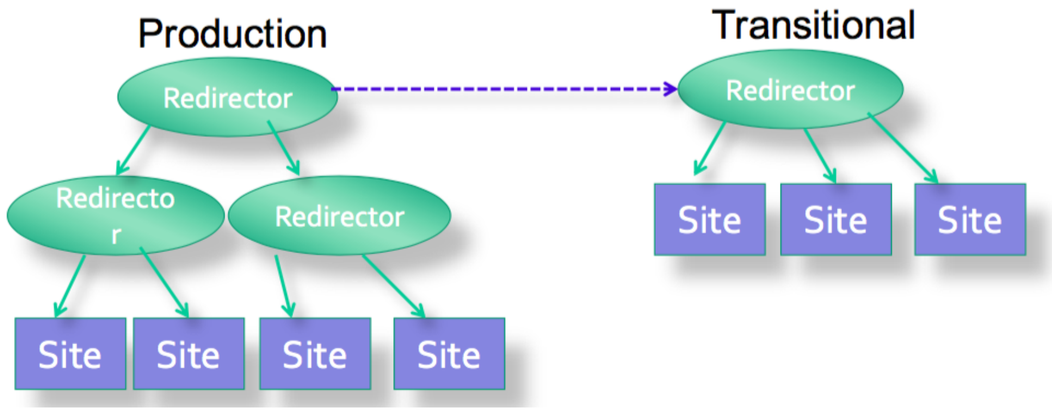

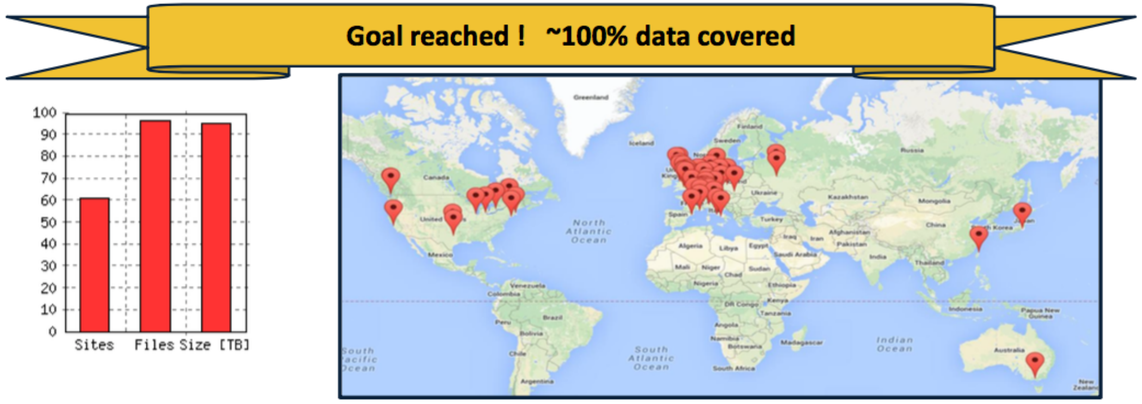

OSG StashCache

- OSG: StashCache
  - Bringing opportunistic storage usage to all users of OSG
  - OSG collaborators provide local disk space
  - OSG is running xrootd cache servers
    - Dynamic population of caches ➜ efficient distributed access to files
      - For users that don't have infrastructures like CMS and Atlas

Stash origin: ★

OSG Caches: ●

Active Archival Facility

- HEP has the tools and experience for the **distributed exabyte scale**
  - We are “best in class” in the field of scientific data management

- We are working with and for the whole science community
  - To bring our expertise to everyone’s science
  - To enable everyone to manage, distribute and access their data, globally

- Example: Fermilab’s Active Archival Facility (AAF)
  - Provide services to other science activities to preserve integrity and availability of important and irreplaceable scientific data
  - Projects:
    - Genomic research community is archiving datasets at Fermilab’s AAF and providing access through Fermilab services to ~300 researchers all over the world
    - University of Nebraska and University of Wisconsin are setting up archival efforts with Fermilab’s AAF

