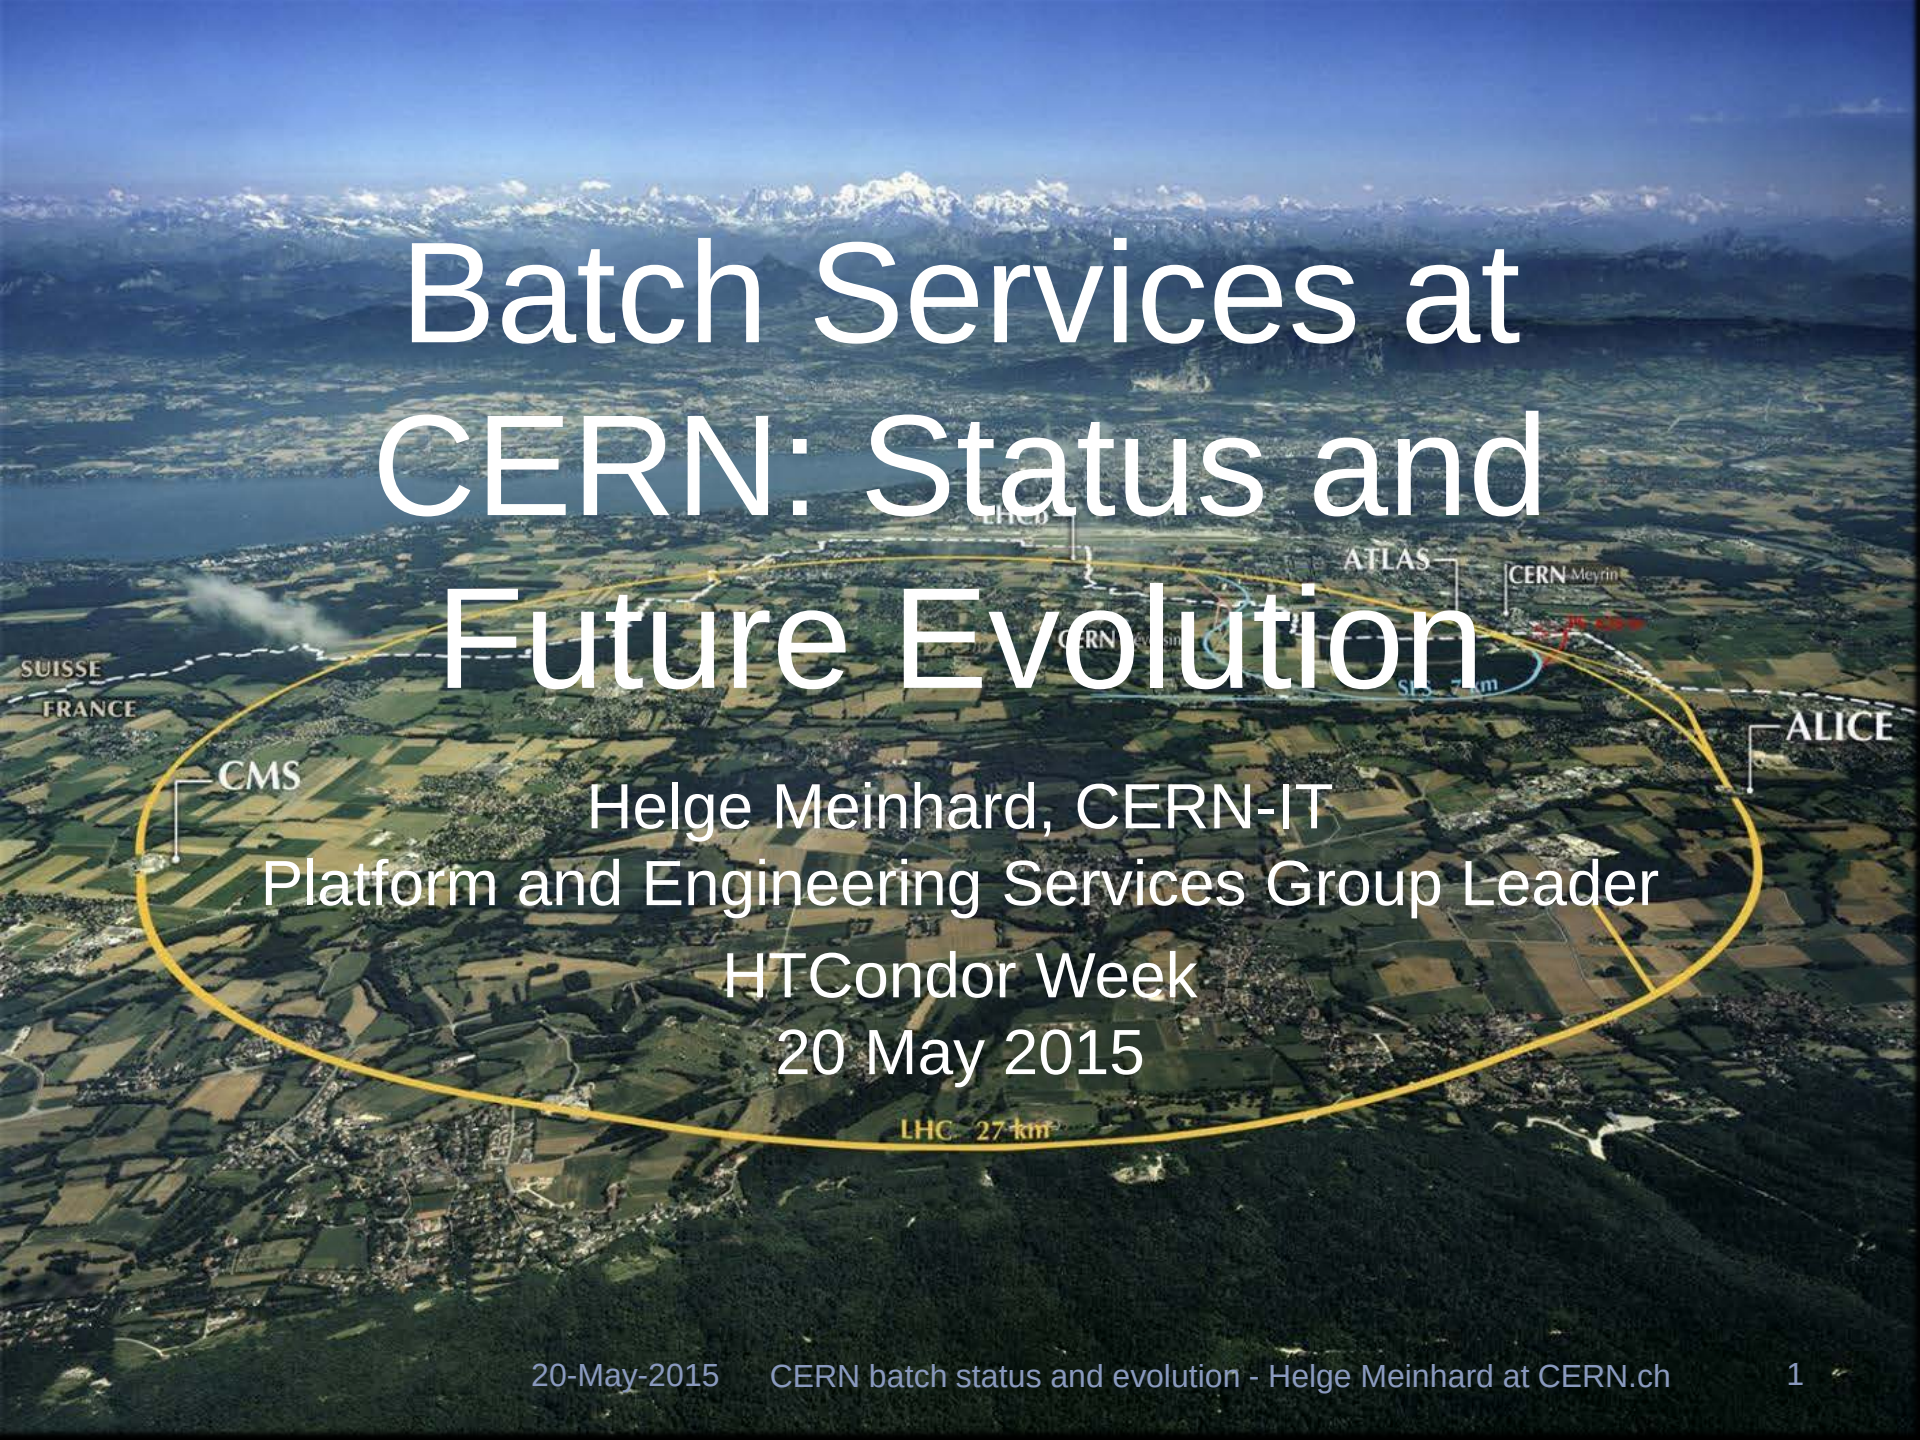

Batch Services at CERN: Status and Future Evolution

Helge Meinhard, CERN-IT
Platform and Engineering Services Group Leader

HTCondor Week
20 May 2015

CERN

CERN

CERN

- International organisation close to Geneva, straddling Swiss-French border, founded 1954

