# The U.S. ATLAS Facility and Agile Infrastructure

Michael Ernst, BNL

US ATLAS Distributed Facilities Meeting, University of Arizona

December 11, 2013

# LHC Long Shutdown 1 (LS1)

**We're about halfway through LS1…**

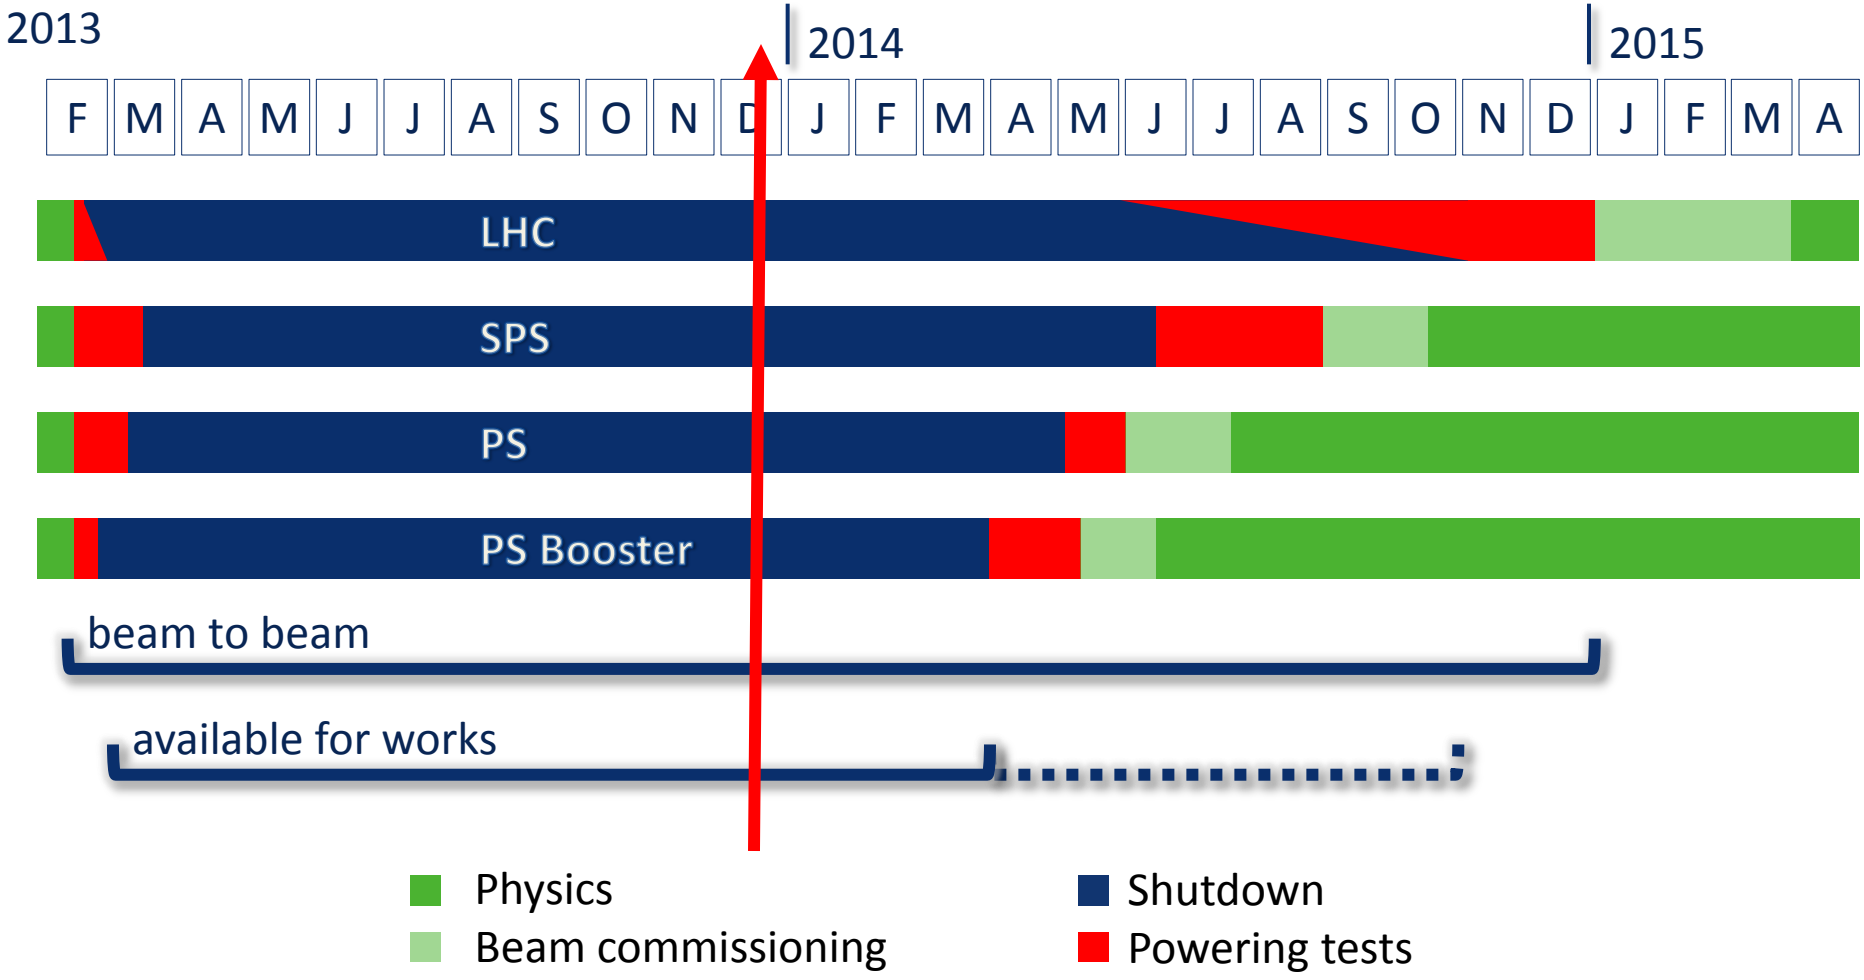


**… beam still foreseen for January 2015 (physics in ~April)**

# LHC schedule beyond LS1

Only EYETS (19 weeks) (no Linac4 connection during Run2)