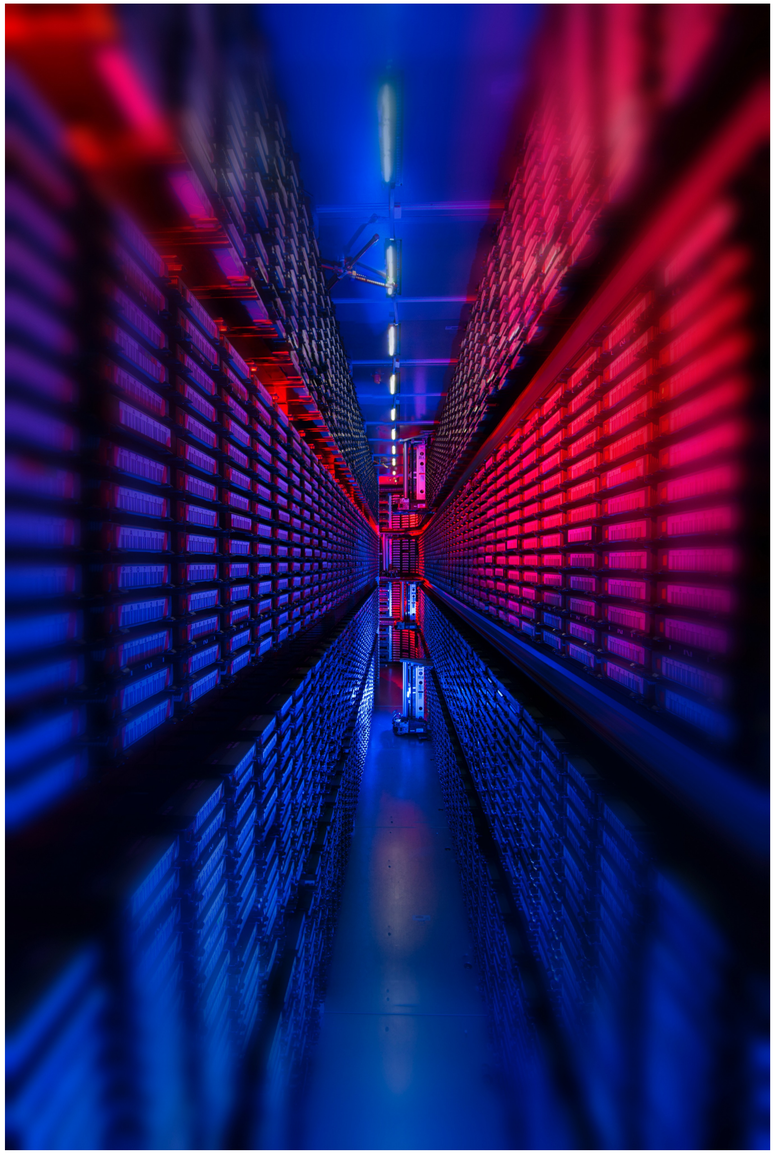

Processing

New resource providers

Grid

- Virtual Organizations (VOs) of users trusted by Grid sites
- VOs get allocations ➜ **Pledges**
  - Unused allocations: opportunistic resources

Trust Federation

Cloud

- Community Clouds - Similar trust federation to Grids
- Commercial Clouds - **Pay-As-You-Go** model
  - Strongly accounted
  - Near-infinite capacity ➜ **Elasticity**
  - Spot price market

Economic Model

HPC

- Researchers granted access to HPC installations
- Peer review committees award **Allocations**
  - Awards model designed for individual PIs rather than large collaborations

