Scaling Up CMS Tier-3 Data Processing

Kevin Lannon, Mike Hildreth
on behalf of ND CMS Group

Last year's Workshop: Data Preservation

2013 Physics Prize

Francois Englert

Peter Higgs

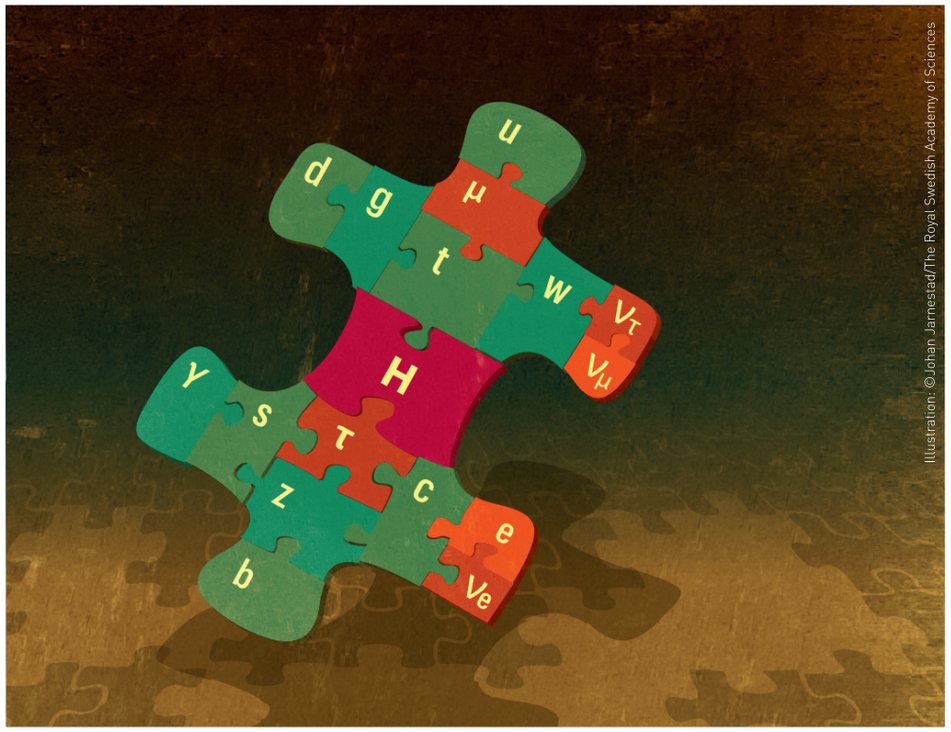

The Higgs Boson (aka "God Particle")

"...for the theoretical discovery of a mechanism that contributes to our understanding of the origin of mass of subatomic particles, and which recently was confirmed through the discovery of the predicted fundamental particle, by the ATLAS and CMS experiments at CERN's Large Hadron Collider"

2013 Physics Prize

Brought to you by the power of the grid...

Illustration: ©Johan Jarnestad/The Royal Swedish Academy of Sciences

Francois Englert

Boson (aka

"... mechanism
that co... ding of the
origin of m... ticles, and which
recently was co... gh the discovery of
the predicted fundam... particle, by the
ATLAS and CMS experiments at CERN's
Large Hadron Collider"


LHCb
LHCb
CMS
CMS
ATLAS
ATLAS
ALICE
ALICE

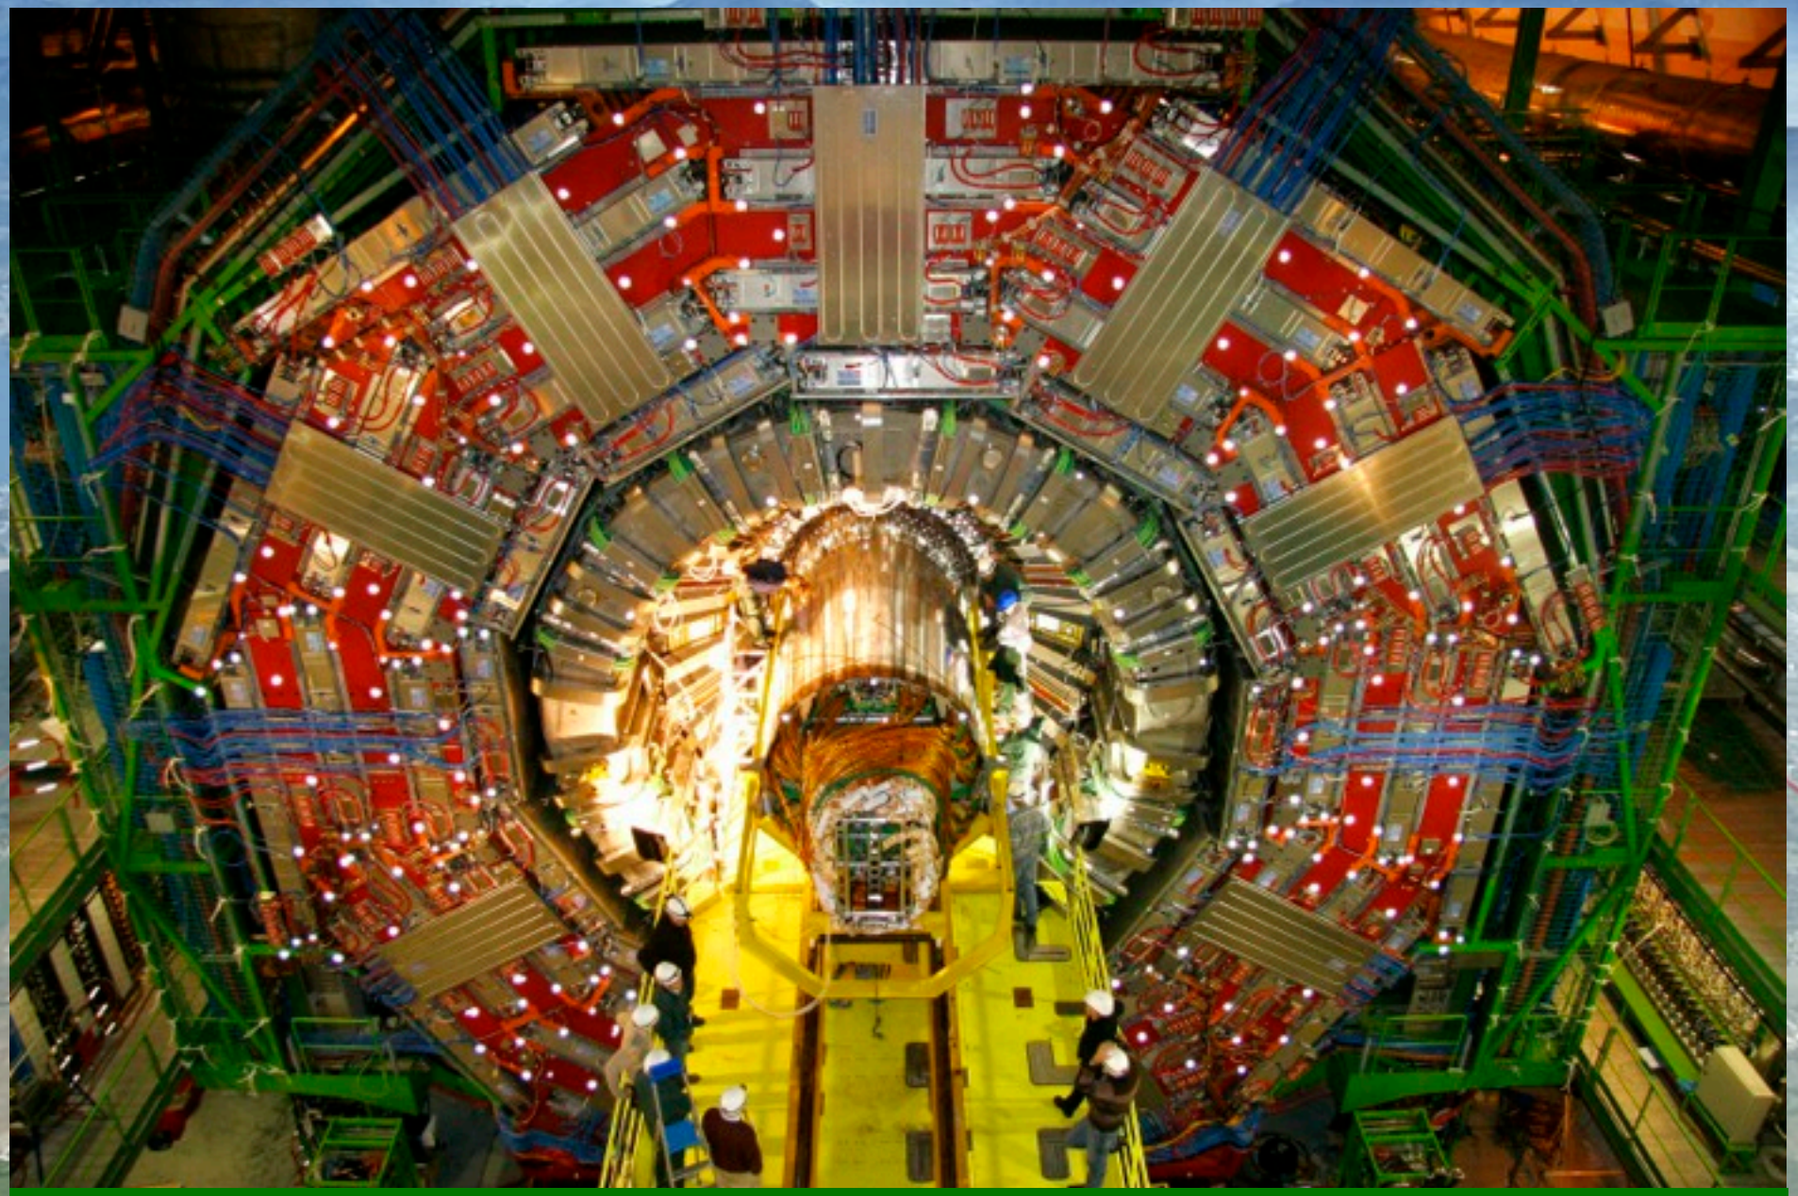

The CMS Experiment at the LHC (ND involvement)

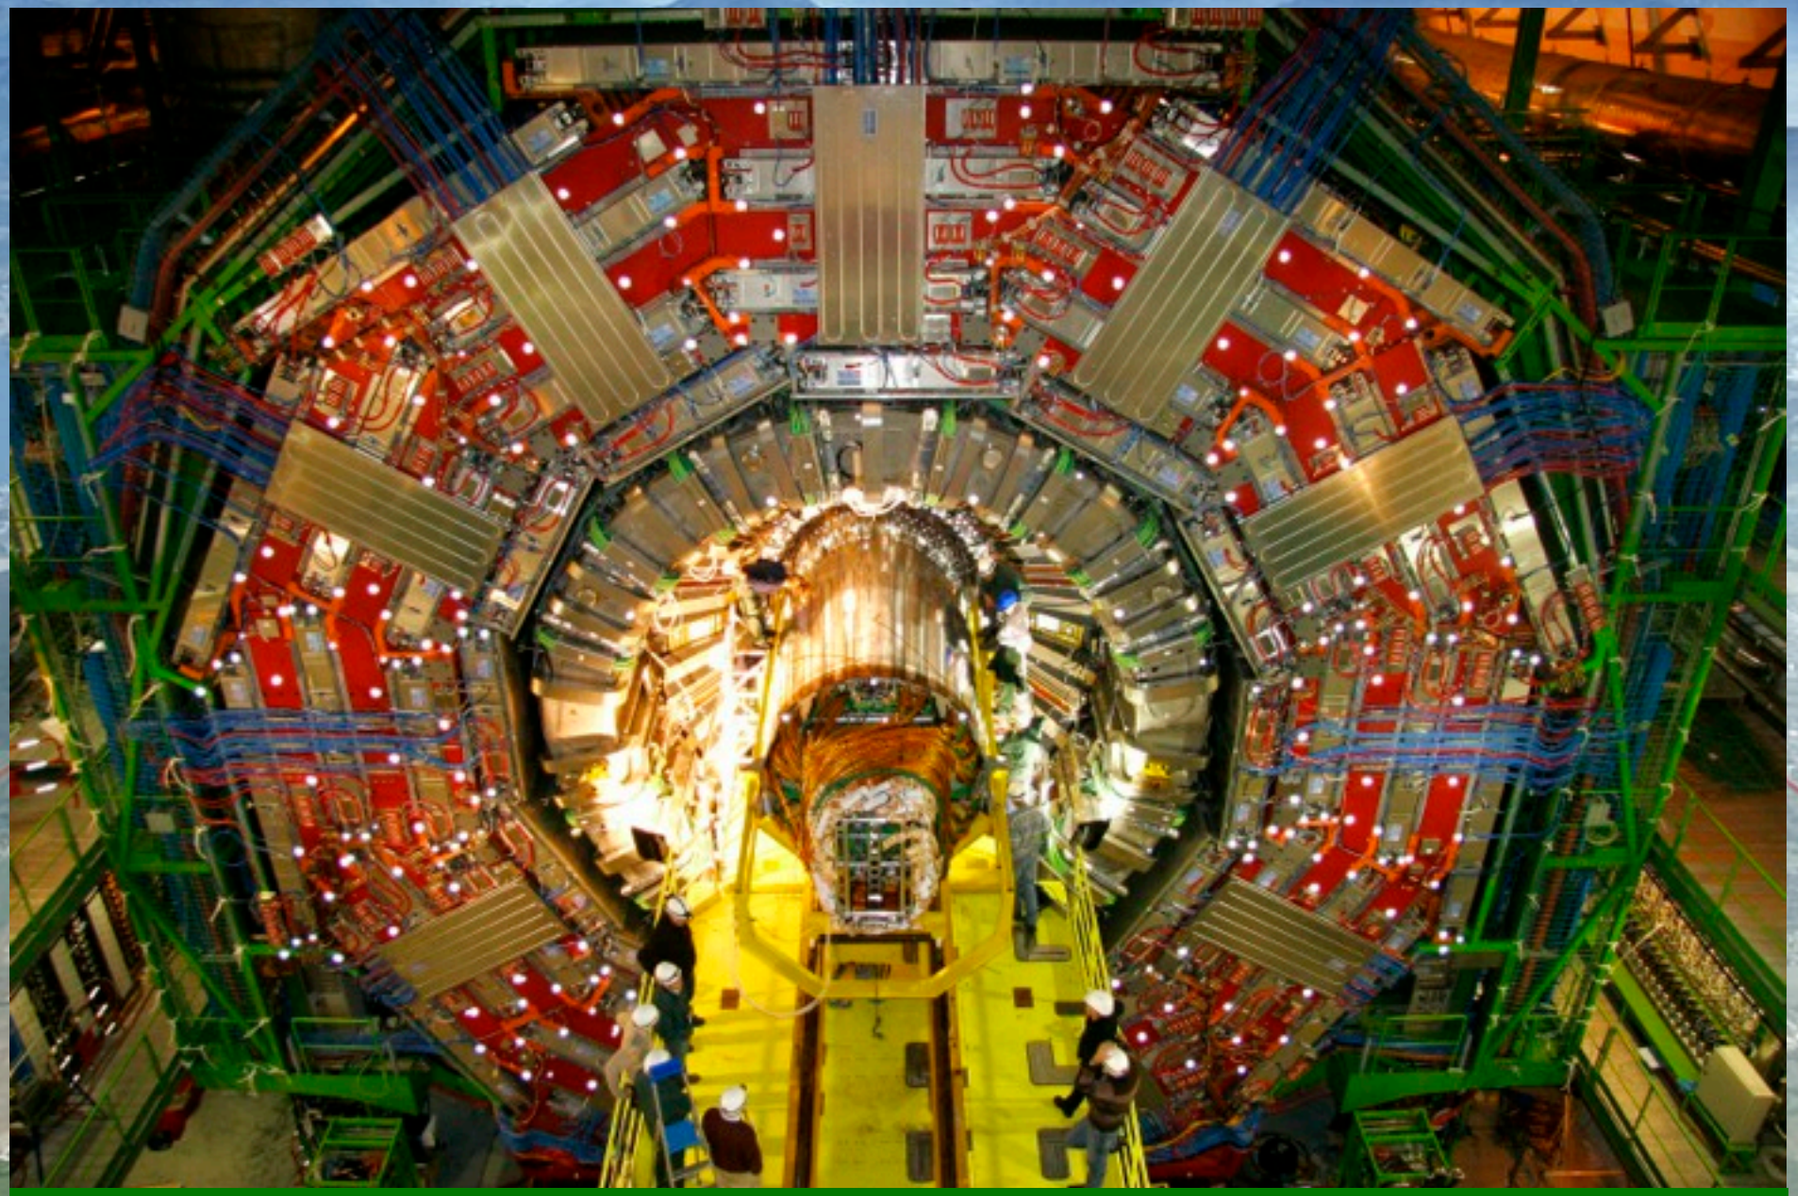

The CMS Experiment at the LHC (ND involvement)





Oceans of Data

Oceans of Data

- One Higgs boson produced every 3 billion collisions

Oceans of Data

- One Higgs boson produced every 3 billion collisions
- Peak rate of 9 Higgs bosons/minute

Oceans of Data

- One Higgs boson produced every 3 billion collisions
- Peak rate of 9 Higgs bosons/minute
- Total number of collisions produced to find Higgs: 690 trillion

Oceans of Data

- One Higgs boson produced every 3 billion collisions
- Peak rate of 9 Higgs bosons/minute
- Total number of collisions produced to find Higgs: 690 trillion
- If each collisions were one grain of sand... 