LS2 starting in 2018 (July) 18 months + 3months BC (Beam Commissioning)

LS3 LHC: starting in 2023 => 30 months + 3 BC

injectors: in 2024 => 13 months + 3 BC


LHC schedule approved by CERN management and LHC experiments spokespersons and technical coordinators
Monday 2nd December 2013

# LHC Run 2

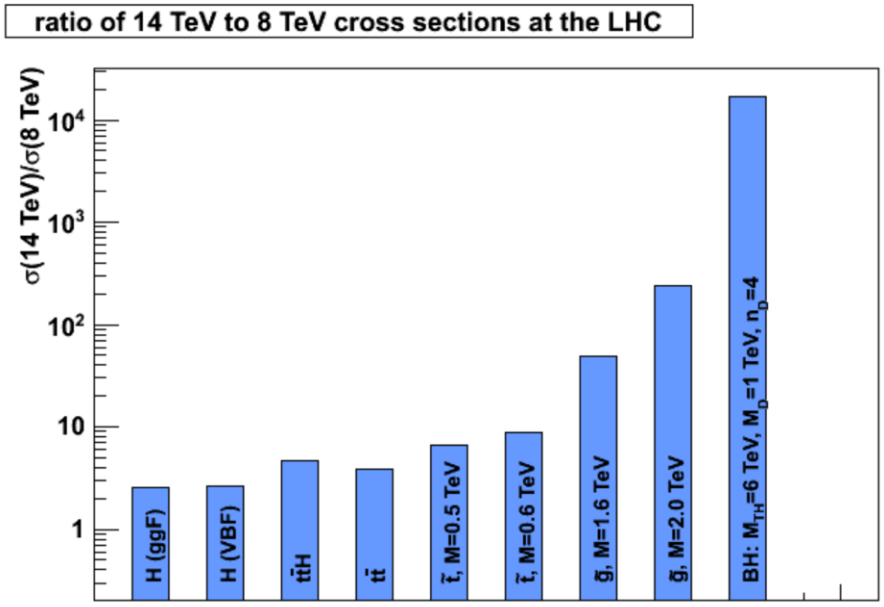


- 2015 is a whole new energy domain
  - Every measurement is new
  - Every search has a new chance
- **Early Run-2 data has discovery potential**
  - With a few $fb^{-1}$ many searches reach or surpass current sensitivity
  - Discovery of TeV scale particles possible
- Center of mass energy ~2x
- Event rate to storage ~2x, pile-up above 30, 25 ns bunch spacing
- Foreseen computing capacity increment much less than extrapolation from 2013 practice would suggest

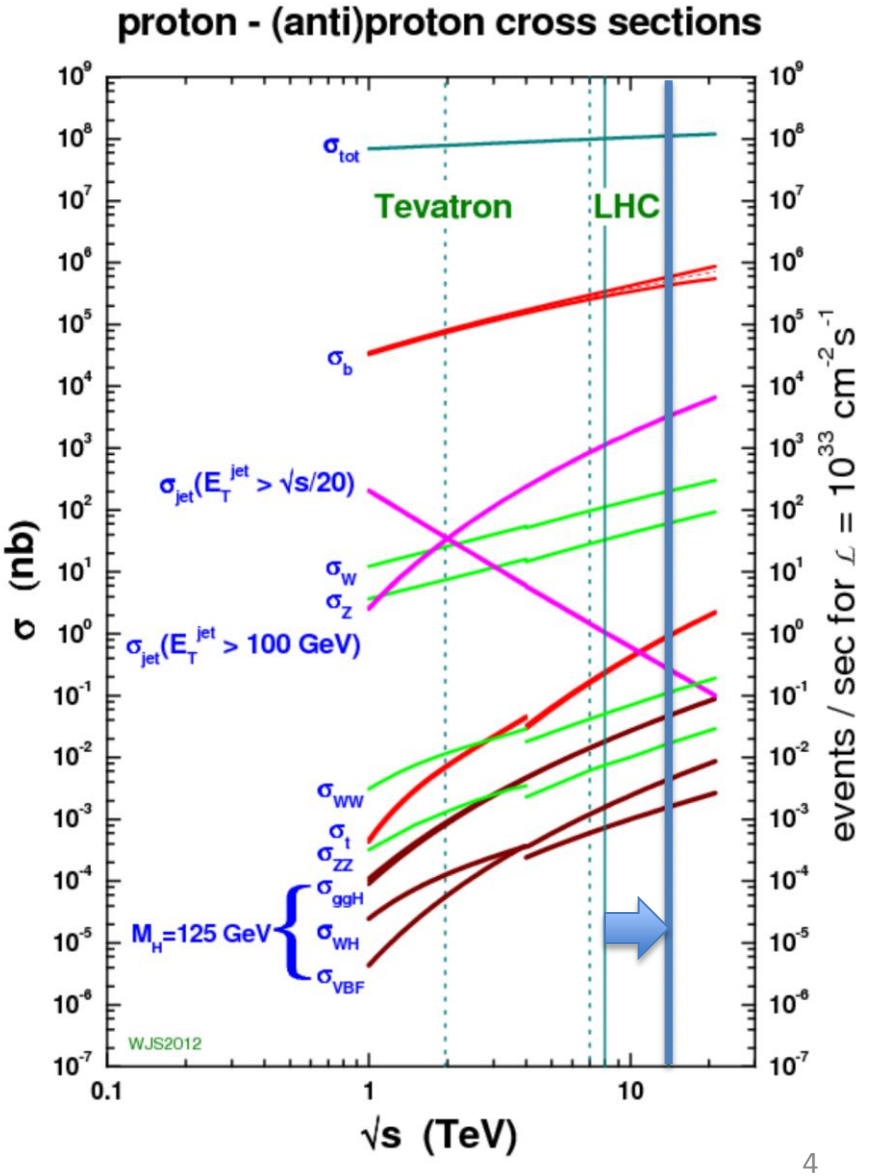

# LS1 S&C Upgrades

- ATLAS S&C following an evolutionary plan towards Run-2
- A sizable, complex and challenging task
  - Tight budgets – maximize event throughput per unit cost
  - Computing requirements grow substantially
  - Computing model evolution leveraging powerful networks will continue, putting new demands on facilities and fabric, e.g.
    - Federated Data Stores w/ direct access across the WAN
    - Potential to (transparently) integrate user analysis stations with resources & capabilities at Tier2s and the Tier-1 (e.g. ATLAS Connect)
  - Technology evolution presents opportunities but also challenges
    - New processors => Sequential to Parallel processing
  - We learned much in Run-1 all along the processing & analysis chain, we want to incorporate the lessons
  - Facilities delivered stable ops and high availability in Run-1, LS1 is an opportunity for major upgrades
- Hence a broad range of upgrades planned and now well underway
  - While sustaining full scale steady state operations