Grant Allocation

Evolving the Grid

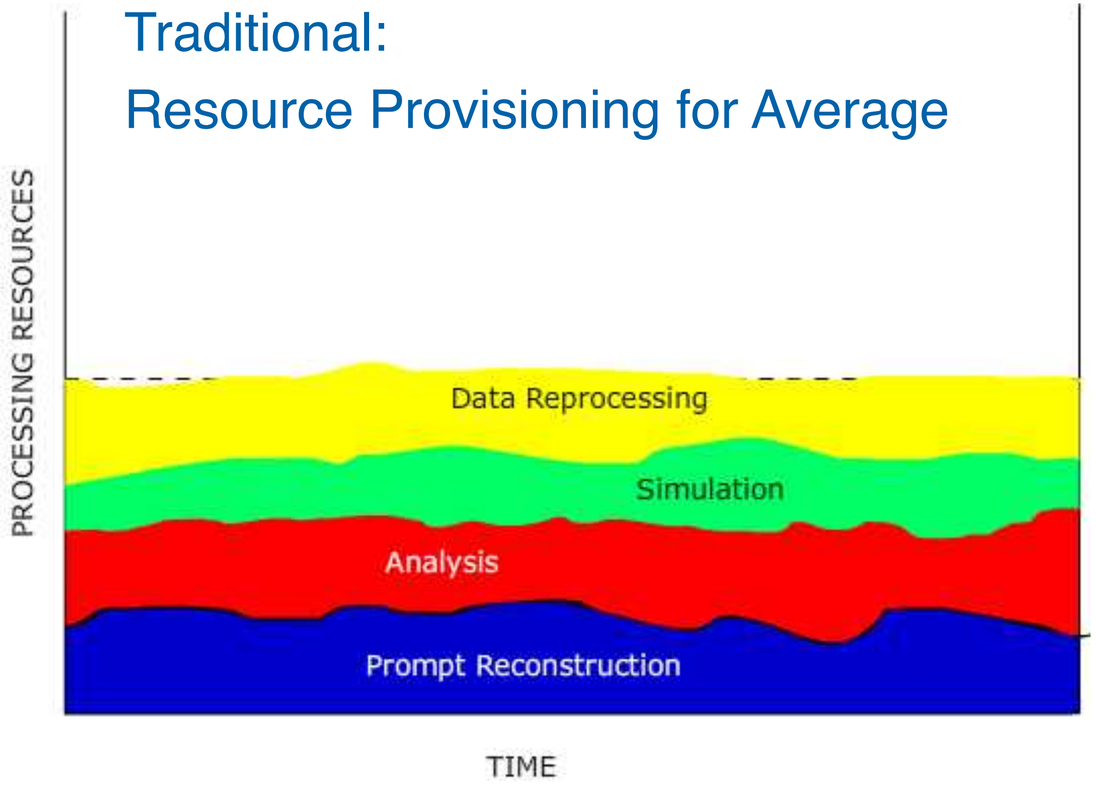

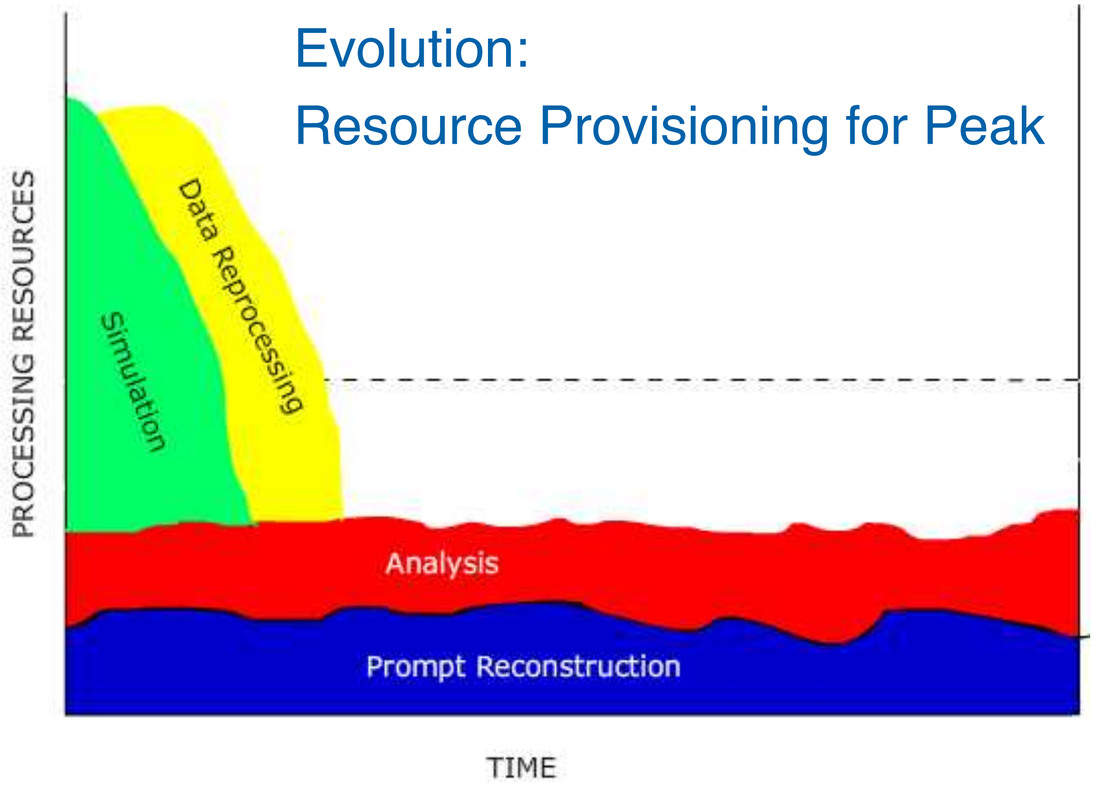

- Experiments don’t need all the resources all the time
  - Conference schedule, holiday seasons, accelerator schedules, etc.
  - Resource needs vary with time ➜ Provisioning needs to adapt

Fermilab's HEPCloud

- Many experiments and facilities are exploring using commercial cloud providers to provision for peak
  - Examples: Atlas, CMS, STAR, NOvA, etc. / BNL, FNAL, CNAF, etc.