Oceans of Data

- Would fill 17 Olympic-sized swimming pools

Oceans of Data

- Would fill 17 Olympic-sized swimming pools

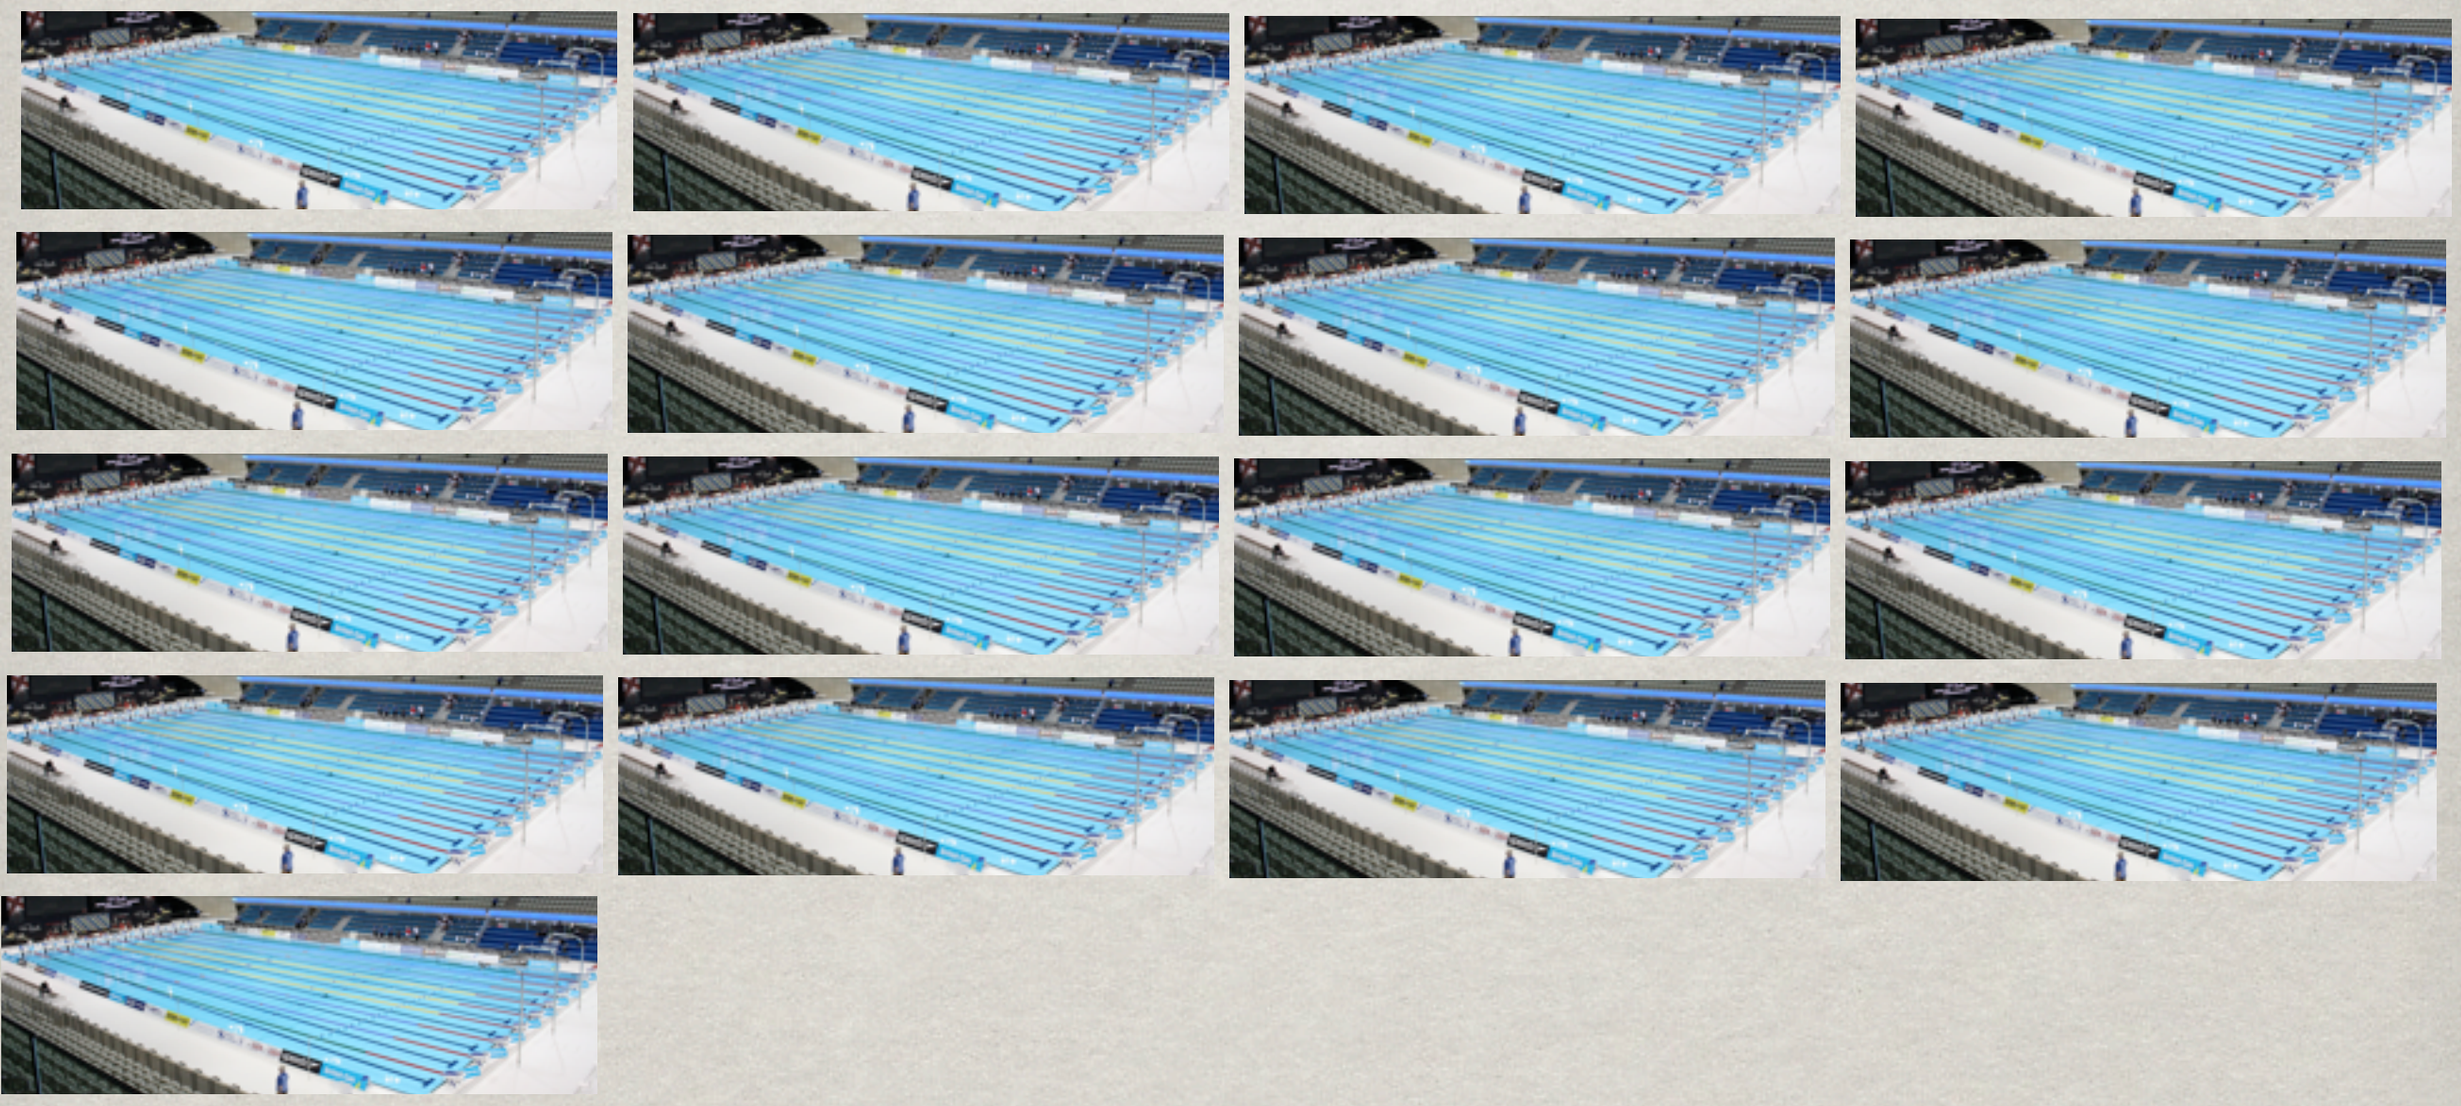

Oceans of Data

- Would fill 17 Olympic-sized swimming pools

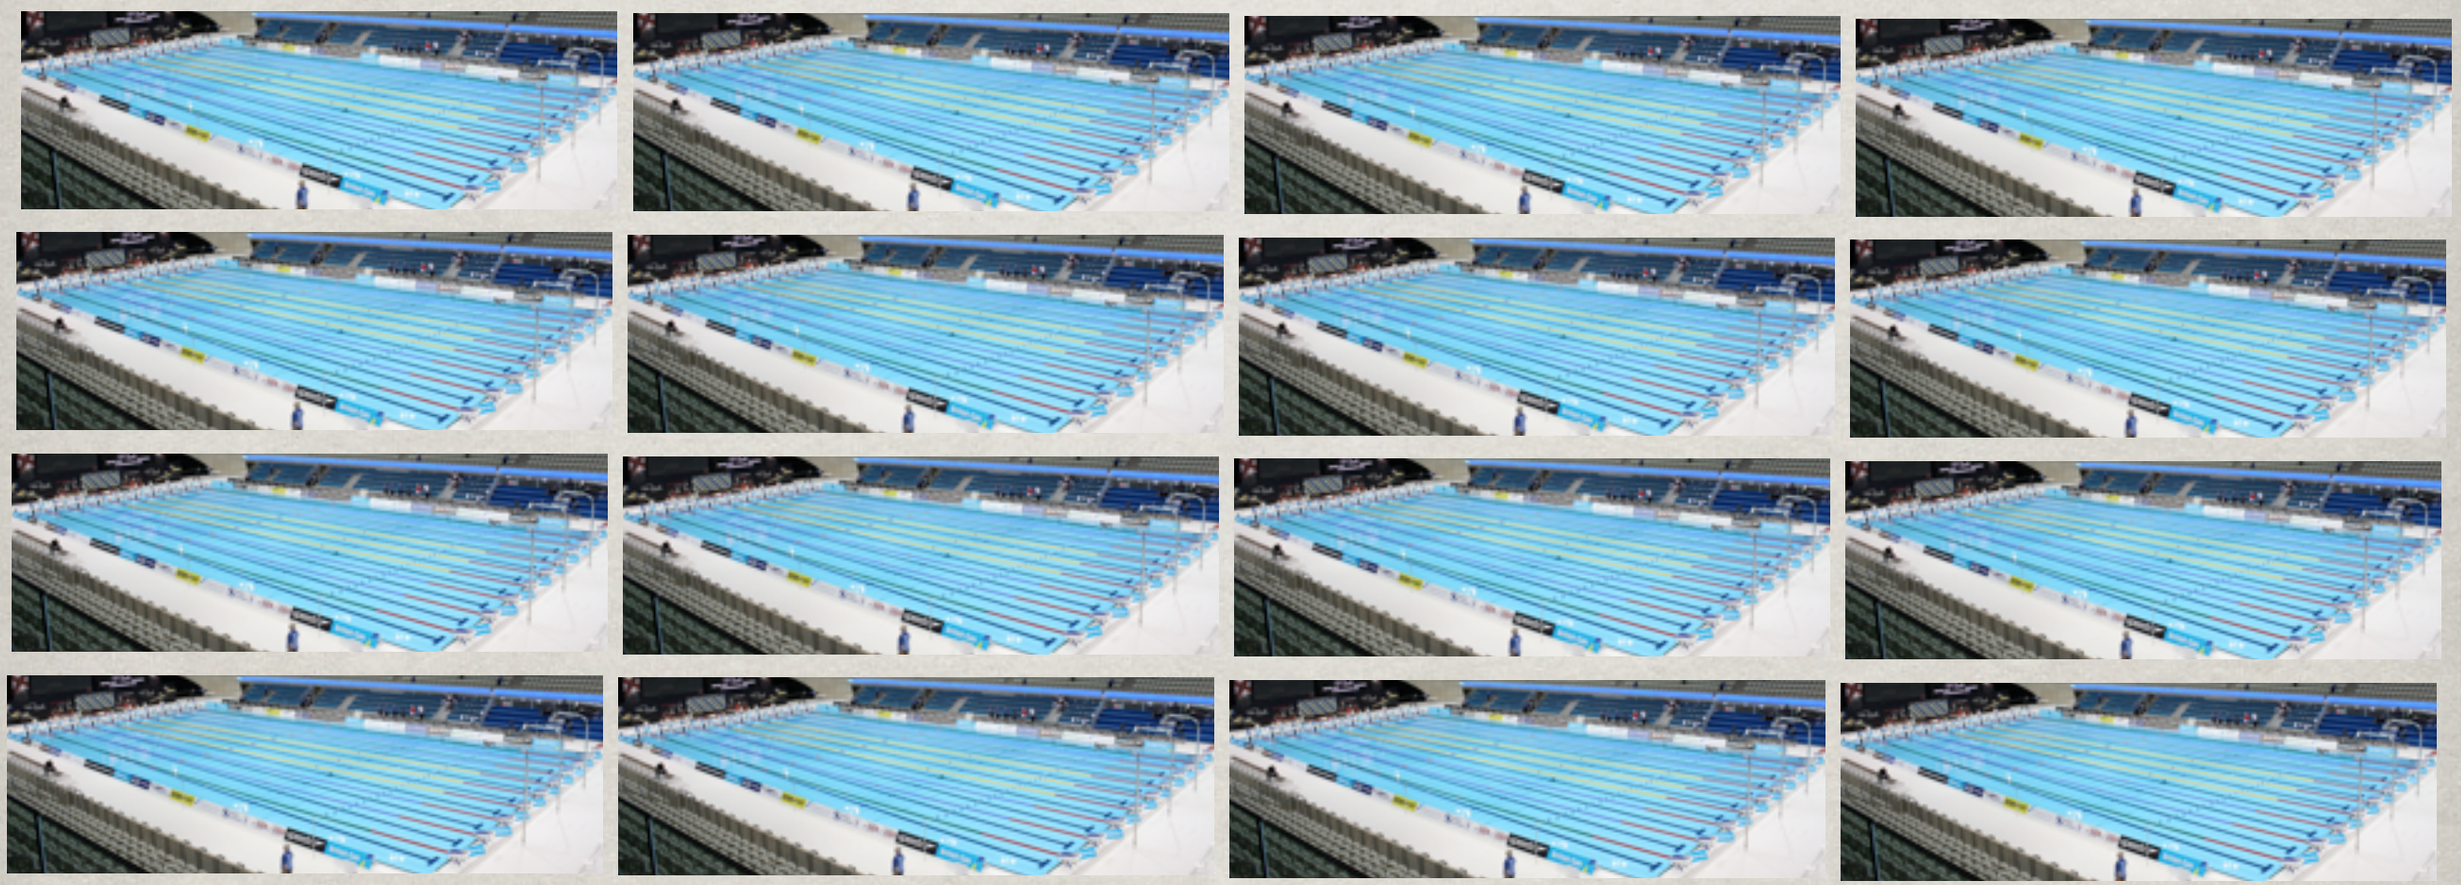

All of the Higgs collisions would fill 1/2 Tbsp

How fast do we need to Go?

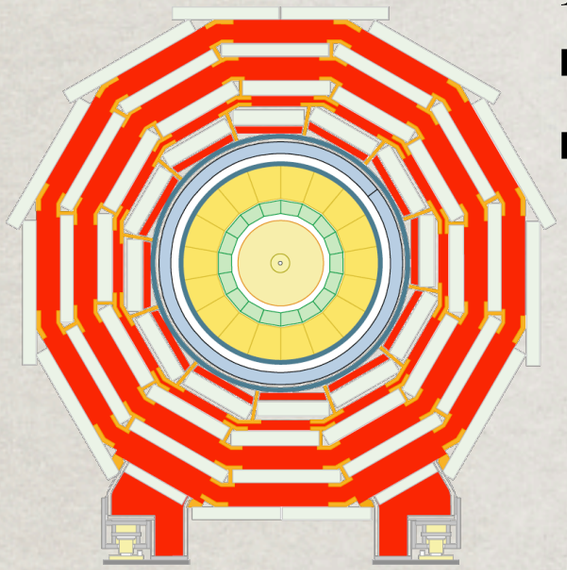

Basic facts:

➡ Data from detector: 200 kB/ collision

➡ Processing time for analysis: 5 sec (basic)

		Proton Collisions in Detector	Level 1 Trigger	High Level Trigger
	Data Rate	16 MHz	60 kHz	300 Hz
For 1 year's worth of data	Data Collected	50 EB	200 PB	1-2 PB
	Processing time	45 Million CPU years!	170 Thousand CPU years	860 CPU years

How fast do we need to Go?

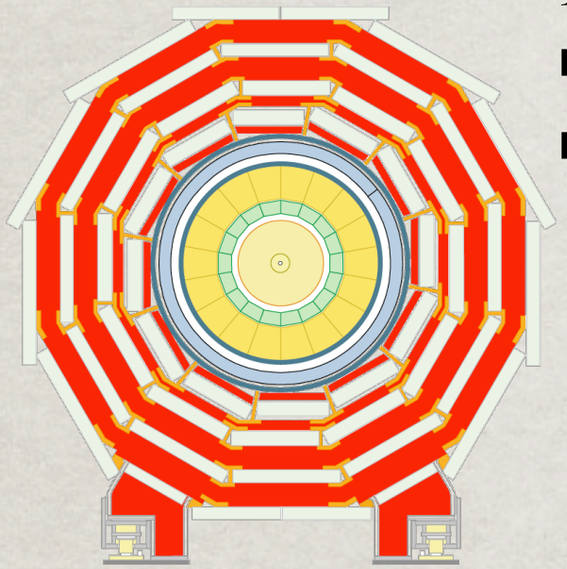

Basic facts:

➡ Data from detector: 200 kB/ collision

➡ Processing time for analysis: 5 sec (basic)

TB:

1 Million ×

	Proton Collisions in Detector	Level 1 Trigger	High Level Trigger
Data Rate	16 MHz	60 kHz	300 Hz
Data Collected	50 EB	~~200 PB~~	1-2 PB
Processing time	45 Million CPU years!	170 Thousand CPU years	860 CPU years

For 1 year's worth of data

How fast do we need to Go?