1954: 12 Member States

CERN

- International organisation close to Geneva, straddling Swiss-French border, founded 1954
- Facilities for fundamental research in particle physics

1954: 12 Member States

CERN

- International organisation close to Geneva, straddling Swiss-French border, founded 1954
- Facilities for fundamental research in particle physics
- 21 member states, 1 B CHF budget

1954: 12 Member States

Members: Austria, Belgium, Bulgaria, Czech republic, Denmark, Finland, France, Germany, Greece, Hungary, Israel, Italy, Netherlands, Norway, Poland, Portugal, Slovak Republic, Spain, Sweden, Switzerland, United Kingdom
Candidate for membership: Romania
Associate member: Serbia
Observers: European Commission, India, Japan, Russia, Turkey, UNESCO, United States of America
Numerous **non-member states with collaboration agreements**

CERN

- International organisation close to Geneva, straddling Swiss-French border, founded 1954
- Facilities for fundamental research in particle physics
- 21 member states, 1 B CHF budget
- 3’581 staff, fellows, students, apprentices, …

1954: 12 Member States

Members: Austria, Belgium, Bulgaria, Czech republic, Denmark, Finland, France, Germany, Greece, Hungary, Israel, Italy, Netherlands, Norway, Poland, Portugal, Slovak Republic, Spain, Sweden, Switzerland, United Kingdom
Candidate for membership: Romania
Associate member: Serbia
Observers: European Commission, India, Japan, Russia, Turkey, UNESCO, United States of America
Numerous **non-member states with collaboration agreements**

2’513 staff members, 566 fellows, 481 students, 21 apprentices

CERN

- International organisation close to Geneva, straddling Swiss-French border, founded 1954
- Facilities for fundamental research in particle physics
- 21 member states, 1 B CHF budget
- 3’581 staff, fellows, students, apprentices, …
- 11’000 users

1954: 12 Member States

Members: Austria, Belgium, Bulgaria, Czech republic, Denmark, Finland, France, Germany, Greece, Hungary, Israel, Italy, Netherlands, Norway, Poland, Portugal, Slovak Republic, Spain, Sweden, Switzerland, United Kingdom
Candidate for membership: Romania
Associate member: Serbia
Observers: European Commission, India, Japan, Russia, Turkey, UNESCO, United States of America
Numerous **non-member states with collaboration agreements**

2’513 staff members, 566 fellows, 481 students, 21 apprentices

6’700 member states, 1’800 USA, 900 Russia, 230 Japan, …

CERN

"Science for peace"

- International organisation close to Geneva, straddling Swiss-French border, founded 1954
- Facilities for fundamental research in particle physics
- 21 member states, 1 B CHF budget
- 3'581 staff, fellows, students, apprentices, …
- 11'000 users

1954: 12 Member States

Members: Austria, Belgium, Bulgaria, Czech republic, Denmark, Finland, France, Germany, Greece, Hungary, Israel, Italy, Netherlands, Norway, Poland, Portugal, Slovak Republic, Spain, Sweden, Switzerland, United Kingdom
Candidate for membership: Romania
Associate member: Serbia
Observers: European Commission, India, Japan, Russia, Turkey, UNESCO, United States of America
Numerous **non-member states with collaboration agreements**