# Multi-Core

- GEANT 4.10 has the capability for event level parallelism to run on multiple cores
  - Saving memory by sharing common elements like geometry
  - Allows to utilize highly parallel machines efficiently
- ATLAS is looking into threaded frameworks
  - Event level parallelism as a start with multi-threaded reconstruction to follow
- I think a good transition system would be through the pilot systems
  - A pilot job would be assigned multiple cores and assign a mix of single and multiple core workflows
  - Production applications will move to multi-core first
    - Analysis may be single core for quite some time

# Transition

- My guess is that we will move to new hardware platforms when they make a significant enough impact that they can replace a workflow nearly completely
  - For example if suddenly a year's worth of Gen-Sim could be done on a co-processor enabled farm in a month
    - This is an example of the kind of resource we might be willing to get an allocation and free up dedicated computing for analysis
  - If reprocessing could be done 10 times faster on a GPU farm, we could imagine getting one farm and freeing the remaining Tier-1 computing for other tasks
- Dedicated hardware will most likely be applied to dedicated tasks
  - Freeing up other computing resources

# Run-2 Readiness

- The LS1 S&C upgrade program is progressing on many fronts
- ATLAS needs to get Run-2 physics out quickly
  - As in Run-1 have to avoid limitations from computing
- How do we ensure S&C is ready?
  - Facility deliverables completed according to the plan
  - Components integrated as a functioning system
    - Important for the Facilities in view of many changes
  - Computing Fabric operational at scale
    - Storage & Network performance is the primary concern
- Data Challenge (DC14) to assess S&C Readiness
  - Will run from early July until September 2014

# Goals for Data Challenge 2014 (DC14)

- Prepare for Run-2 physics analyses
- Test the Run-2 analysis model
  - xAOD, reduction framework, common analysis tools
- Gain experience with Run-2 analysis model
  - Possibly adjust model based on that experience
- Commission the Integrated Simulation Framework (ISF) in context of physics analysis
- Test updated reconstruction algorithms (Rel19)
- Provide large scale test of upgraded distributed computing environment