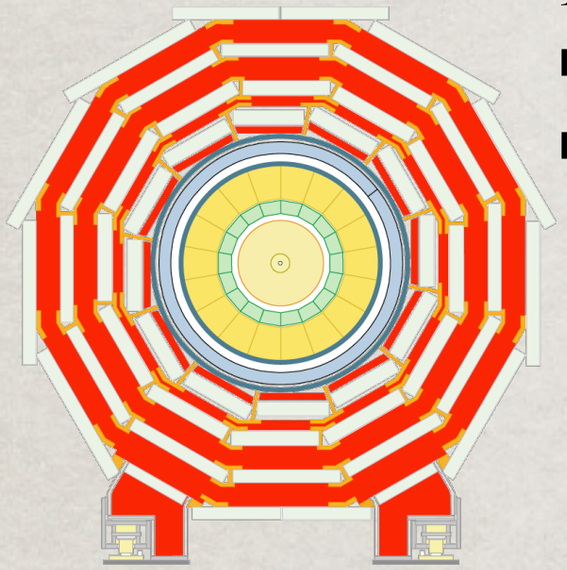

Basic facts:

➡ Data from detector: 200 kB/ collision

➡ Processing time for analysis: 5 sec (basic)

	Proton Collisions in Detector	Level 1 Trigger
Data Rate	16 MHz	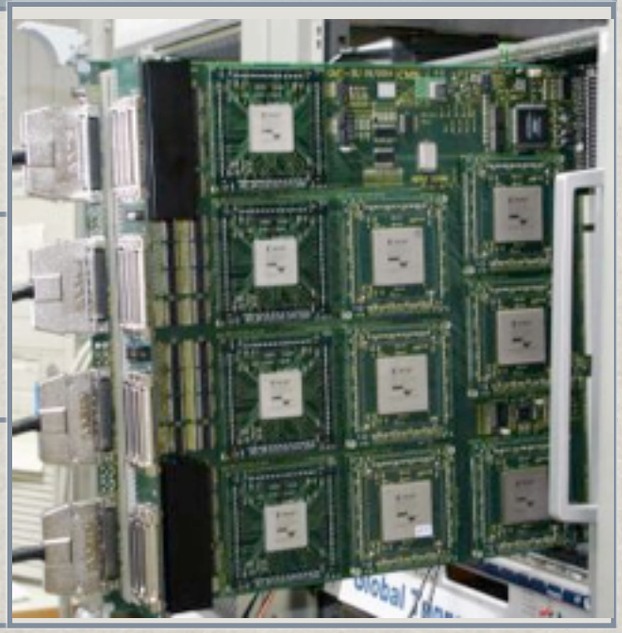
Data Collected	50 EB	
Processing time	45 Million CPU years!	

For 1 year's worth of data

How fast do we need to Go?

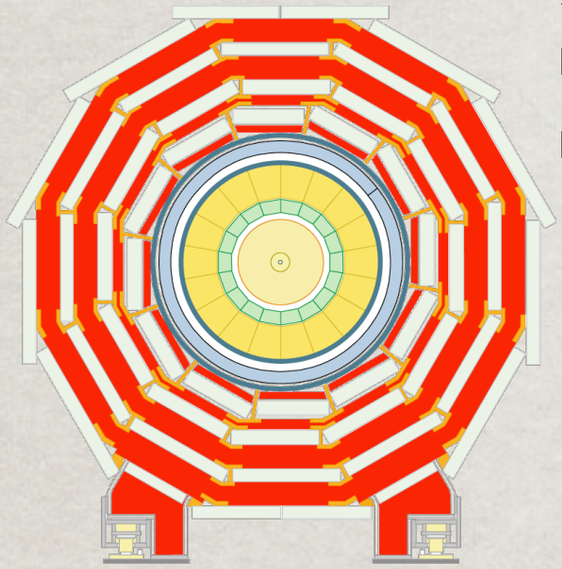

Basic facts:

➡ Data from detector: 200 kB/ collision

➡ Processing time for analysis: 5 sec (basic)

FPGA Chips do very simple analysis ~ μs to analyze data

	Proton Collisions in Detector	Level 1 Trigger
Data Rate	16 MHz	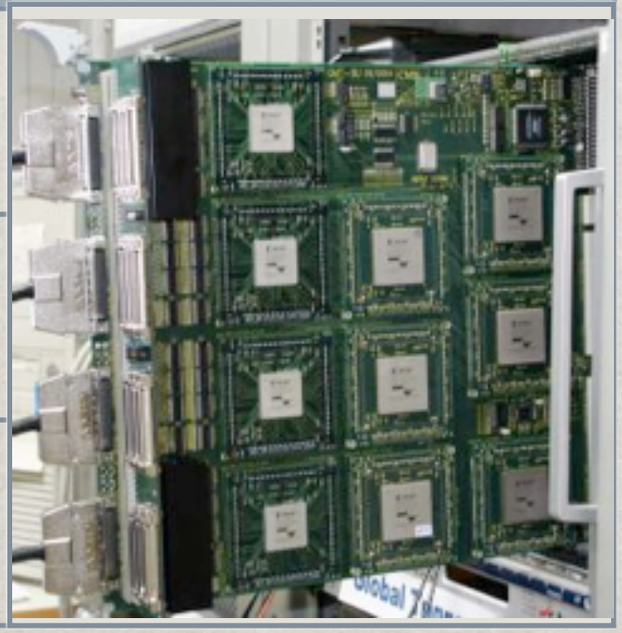
Data Collected	50 EB	
Processing time	45 Million CPU years!	

For 1 year's worth of data

How fast do we need to Go?

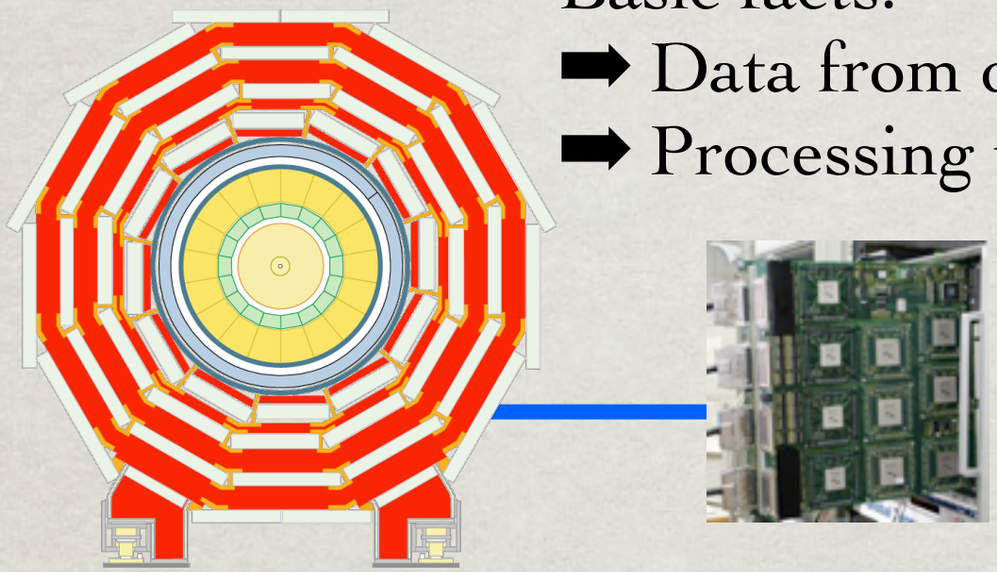

Basic facts:

➡ Data from detector: 200 kB/ collision

➡ Processing time for analysis: 5 sec (basic)

	Proton Collisions in Detector	Level 1 Trigger
Data Rate	16 MHz	60 kHz
Data Collected	50 EB	200 PB
Processing time	45 Million CPU years!	170 Thousand CPU years

For 1 year's worth of data

How fast do we need to Go?

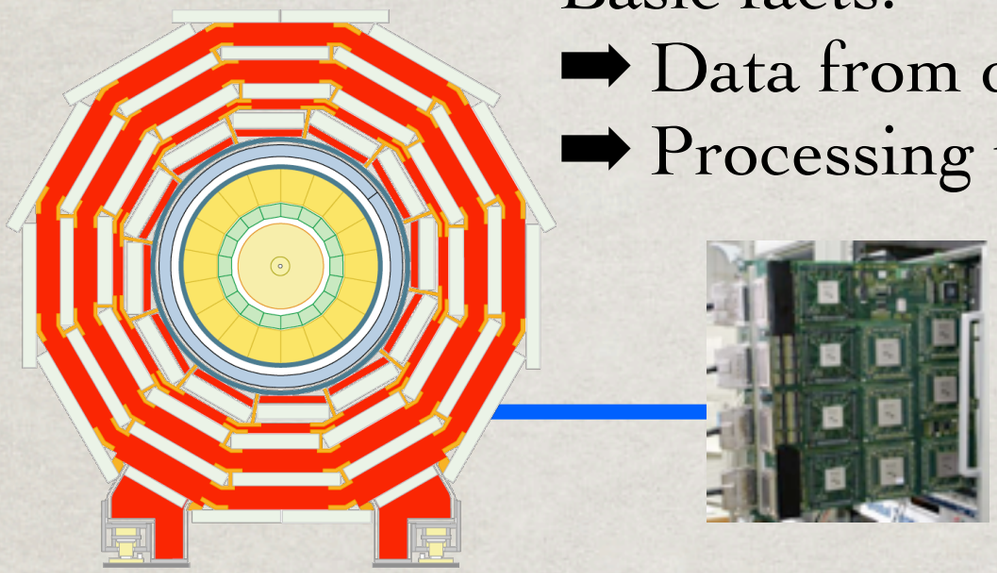

Basic facts:

➡ Data from detector: 200 kB/ collision

➡ Processing time for analysis: 5 sec (basic)

	Proton Collisions in Detector	Level 1 Trigger	High Level Trigger
Data Rate	16 MHz	60 kHz	Computer farm with ~ 3000 CPUs High speed network/ switch
Data Collected	50 EB	200 PB	
Processing time	45 Million CPU years!	170 Thousand CPU years	

For 1 year's worth of data

How fast do we need to Go?

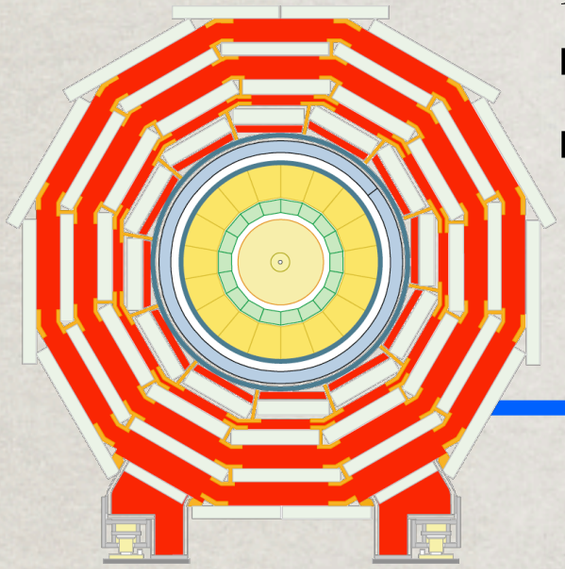

Basic facts:

➡ Data from detector: 200 kB/ collision

➡ Processing time for analysis: 5 sec (basic)

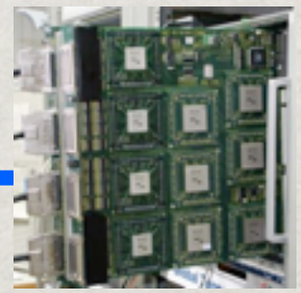

Simplified analysis code

~ ms to analyze data

	Proton Collisions in Detector	Level 1 Trigger	High Level Trigger
Data Rate	16 MHz	60 kHz	Computer farm with ~ 3000 CPUs High speed network/ switch
Data Collected	50 EB	200 PB	
Processing time	45 Million CPU years!	170 Thousand CPU years	

For 1 year's worth of data

How fast do we need to Go?

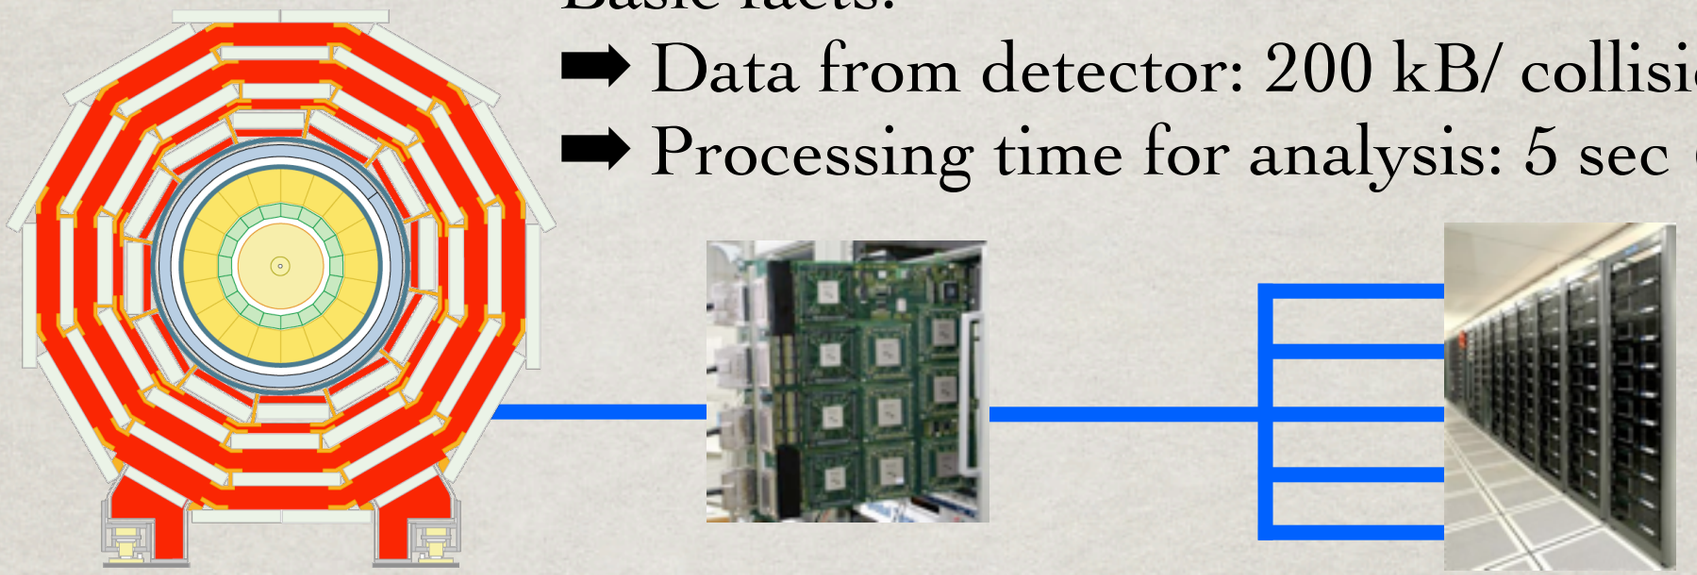

Basic facts:

➡ Data from detector: 200 kB/ collision

➡ Processing time for analysis: 5 sec (basic)

	Proton Collisions in Detector	Level 1 Trigger	High Level Trigger
Data Rate	16 MHz	60 kHz	300 Hz
For 1 year's worth of data: Data Collected	50 EB	200 PB	1-2 PB
For 1 year's worth of data: Processing time	45 Million CPU years!	170 Thousand CPU years	860 CPU years

How fast do we need to Go?

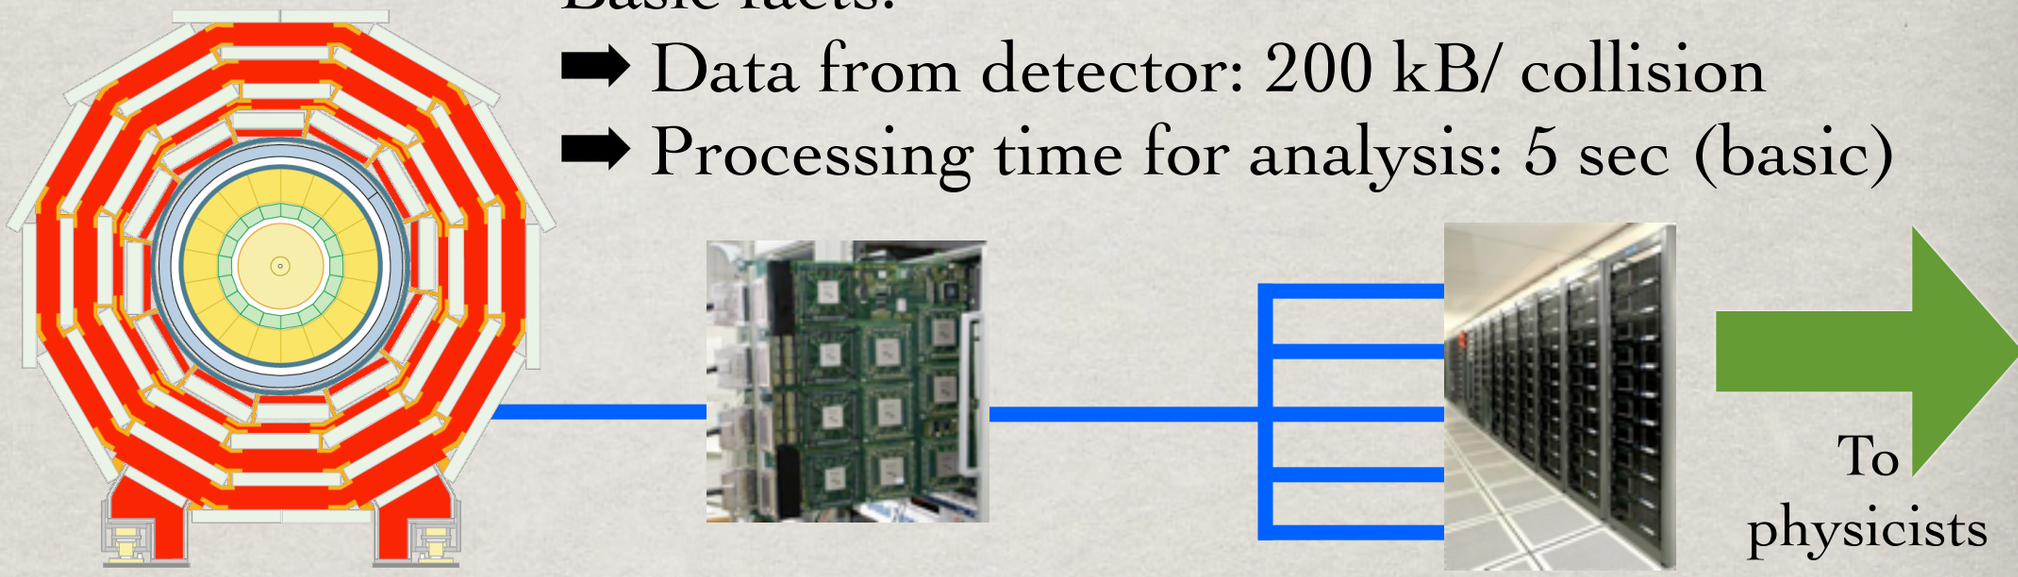

Basic facts:

➡ Data from detector: 200 kB/ collision

➡ Processing time for analysis: 5 sec (basic)

For 1 year's worth of data

	Proton Collisions in Detector	Level 1 Trigger	High Level Trigger
Data Rate	16 MHz	60 kHz	300 Hz
Data Collected	50 EB	200 PB	1-2 PB
Processing time	45 Million CPU years!	170 Thousand CPU years	860 CPU years

Processing CMS Data

~ 5-10 PB/year

Simulated Data

~ 2 PB/year

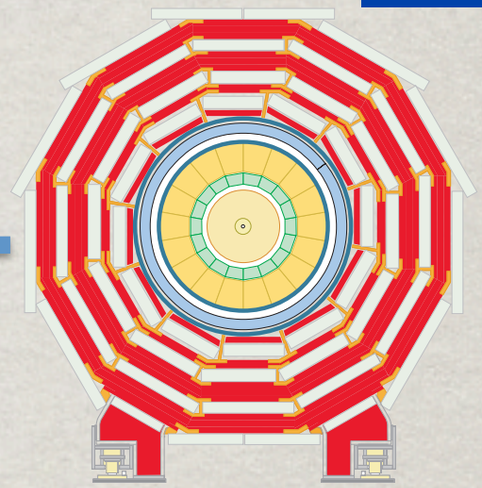

CMS detector

Prompt Reco.

