Lightweight sites in WLCG

Maarten Litmaath, CERN

Rationale

One of the goals of WLCG Operations Coordination activities is to help simplify what the majority of the sites, i.e. the smaller ones, need to do to be able to contribute resources in a useful manner, i.e. with large benefits compared to efforts invested.

Classic grid sites may profit from simpler mechanisms to deploy and manage services. Moreover, we may be able to get rid of some service types in the end.

New and existing sites may rather want to go into one of the cloud directions that we will collect and document.

There may be different options also depending on the experiment(s) that the site supports.

There is no one-size-fits-all solution. We will rather have a matrix of possible approaches, allowing any site to check which ones could work in its situation, and then pick the best.

Boundaries – storage & data access

- Under the aegis of the WLCG Data Steering Group
  - Data federations
  - Multi-site storage
  - Caches
  - Diskless sites
  - Big data technologies

Potential for paradigm changes

- A number of these areas will be covered by other presentations in today's session

- Further information
  - WLCG workshop June 2017 Data session
  - May 2017 GDB

- Here we focus on computing resources instead

Boundaries – EGI and OSG

- In OSG every WLCG site mainly supports just a single LHC experiment
- The sites are managed in close collaboration with the US project in each experiment
  - US-ATLAS, US-CMS, US-ALICE
- Both US-ATLAS and US-CMS have already been working on lighter ways to provision their resources
  - Ubiquitous Cyberinfrastructure, Virtual Clusters
  - Tier-3 in a box, Pacific Research Platform
- In EGI the situation is a lot more complex
  - Multi-experiment sites, many countries/cultures/projects/…, more MW diversity, experiments have less influence, …
- Here we should focus on the EGI sites then
  - While learning from the OSG sites

T2 vs. T3 sites

- T3 sites have not signed the WLCG MoU
  - Typically dedicated to a single experiment → can take advantage of shortcuts
- T2 sites have rules that apply
  - Availability / Reliability targets
  - Accounting into EGI / OSG / WLCG repository
  - EGI: presence in the info system for the Ops VO
  - Security regulations
    - Mandatory OS and MW updates and upgrades
    - Isolation
    - Traceability
    - Security tests and challenges
- Evolution is possible
  - Some rules could be adjusted
  - The infrastructure machinery can evolve

How to enable computing

- Services that currently are or may be needed to enable computing at a grid site:
  - Computing Element
  - Batch system
  - Cloud setups
  - Authorization system
  - Info system
  - Accounting
  - CVMFS Squid
  - Monitoring

How to enable computing

- Services that currently are or may be needed to enable computing at a grid site:
  - Computing Element
  - Batch system
  - Cloud setups
  - Authorization system
  - Info system
  - Accounting
  - CVMFS Squid
  - Monitoring

- Reduce the catalog of required services, where possible?
- Replace classic, complex portfolio with alternative, more widespread technologies?
- Simplify deployment, maintenance and operation of what needs to remain?

Less diversity would help

- Batch systems on the rise
  - HTCondor
  - Slurm
- CE implementations on the rise
  - HTCondor
  - ARC
- Configuration systems on the rise
  - Puppet
  - Ansible

Tap into popular technologies?

- Cloud systems on the rise
  - OpenStack
  - …
- Container systems on the rise
  - Docker
  - Singularity
  - Kubernetes
  - Mesos
  - OpenStack Magnum
  - OpenShift
  - ...

Many winners for now?

Lightweight sites – classic view

- How to provide resources with less effort?
  - Keep things basically the same, but easier

- Site responses to a questionnaire show the potential benefits of shared repositories
  - OpenStack images
    - Pre-built services, pre-configured where possible
  - Docker containers
    - Ditto
  - Puppet modules
    - For site-specific configuration

Lightweight sites – alternative view

- CE + batch system *not strictly needed*
- Cloud VMs or containers could be sufficient
- They can be managed e.g. with Vac or Vcycle
  - Several GridPP sites are doing that already
    - All 4 experiments are covered
  - The resources are properly accounted
- They can directly receive work from an experiment's central task queue
- Or they can rather join a regional or global *HTCondor pool* to which an experiment submits work
  - Proof of concept used by GridPP sites for ALICE
  - Cf. the CMS global GlideinWMS pool → scalable to O(100k)