# S&C Timeline for LS1

Details of this schedule are evolving…

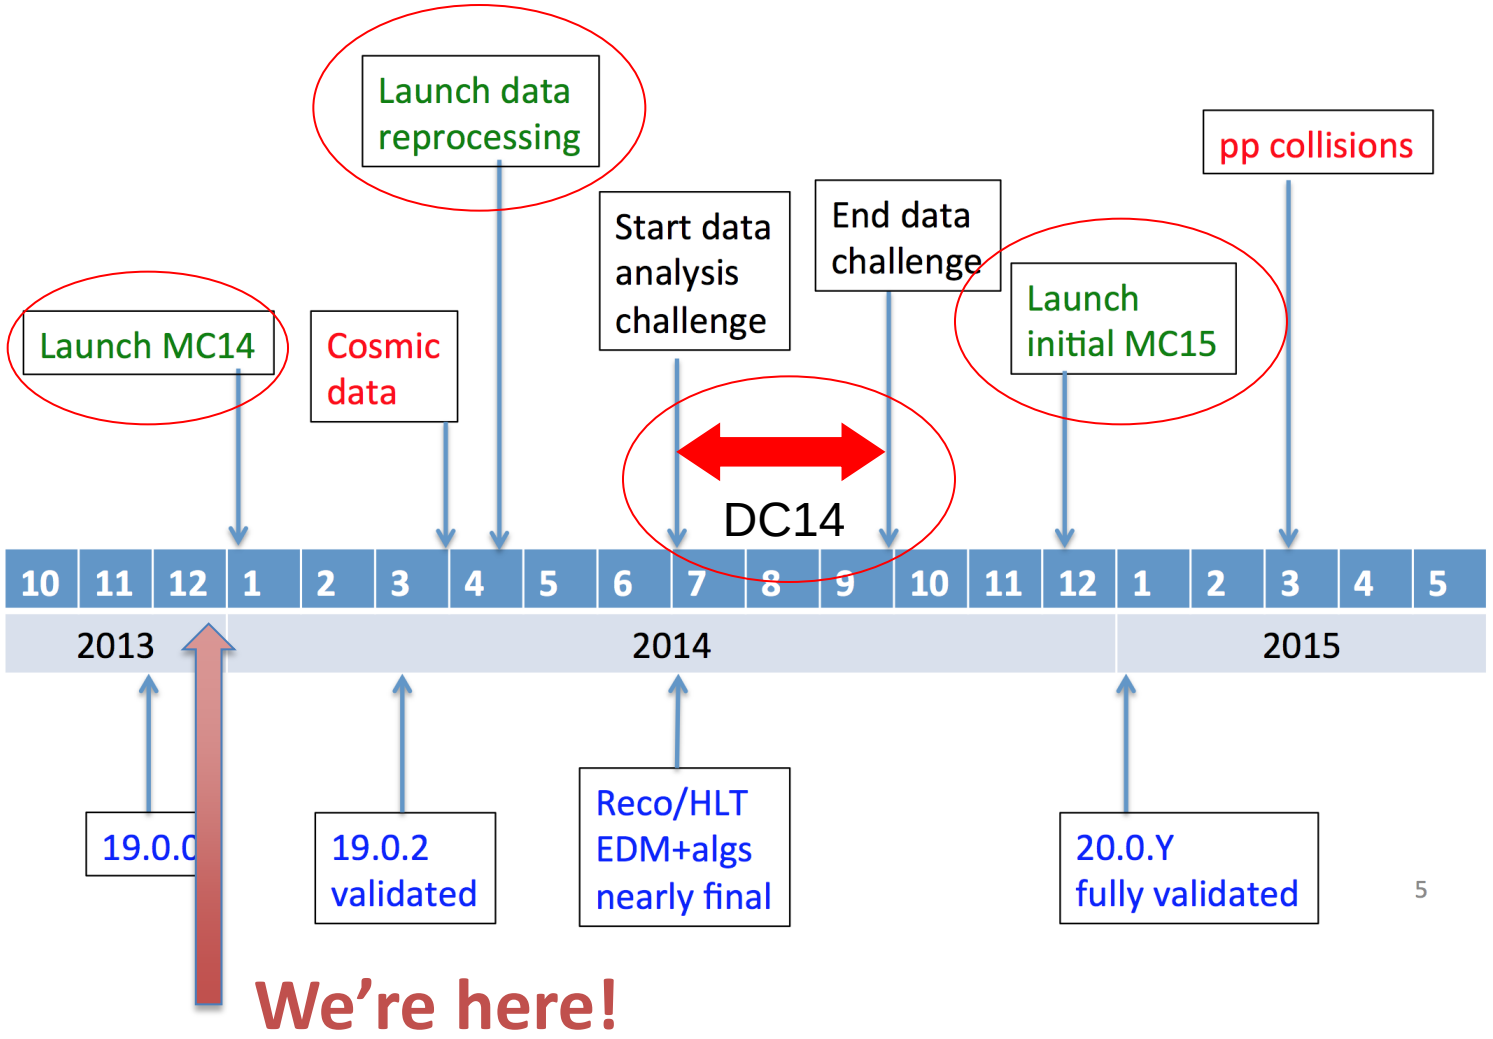


5

# The value of DC14 (1/2)

- DC14 provides a milestone with inter-dependent deliverables coming together
  - Likelihood that lessons from DC14 will drive a round of adjustment/development, so there's a premium – or necessity – to being ready for DC14 and making the most of it
- For example…
  - xAOD – including its impact on workflows, data access performance, …
  - Federated Storage – its impact on production and analysis efficiency
  - ISF – all of DC14 to be based on ISF and G4
  - Reduction framework – the new train model
  - Massive use of athenaMP – else lose 30-50% of cores due to insufficient memory
  - Prodsys2 – particularly DEFT which is green-field new and a dependency for reduction framework. Prodsys1 is not a fallback
  - Rucio – will benefit from very nice adiabatic completion/migration plan
  - Analysis releases and analysis framework – users should have migrated to analysis release 19 by DC14

# The value of DC14 (2/2)

- New/increased demands on facilities and networking – are we on track to meet the demand in DC14?
  - Readiness of multicore queues for athenaMP bulk production
    - Need to move from hard partitioning to something more dynamic for optimal resource utilization for parallel/sequential job env
  - Greater role for tape as disk replica counts fall
    - Readiness of tape facilities and the workflows that utilize them
  - Increased usage of WAN data access
    - Supporting, managing, monitoring WAN as well as LAN loads
    - Requires special attention to facility/site configuration
      - Brian will address those in his talk tomorrow

# Flexibility

- There is some skepticism because so much effort was invested in getting the data close
- Computing intensive tasks like reprocessing can be sustained reading data from remote storage
  - Input size is small compared to the time of the application
  - ~100kB/s is enough to sustain the ATLAS application per job slot
  - Even thousands of cores can be reasonably fed with Gb/s
- Analysis as is today can be supported but it would be better to have data formats that allow to read the needed objects only
- Intelligent IO and high capacity networks have changed the rules

# Networks

- 100Gb/s links are in production for OSG/WLCG
  - The number is still growing fast, i.e. site connectivity
  - When reaching 100Gb/s WAN network, there is little difference between WAN and LAN for capacity (though re latency)
- We have reduced the differences in site functionality
- Then reduced the difference in data accessibility
- We are in the process of reducing the difference in the perception that two *logical* sites are separate
  - We can begin to think of the facility as a big center and not a cluster of centers

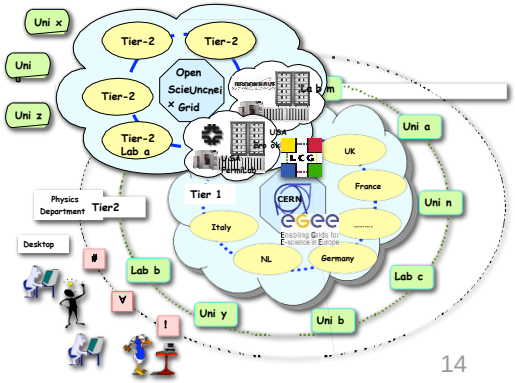

# Changing the Services

- The OSG/WLCG service architecture has been reasonably stable for over a decade
  - This is beginning to change with new Middleware for resource provisioning
- A variety of places are opening their resources to “Cloud” type of provisioning
  - From a site perspective this is often chosen for cluster management and flexibility reasons
  - Everything is virtualized and services are put on top
- There is nothing that prevents a site from bringing up exactly the same environment currently deployed for the WLCG, but maybe it’s not needed

# Flexibility

•We have ways of packaging the local environment and deploying it even in user space (cvmfs local mounts).
•We can make virtual machines if we have to. We can take very diverse systems and make them look the same to our application

•Sending data over the wide areas means not having to care about data location and combined with controlling the environment it is a transformative change

- Anything you bring to the table can make a contribution
- Local, opportunistic, and temporarily allocated contributions can be a huge increase in capacity