Physics Analysis

Re-reco.

Analysis improvements

Physic Results

Worldwide LHC Computing Grid

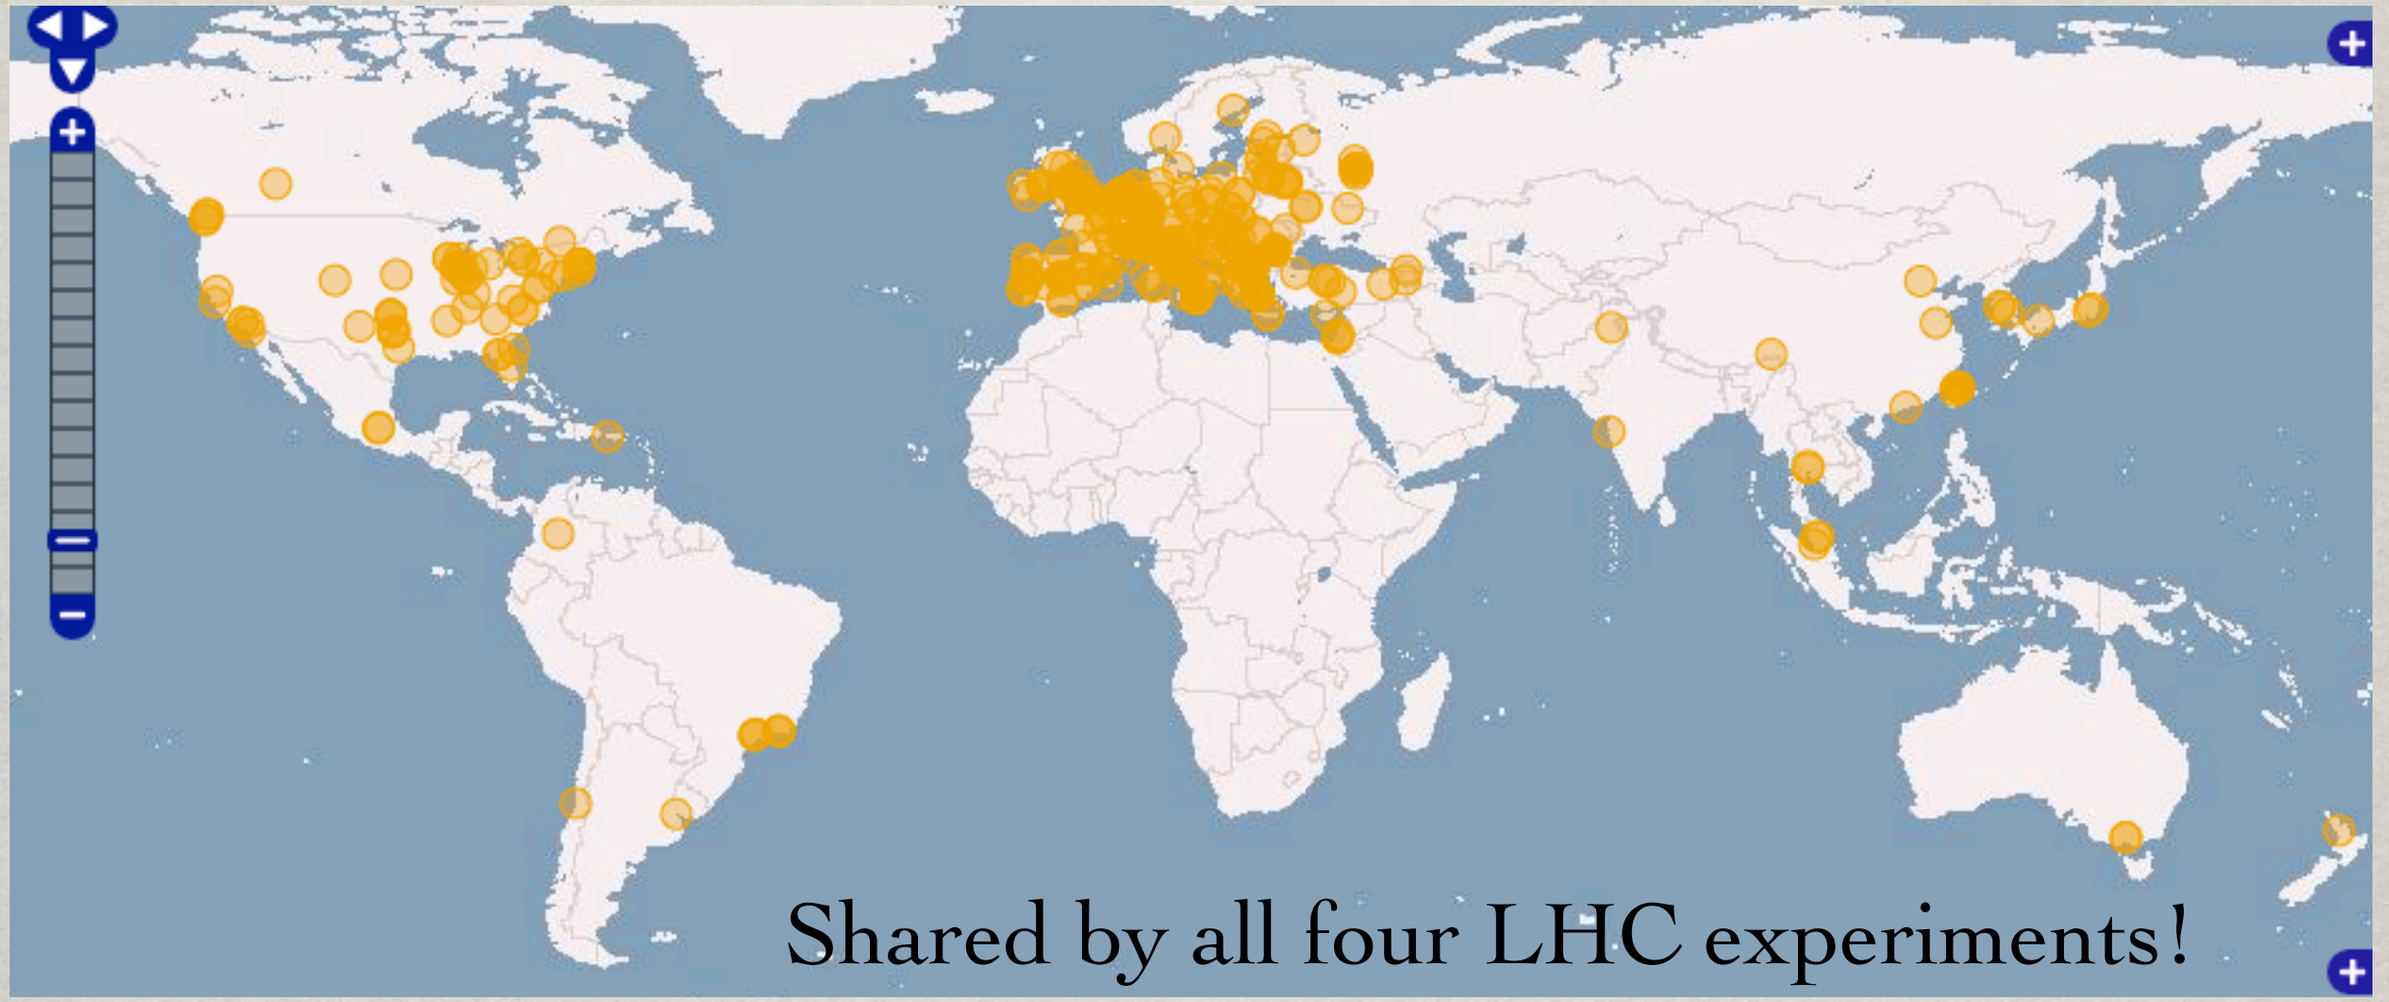

Worldwide LHC Computing Grid

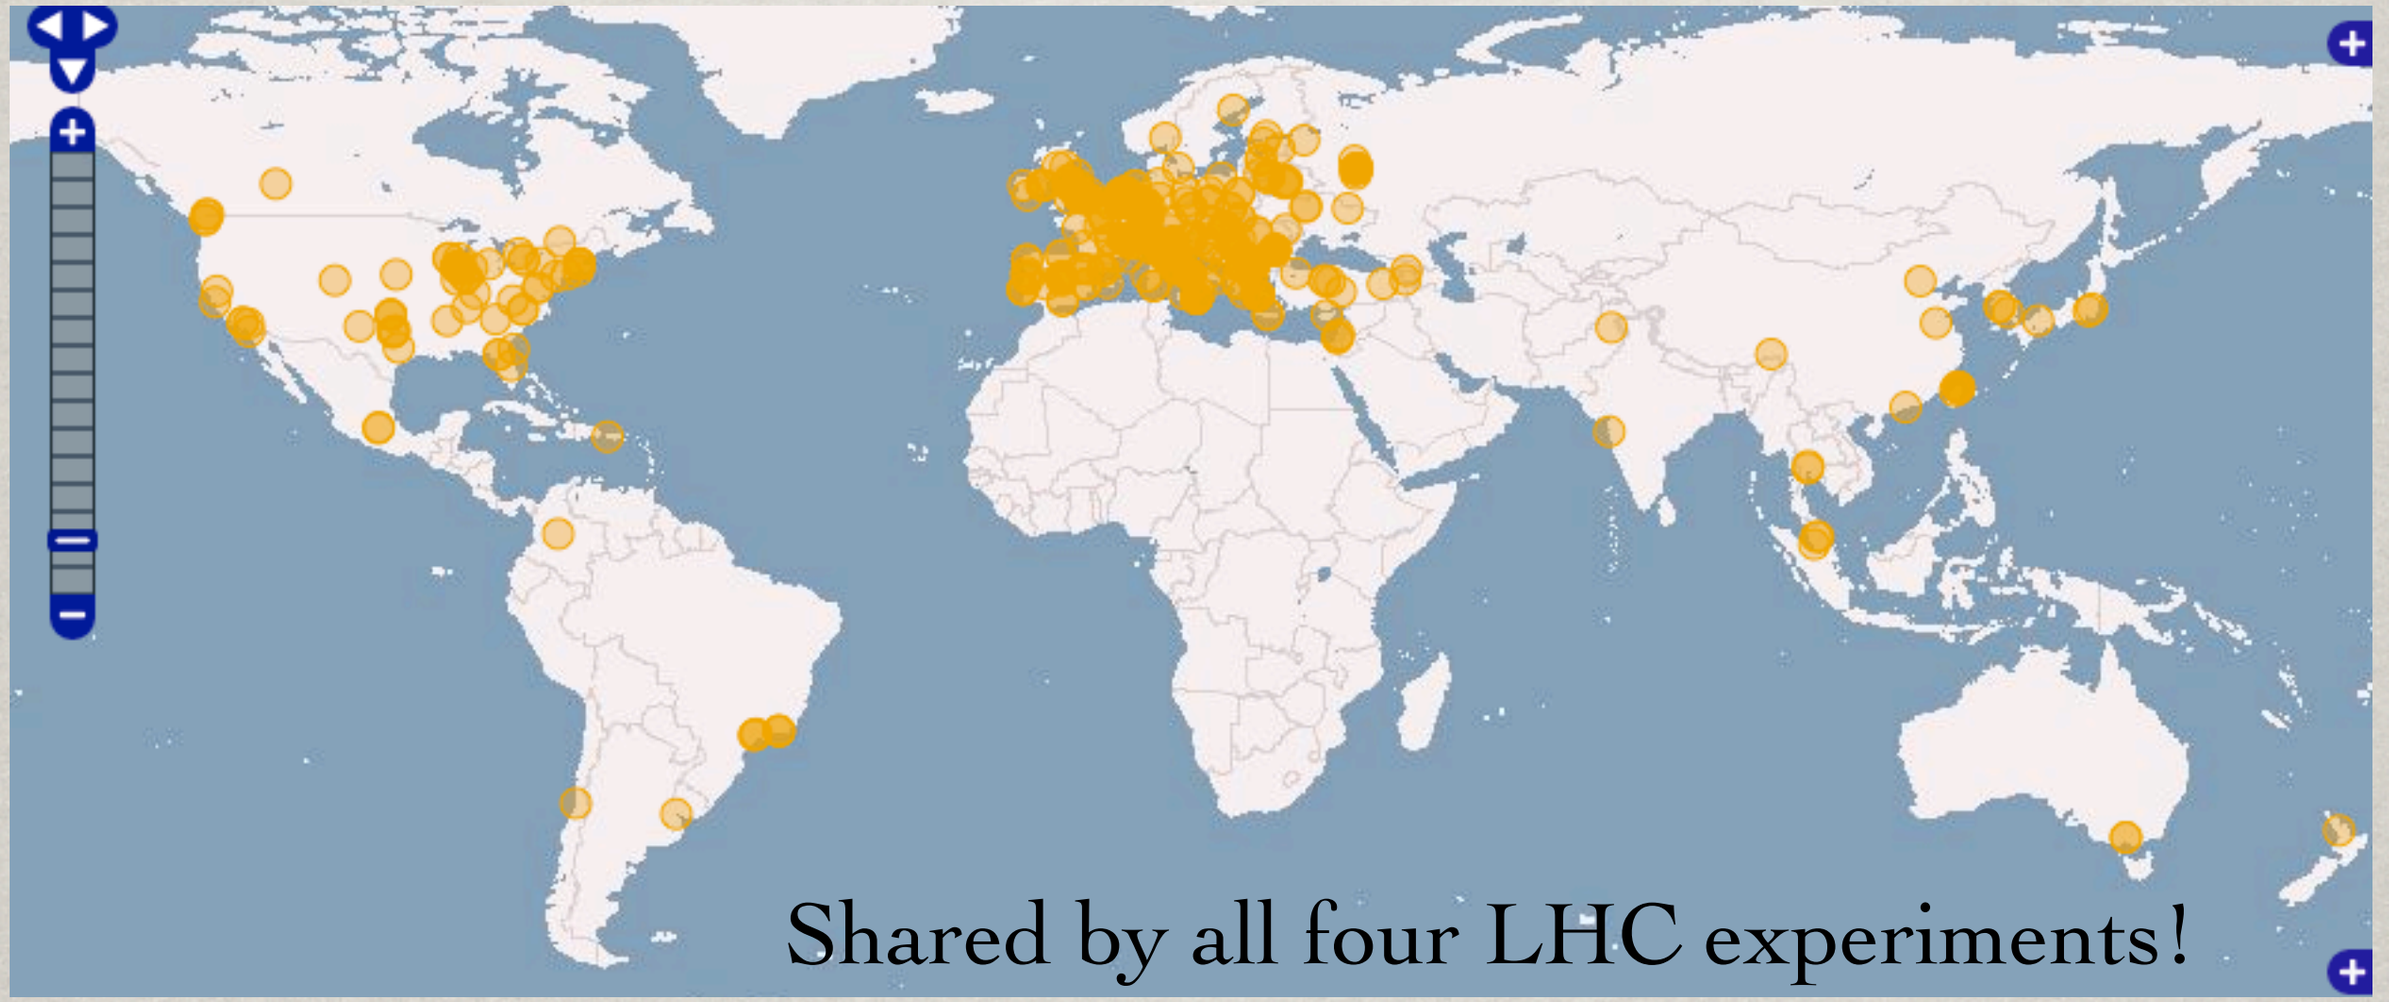

- Over 160 sites around world (including OSG sites in US)

