IT-CM-IS
Introduction and Plans

Gavin McCance
IT-CM Group Meeting
26 January 2016

Who are we?

BELLEMAN, Jerome
BORKOVEC, Krystof
BOZEK, Karolina Maria
DIAS MEIRELES, Fernando Miguel
FIELD, Laurence
HOIMYR, Nils-Joar
JONES, Ben Dylan
LEE, Hung-Te
MARQUES COELHO DOS SANTOS, Miguel
MCCANCE, Gavin John
SANDALL, Richard
STEERS, Iain Bradford
STIVANIN, Paolo

Who a

BELLE
BORK
BOZEK
DIAS M
FIELD,
HOIMY
JONES
LEE, H
MARQ
MCCA
SANDA
STEER
STIVA

Section mandate

Provide high-level compute services to the CERN Tier-0 and WLCG

Extension of those services into the public cloud / external hosting

Compute Workloads

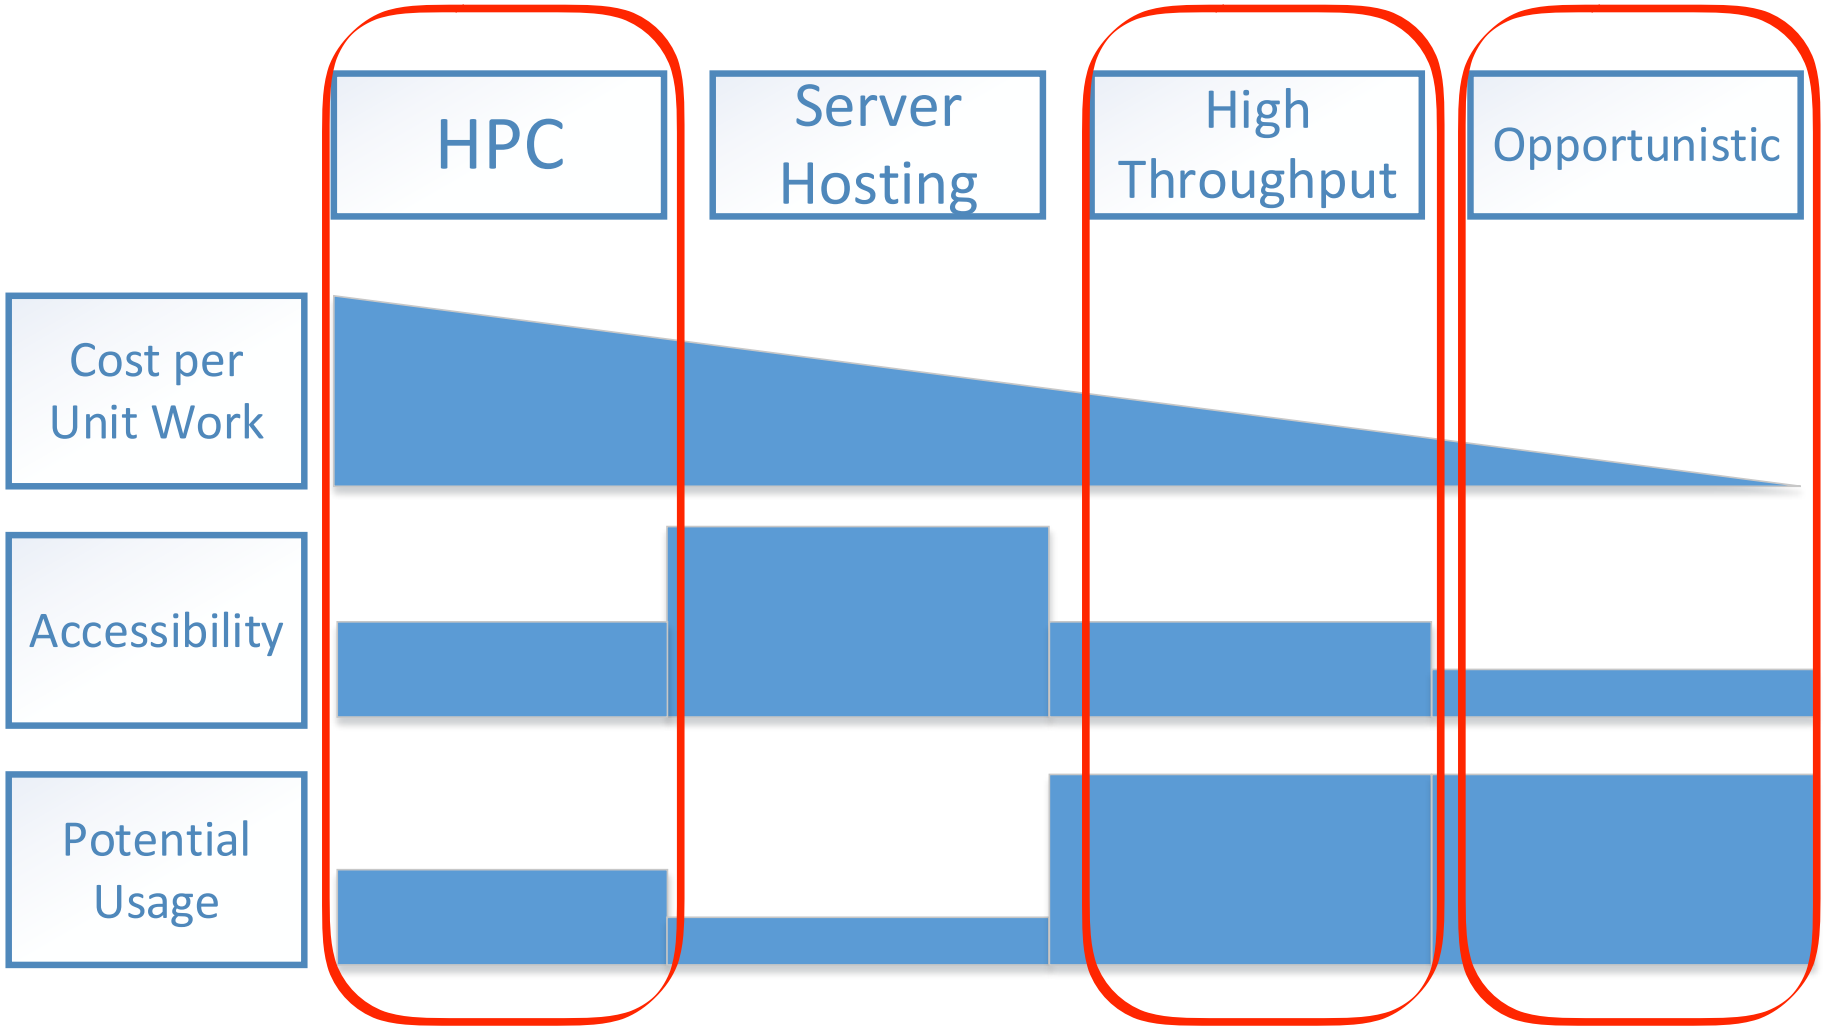

Compute Workloads

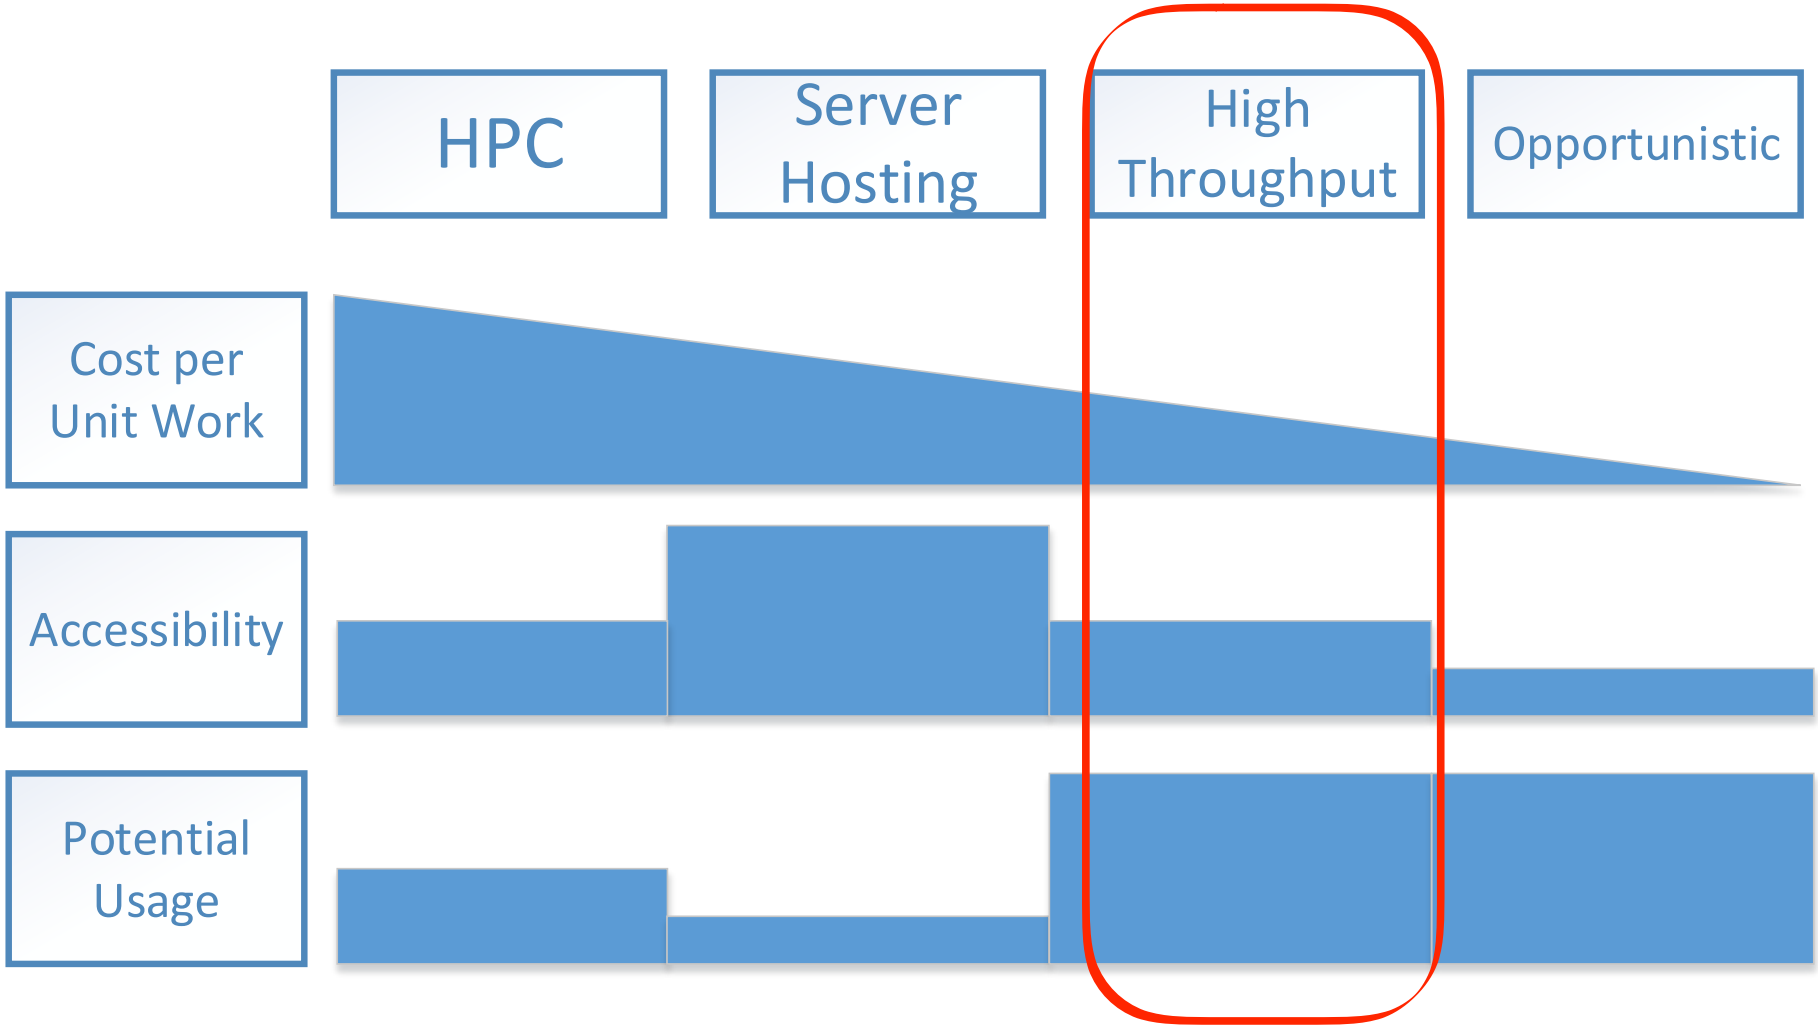

High-throughput batch service

- 90k CPU cores and increasing
  - LHC: Prompt Tier-0 calibration / hot events
  - LHC: Tier-0 bulk reconstruction
  - LHC: Our share of WLCG Grid quota
  - All other CERN experiments (e.g. Compass)
  - Associated experiments (e.g. AMS)
  - Various local CERN groups in EP (ATLAS, LHCb, etc)
  - Typically delivering around 480k jobs per day

High-throughput batch service

- Batch service balances the fair-share across all competing applications according to CERN resource policies
- Ensures unused shares are not wasted but given to other users of the system
- Users interaction pattern: “submit a job, sits in queue, runs, get result back”

- Currently commercial product (IBM LSF)
- We’re in the process of migrating to open-source product (HTCondor)
- We also run misc Griddy services (Compute Element and Argus) interfacing WLCG Grid to local site resources

Compute challenges

The outline LHC schedule out to 2035 presented by Frederick Bordry to the SPC and FC June 2015 can be found here