- Example: Fermilab's HEPCloud
  - Provision commercial cloud resources in addition to physically owned resources
  - Transparent to the user

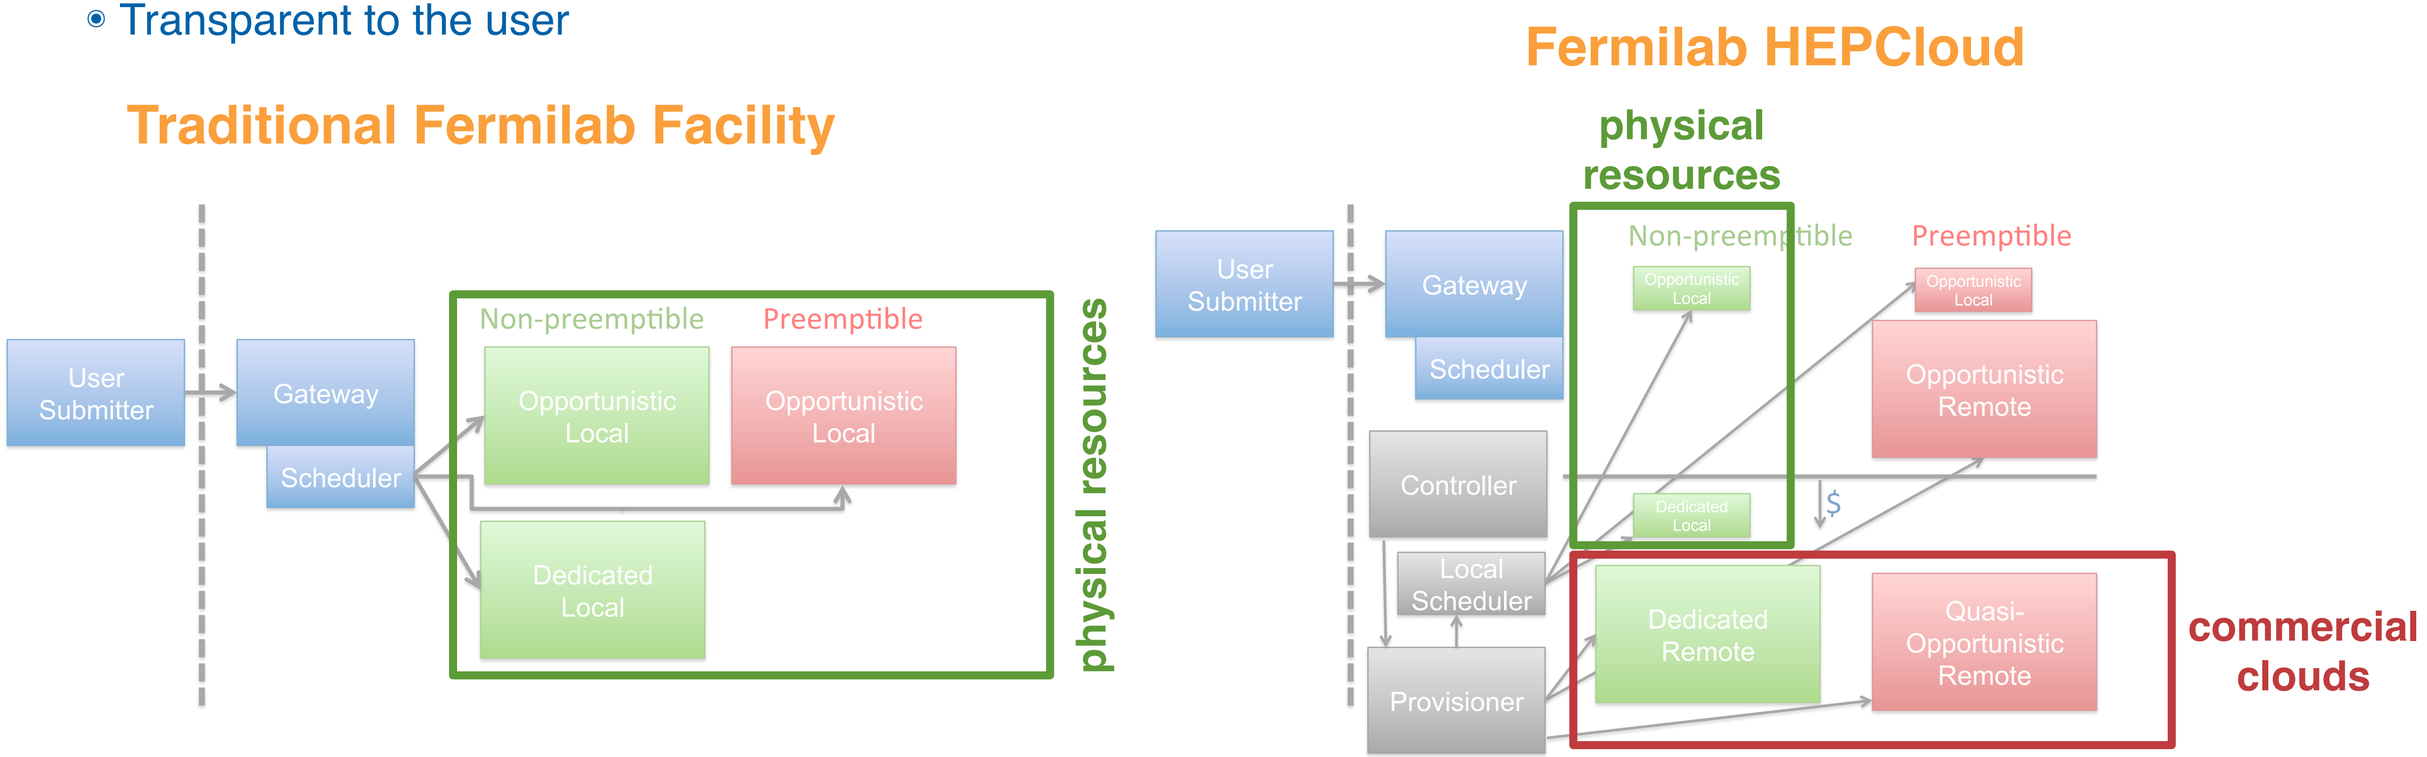

Open Science Grid ➜ Facilitating shared access

- Researcher use a single interface to use resources …
  - … they own
  - … others are willing to share
  - … they have an allocation on
  - … they buy from a commercial (cloud) provider