Worldwide LHC Computing Grid

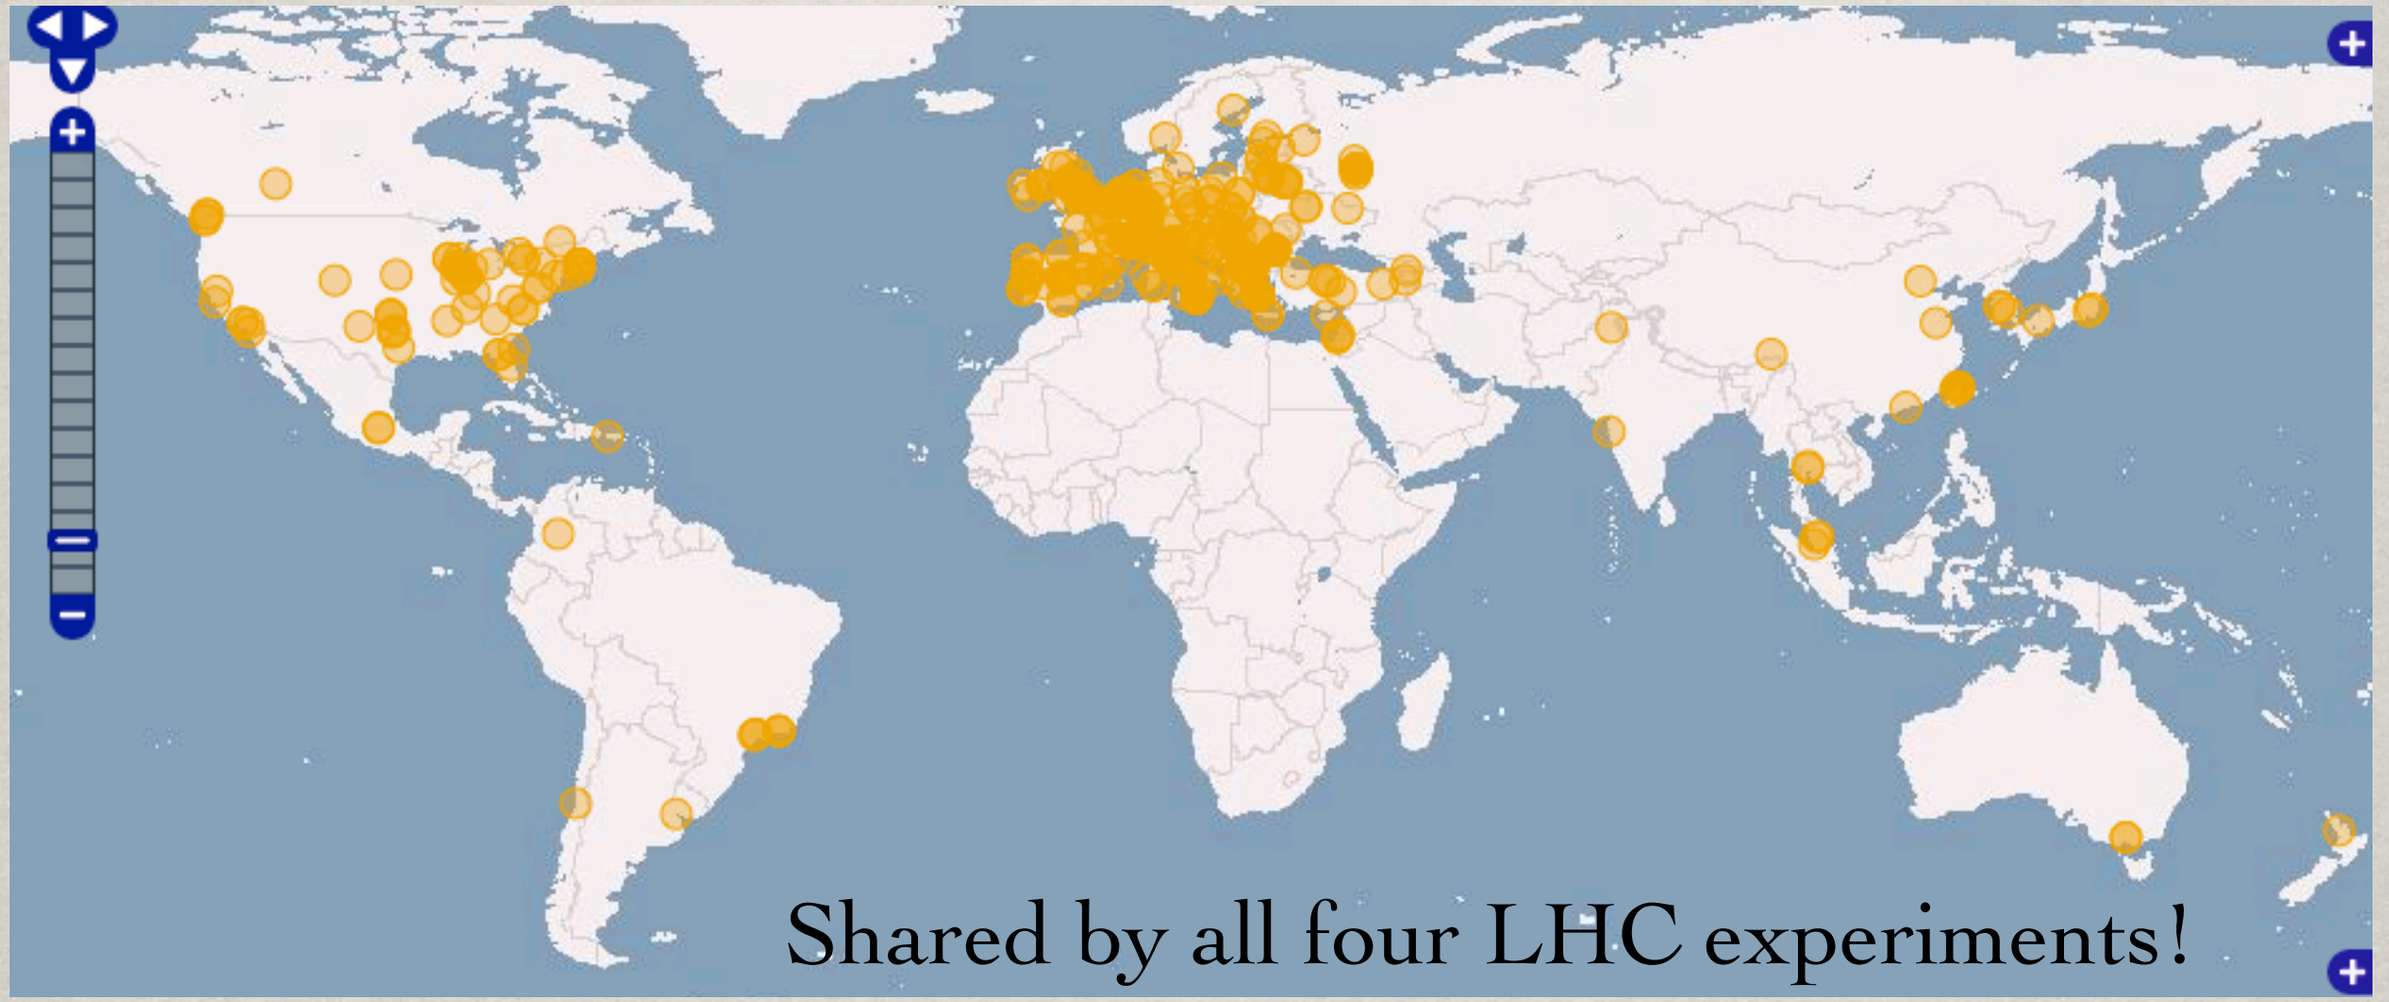

- Over 160 sites around world (including OSG sites in US)
- > 200k CPU cores available

Worldwide LHC Computing Grid

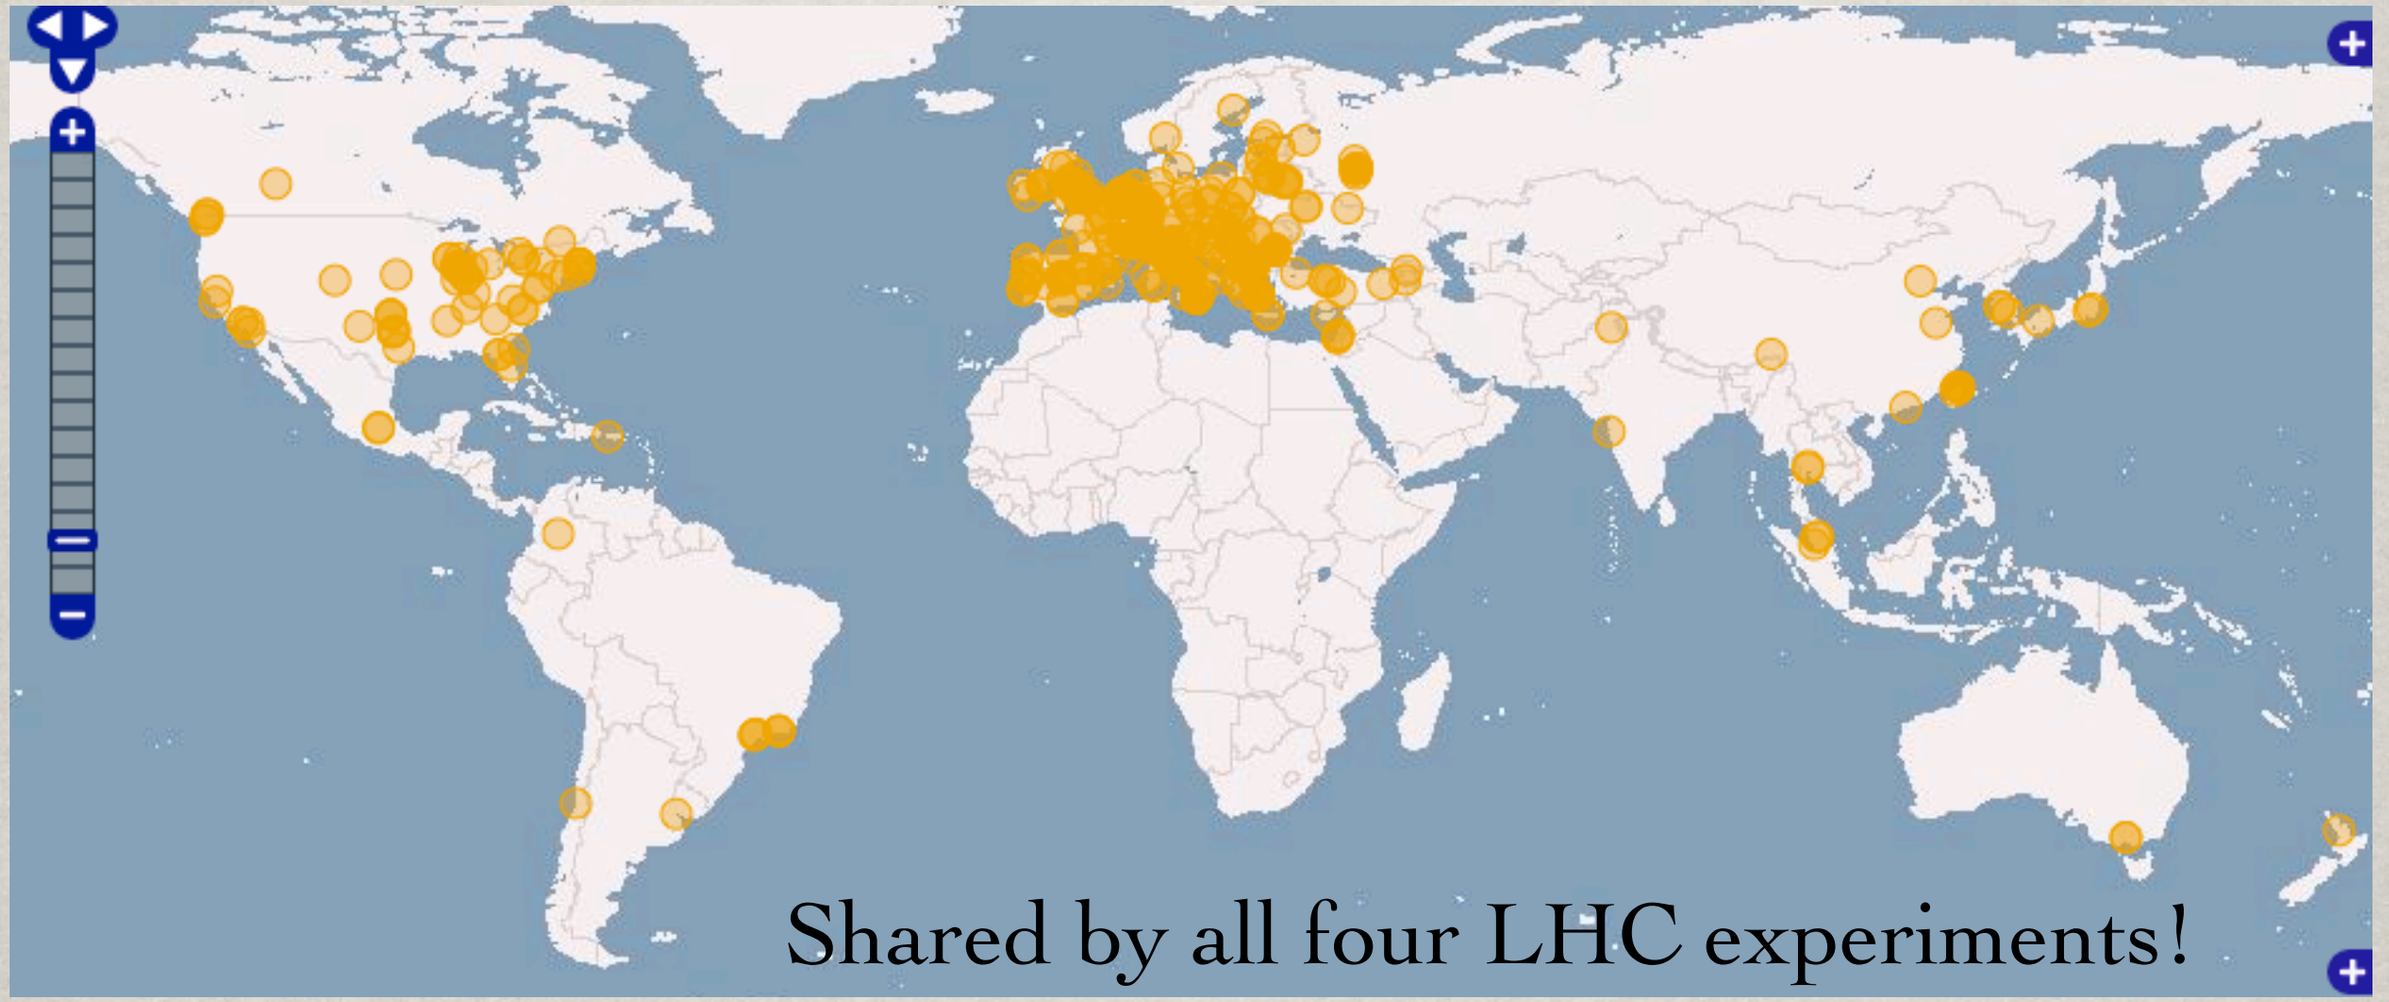

- Over 160 sites around world (including OSG sites in US)
- > 200k CPU cores available
- Has gone as high as ~ 1 million jobs submitted in a single day