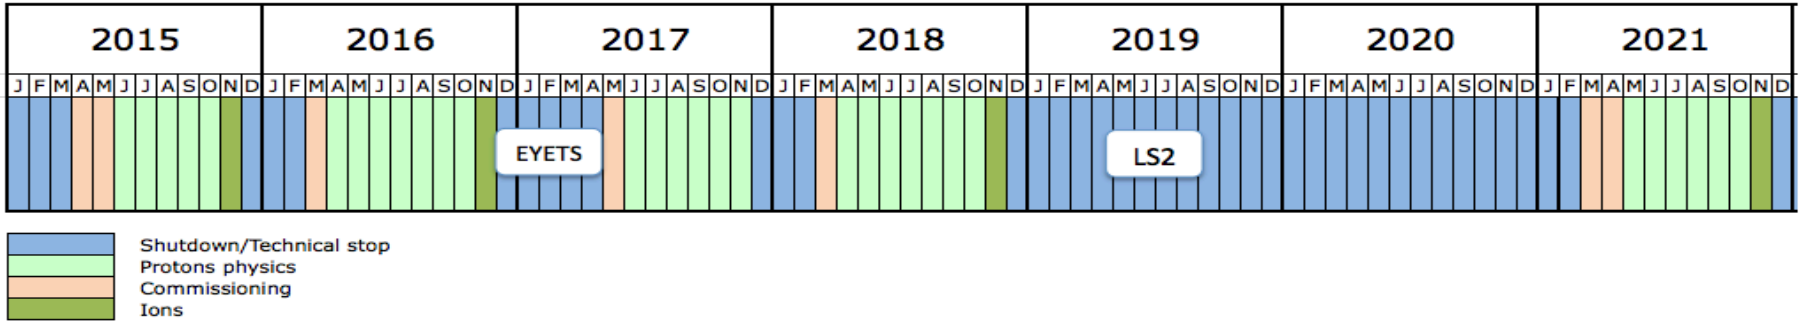

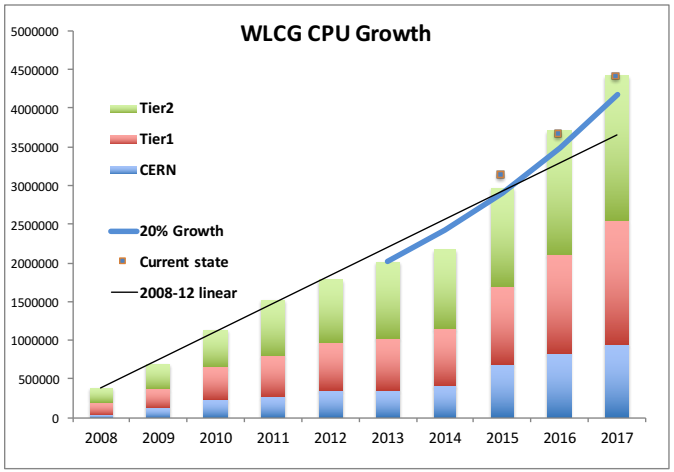

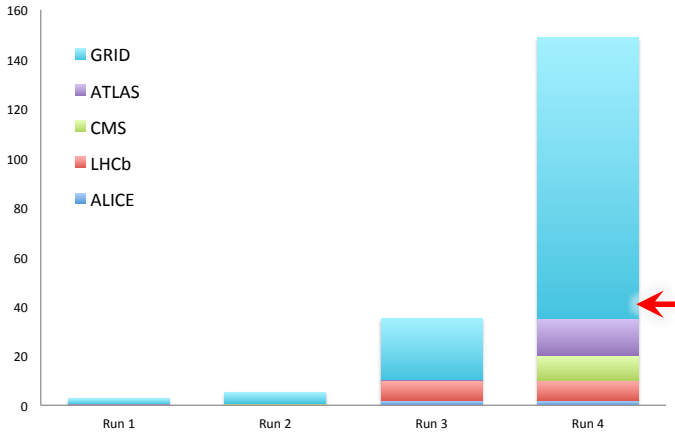

Compute: Growth > x50

What we think is affordable unless we do something differently

Batch HTC challenges and plans

- Scaling the service for Run3 and beyond
- Understanding and improving overall CPU efficiency (currently 60%)
- Moving 1300 users to Condor (~end of Run2) and retiring LSF
- Expanding into the external cloud / hosting
  - Minimise the ops cost by using the same tools as we use here; run as a single tier-0
  - Work with experiments to push their workloads there
    - Simulation workloads proven, data-intensive workloads are now the challenge