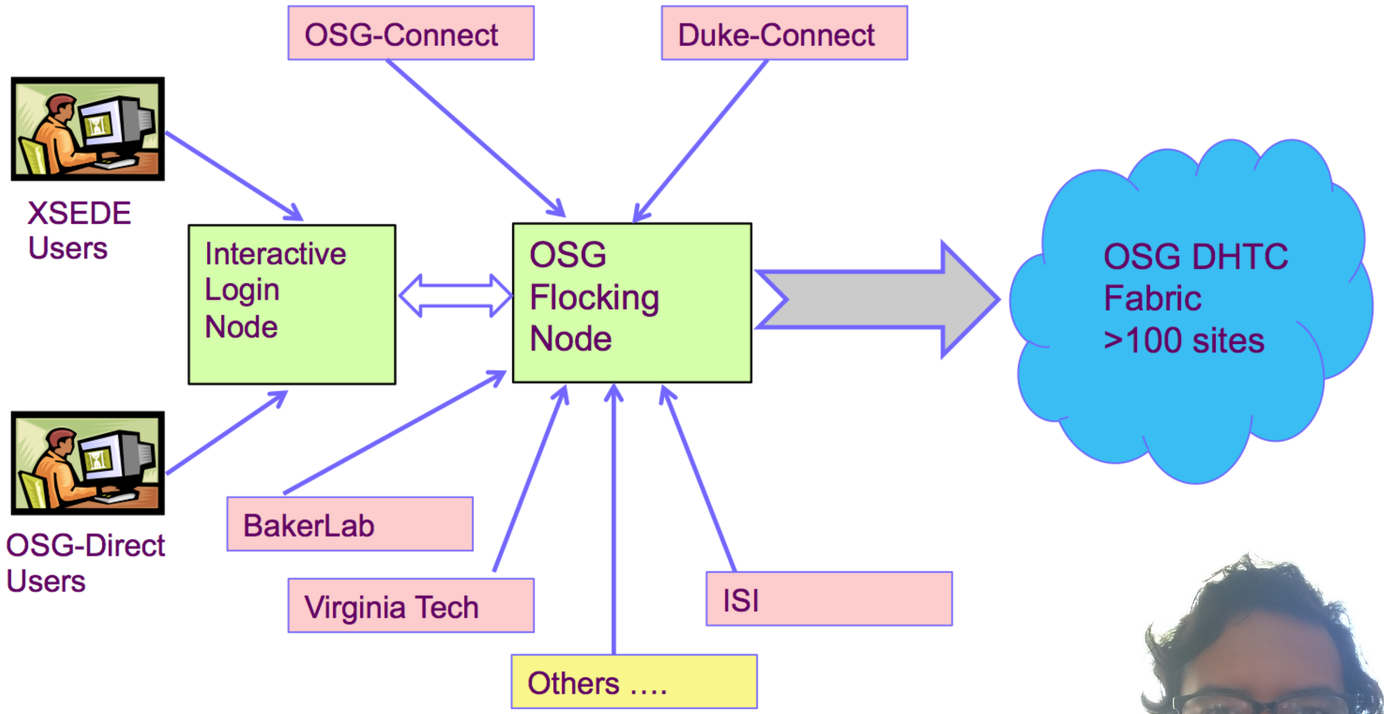

- OSG focuses on making this technically possible for Distributed High Throughput Computing
  - Operate a shared Production Infrastructure ➜ Open Facility (glideinWMS)
  - Advance a shared Software Infrastructure ➜ Open Software Stack
  - Spread knowledge across Researchers, IT professionals & Software developers ➜ Open Ecosystem