Worldwide LHC Computing Grid

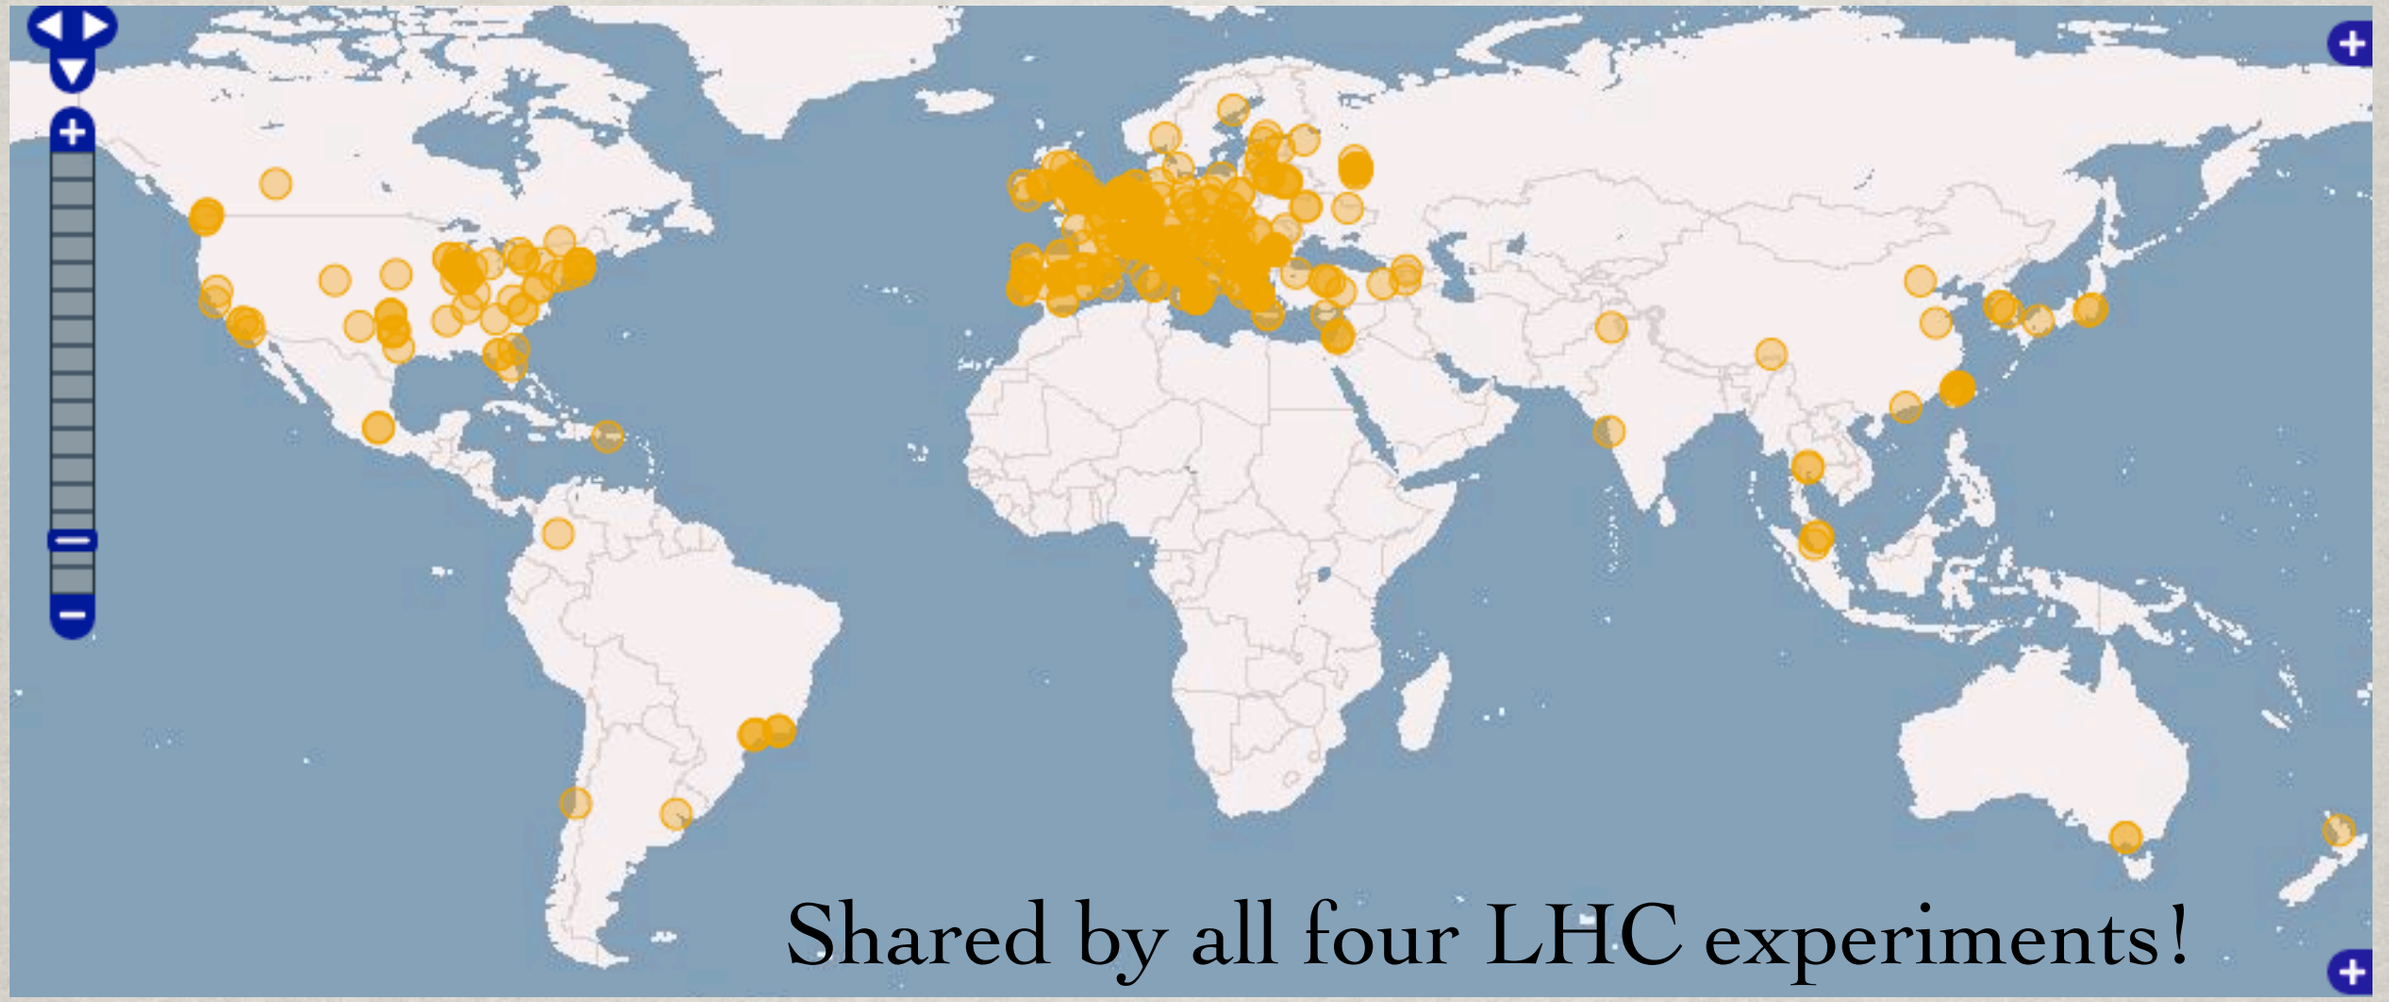

- Over 160 sites around world (including OSG sites in US)
- > 200k CPU cores available
- Has gone as high as ~ 1 million jobs submitted in a single day
- > 300 PB of total storage available

Organization

Tier 0

CERN

Organization

CERN

- All LHC data passes through T0 for initial processing
- Provides less than 20% of total CPU resources for LHC experiments
- Basic data processing common to all analyses

Organization

Tier 1
Taipei

Tier 1
Canada

Tier 1
Germany

Tier 1
Spain

Tier 1
Netherlands

Tier 0

Tier 1
France

Tier 1
United Kingdom

Tier 1
USA (FNAL)

Tier 1
USA (BNL)

Tier 1
Nordic Countries

Tier 1
Italy

Organization

Tier 1 — Taipei

Tier 1 — Canada

Tier 1 — Germany

Tier 1 — Spain

Tier 1 — Netherlands

Tier 1 — France

- Distribute data around the world
- Provide CPU for central reprocessing
- Generate simulated data

Tier 0

Tier 1 — United Kingdom

Tier 1 — Italy

Tier 1 — USA (FNAL)

Tier 1 — USA (BNL)

Tier 1 — Nordic Countries

Organization

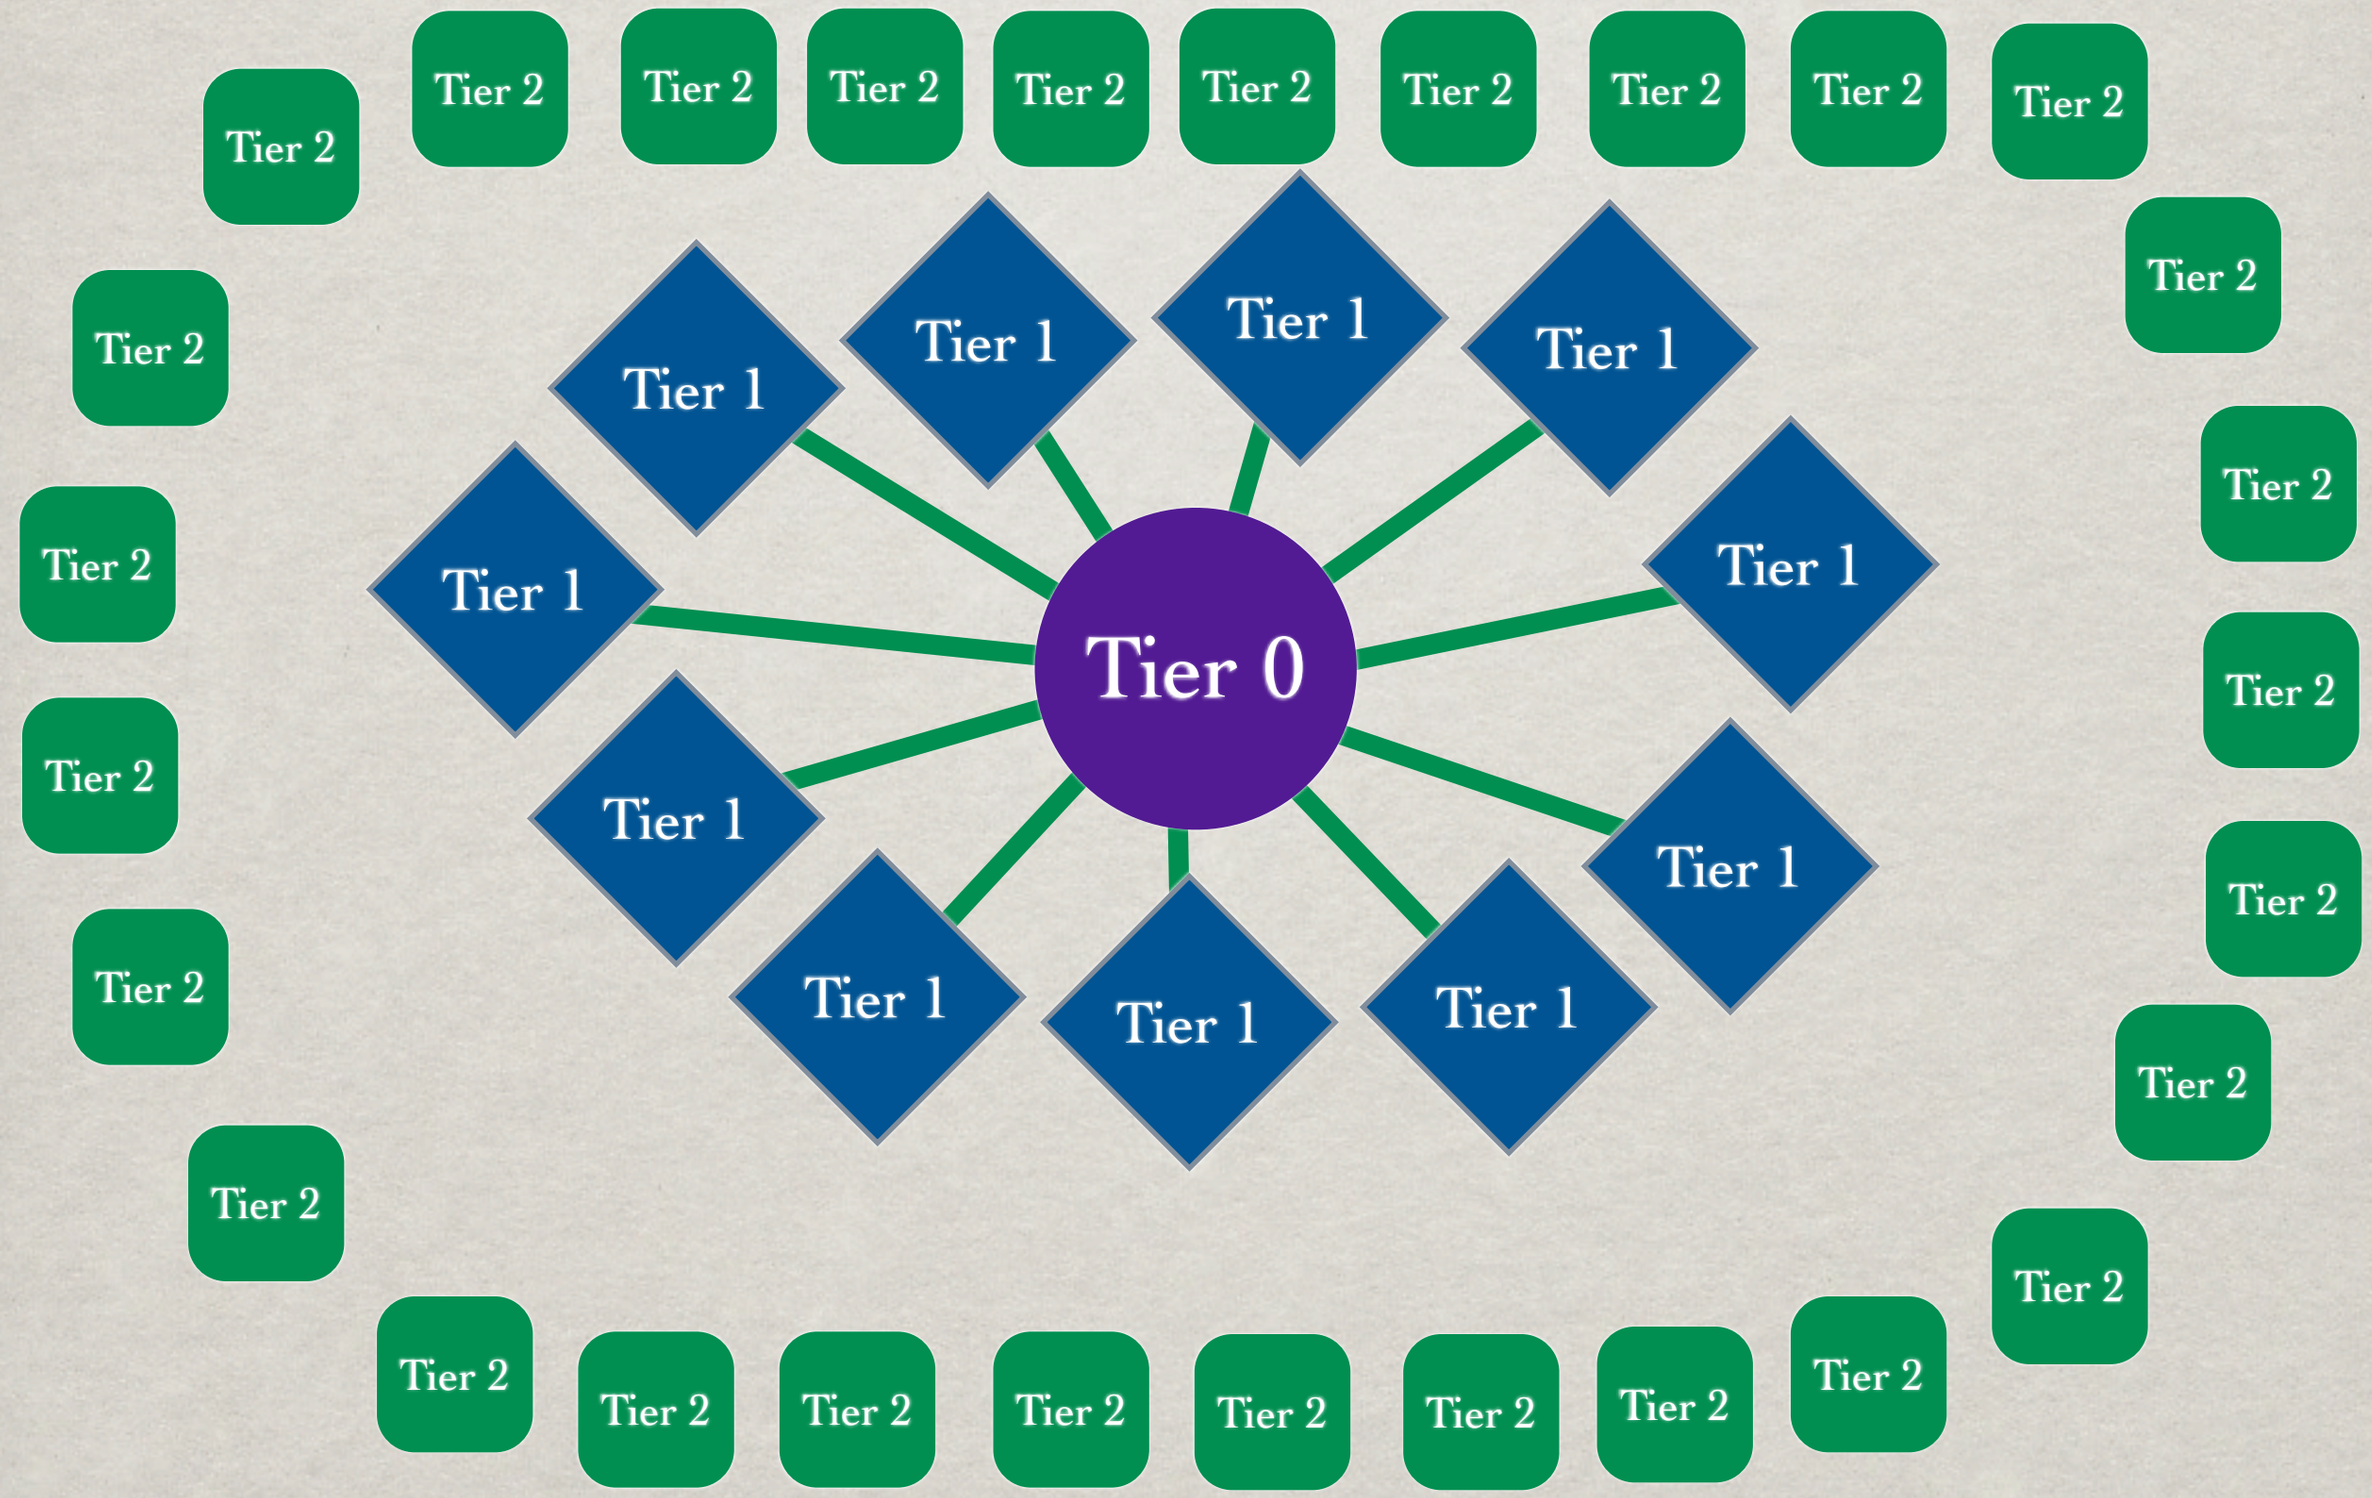

Organization

- Over 140 T2 sites throughout the world
- Average site has ~800 CPU's and 300 TB storage
  - Some have much more: > 1 PB storage!
- Provide CPU and storage for analysis of data, plus some simulation
- This is where "average user" runs analysis
- Connected via regional internet links

Where ND Fits In

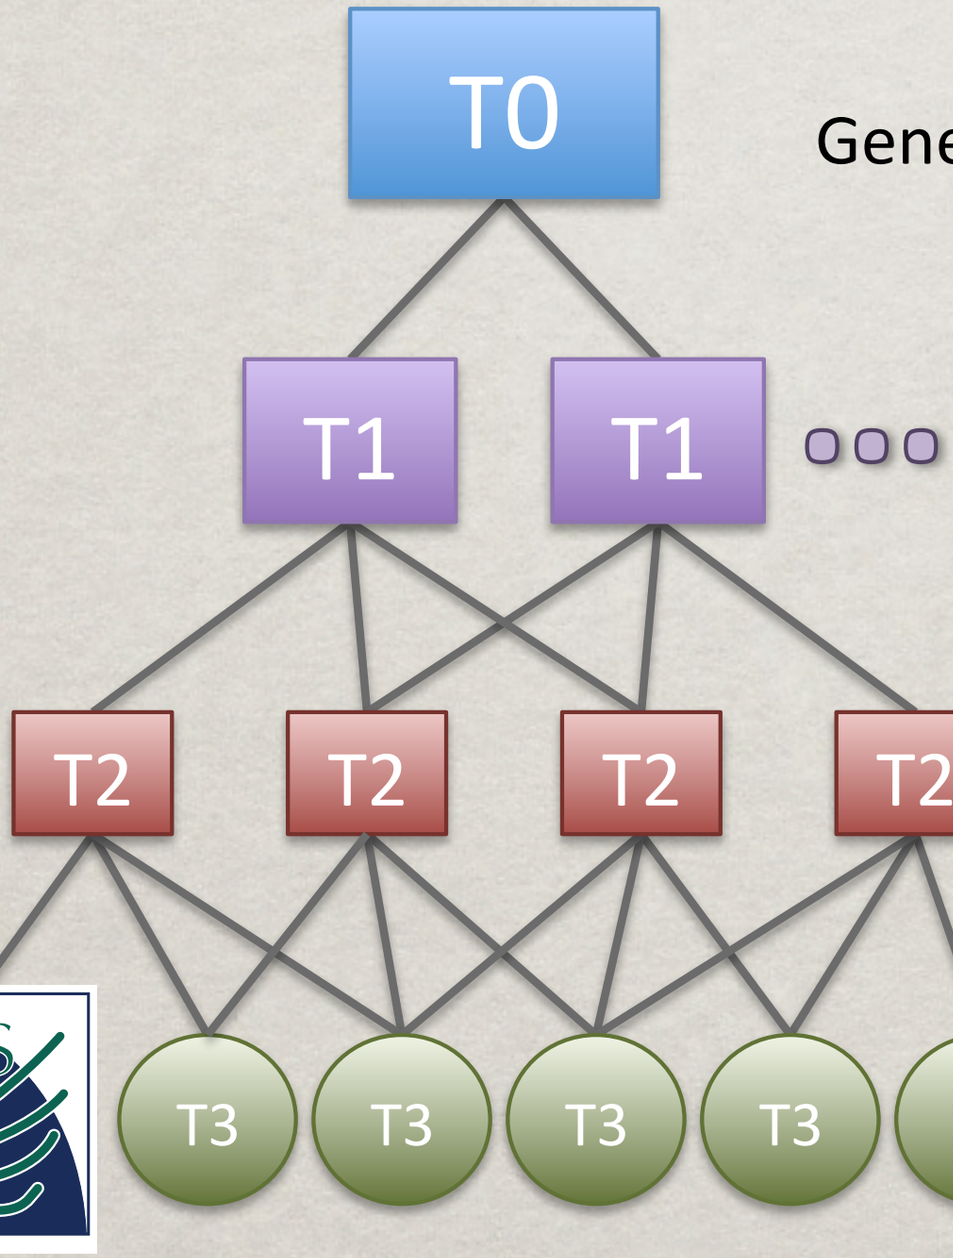

CERN Lab
Geneva, Switzerland

National Labs
(Fermilab, etc.)