2'513 staff members, 566 fellows, 481 students, 21 apprentices

6'700 member states, 1'800 USA, 900 Russia, 230 Japan, …

Tools: LHC and Detectors

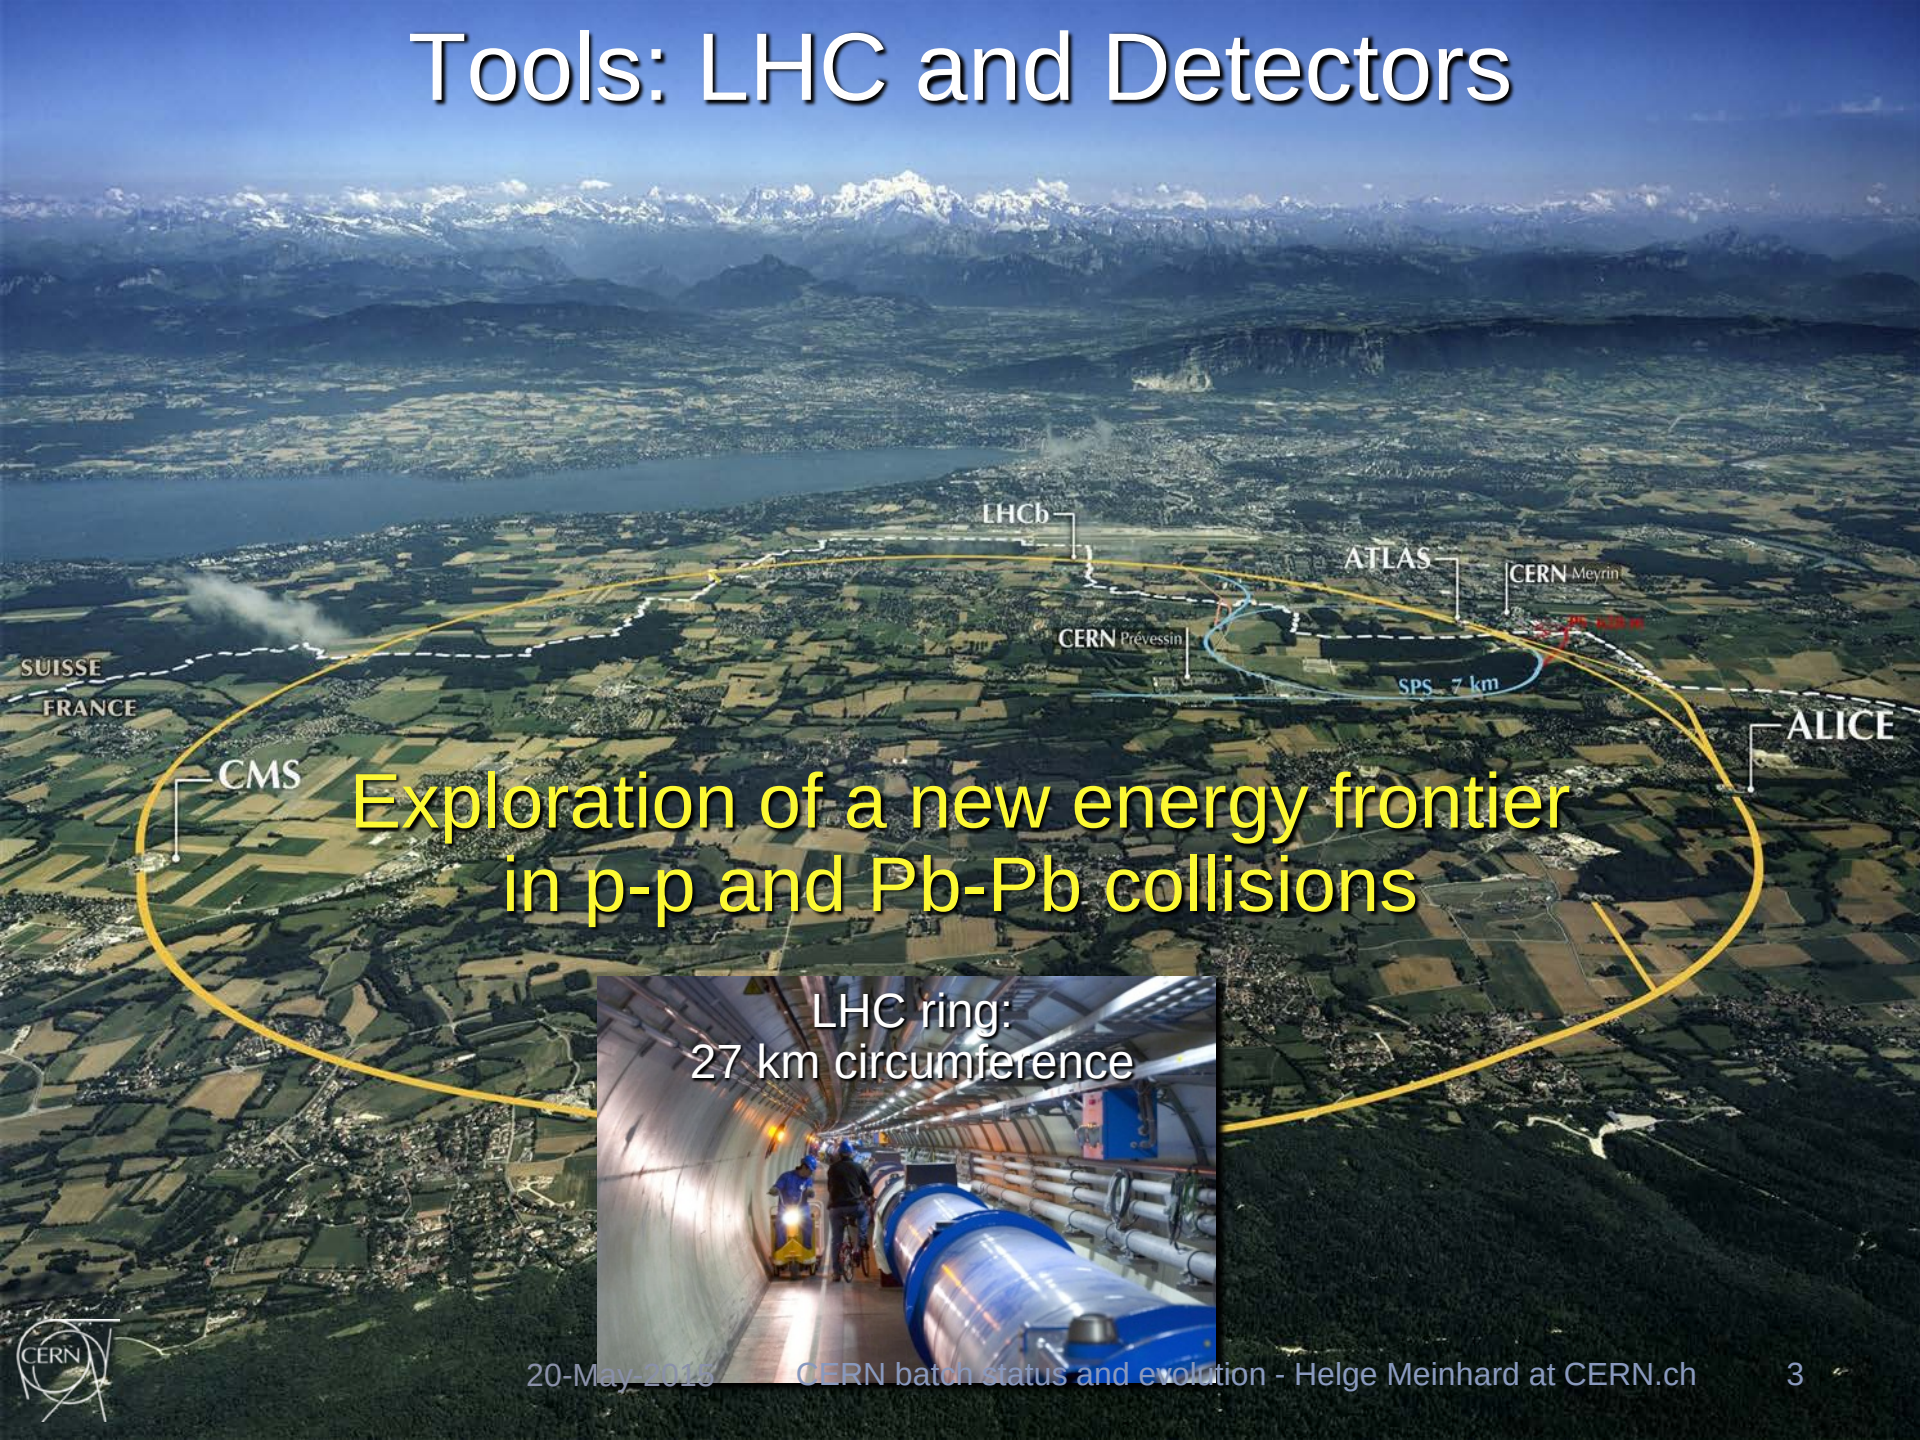

Exploration of a new energy frontier
in p-p and Pb-Pb collisions

Tools: LHC and Detectors

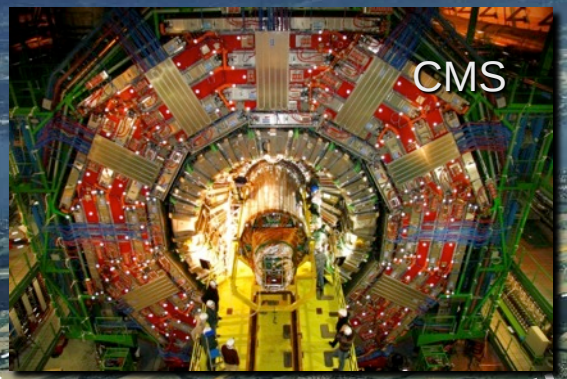

LHCb

ATLAS

CERN Meyrin

SPS 7 km

ALICE

CMS

General Purpose,
proton-proton, heavy ions
Discovery of new physics:
Higgs, SuperSymmetry

Exploration of a new energy frontier
in p-p and Pb-Pb collisions

Tools: LHC and Detectors

pp, B-Physics, CP Violation
(matter-antimatter symmetry)

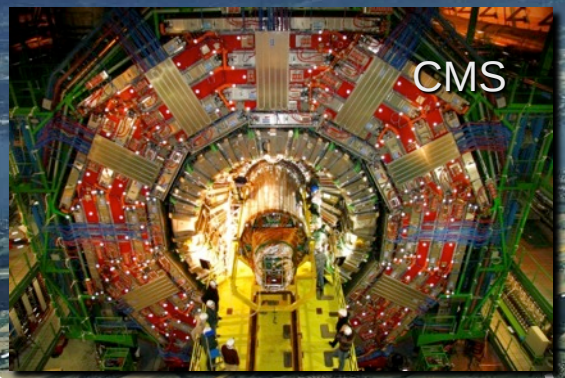

LHCb
ATLAS
CERN Meyrin
SPS 7 km
ALICE
CMS

General Purpose,
proton-proton, heavy ions
Discovery of new physics:
Higgs, SuperSymmetry

Exploration of a new energy frontier
in p-p and Pb-Pb collisions

Tools: LHC and Detectors

pp, B-Physics, CP Violation
(matter-antimatter symmetry)

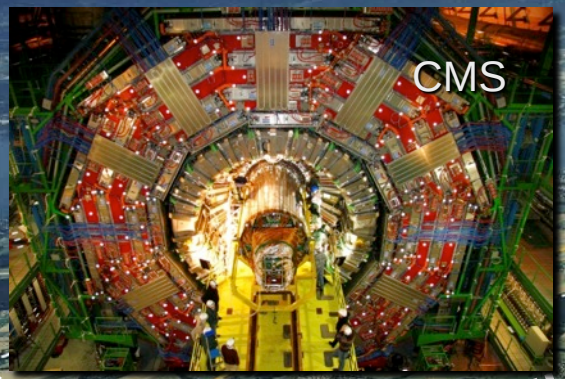

General Purpose,
proton-proton, heavy ions
Discovery of new physics:
Higgs, SuperSymmetry

Exploration of a new energy frontier
in p-p and Pb-Pb collisions

Heavy ions, pp
(state of matter of early universe)

Results so far

- Many… the most spectacular one being

Results so far

- Many… the most spectacular one being
- 04 July 2012: Discovery of a “Higgs-like particle”

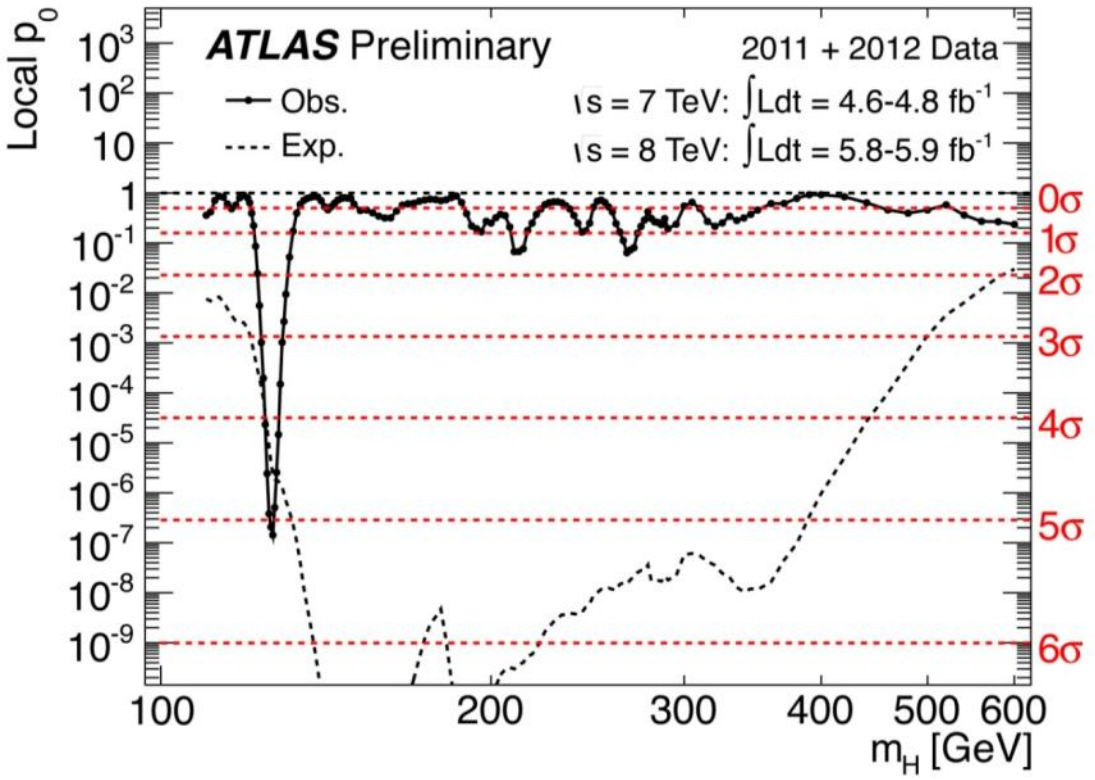

Results so far

- Many… the most spectacular one being
- 04 July 2012: Discovery of a “Higgs-like particle”