HPC & HEP

- HTC: High Throughput Computing
  - Independent, sequential jobs that can be individually scheduled on many different computing resources across multiple administrative boundaries(*)
- HPC: High Performance Computing
  - Tightly coupled parallel jobs, must execute within a particular site with low-latency interconnects(*)

- Long history in HEP in using HPC installations
  - Lattice QCD and Accelerator Modeling exploit the low latency interconnects successfully for a long time

- Community effort: enable traditional HEP framework applications to run on HPC installations
  - Example: Mira at Argonne (PowerPC, ~49k nodes each 16 cores, almost 800k cores)
  - Generating Atlas LHC Events with Algren

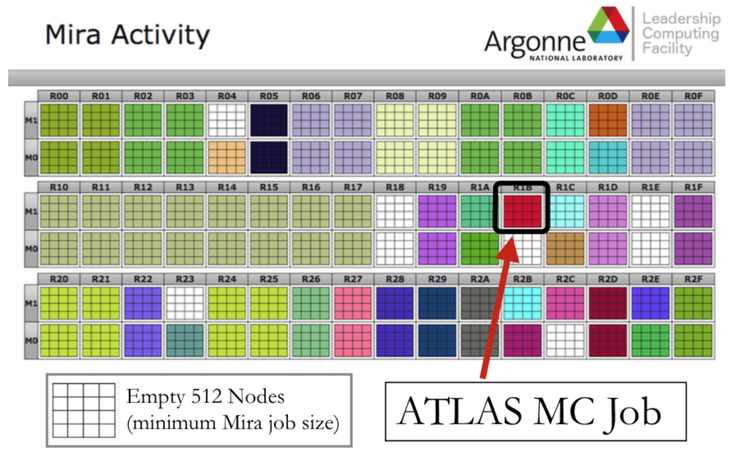

The Future: Exascale ➜ more cores!

System attributes	NERSC Now	OLCF Now	ALCF Now	NERSC Upgrade	OLCF Upgrade	ALCF Upgrades	
Name Planned Installation	Edison	TITAN	MIRA	Cori 2016	Summit 2017-2018	Theta 2016	Aurora 2018-2019
System peak (PF)	2.6	27	10	> 30	150	>8.5	180
Peak Power (MW)	2	9	4.8	< 3.7	10	1.7	13
Total system memory	357 TB	710TB	768TB	~1 PB DDR4 + High Bandwidth Memory (HBM) +1.5PB persistent memory	> 1.74 PB DDR4 + HBM + 2.8 PB persistent memory	>480 TB DDR4 + High Bandwidth Memory (HBM)	> 7 PB High Bandwidth On-Package Memory Local Memory and Persistent Memory
Node performance (TF)	0.460	1.452	0.204	> 3	> 40	> 3	> 17 times Mira
Node processors	Intel Ivy Bridge	AMD Opteron Nvidia Kepler	64-bit PowerPC A2	Intel Knights Landing many core CPUs Intel Haswell CPU in data partition	Multiple IBM Power9 CPUs & multiple Nvidia Voltas GPUS	Intel Knights Landing Xeon Phi many core CPUs	Knights Hill Xeon Phi many core CPUs
System size (nodes)	5,600 nodes	18,688 nodes	49,152	9,300 nodes 1,900 nodes in data partition	~3,500 nodes	>2,500 nodes	>50,000 nodes
System Interconnect	Aries	Gemini	5D Torus	Aries	Dual Rail EDR-IB	Aries	2nd Generation Intel Omni-Path Architecture
File System	7.6 PB 168 GB/s, Lustre®	32 PB 1 TB/s, Lustre®	26 PB 300 GB/s GPFS™	28 PB 744 GB/s Lustre®	120 PB 1 TB/s GPFS™	10PB, 210 GB/s Lustre initial	150 PB 1 TB/s Lustre®