# Evolving the Infrastructure

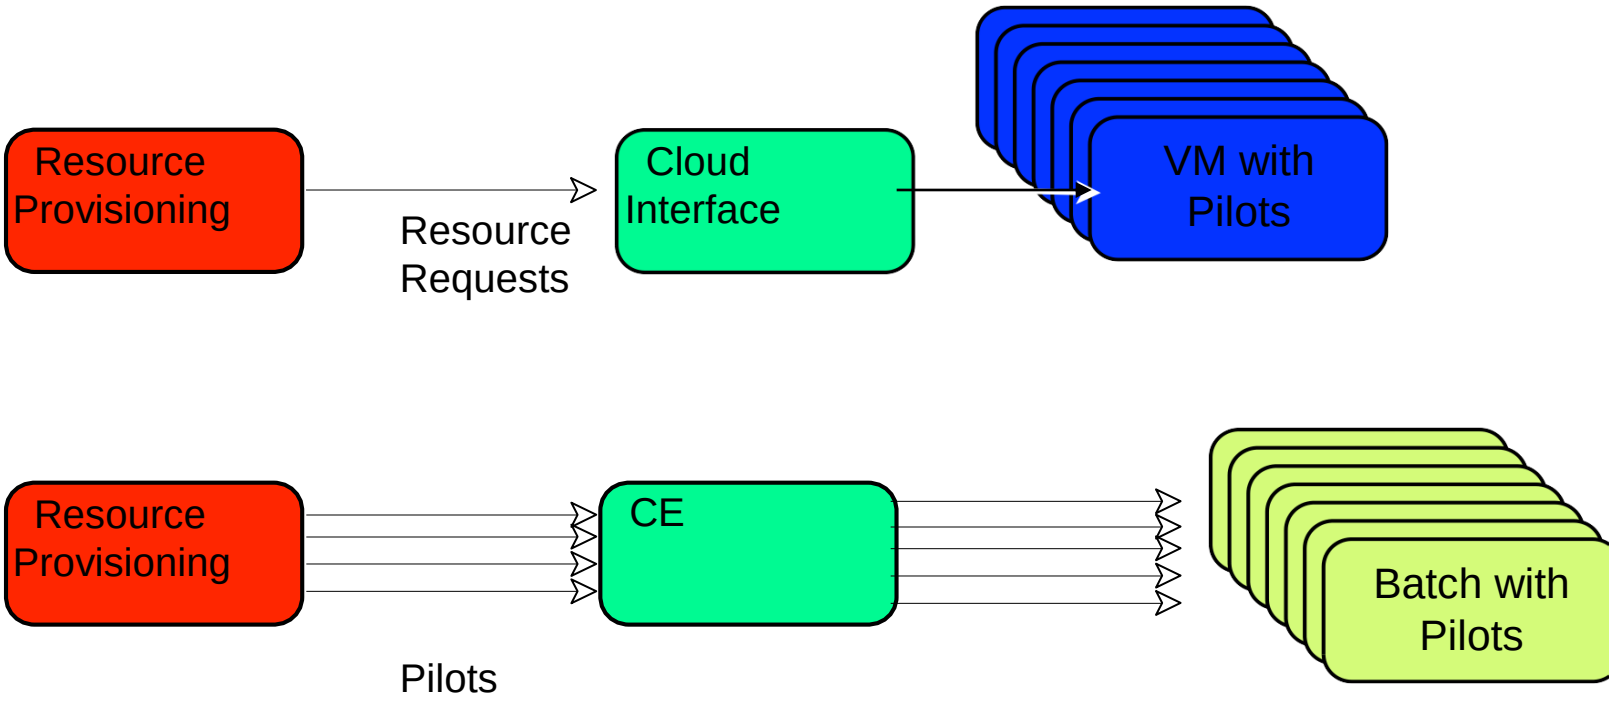


- In the new resource provisioning model the pilot infrastructure communicates with the resource provisioning tools directly
  - Requesting groups of machines for periods of time

# Impact of Resource Provisioning

- The resource provisioning layer becomes a potentially common, but complicated interface
  - Goal is to get pilots into a diverse set of computers
    - Dedicated resources, opportunistic computing, local resources, clouds, super computer allocations, etc.