Compute Workloads

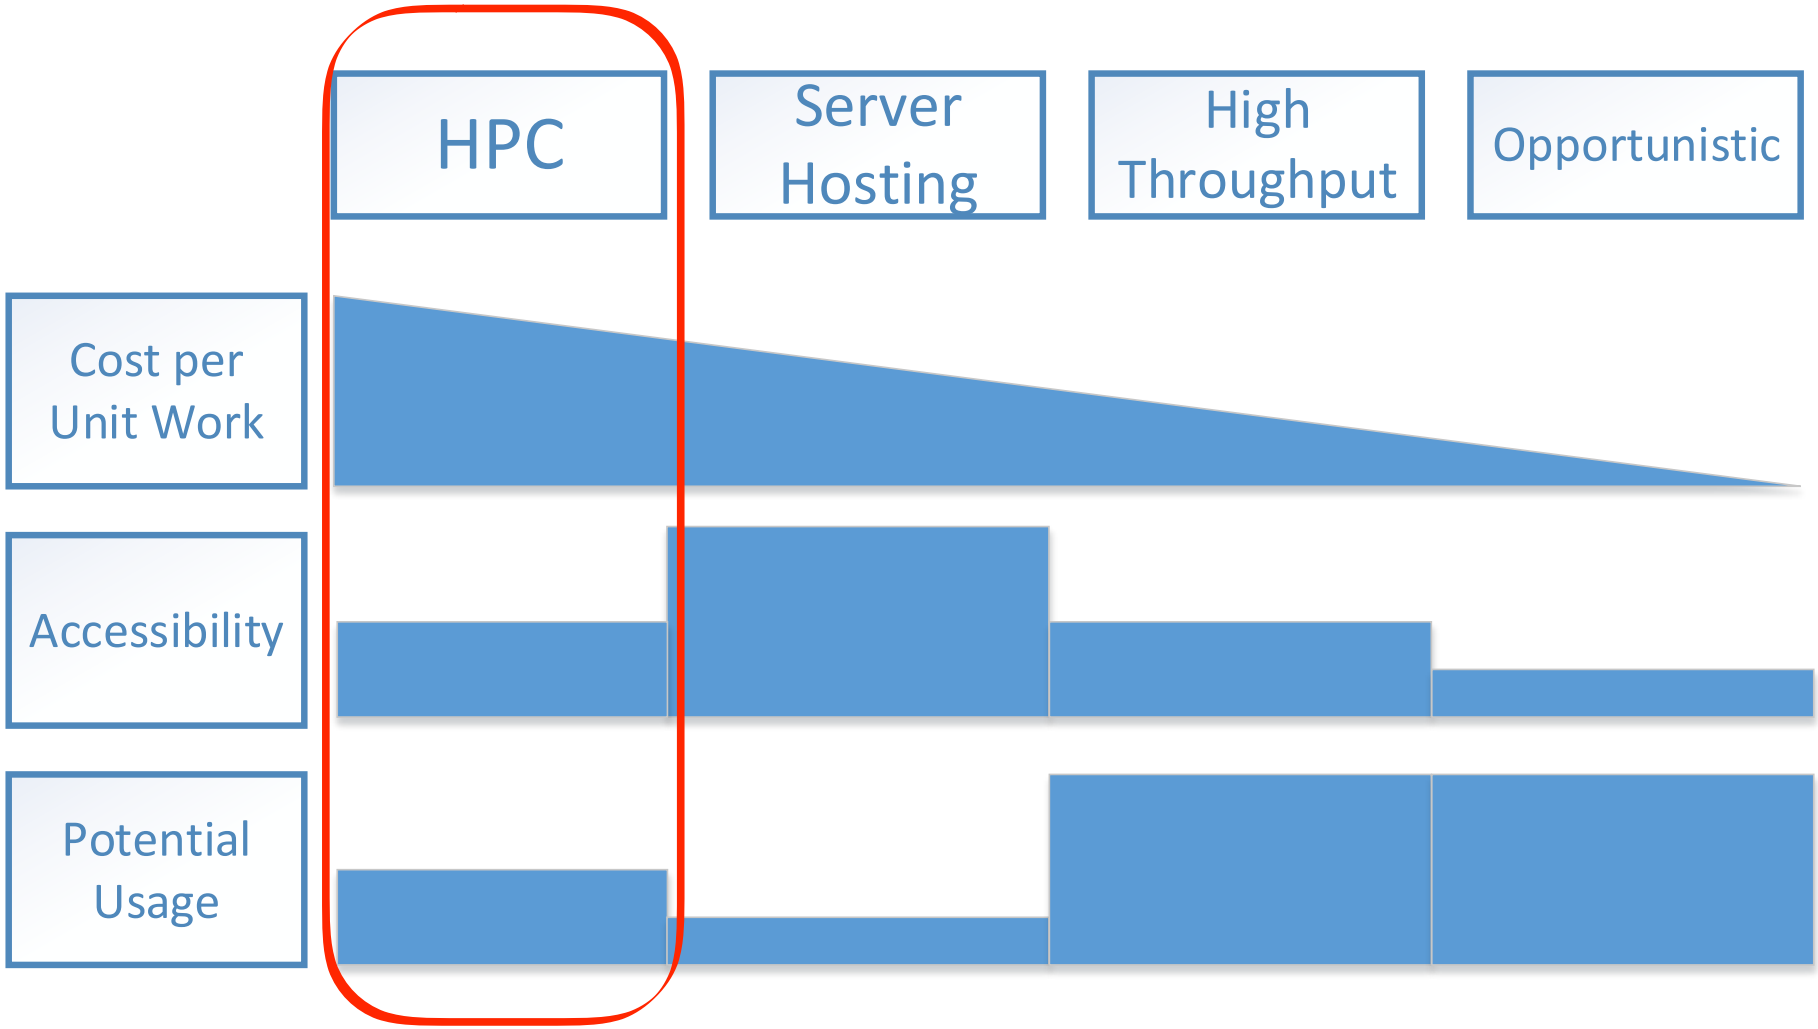

High-performance Compute

- Linux, scheduled by the batch service, but applications that don't fit the standard HTC pattern
- MPI, shared memory across nodes, fast low-latency interconnects needed
- A single job typically runs on multiple nodes for a long-time
  - Lattice QCD Theory simulations
  - Beam / plasma cell simulations
  - Fluid dynamics applications