Projected Parallelism for Exascale

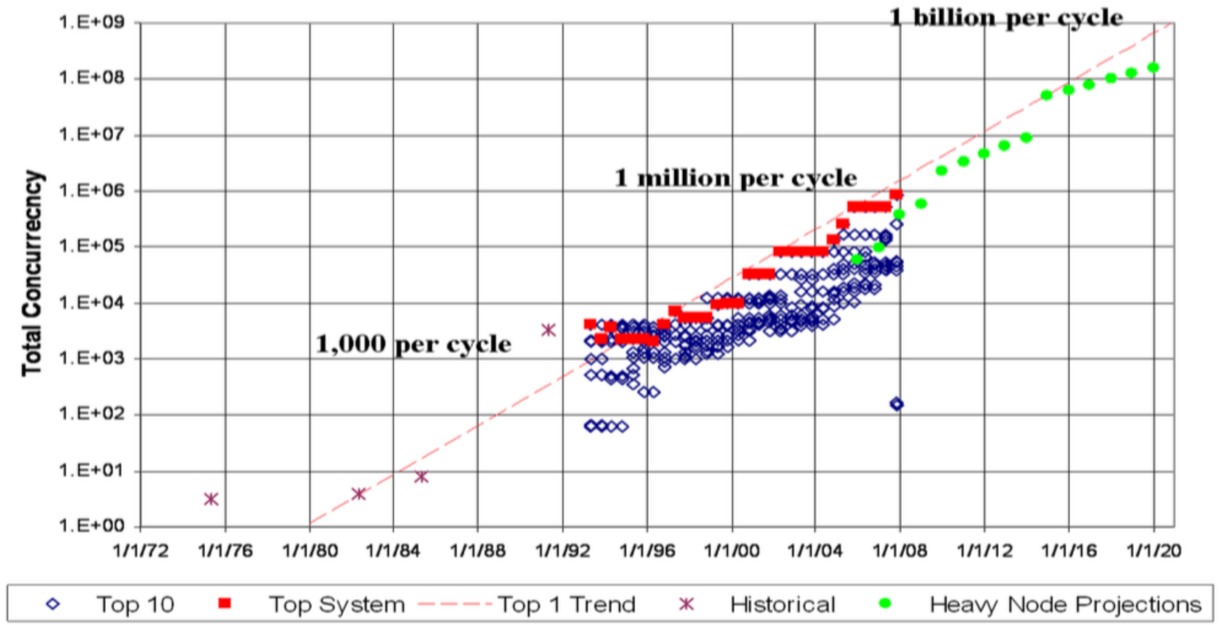

- Department of Energy's (DOE) Advanced Scientific Computing Research (ASCR) program plans for Exascale Era ➜ **"A lot more cores!"**
- Opens up exciting possibilities for HEP: in the light of significantly increasing resource needs (for example for the High Luminosity LHC)