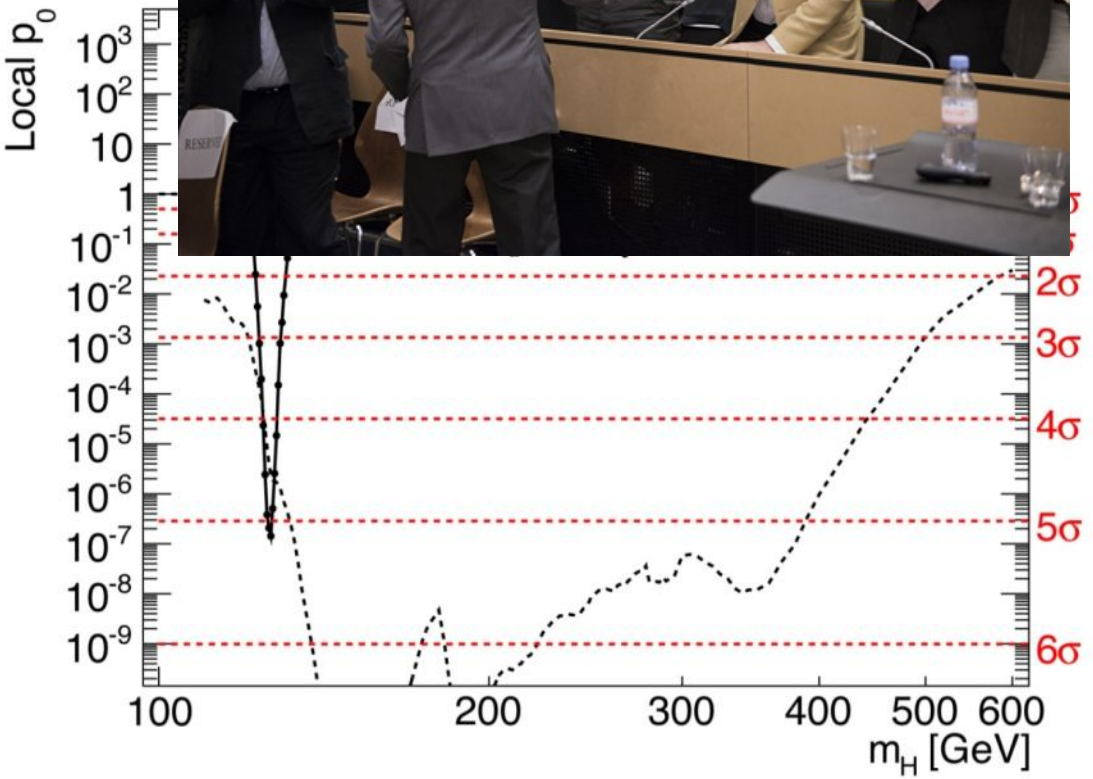

Results so far

- Many… the most spectacular one being
- 04 July 2012: Discovery of a “Higgs-like particle”

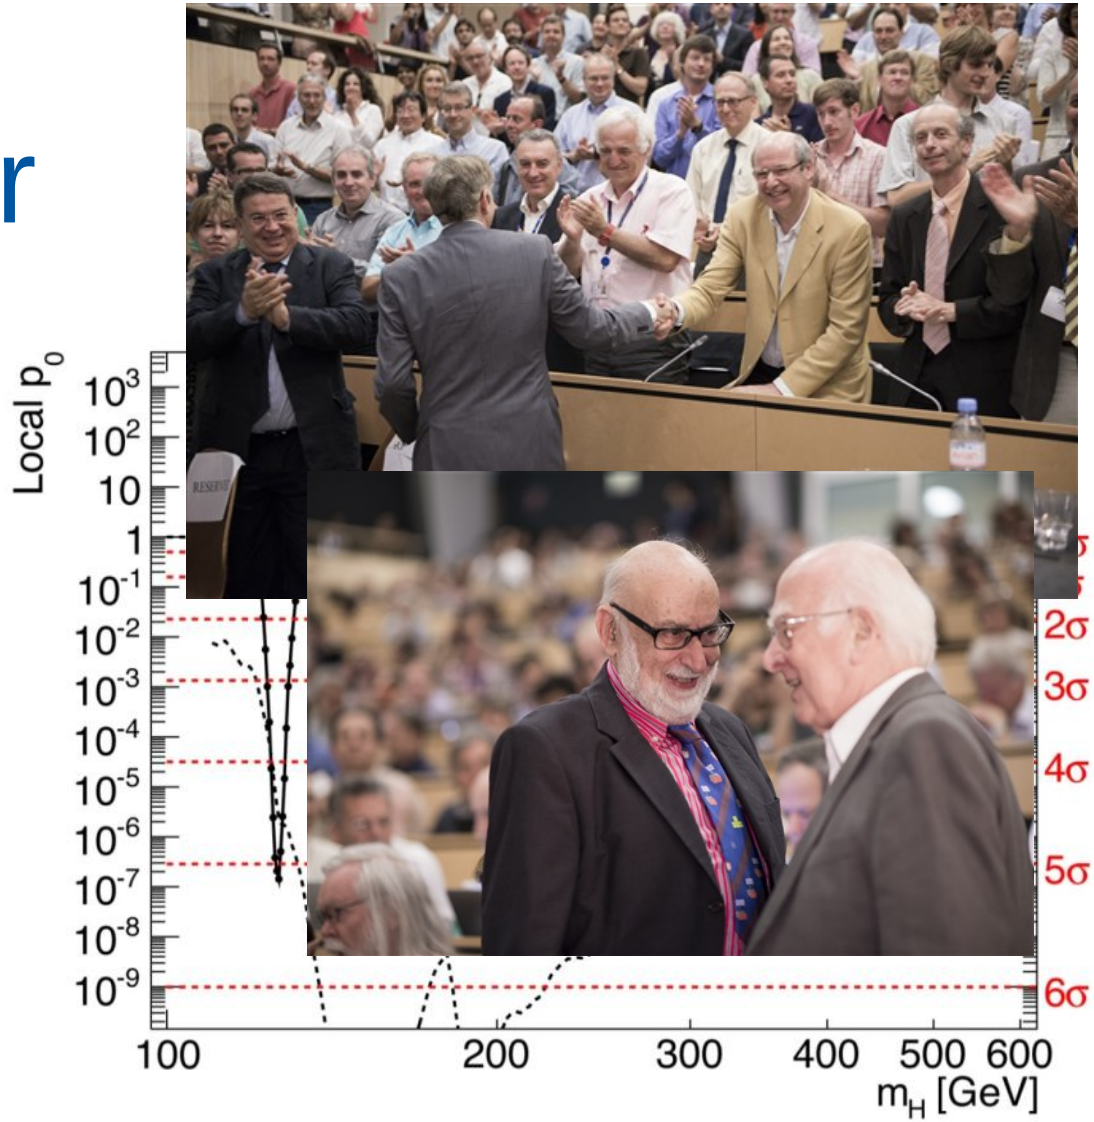

Results so far

- Many… the most spectacular one being
- 04 July 2012: Discovery of a “Higgs-like particle”
- March 2013: The particle is indeed a Higgs boson