Universities
(MIT, Wisconsin,
Nebraska, Purdue, etc.)

Universities
(ND, Colorado,
UMD, OSU, etc.)

ND T3


© Nil Valls - www.nilvalls.com


© Nil Valls - www.nilvalls.com


© Nil Valls - www.nilvalls.com


© Nil Valls - www.nilvalls.com

Basic Idea In pictures

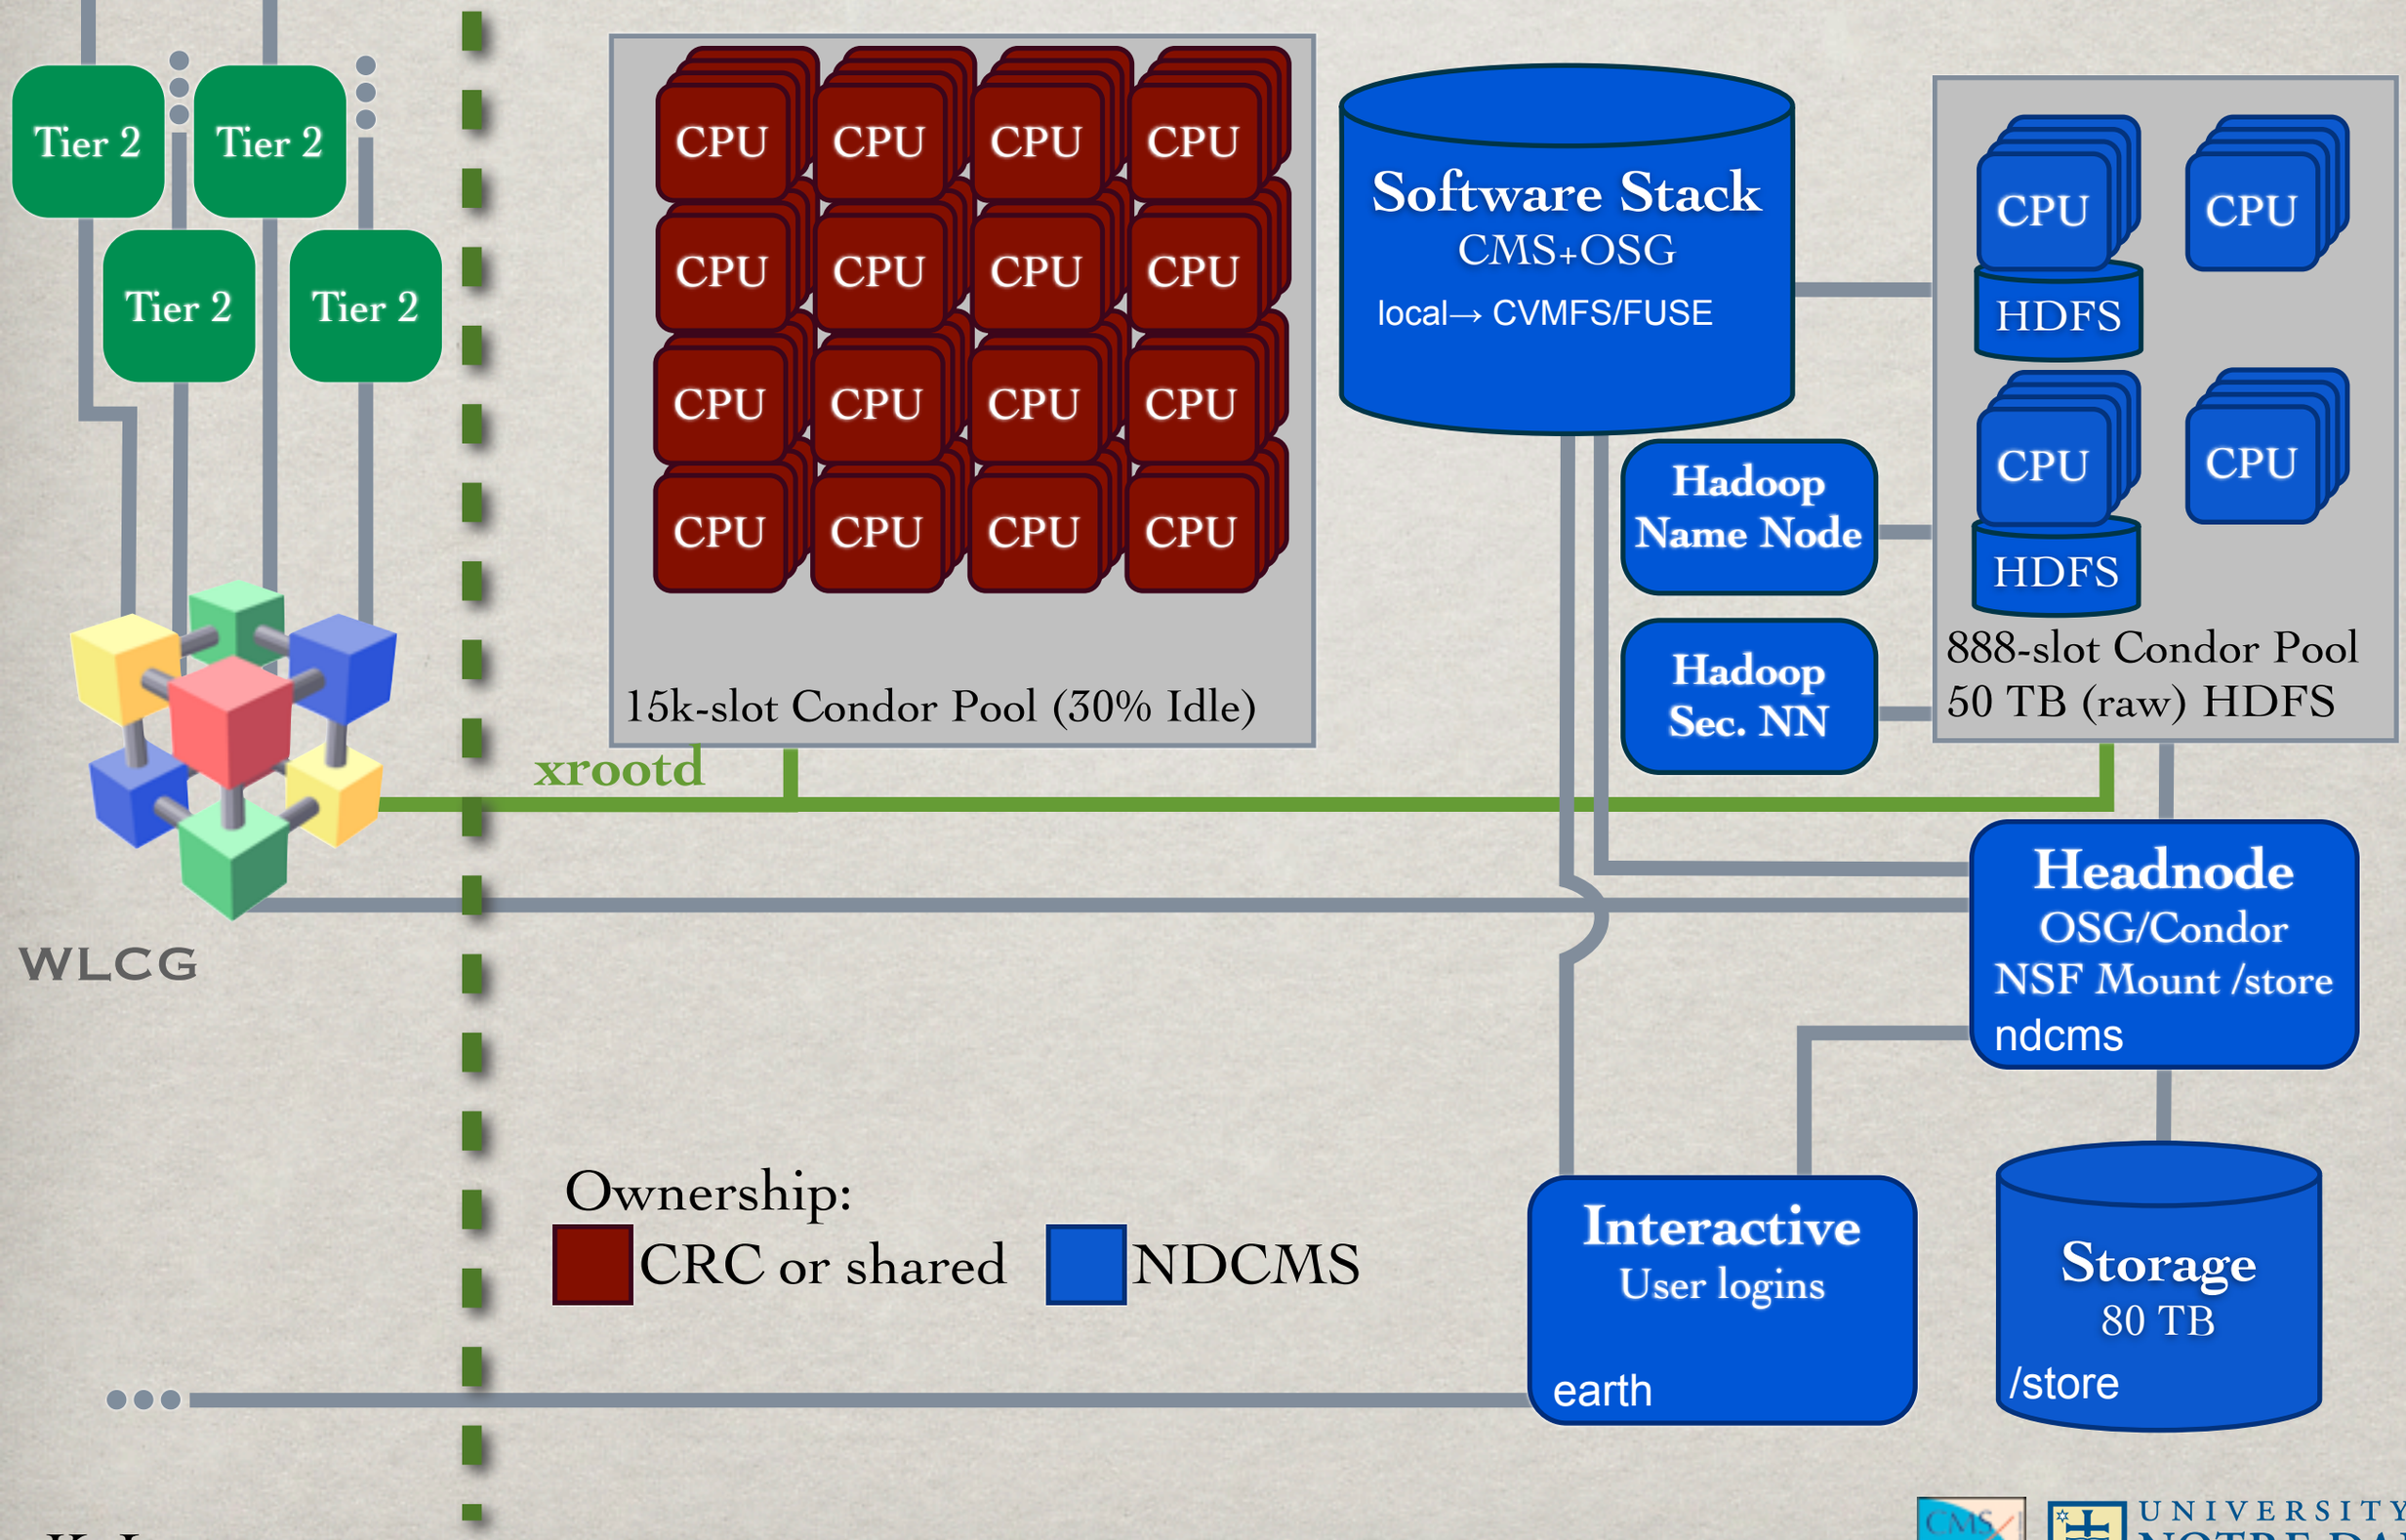

What do we use it for?

- Short answer: physics analysis
- More details
  - Start with large general purpose dataset (simulations + real data): ~400TB
  - Reduce by keeping only essential information for an analysis: ~20 TB
  - Select most interesting subset of reduced data: ~100 GB
  - Make plots
- First step done by reading data stored at T2 over network (xrootd)
  - Goes against trend of "moving code to data" (I think)
  - Solves two problems
    - Lack of storage (and need to manage large amounts of data needed only briefly)
    - Transfer time (overlaps with processing time)
- Rest done with local data access

Talk by Brian Bockelman (last year)