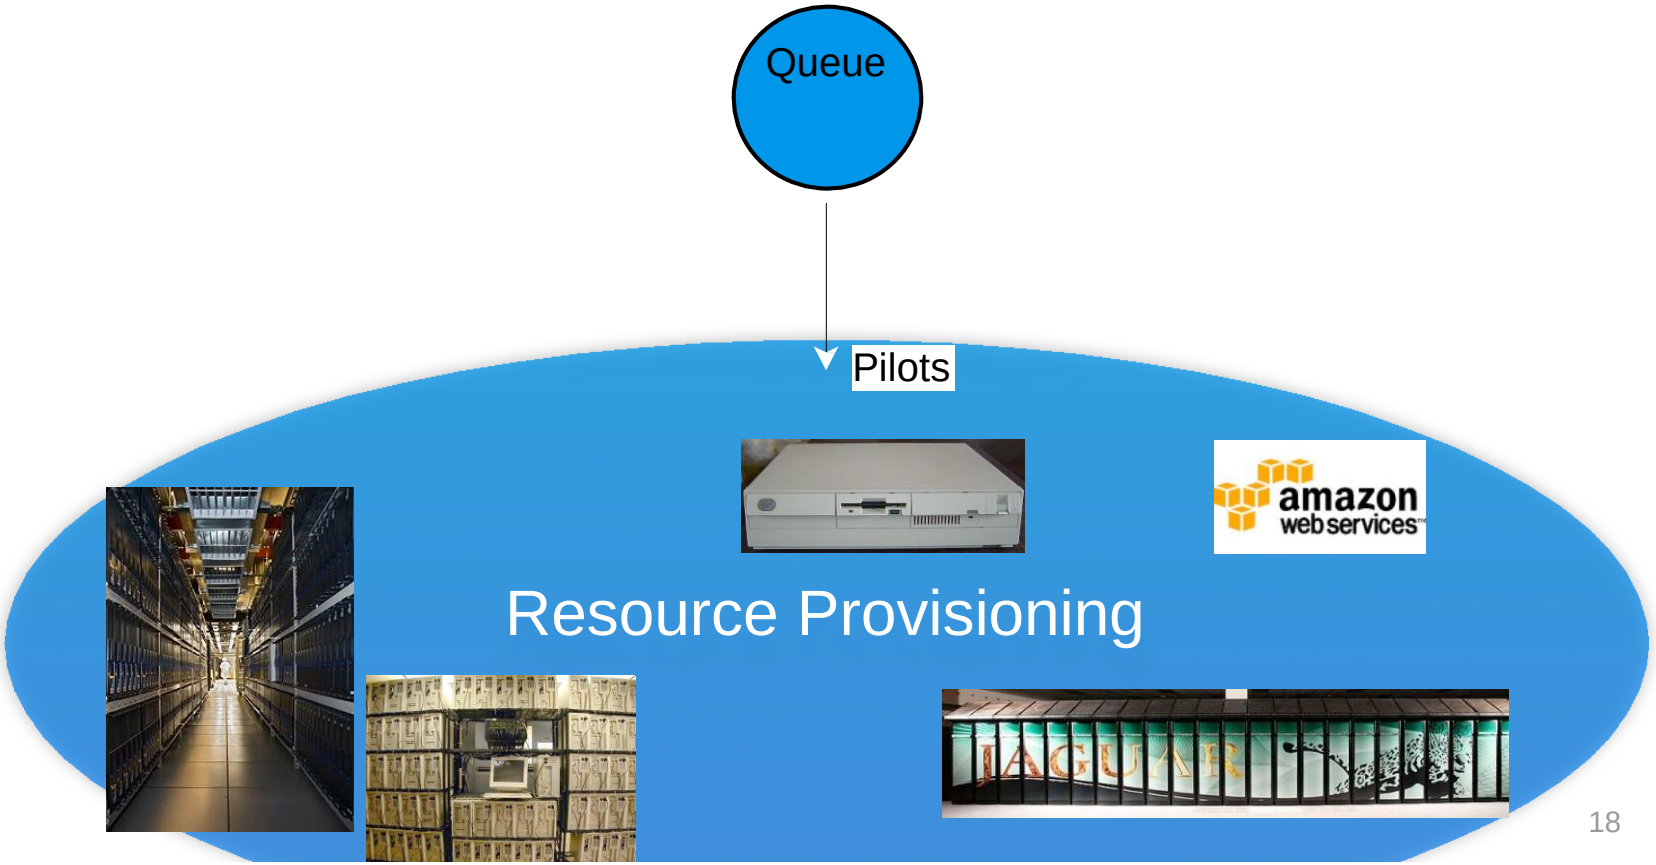

# Evolution

- After many years of operations our system is evolving

– We are reducing how strictly we define the functionality in each tier
  - Lines and capabilities are blurring together

– We have much better access to data
  - With it we begin to eliminate the boundaries between sites

– We will have a much more diverse set of resources
  - Cloud and opportunistic access will be big areas of growth

– The availability of capability computing is likely to change the landscape

# Testing Interfaces

- At the T1 we are bringing up a new agile infrastructure of services that will be configurable through a cloud interface
  - Similar to how we interacts with the HLT
  - We should consider using similar resources at more sites.
- It is potentially a big change in how we handle prioritization and release of resources
  - Needs to be tested, but could be more efficient for us and the experiment

# Agile Infrastructure - IaaS

- Rationale
  - Need to manage a lot of servers
  - No increase in staff effort
  - …we have to manage the infrastructure efficiently
- Infrastructure as a Service (IaaS)

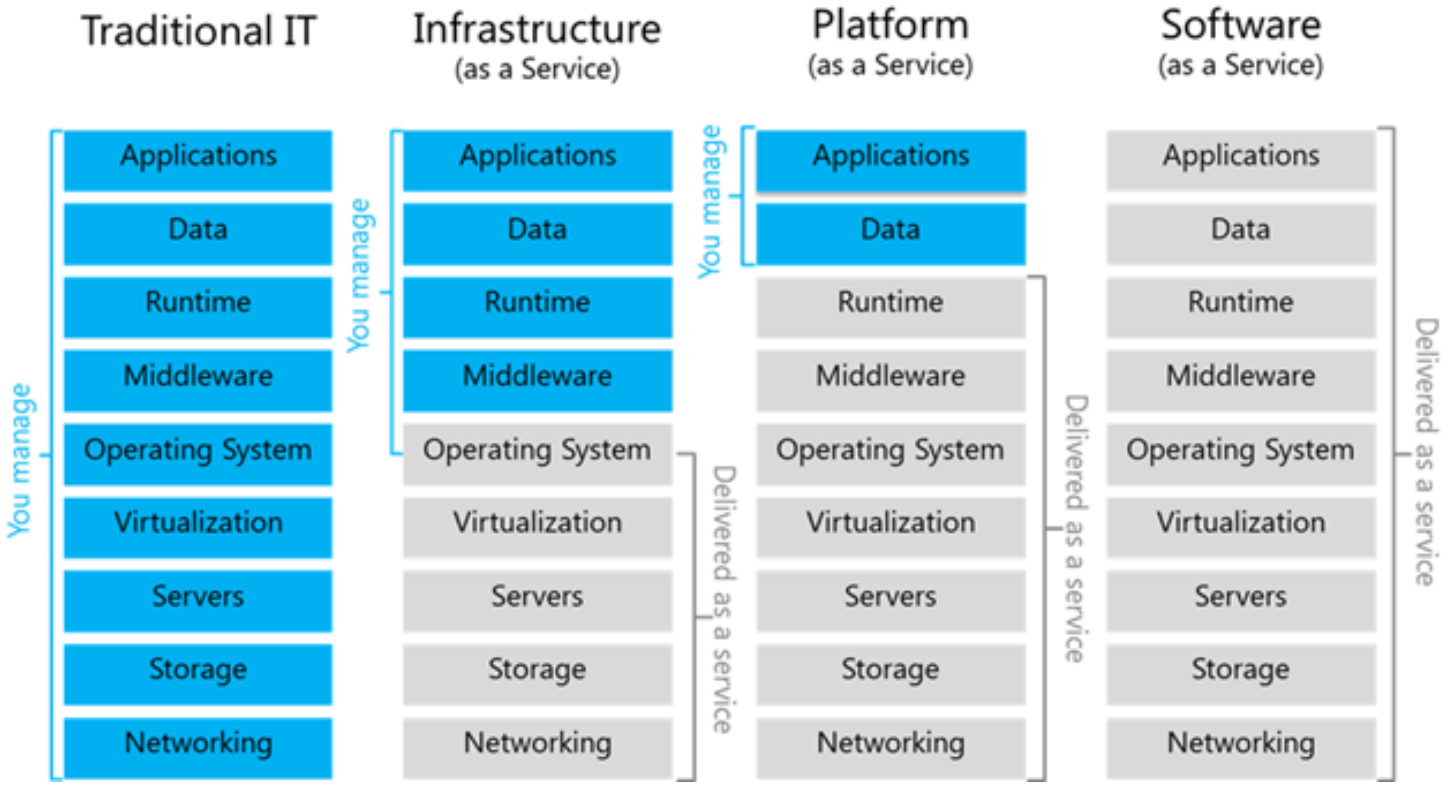

# U.S. ATLAS Facilities as a Cloud Service Provider ?

- Following industry reference structure
  - Infrastructure, Platform and Applications
- Start first with Infrastructure-as-a-Service at scale
  - Just provide Amazon functionality
- Value increases up the chain so we need a vibrant ecosystem for PaaS and SaaS