Distributed site operations model

- A site needs to provide resources at an agreed QoS level
- HW needs to be administered by the site
- Other admin operations could be done by a remote, possibly distributed team of experts
- Site resources within a region could be integrated into a regional cloud
  - Example: **JINR cloud** extending to partner sites
- Or they could be integrated by a regional *virtual* HTCondor batch system
  - VMs/containers of willing sites may join the pool directly
  - CEs and batch systems of other sites can be addressed through Condor-G
  - The virtual site exposes an HTCondor CE interface through which customers submit jobs *to the region*
  - HTCondor then routes the jobs according to fair-share etc.

Volunteer computing …

- The LHC@home project coordinates volunteer computing activities across the experiments
- ATLAS have benefited from 1-2% extra resources for simulation workloads
- See this recent talk by David Cameron
- It could become a way for a *computing-only lightweight site* to provide its resources
  - The central infrastructure can scale at least for simulation jobs
  - The resources can be properly accounted in APEL

… and lightweight sites

- Real sites can be trusted
  - No need for volunteer CA or data bridge
  - A separate, easier infrastructure would be set up

- BOINC can even coexist with a batch system on the same WN
  - Successfully demonstrated at IHEP, Beijing

- Also here HTCondor is used under the hood
  - Standard for experiments and service managers