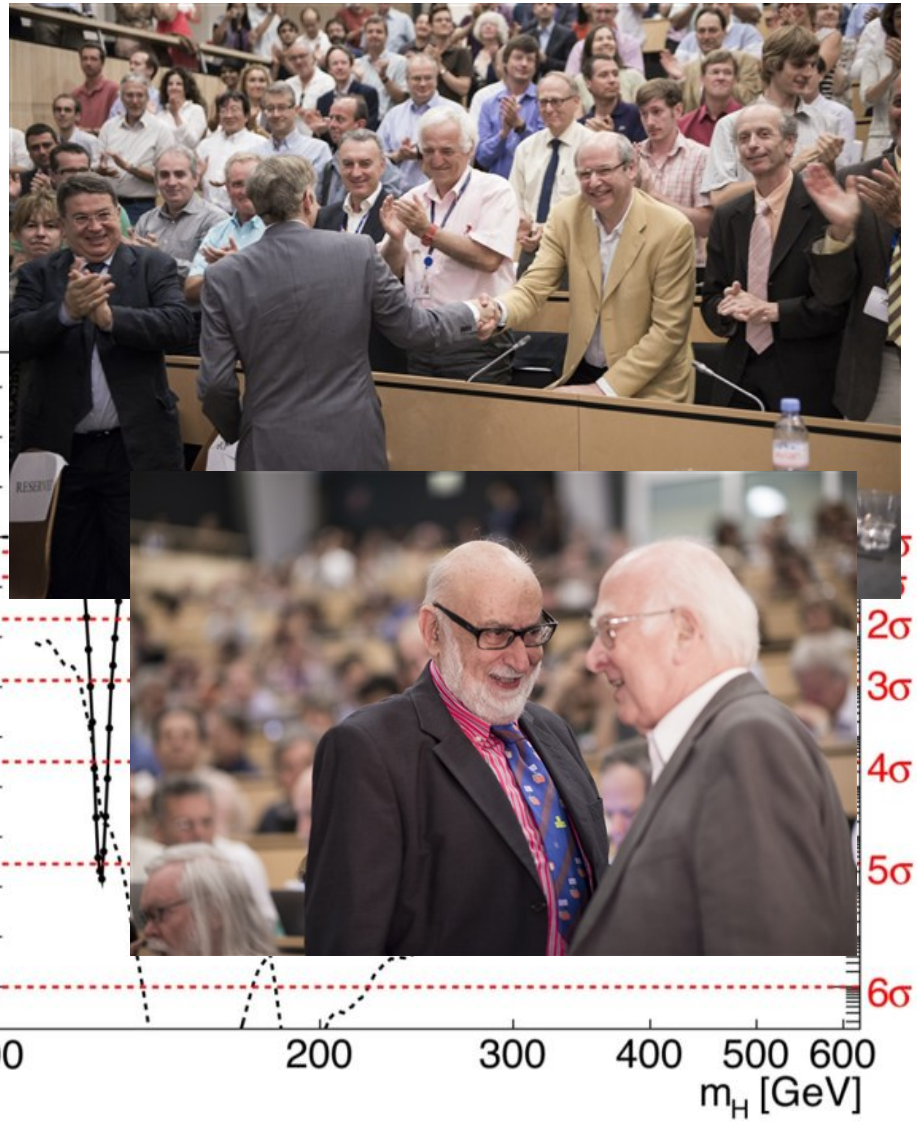

Results so far

- Many… the most spectacular one being
- 04 July 2012: Discovery of a “Higgs-like particle”
- March 2013: The particle is indeed a Higgs boson
- 08 Oct 2013 / 10 Dec 2013: Nobel price to Peter Higgs and François Englert
  - CERN, ATLAS and CMS explicitly mentioned

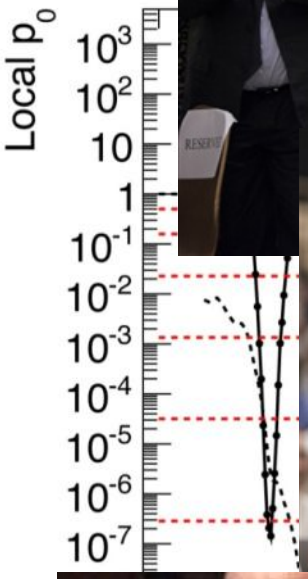

What is the data?

- 150 million sensors deliver data …40 million times per second

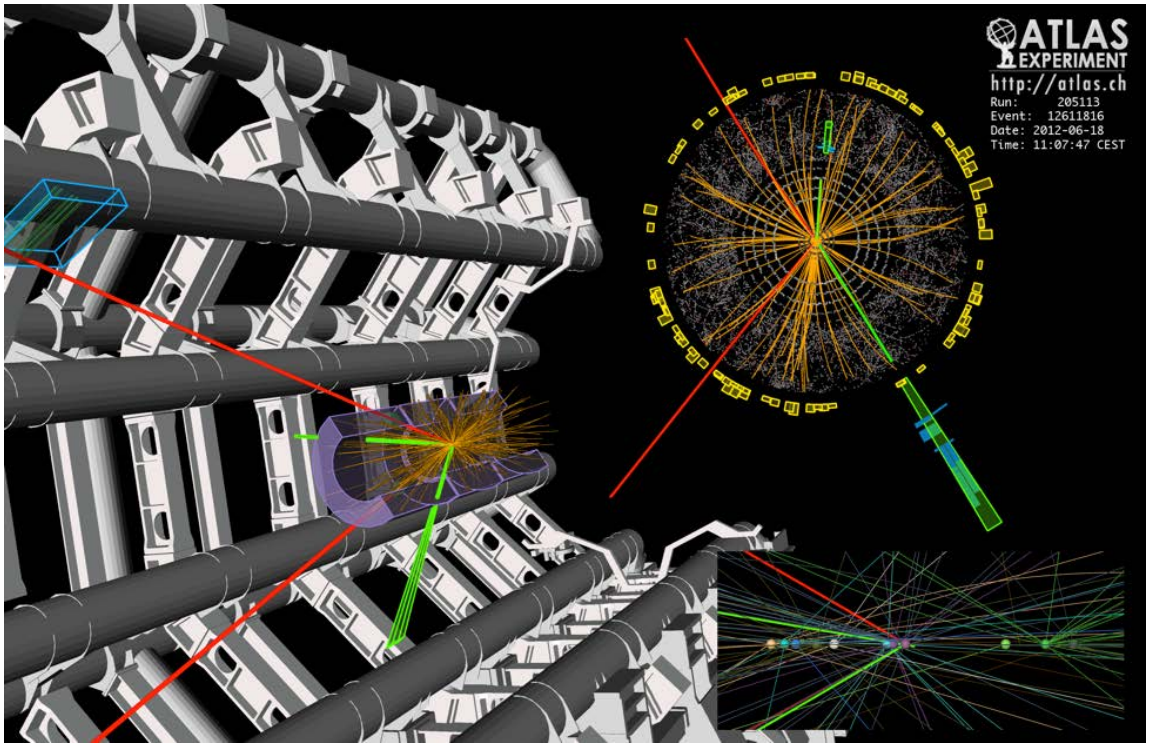

- Up to 6 GB/s to be permanently stored after filtering
- Almost 30 PB/y in Run 1
- Expect ~50 PB/y in Run 2

The Worldwide LHC Computing Grid

An International collaboration to distribute and analyse LHC data

The Worldwide LHC Computing Grid

An International collaboration to distribute and analyse LHC data

Integrates computer centres worldwide that provide computing and storage resource into a single infrastructure accessible by all LHC physicists

The Worldwide LHC Computing Grid

An International collaboration to distribute and analyse LHC data

Integrates computer centres worldwide that provide computing and storage resource into a single infrastructure accessible by all LHC physicists

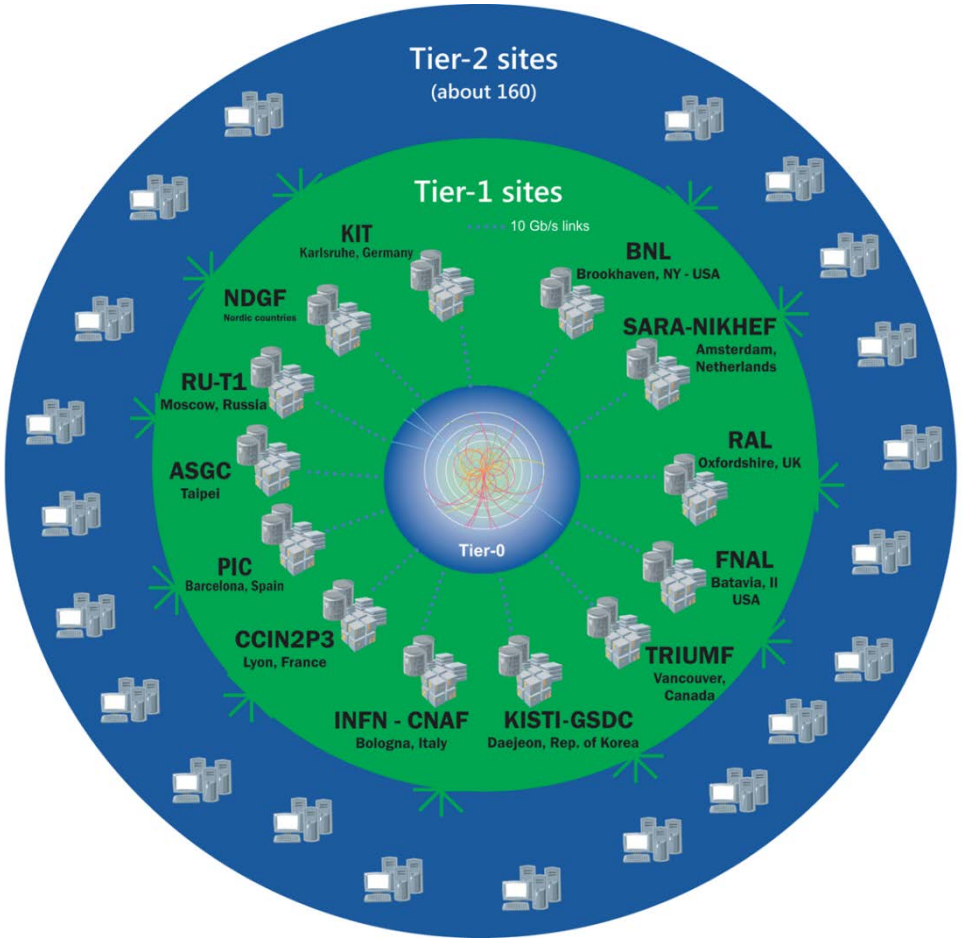

The Worldwide LHC Computing Grid

An International collaboration to distribute and analyse LHC data

Integrates computer centres worldwide that provide computing and storage resource into a single infrastructure accessible by all LHC physicists