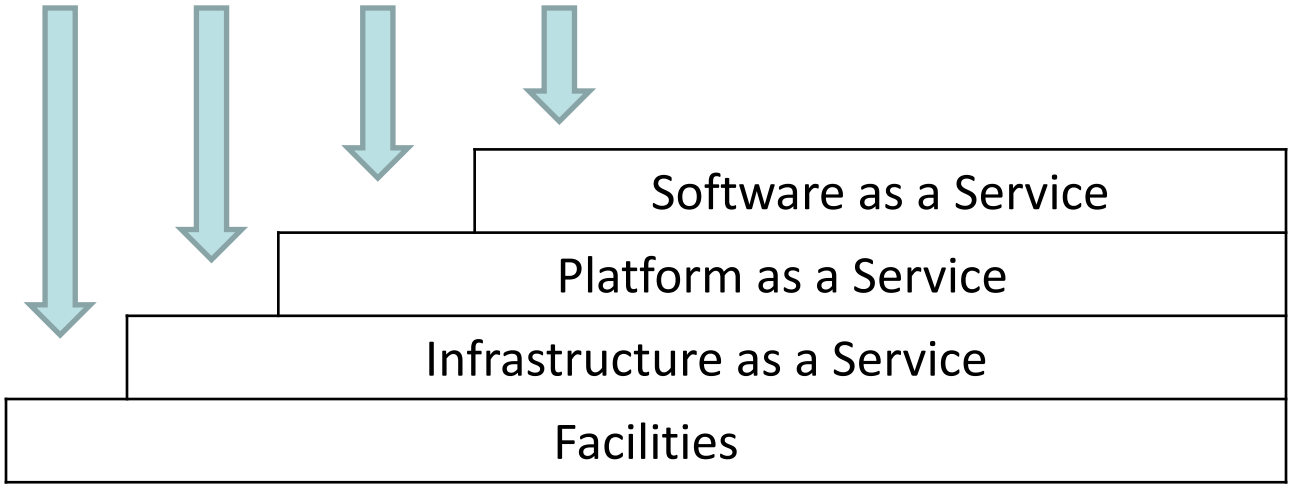

# Goals

- Objectives
  - 90% of hardware virtualized
  - ~3-5k hypervisors needed over next two years
- Following industry reference structure
  - Infrastructure, Platform and Applications
- Deploy multi-site
  - Extend to Tier-2 sites
  - Enable remote management
- More efficient use of our hardware
- Better tracking of usage
- Consolidate support to a single device
- Support potential new use cases (PaaS)

# Timeline

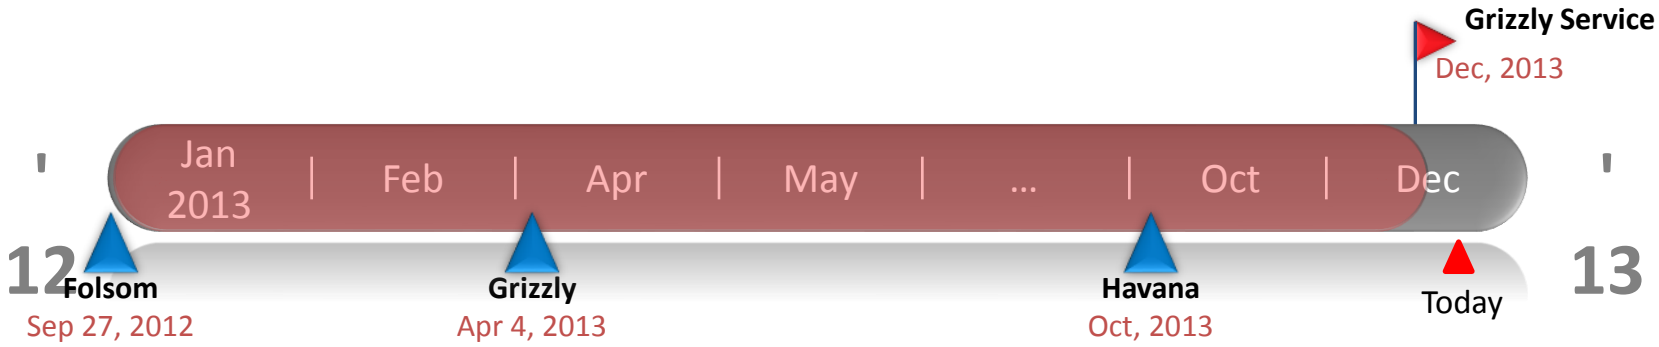


- Maintain the project linked to upstream releases
- Deployment delay:
  - Time to understand (challenging learning curve) and for testing
  - Checking suitability of new features
- We have been delivering a series of pre-production services…
  - Adding functionality…
  - Scaling out…
  - Integrating into infrastructure at the T1 …

# Interfaces

- Multiple interfaces available
  - Supported OpenStack APIs for python
  - REST APIs for many public domain libraries available
  - Same API for SIM@P1 cloud
- We have demonstrated that PanDA can use cloud resources transparently
- Foreman GUI/CLI to create a server with Puppet
  - Create me a new (e.g.) analysis server instance with 1 command

# Agile Infrastructure - Outlook (I)

- Going to production for managed servers
  - End of Jan/Feb, 2014
  - Based on OpenStack Grizzly release
  - ~200 Hypervisors

# Agile Infrastructure - Outlook (II)

- Non-managed machines, missing requirements:
  - Shared storage for volume service and live-migration
- Investigating alternatives for block storage:
  - Ceph, others
- Special attention to cells
  - Exploit cells functionality for scalability
  - Multi-site operations (T2s willing to join?)
- And more:
  - Kiosk to create VMs on demand (user-initiated)