Recent ATLAS T3 stats

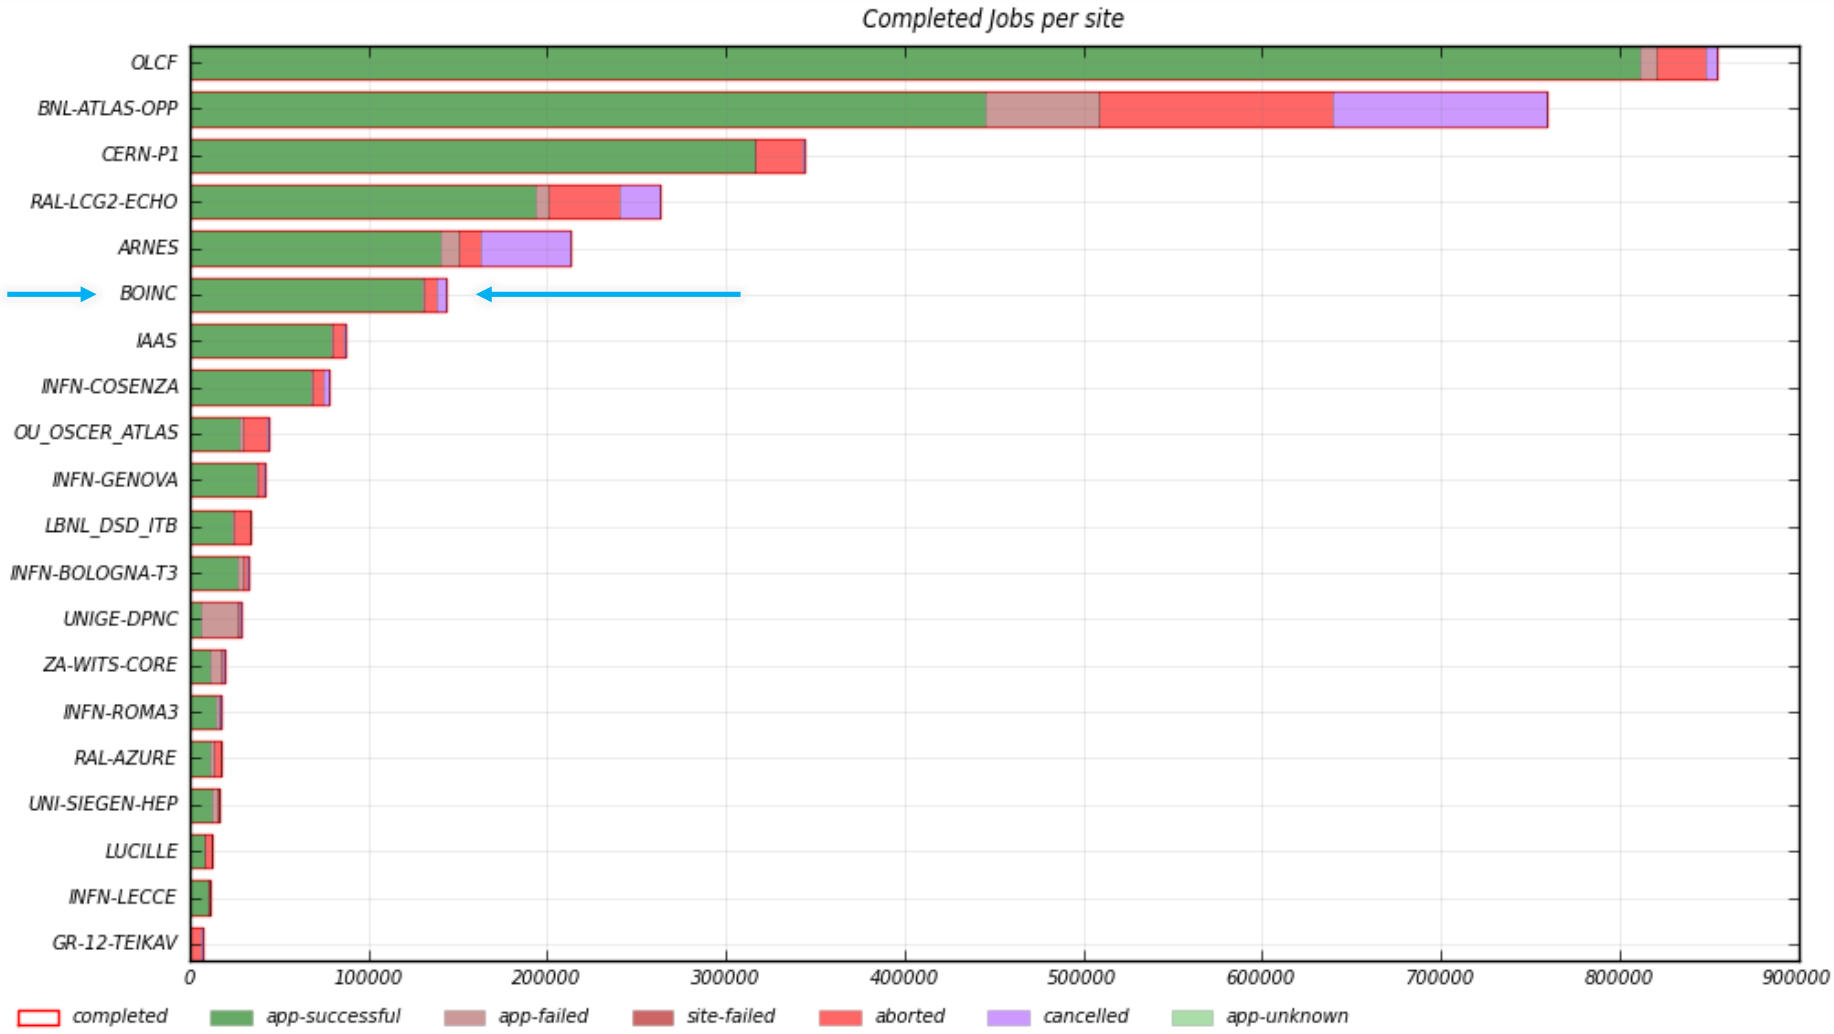

Recent CMS T3 stats

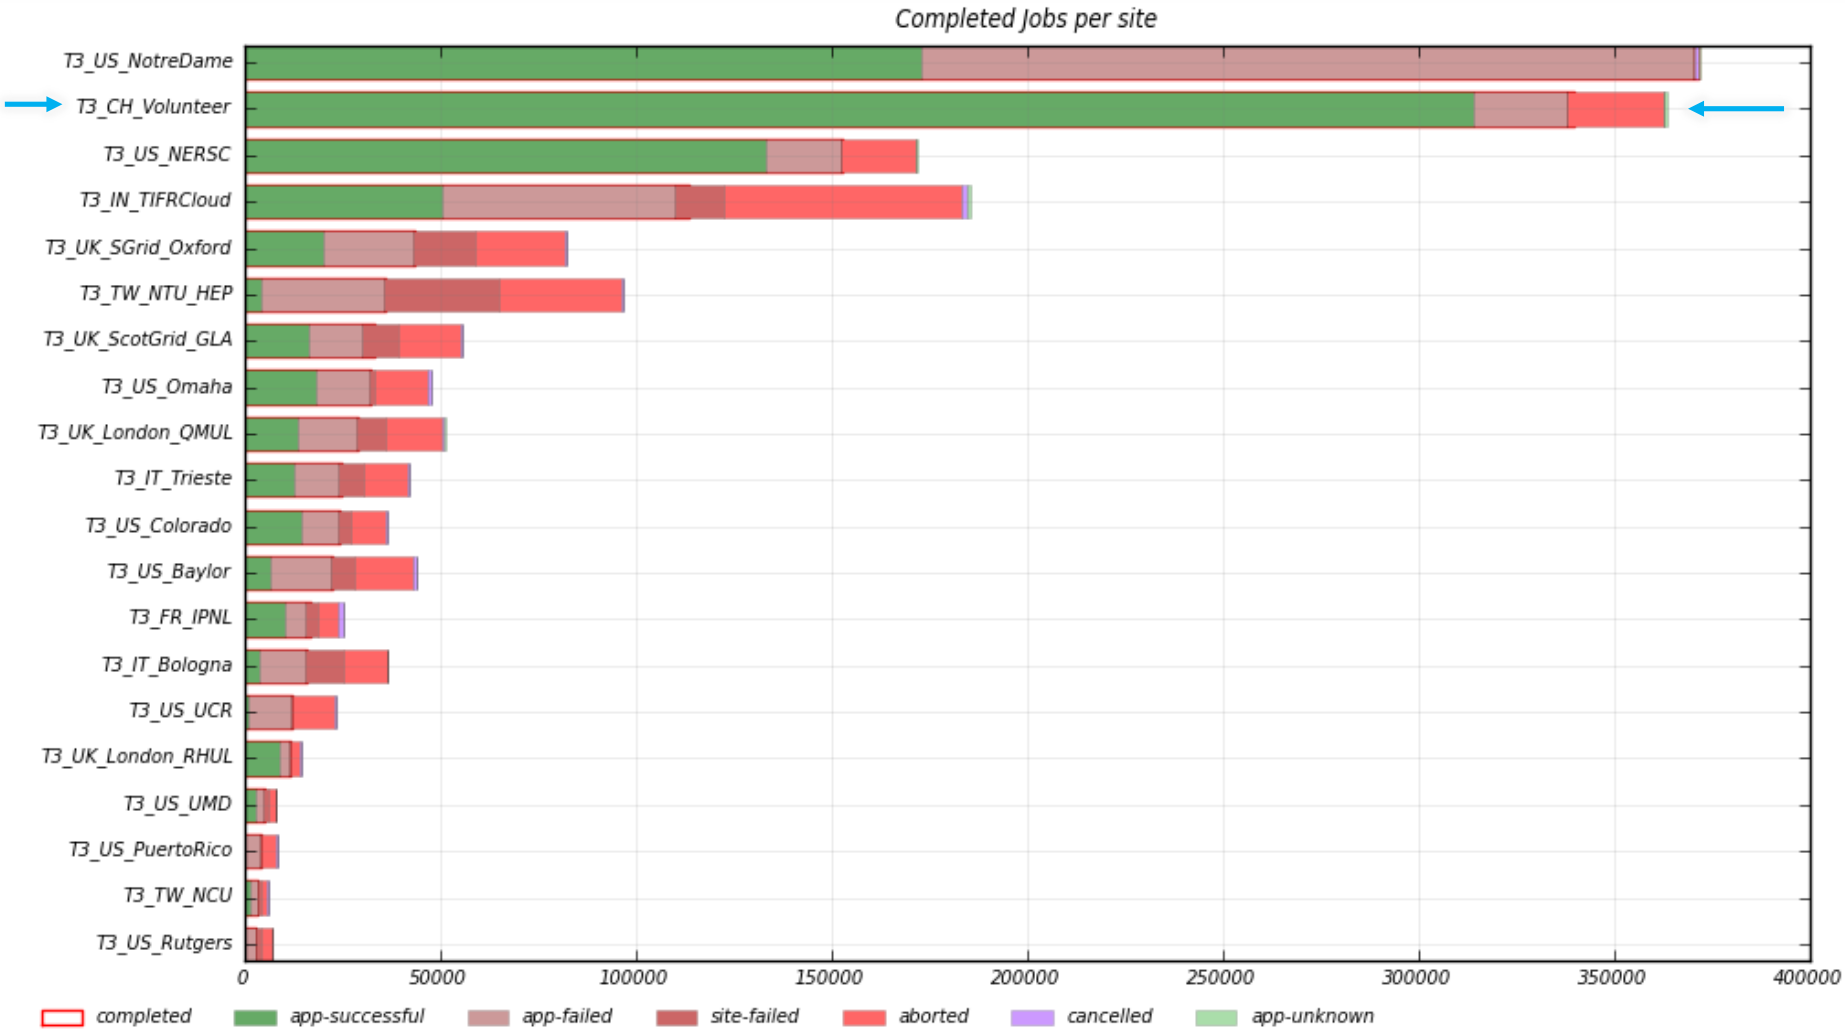

Volunteer potential

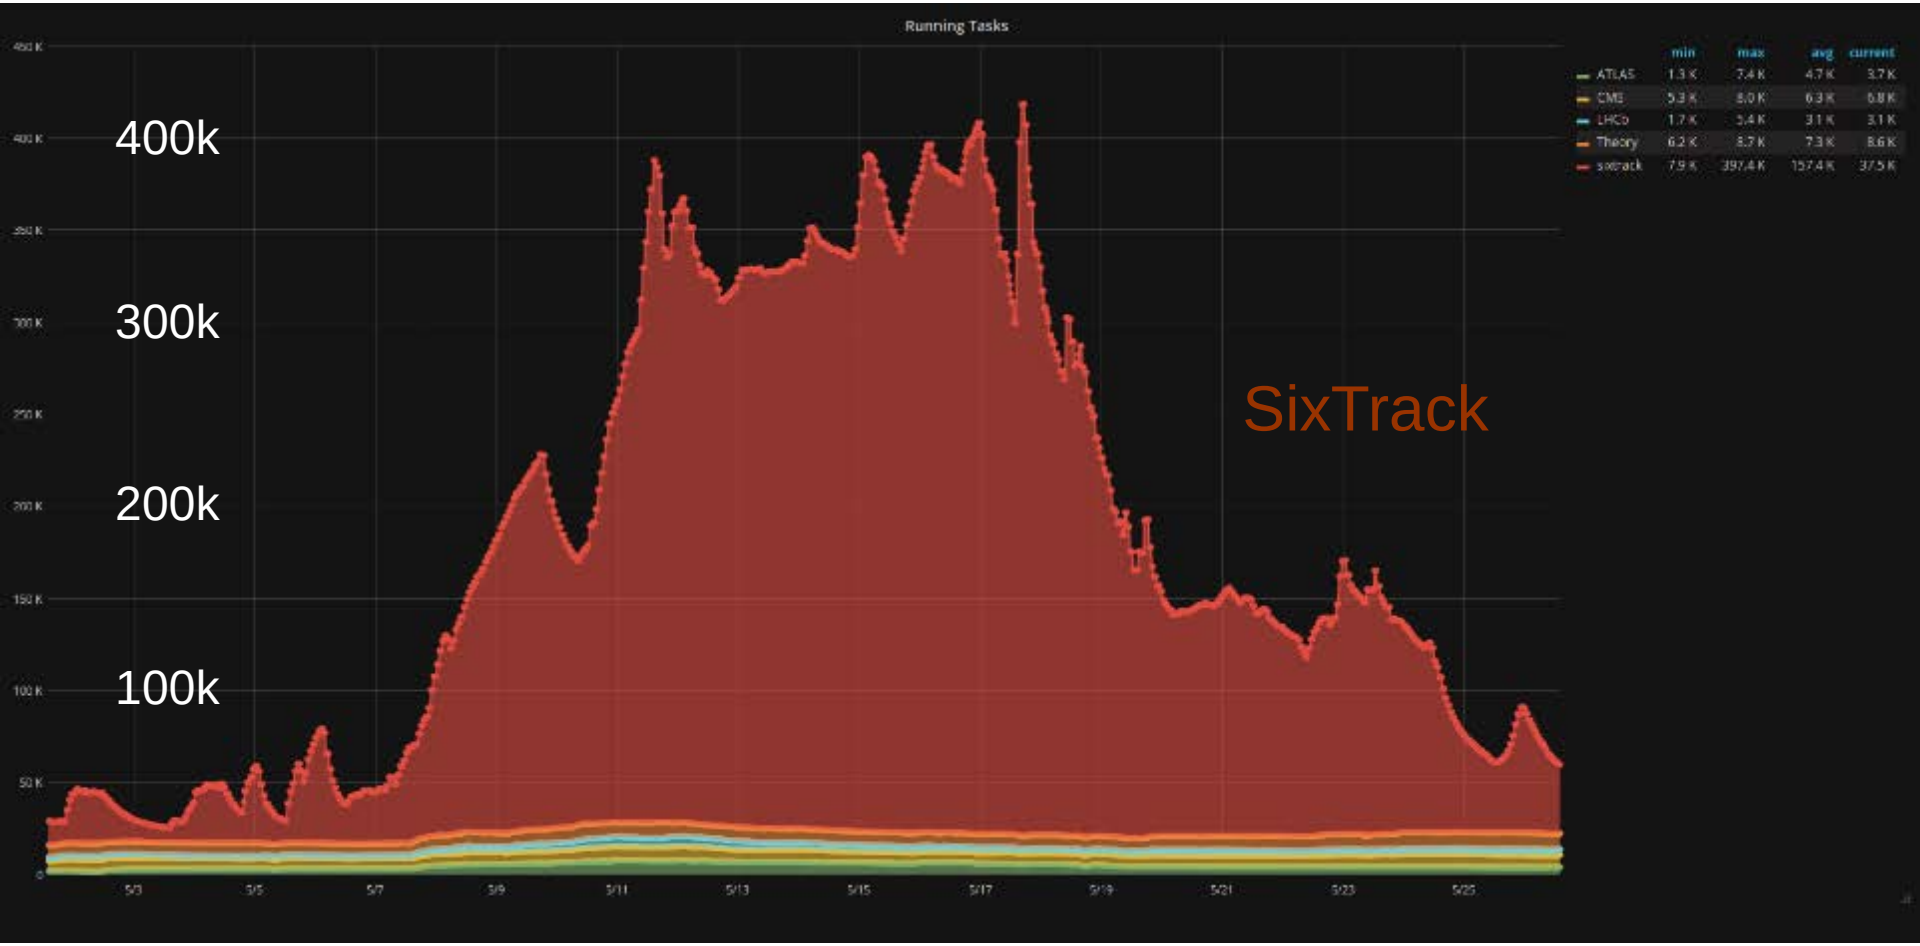

8th BOINC Pentathlon 2017

Computing resource SLAs

- The resources themselves can also be “lightweight”
- Please see this recent talk by Gavin McCance
- Extra computing resources could be made available at a lower QoS than usual
  - Disk server CPU cycles, spot market, HPC backfill, intervention draining, …
  - Jobs might e.g. get lower IOPS and would typically be pre-emptible
  - Machine-Job Features (MJF) functionality can help smooth the use
- They would have an SLA *between* those of *standard* and *volunteer* resources → a *mid-SLA*

Lightweight operations

- We would like to have sites which can run almost “by themselves”
  - With minimal oversight and operational efforts from people at the site
- Could we make use of Machine Learning algorithms to improve our monitoring?
  - Automatic classification and filtering of log messages
  - Definition of metrics that characterize the state of operations
  - Early identification of remarkable trends

Conclusions and outlook

- Many small sites currently need to invest efforts that are not commensurate with their size nor available funding
- Multiple areas are being investigated to allow small sites to become more lightweight
- Sites are envisaged to be able to pick the best choice from a matrix of solutions
- WLCG thus may evolve toward increased flexibility and sustainability !