Usage Statistics

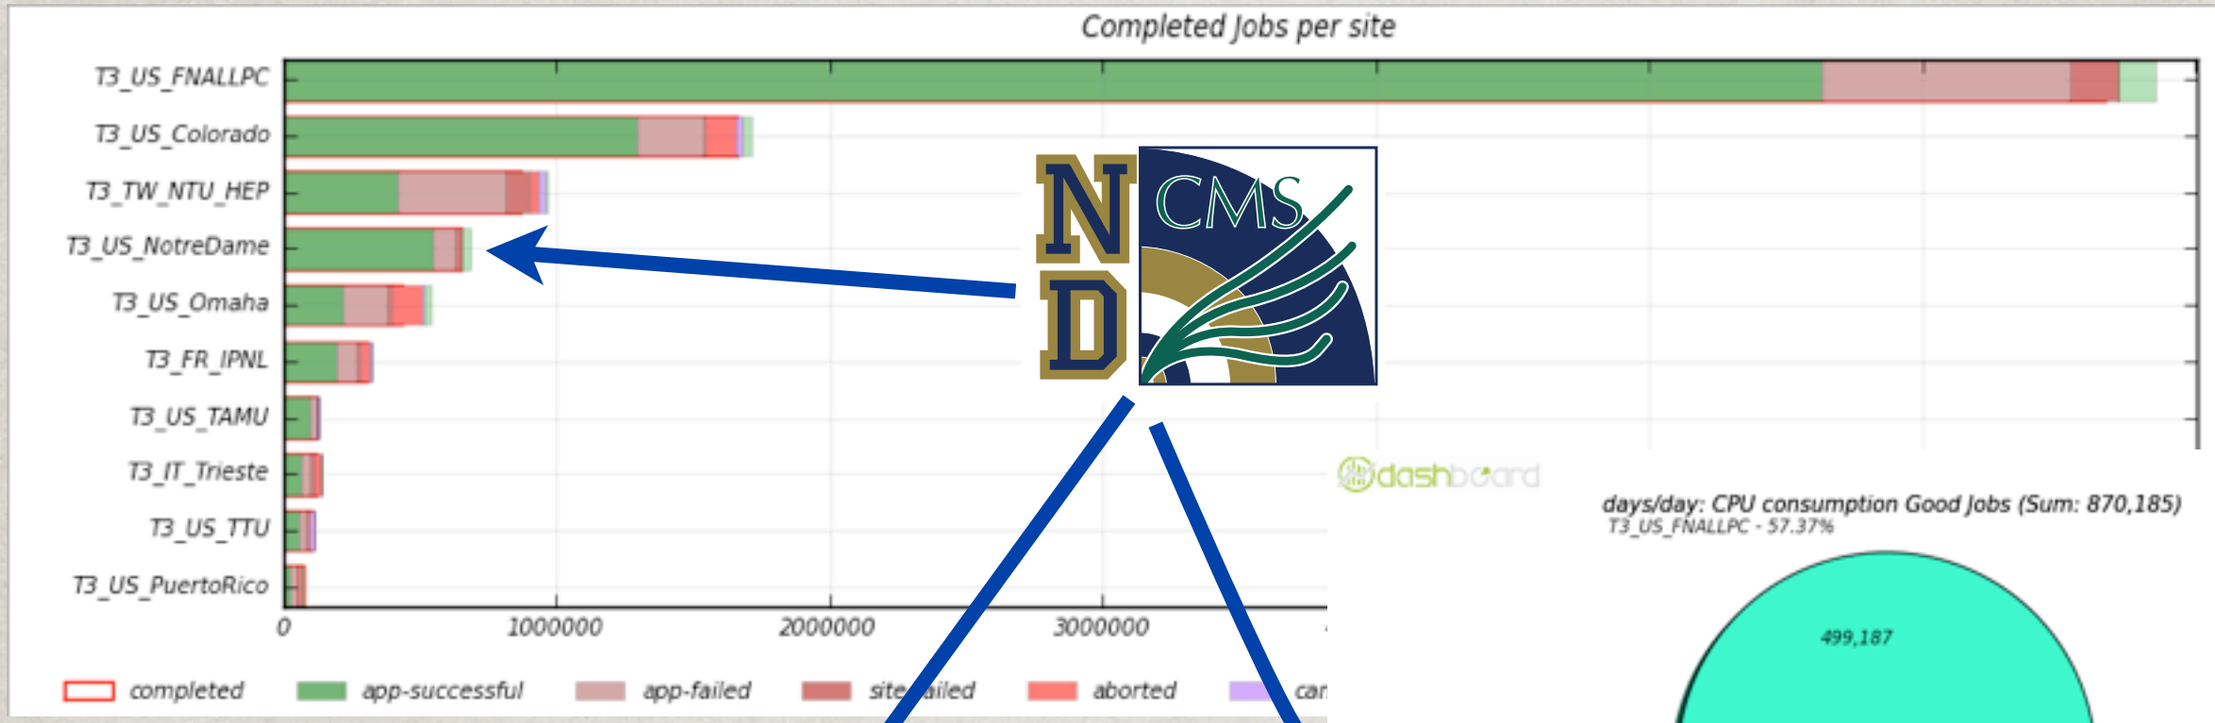

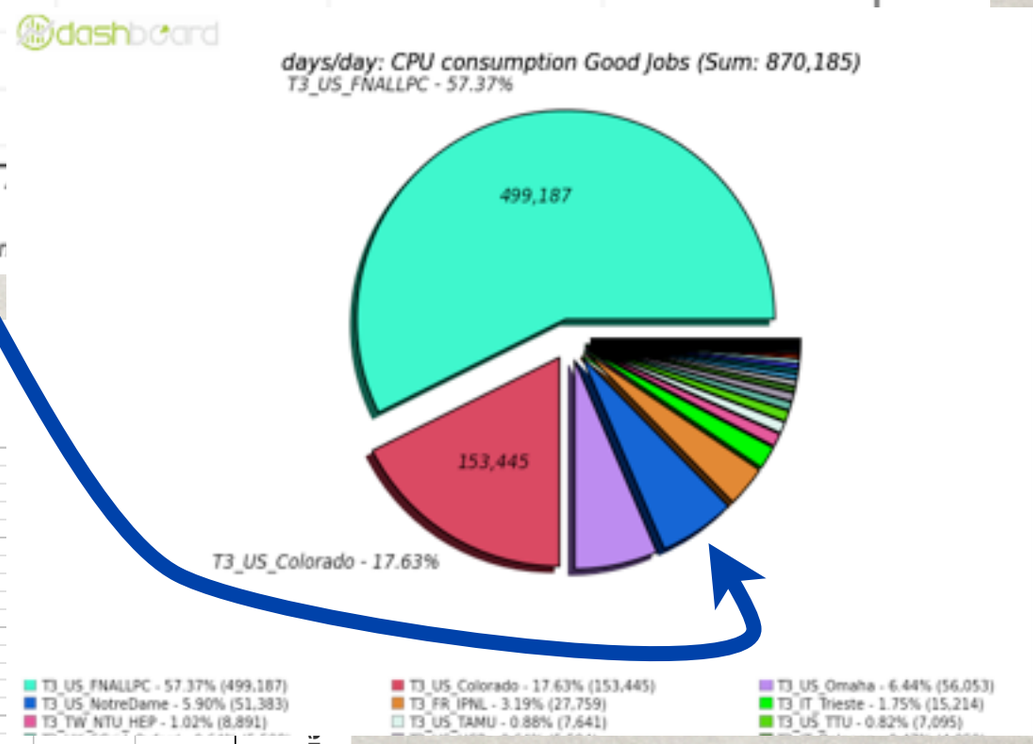

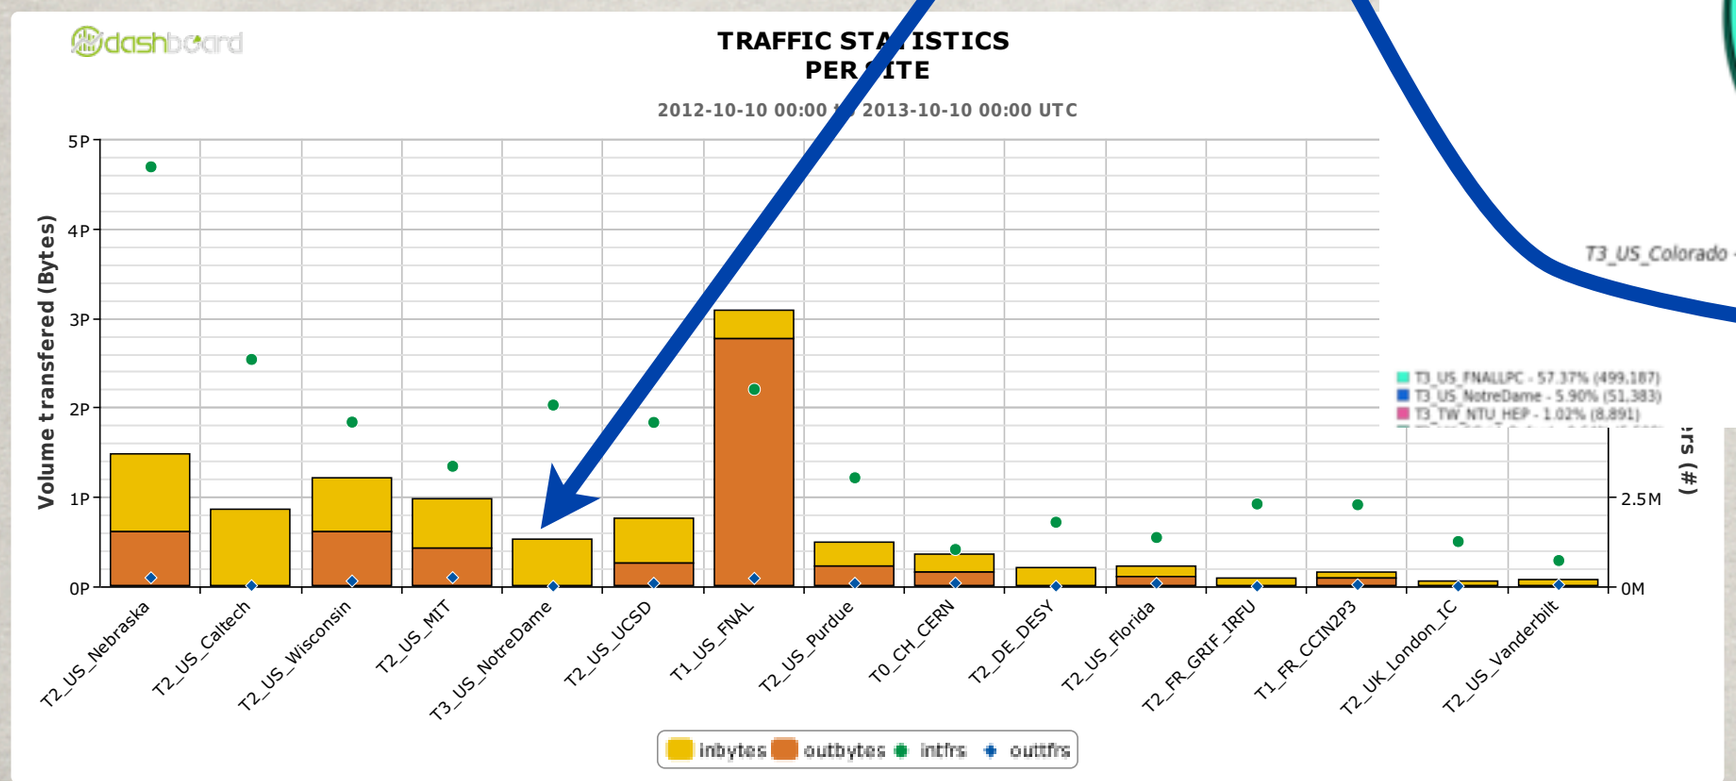

UNIVERSITY OF NOTRE DAME

Bandwidth Usage

Haven't saturated campus bandwidth yet!

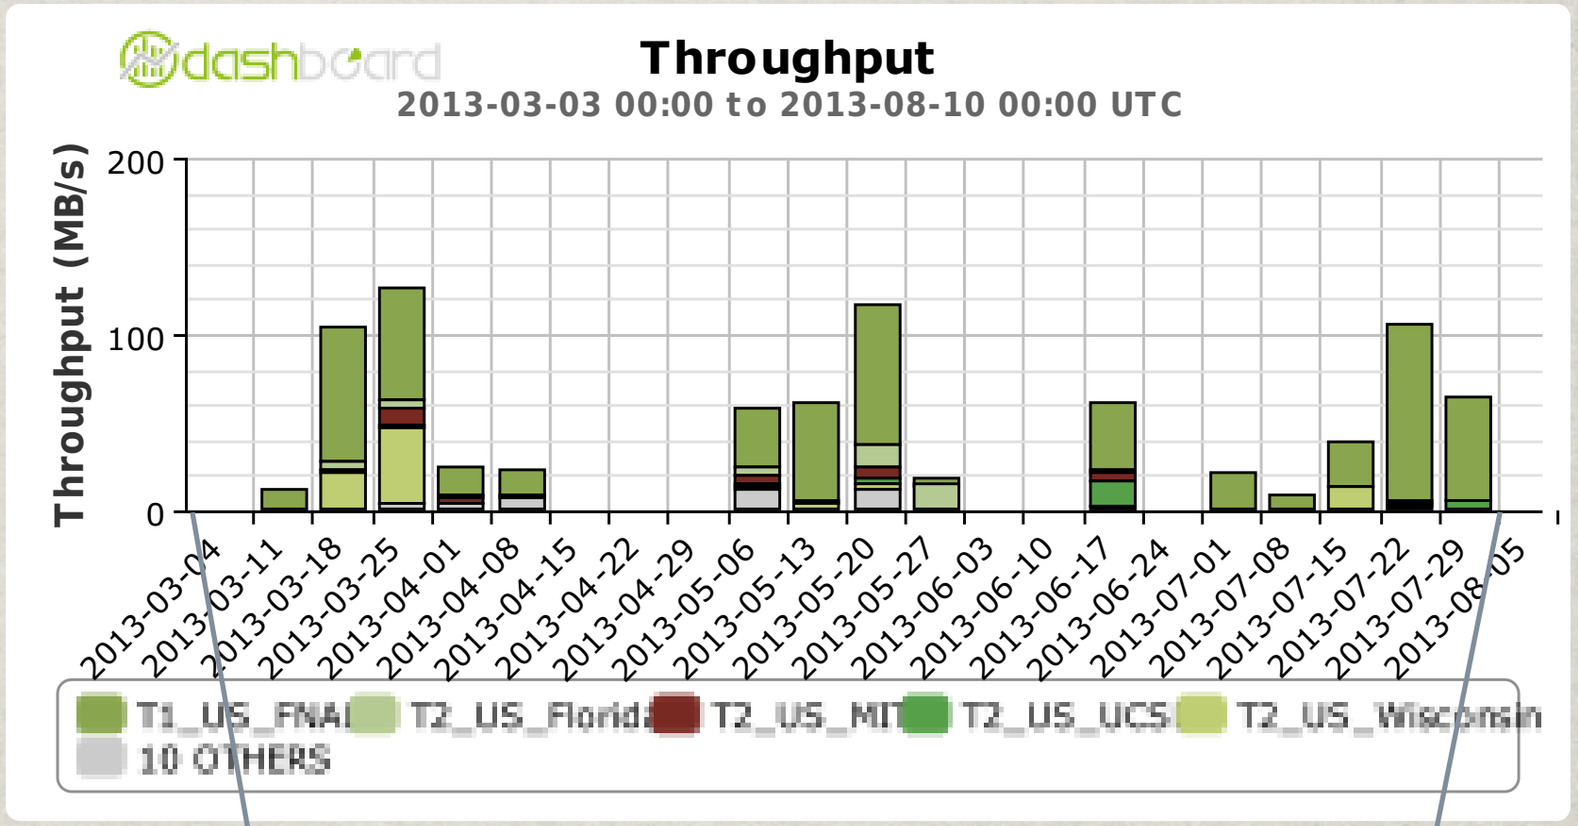

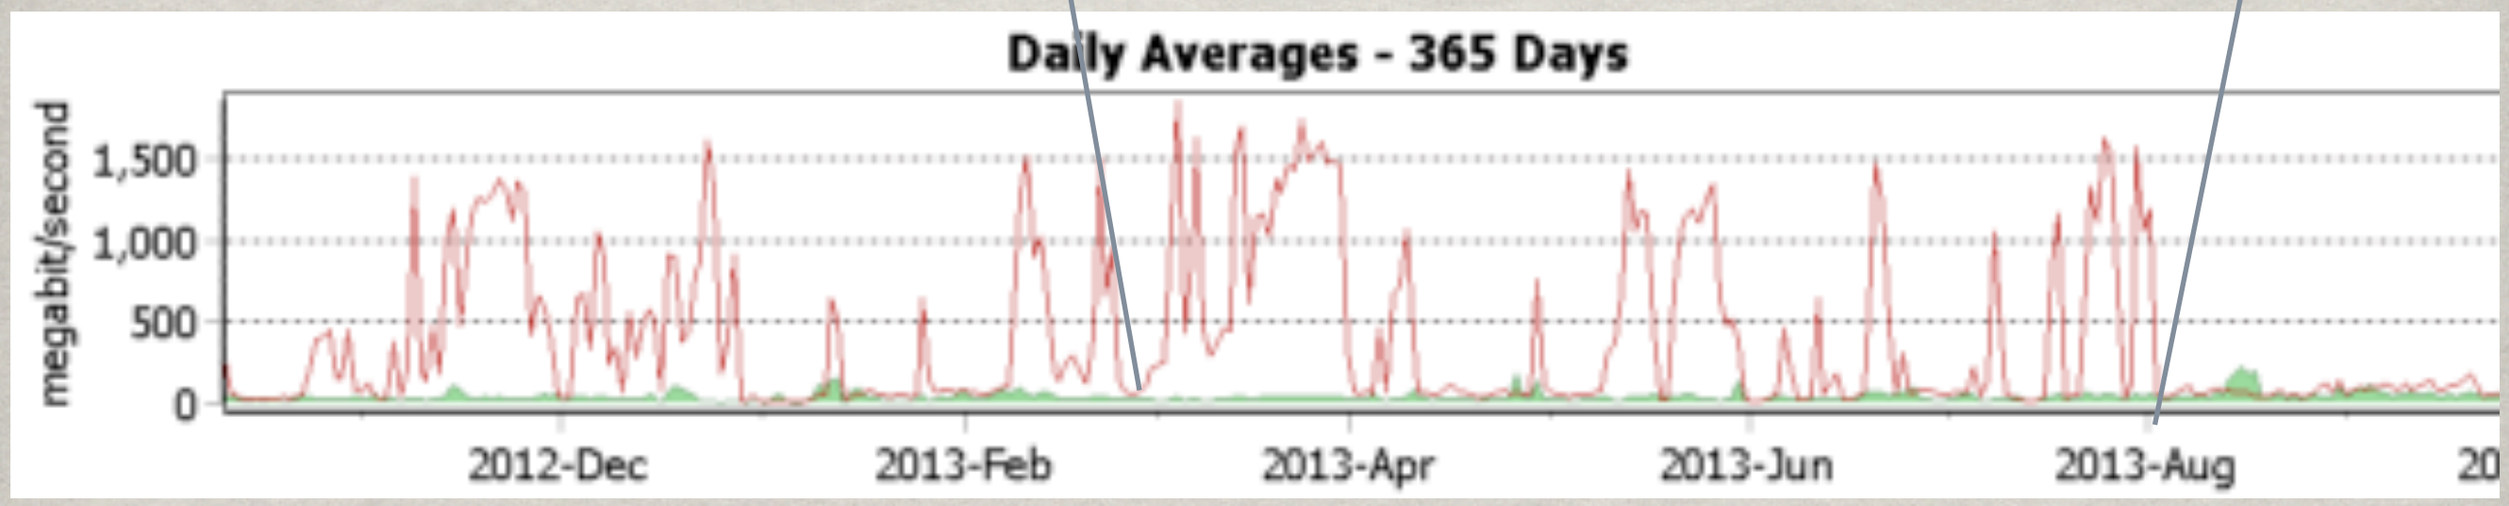

Scaling from Here

- Main bottleneck (right now): CPU
  - Most intensive part of the processing takes ≥ 3 weeks using full T3 resources
  - Would like it to be faster
- Plan: top into opportunistic resources on campus
  - CRC pool
    - Potentially 15k cores available
    - Only ~30% idle at any moment: still 5k cores!
    - Comparable resources to T2

Challenges

- Problem: CRC condor machines don't have CMS/ OSG software
  - Solution: Use CVMFS + Parrot
  - Status: Successful small scale tests: working to scale to 1000's of cores

Good talk last year by Dan Bradley

CRC REU: Dillon Skeehan (+Paul Brenner)

- Problem: CRC machines need access to data at T2 sites
  - Solution: xrootd (just like T3 workers)
  - Status: After reconfiguration of CRC network to allow outside access on nodes, successful small scale tests

Challenges

- Problem: Data stored at ND T3 not accessible to CRC nodes
  - Solution: Still TBD
    - Run xrootd server for ND storage?
    - Chirp?
- Problem: Resource management
  - Ideally jobs would overflow automatically from dedicated T3 resources to opportunistic when necessary
  - Right now, need to manually decide where to run
  - Solution: TBD
    - Natural sort of approach: Condor glidein
    - Could WorkQueue be an interesting alternative?

Challenges

- Problem: Preemption
  - Serious problem:
    - Jobs don't own slot. Will be evicted if owner wants machine
    - Current CMS workflow control doesn't handle preemption well: no checkpointing or automatic restarting
  - Solution: ???
    - In principle: address no restart issue with tuning to condor submit file generated by CMS workflow tool
    - Doesn't handle issue of wasted resources when evicted jobs progress lost
    - Need solution to handle: checkpointing? Run in VM (stop/restart on eviction)?
    - Can imagine building solution using Chirp and/or Parrot to store checkpoint information across network to allow for graceful restarts

Conclusions

- ND T3 has transformed our group's ability to do data anlaysis
- Entering new territory in terms of scaling T3 resources (not counting FNAL)
- CCL tools critical to successes so far
  - Everything rides on CVMFS + Parrot
  - See opportunities for other tools to play important role(s)
- Several opportunities to innovate (network data access, workflow management, preemption, etc.)
  - Innovations can be fed back to larger T1/2/3 community