New architectures

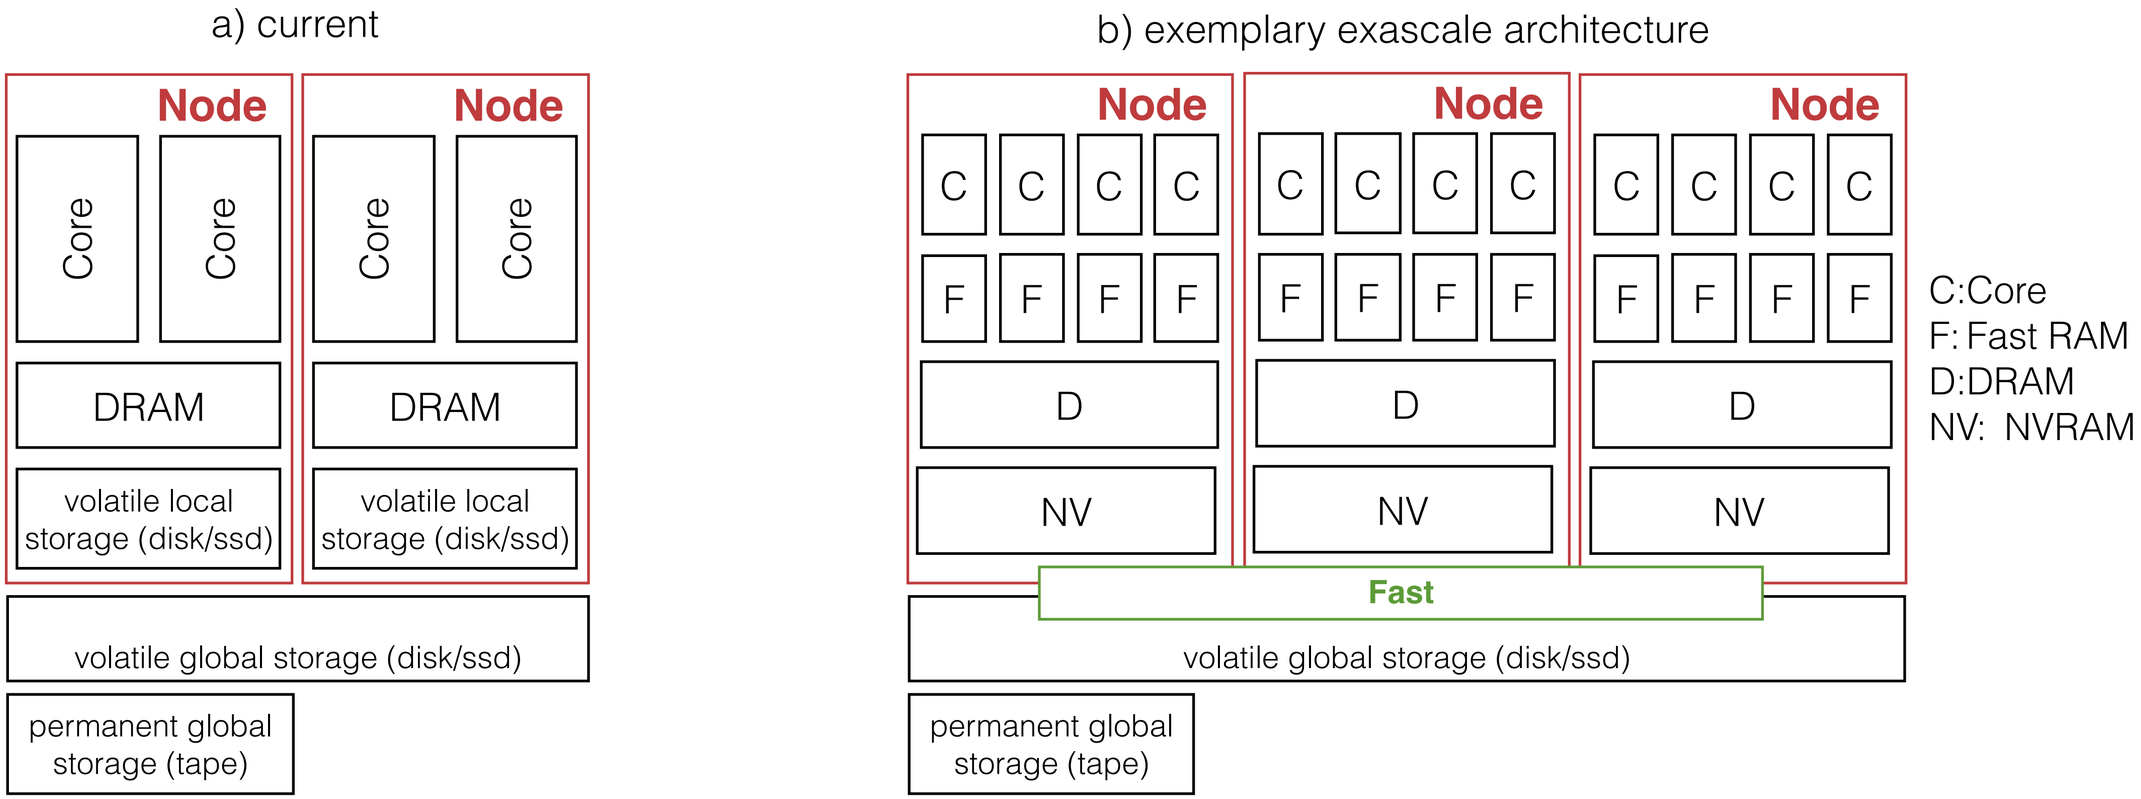

- HEP applications need a lot of memory and memory bandwidth
  - Cannot have both in Exascale machines ➜ **new architectures**
  - Requires to rethink how we design HEP applications!

Summary & Outlook

Summary & Outlook

- Software and Computing are integral parts of the HEP science process
  - Fermilab is home for a tremendous amount of software and computing expertise
  - We are supporting a wide variety of experiments from many corners of the HEP phase space

- For the future:
  - Learn multi-threaded programming!!!
  - Having to handle Exabytes of data is not that far off
    - Many new tools help you, both if you are working for a LHC collaboration, the Neutrino and Muon Experiment Community or any other HEP or non-HEP experiments
  - Science will look different in the Exascale era
    - Commercial clouds and Exascale HPC machines will change the way when and how we do computing

Acknowledgements

- Many thanks to Carlos Avila and his team and all at the Universidad de los Andes for the invitation

- Thanks to
  - All my colleagues who make running science software at unprecedented scales possible
  - All my colleagues who helped preparing this talk