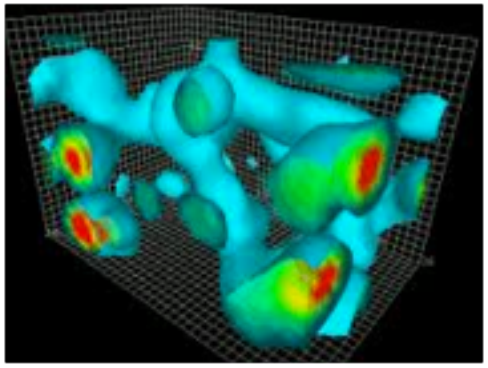

High-performance Compute

- Current facilities:
  - ..are all a bit dusty
  - Somewhat old servers with lowish-latency interconnects (10GB ethernet with RDMA)
  - Really old Theory infiniband cluster, due to be retired

HPC challenges and plans

- Take a step back and build a new general purpose HPC facility as part of batch service
  - New Theory cluster ordered; arriving in spring
  - Possible expansion for other HPC applications depending on funding

- Challenges:
  - Special hardware (interconnects)
  - Limited MPI and HPC competence - work with other teams to learn

Compute Workloads

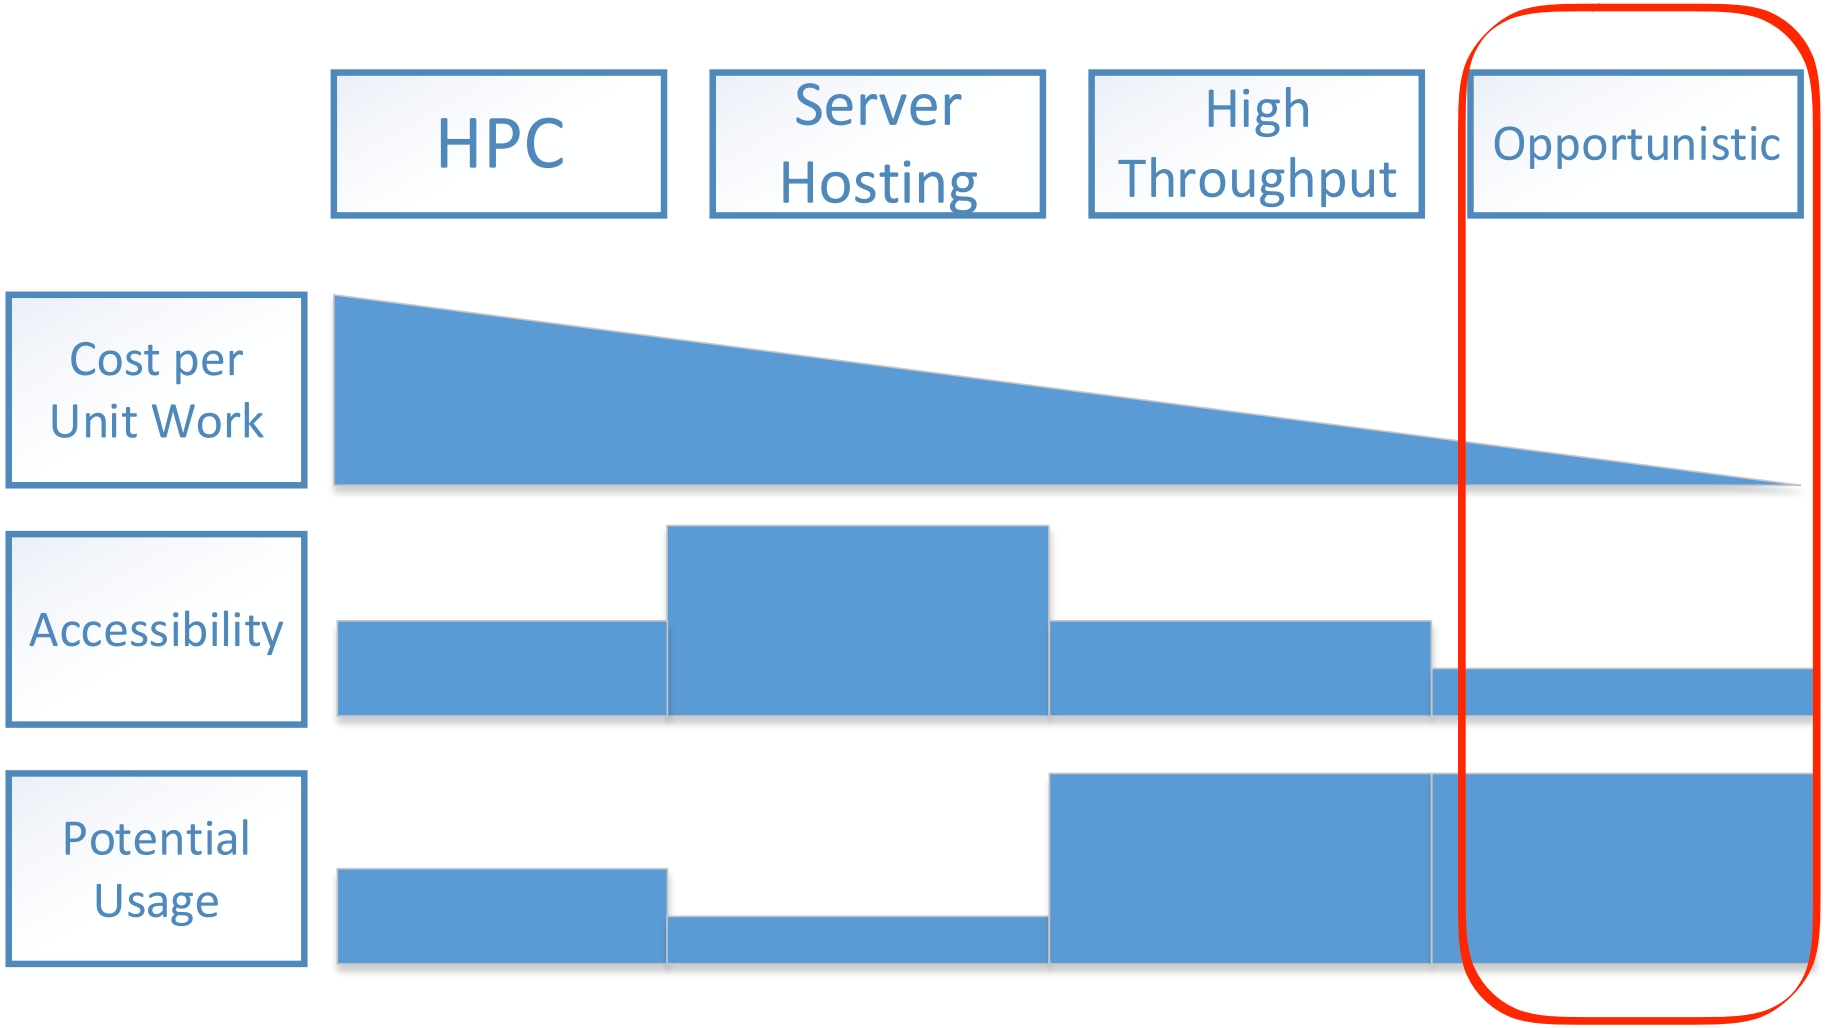

Opportunistic compute

- Aim is to be able to make use of free resources, quickly and efficiently as they come up
  - LHC/ATLAS/CMS/…@HOME volunteer
  - Supercomputers slots
  - Tests / donations / cheap backfill from large cloud providers

- Service Level: Not guaranteed, transient, low network bandwidth and volatile
- Working with experiments to find workloads which suit this environment (e.g simulation)

Opportunistic compute

- Significant resources avail. O(100k) Sixtrack beam sims
- ATLAS@Home (10k) - T1 site or 2.6MEUR on cloud