Tier-0 (CERN): data recording, reconstruction and distribution

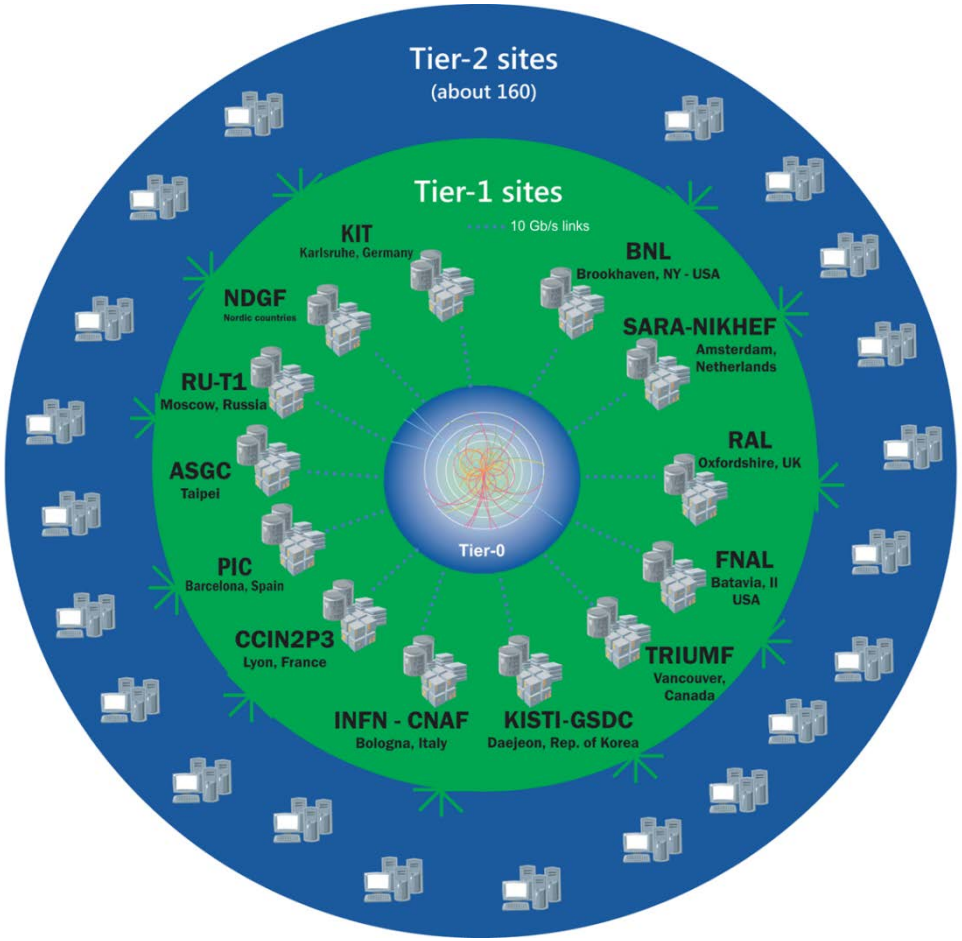

The Worldwide LHC Computing Grid

An International collaboration to distribute and analyse LHC data

Integrates computer centres worldwide that provide computing and storage resource into a single infrastructure accessible by all LHC physicists

Tier-0 (CERN): data recording, reconstruction and distribution

Tier-1: permanent storage, re-processing, analysis

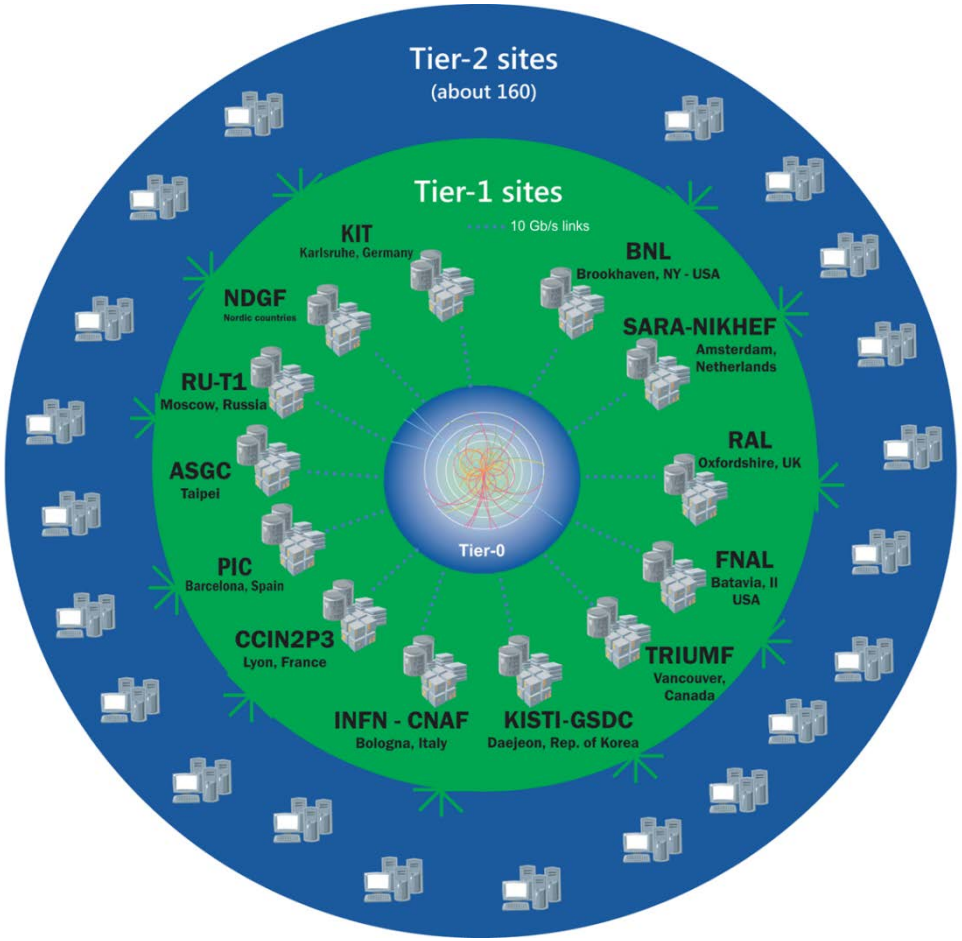

The Worldwide LHC Computing Grid

An International collaboration to distribute and analyse LHC data

Integrates computer centres worldwide that provide computing and storage resource into a single infrastructure accessible by all LHC physicists

Tier-0 (CERN): data recording, reconstruction and distribution

Tier-1: permanent storage, re-processing, analysis

Tier-2: Simulation, end-user analysis

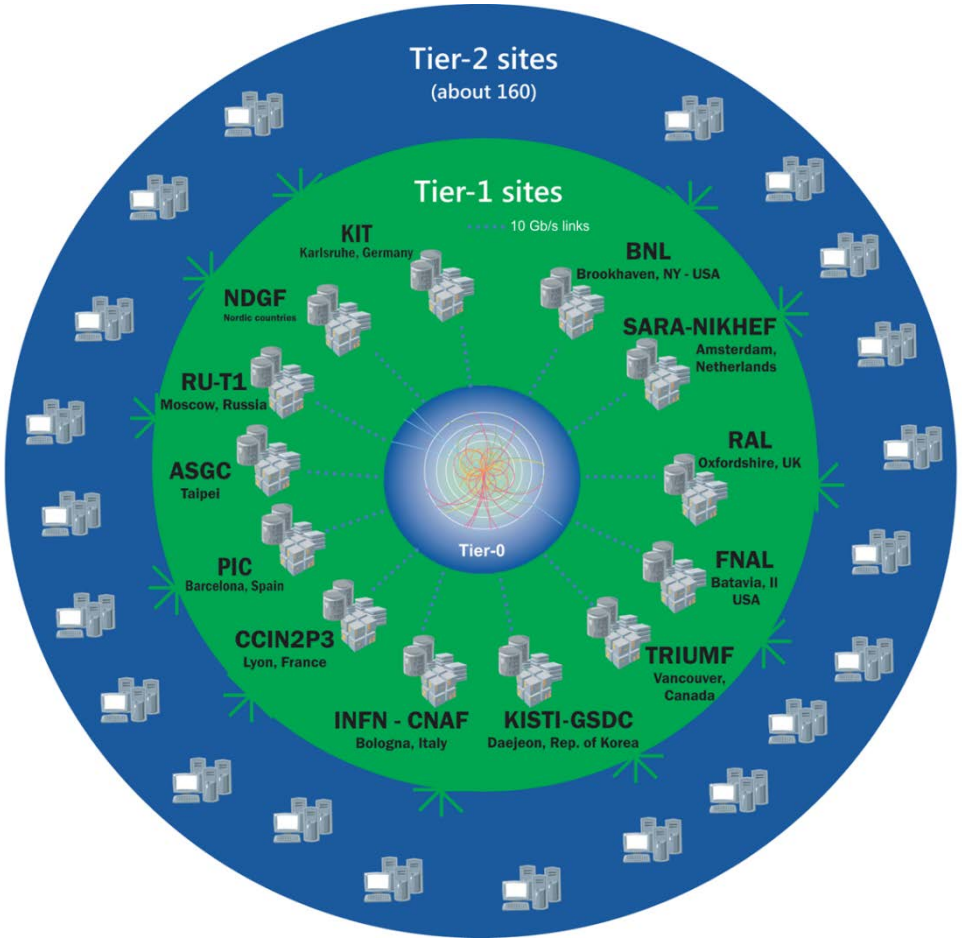

The Worldwide LHC Computing Grid

An International collaboration to distribute and analyse LHC data

Integrates computer centres worldwide that provide computing and storage resource into a single infrastructure accessible by all LHC physicists