# Planning the Challenges

- This meeting is the opportunity to start understanding and defining the goals and timelines for near/mid-term computing aspects
  - Understanding and mitigating performance issues
  - Agree on areas where R&D is needed
- Planning, in collaboration with others (e.g. T3 implementation TF), can then proceed in the weeks & months to come
- Integrating services currently under development and participating in new activities
  - I.e. making our beyond-pledge resources available to a broad user community in the U.S.
  - This is a huge and hugely important topic we will spend a lot of time on at this meeting
- Tracking critical path deliverables and taking remedial actions, launching backup plans where needed

# Completing a table like this may be useful for initial DC14 / Late LS1 / Run-2 planning…

| | DC14 essential | DC14 best effort | LS1, not DC14 | Run-2 |
|---|---|---|---|---|
| xAOD – including its impact on workflows, data access performance, … | | | | |
| ISF bulk production – all of DC14 to be based on ISF and G4 | | | | |
| Reduction framework and the new train model | | | | |
| Analysis releases and analysis framework tools | | | | |
| Tracking migrated to Eigen | | | | |
| Massive use of athenaMP | | | | |
| Prodsys2 | | | | |
| Rucio | | | | |
| Metadata task force outcomes integrated | | | | |
| | | | | |
| Multicore queues for athenaMP bulk production | | | | |
| Validation of large(r) scale tape usage -- facilities and workflows | | | | |
| WAN data access stress test | | | | |
| | | | | |
| PanDA analysis (in limited region) with loosened data locality | | | | |
| Event service for simulation | | | | |
| HPC site(s) in DC14 production | | | | |
| Volunteer computing (BOINC or…) | | | | |
| (Short) jobs to Amazon EC2 spot market | | | | |

Relevant to Facilities

Torre at Oct S&C week

# Agenda of First Block

09:00 - 10:30 Introduction, Overview, Goals

Convener: Michael Ernst

•09:00 The US ATLAS Computing Facility and Agile Infrastructure *30'*

Speaker: Michael Ernst

•09:30 US ATLAS Facilities Projects: AtlasConnect, FAXbox, T3 Flocking *30'*

Speaker: Robert William Gardner Jr (University of Chicago (US))

•10:00 Overview of the Analysis Support group activities in US ATLAS *30'*

Speaker: Erich Ward Varnes (University of Arizona (US))

# Offline Schedule in LS1

- Data challenge based on 8 TeV data and new 8 TeV and 13 TeV MC in summer 2014 based on release 19
  - six major goals defined (see next slide)
  - release 19 ready for data challenge and cosmic data taking by February 2014
- Use release 20 for 2015 data taking
  - First deadline will be in July 2014 for algorithms and EDM for objects used in the HLT
    - Ensure excellent correspondence between HLT and offline => sharp turnon
  - Final release 20 available in January 2015
- Start MC15 production for 13 TeV data analysis in late 2014
  - Not yet clear how much and which should be started when exactly
    - Need some MC early for detector commissioning and preparation of discovery physics
    - Launch most MC later after beamspot (and energy) really known

Beate Heinemann, June 2013 ATLAS Week

# DC14 Processing

- Run-1 8 TeV data
  - reprocessed with release 19 and updated conditions
- 8 TeV MC14_G4 HITS
  - produced with ISF-based on r17.7 G4 simulation reconstructed with r19 and new conditions
- 8 TeV MC14_AF and/or MC14_FF HITS
  - produced with fast aspects of ISF for commissioning purposes
- 13 TeV MC14 HITS
  - produced with ISF-based simulation (G4, AF and FF) reconstructed with r19 and conditions

G4: Geant4 simulation
AF: AtlFast2 simulation
FF: faster flavors of simulation

Not clear how much MC will be needed => depends on how many and which analyses will be done (~ 1 billion G4 + many billion AF/FF?)

Beate Heinemann, June 2013 ATLAS Week

# What is OpenStack?

"OpenStack is a cloud operating system that controls large pools of compute, storage, and networking resources throughout a datacenter, all managed through a dashboard that gives administrators control while empowering their users to provision resources through a web interface"

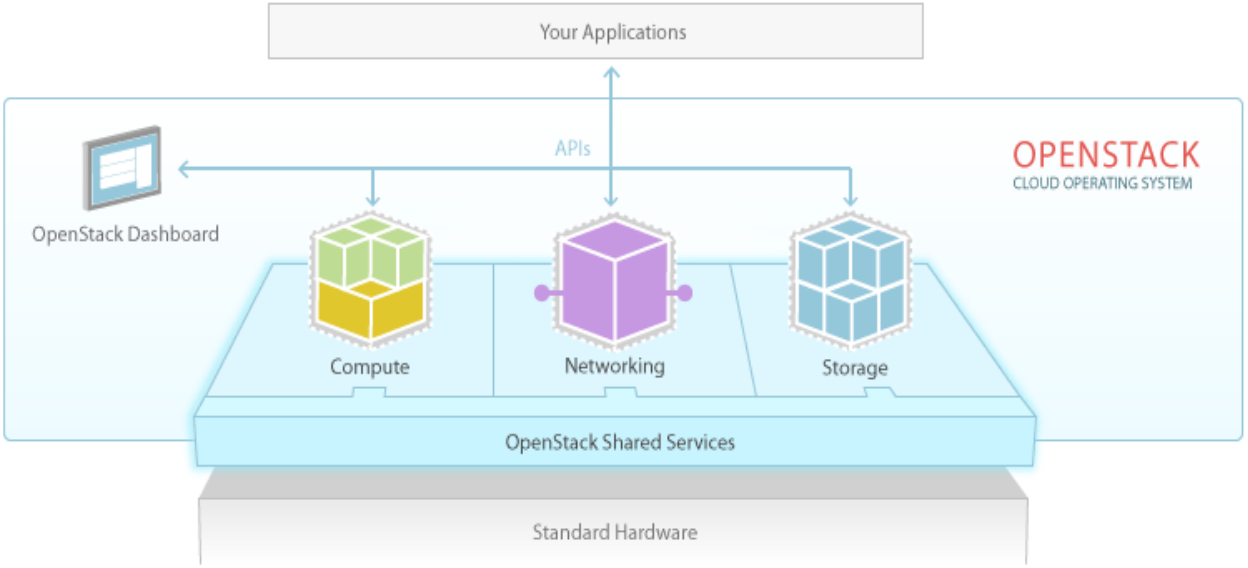


Swift

**Nova**

Quantum