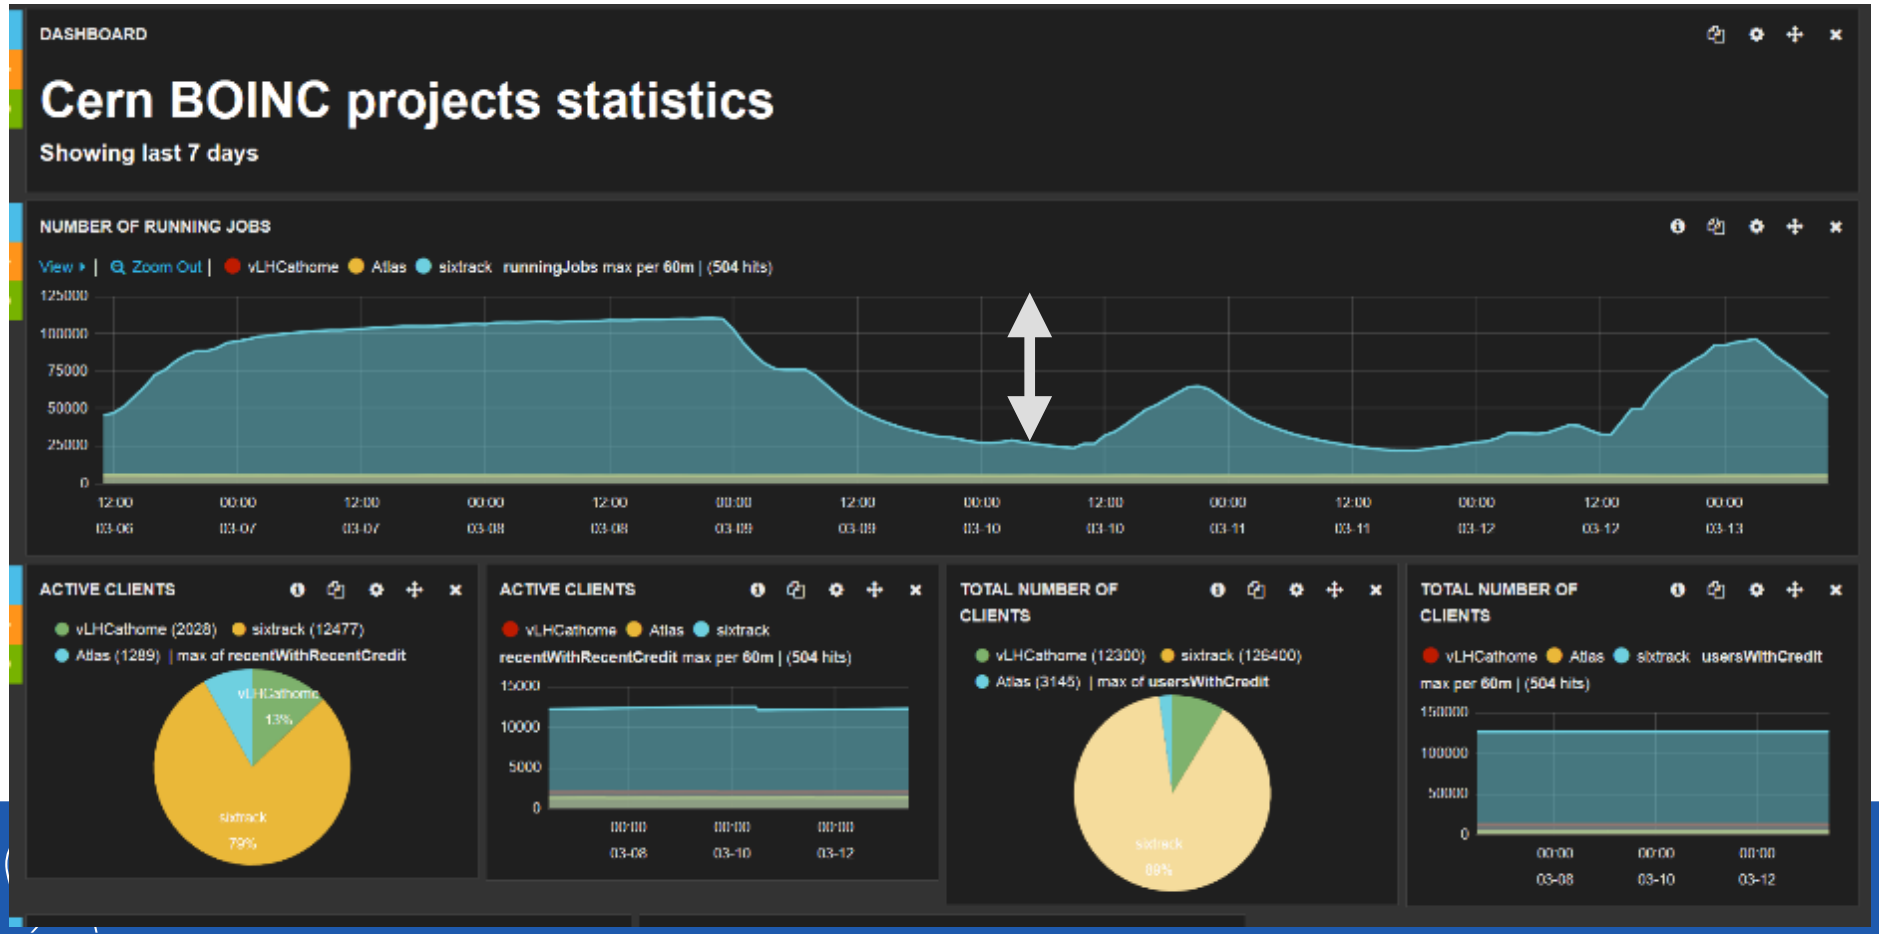

Opportunistic challenges and plans

- Resources free but ops cost is not
- Reduce ops overhead - leverage same technologies as we're using to expand into the cloud (e.g. Condor)
- Ramping up capacity and ensure the BOINC infrastructure can scale

- Longer term - can also we expand opportunistic onto Cloud backfill (internal and external)?
- Continue to work with experiments to find workload that's suitable for this use-case

WLCG support services

- Critical services used world-wide by all running jobs on O(200) sites

- *MyProxy*: allows trusted compute services to renew your x509 user proxy to interact with the Grid
- *VOMS*: decorates your x509 user proxy with trusted group / role assertions e.g. “Bob is the production manager of ATLAS”
- *BDII*: An LDAP information system for WLCG service discovery and critical site information

- Stability and some evolution within WLCG

Accounting

- Cross-section activity
- Significant money is spent on LHC computing, so we need to show how that money is spent
  - to CERN internally;
  - to WLCG as a whole, via our experiment pledges
- Challenges: get, better, clearer numbers
- Understand and maintain those numbers as we move into the cloud / hosted environments
- Be able to verify cloud provider bills

Automation and tooling

- Investigating and investing in automation tools to bring down the overall ops cost

- Challenge:
  - Work out the single set of tools we need to run our internal and externally hosted tier-0 with minimal ops cost

Summary

- We're expanding CERN's compute capacity across HTC, HPC and opportunistic

- We're expanding the ways in which we deliver the compute capacity
  - external cloud / hosting / opportunistic
  - aim to consolidate on a single set of ops tools

- Making as much use as we can of industry standard tooling to lower the ops cost and improve the overall efficiency