Tier-0 (CERN): data recording, reconstruction and distribution

Tier-1: permanent storage, re-processing, analysis

Tier-2: Simulation, end-user analysis

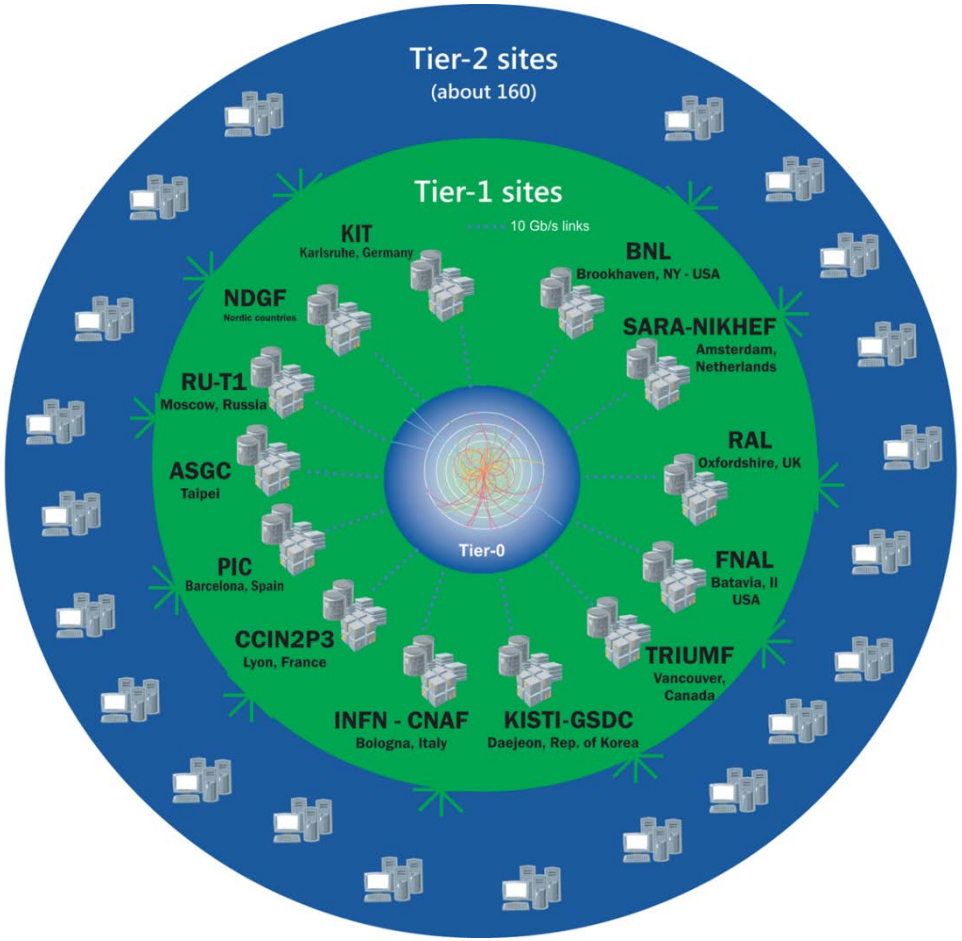

nearly 170 sites, 40 countries

~350’000 cores

500 PB of storage

> 2 million jobs/day

10-100 Gb links

WLCG Resources

[kHS06]	2014		2015		2016	
	Tier-0	All	Tier-0	All	Tier-0	All
ALICE	90	366	175	495	215	609
ATLAS	111	856	205	1’175	257	1’343
CMS	121	738	271	1’071	317	1’417
LHCb	34	218	36	240	51	315
Total	**356**	**2’178**	**687**	**2’981**	**840**	**3’684**

WLCG Resources

[kHS06]	2014		2015		2016	
	Tier-0	All	Tier-0	All	Tier-0	All
ALICE	90	366	175	495	215	609
ATLAS	111	856	205	1’175	257	1’343
CMS	121	738	271	1’071	317	1’417
LHCb	34	218	36	240	51	315
Total	**356**	**2’178**	**687**	**2’981**	**840**	**3’684**

- One x86 core: 6…15 HS06

WLCG Resources

[kHS06]	2014		2015		2016	
	Tier-0	All	Tier-0	All	Tier-0	All
ALICE	90	366	175	495	215	609
ATLAS	111	856	205	1’175	257	1’343
CMS	121	738	271	1’071	317	1’417
LHCb	34	218	36	240	51	315
Total	**356**	**2’178**	**687**	**2’981**	**840**	**3’684**

- One x86 core: 6…15 HS06
- At CERN:
  - Some capacity provided in addition for analysis (Tier-3)
  - Experiments choose to split pledge across batch, cloud, and service nodes

Current Situation – Batch at CERN

- Currently (08 May) deployed:
  - 4’058 worker nodes (of which 3’669 virtual)
  - 58’488 cores
  - 530 kHS06

Current Situation – Batch at CERN

- Currently (08 May) deployed:
  - 4’058 worker nodes (of which 3’669 virtual)
  - 58’488 cores
  - 530 kHS06
- Some 400’000 jobs per day, mostly single-threaded (one core)

Current Situation – Batch at CERN

- Currently (08 May) deployed:
  - 4’058 worker nodes (of which 3’669 virtual)
  - 58’488 cores
  - 530 kHS06
- Some 400’000 jobs per day, mostly single-threaded (one core)
- Mix of local and Grid submission
  - Grid: Experiment frameworks submit to Cream CEs
  - Grid amounts to 20…40% of submissions at CERN

Current Situation – Batch at CERN

- Currently (08 May) deployed:
  - 4’058 worker nodes (of which 3’669 virtual)
  - 58’488 cores
  - 530 kHS06
- Some 400’000 jobs per day, mostly single-threaded (one core)
- Mix of local and Grid submission
  - Grid: Experiment frameworks submit to Cream CEs
  - Grid amounts to 20…40% of submissions at CERN
- Some 25’000 more cores to come before Run 2 physics

Workload Management

- Since the late 1990s, CERN has been using a commercial product: Platform Inc.’s Load Sharing Facility LSF

Workload Management

- Since the late 1990s, CERN has been using a commercial product: Platform Inc.'s Load Sharing Facility LSF
- Platform Inc. was acquired by IBM in 2011/2012

Workload Management

- Since the late 1990s, CERN has been using a commercial product: Platform Inc.'s Load Sharing Facility LSF
- Platform Inc. was acquired by IBM in 2011/2012
- CERN's licence is perpetual, maintenance is currently covered until November 2017

Workload Management

- Since the late 1990s, CERN has been using a commercial product: Platform Inc.'s Load Sharing Facility LSF
- Platform Inc. was acquired by IBM in 2011/2012
- CERN's licence is perpetual, maintenance is currently covered until November 2017
- We are running release 7.0.6
  - Releases 8 and 9 are out; no significant advantages for CERN

Pain Points with LSF (1)

Goal	LSF constraint
30’000…50’000 worker nodes	Max. ~ 6’500 worker nodes
Dynamic cluster	Adding/removing worker nodes requires cluster reconfiguration
10…100 Hz dispatch rate	Transient dispatch problems – sometimes difficult to ensure 1 Hz
100 Hz query scaling	Slow query / submission response times, queries affect submissions
Licence-free system	Licensed product

Pain Points with LSF (2)

- Worker node scaling:
  - Needed as resources grow by more than 100% from 2014 to 2016; unclear what future distribution of batch vs. cloud resources will be

Pain Points with LSF (2)

- Worker node scaling:
  - Needed as resources grow by more than 100% from 2014 to 2016; unclear what future distribution of batch vs. cloud resources will be
  - Limit appears architecture-related (some central processes single-threaded)

Pain Points with LSF (2)

- Worker node scaling:
  - Needed as resources grow by more than 100% from 2014 to 2016; unclear what future distribution of batch vs. cloud resources will be
  - Limit appears architecture-related (some central processes single-threaded)
  - Limit already constrains us to use unnaturally large VMs (whole hypervisor)

Pain Points with LSF (2)

- Worker node scaling:
  - Needed as resources grow by more than 100% from 2014 to 2016; unclear what future distribution of batch vs. cloud resources will be
  - Limit appears architecture-related (some central processes single-threaded)
  - Limit already constrains us to use unnaturally large VMs (whole hypervisor)
  - Limit not changed significantly with LSF 8/9
    - Can set up multiple instances that can submit to each other

Pain Points with LSF (3)

- Cluster dynamism:
  - LSF reconfigurations are expensive – at least some 10 minutes of unresponsiveness

Pain Points with LSF (3)

- Cluster dynamism:
  - LSF reconfigurations are expensive – at least some 10 minutes of unresponsiveness
  - We are running it once per day
    - Sometimes reconfiguration fails, leading to loss of queues etc.

Pain Points with LSF (3)

- Cluster dynamism:
  - LSF reconfigurations are expensive – at least some 10 minutes of unresponsiveness
  - We are running it once per day
    - Sometimes reconfiguration fails, leading to loss of queues etc.
  - Some operations require two reconfigurations, hence up to 48 hours of delay to become effective

Pain Points with LSF (4)

- Query rate:
  - LSF is not (cannot be) protected against users hammering the system with expensive queries

Pain Points with LSF (4)

- Query rate:
  - LSF is not (cannot be) protected against users hammering the system with expensive queries
  - Number of cases in the past where submissions and job dispatch were seriously affected by query activity

Pain Points with LSF (4)

- Query rate:
  - LSF is not (cannot be) protected against users hammering the system with expensive queries
  - Number of cases in the past where submissions and job dispatch were seriously affected by query activity
  - For ATLAS Tier-0 processing for Run 2, separate LSF instance established

Alternatives to LSF 7 (1)

- LSF 8 or 9
  - Not really addressing any one of our pain points

Alternatives to LSF 7 (1)

- LSF 8 or 9
  - Not really addressing any one of our pain points
- PBS offsprings
  - Way too much trouble reported by other LCG sites

Alternatives to LSF 7 (1)

- LSF 8 or 9
  - Not really addressing any one of our pain points
- PBS offsprings
  - Way too much trouble reported by other LCG sites
- SLURM
  - Considered because of claimed scalability
  - Good for many cores for massively parallel computing, serious scaling limits on worker nodes and job slots

Alternatives to LSF 7 (1)

- LSF 8 or 9
  - Not really addressing any one of our pain points
- PBS offsprings
  - Way too much trouble reported by other LCG sites
- SLURM
  - Considered because of claimed scalability
  - Good for many cores for massively parallel computing, serious scaling limits on worker nodes and job slots
- Grid Engine
  - Univa Grid Engine is the only serious contender left
  - Commercial, similar architecture to LSF

Alternatives to LSF 7: HTCondor

- Open-source, academic environment

Alternatives to LSF 7: HTCondor

- Open-source, academic environment
- Already in widespread use in WLCG, e.g. FNAL, BNL, RAL – good experience
  - CERN's requirements are different: CERN cluster already largest and growing; CERN needs to also support local job submission with AFS token passing/extension

Alternatives to LSF 7: HTCondor

- Open-source, academic environment
- Already in widespread use in WLCG, e.g. FNAL, BNL, RAL – good experience
  - CERN's requirements are different: CERN cluster already largest and growing; CERN needs to also support local job submission with AFS token passing/extension
- HTCondor also used in experiment frameworks (and even as a CE…), can be used as cloud scheduler
  - Potential for future further integration

Alternatives to LSF 7: HTCondor

- Open-source, academic environment
- Already in widespread use in WLCG, e.g. FNAL, BNL, RAL – good experience
  - CERN's requirements are different: CERN cluster already largest and growing; CERN needs to also support local job submission with AFS token passing/extension
- HTCondor also used in experiment frameworks (and even as a CE…), can be used as cloud scheduler
  - Potential for future further integration
- Tests so far very successful
  - Adding/removing worker nodes
  - Failing central manager/submission nodes unproblematic
  - Query scaling revealed an issue, fixed by developers very soon after

Alternatives to LSF 7: HTCondor

- Open-source, academic environment
- Already in widespread use in WLCG, e.g. FNAL, BNL, RAL – good experience
  - CERN's requirements are different: CERN cluster already largest and growing; CERN needs to also support local job submission with AFS token passing/extension
- HTCondor also used in experiment frameworks (and even as a CE…), can be used as cloud scheduler
  - Potential for future further integration
- Tests so far very successful
  - Adding/removing worker nodes
  - Failing central manager/submission nodes unproblematic
  - Query scaling revealed an issue, fixed by developers very soon after
- Scaling test (shadows on LSF worker nodes) looked promising
  - 2 central managers, 20 schedulers/submission nodes, 1'300 worker nodes with 62'500 job slots
  - Architecture promises to support further scale-out (unlike LSF, GE, SLURM etc.)

HTCondor Scaling Behaviour

- Job submission time as function of number of worker nodes and total number of jobs

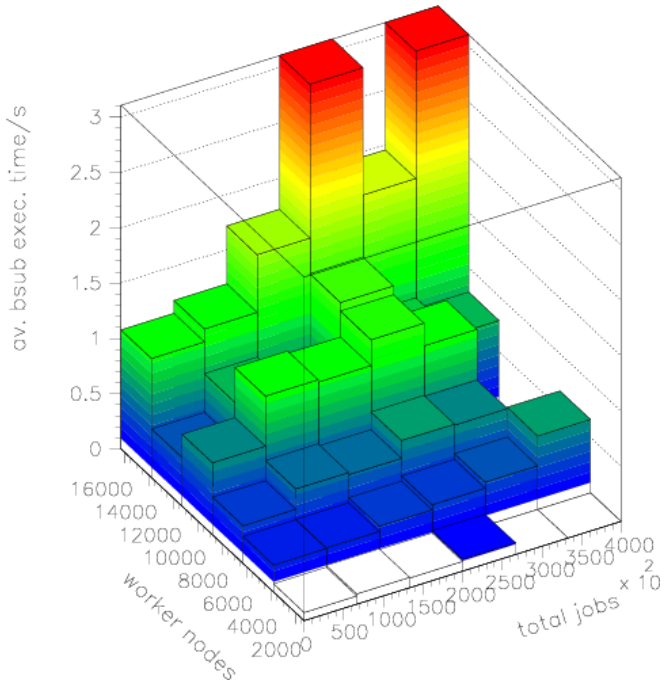

LSF

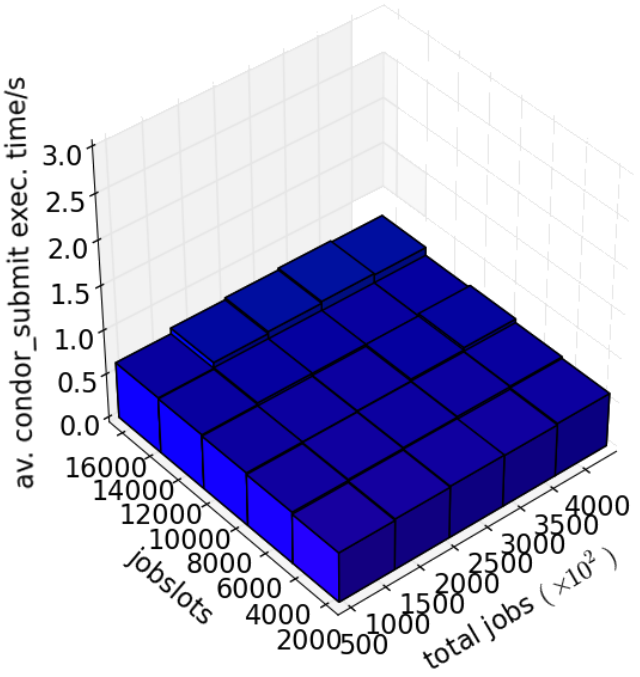

HTCondor

HTCondor Deployment Steps (1)

- Start with a (small) service offering Grid submission only
  - Mostly transparent to users
  - Doesn’t require AFS token passing and extension

HTCondor Deployment Steps (1)

- Start with a (small) service offering Grid submission only
  - Mostly transparent to users
  - Doesn't require AFS token passing and extension

Done – see following talk by Iain Steers

HTCondor Deployment Steps (1)

- Start with a (small) service offering Grid submission only
  - Mostly transparent to users
  - Doesn't require AFS token passing and extension
- Grow that service (up to taking all Grid submissions)
  - Overflowing into LSF part via condor_glidein possible

Done – see following talk by Iain Steers

HTCondor Deployment Steps (2)

- Once necessary developments done, open small service for local job submissions
  - Still to be seen to what extent we can (and wish!) to make condor submission look like LSF submission, idem for queries
  - User support (documentation, handholding, tutorials etc.) will be integral part of deployment (and take significant resources!)

HTCondor Deployment Steps (2)

- Once necessary developments done, open small service for local job submissions
  - Still to be seen to what extent we can (and wish!) to make condor submission look like LSF submission, idem for queries
  - User support (documentation, handholding, tutorials etc.) will be integral part of deployment (and take significant resources!)
- Grow to full size, reducing LSF capacity
  - Close interaction with user community

HTCondor Deployment Timescale

- Grid submissions: see Iain's talk

HTCondor Deployment Timescale

- Grid submissions: see Iain's talk
- Timescale for local submission developments and service to be defined
  - Hoping for pilot by end 2015, but…
  - Priority is on full scale and production quality service for Grid submissions

HTCondor Deployment Timescale

- Grid submissions: see Iain's talk
- Timescale for local submission developments and service to be defined
  - Hoping for pilot by end 2015, but…
  - Priority is on full scale and production quality service for Grid submissions
- Target: Terminate LSF service by end of Run 2