


Experience with dynamic resource provisioning of the CMS online cluster using a cloud overlay




Diego da Silva Gomes - CERN
diego@cern.ch

on behalf of the CMS Data Acquisition Group

CHEP 2018
23rd International Conference on Computing in High Energy and Nuclear Physics
9-13 July 2018, Sofia, Bulgaria


The CMS Online Cluster



The HLT cluster capacity:

- ~1500 nodes (prod + old)
- ~850k HEP-SPEC06
- ~37k cores (74k with hyper-threading)
- 160Gbps connectivity to CERN main facility

The idea: repurpose this computing resource for Grid processing when there is no data taking!

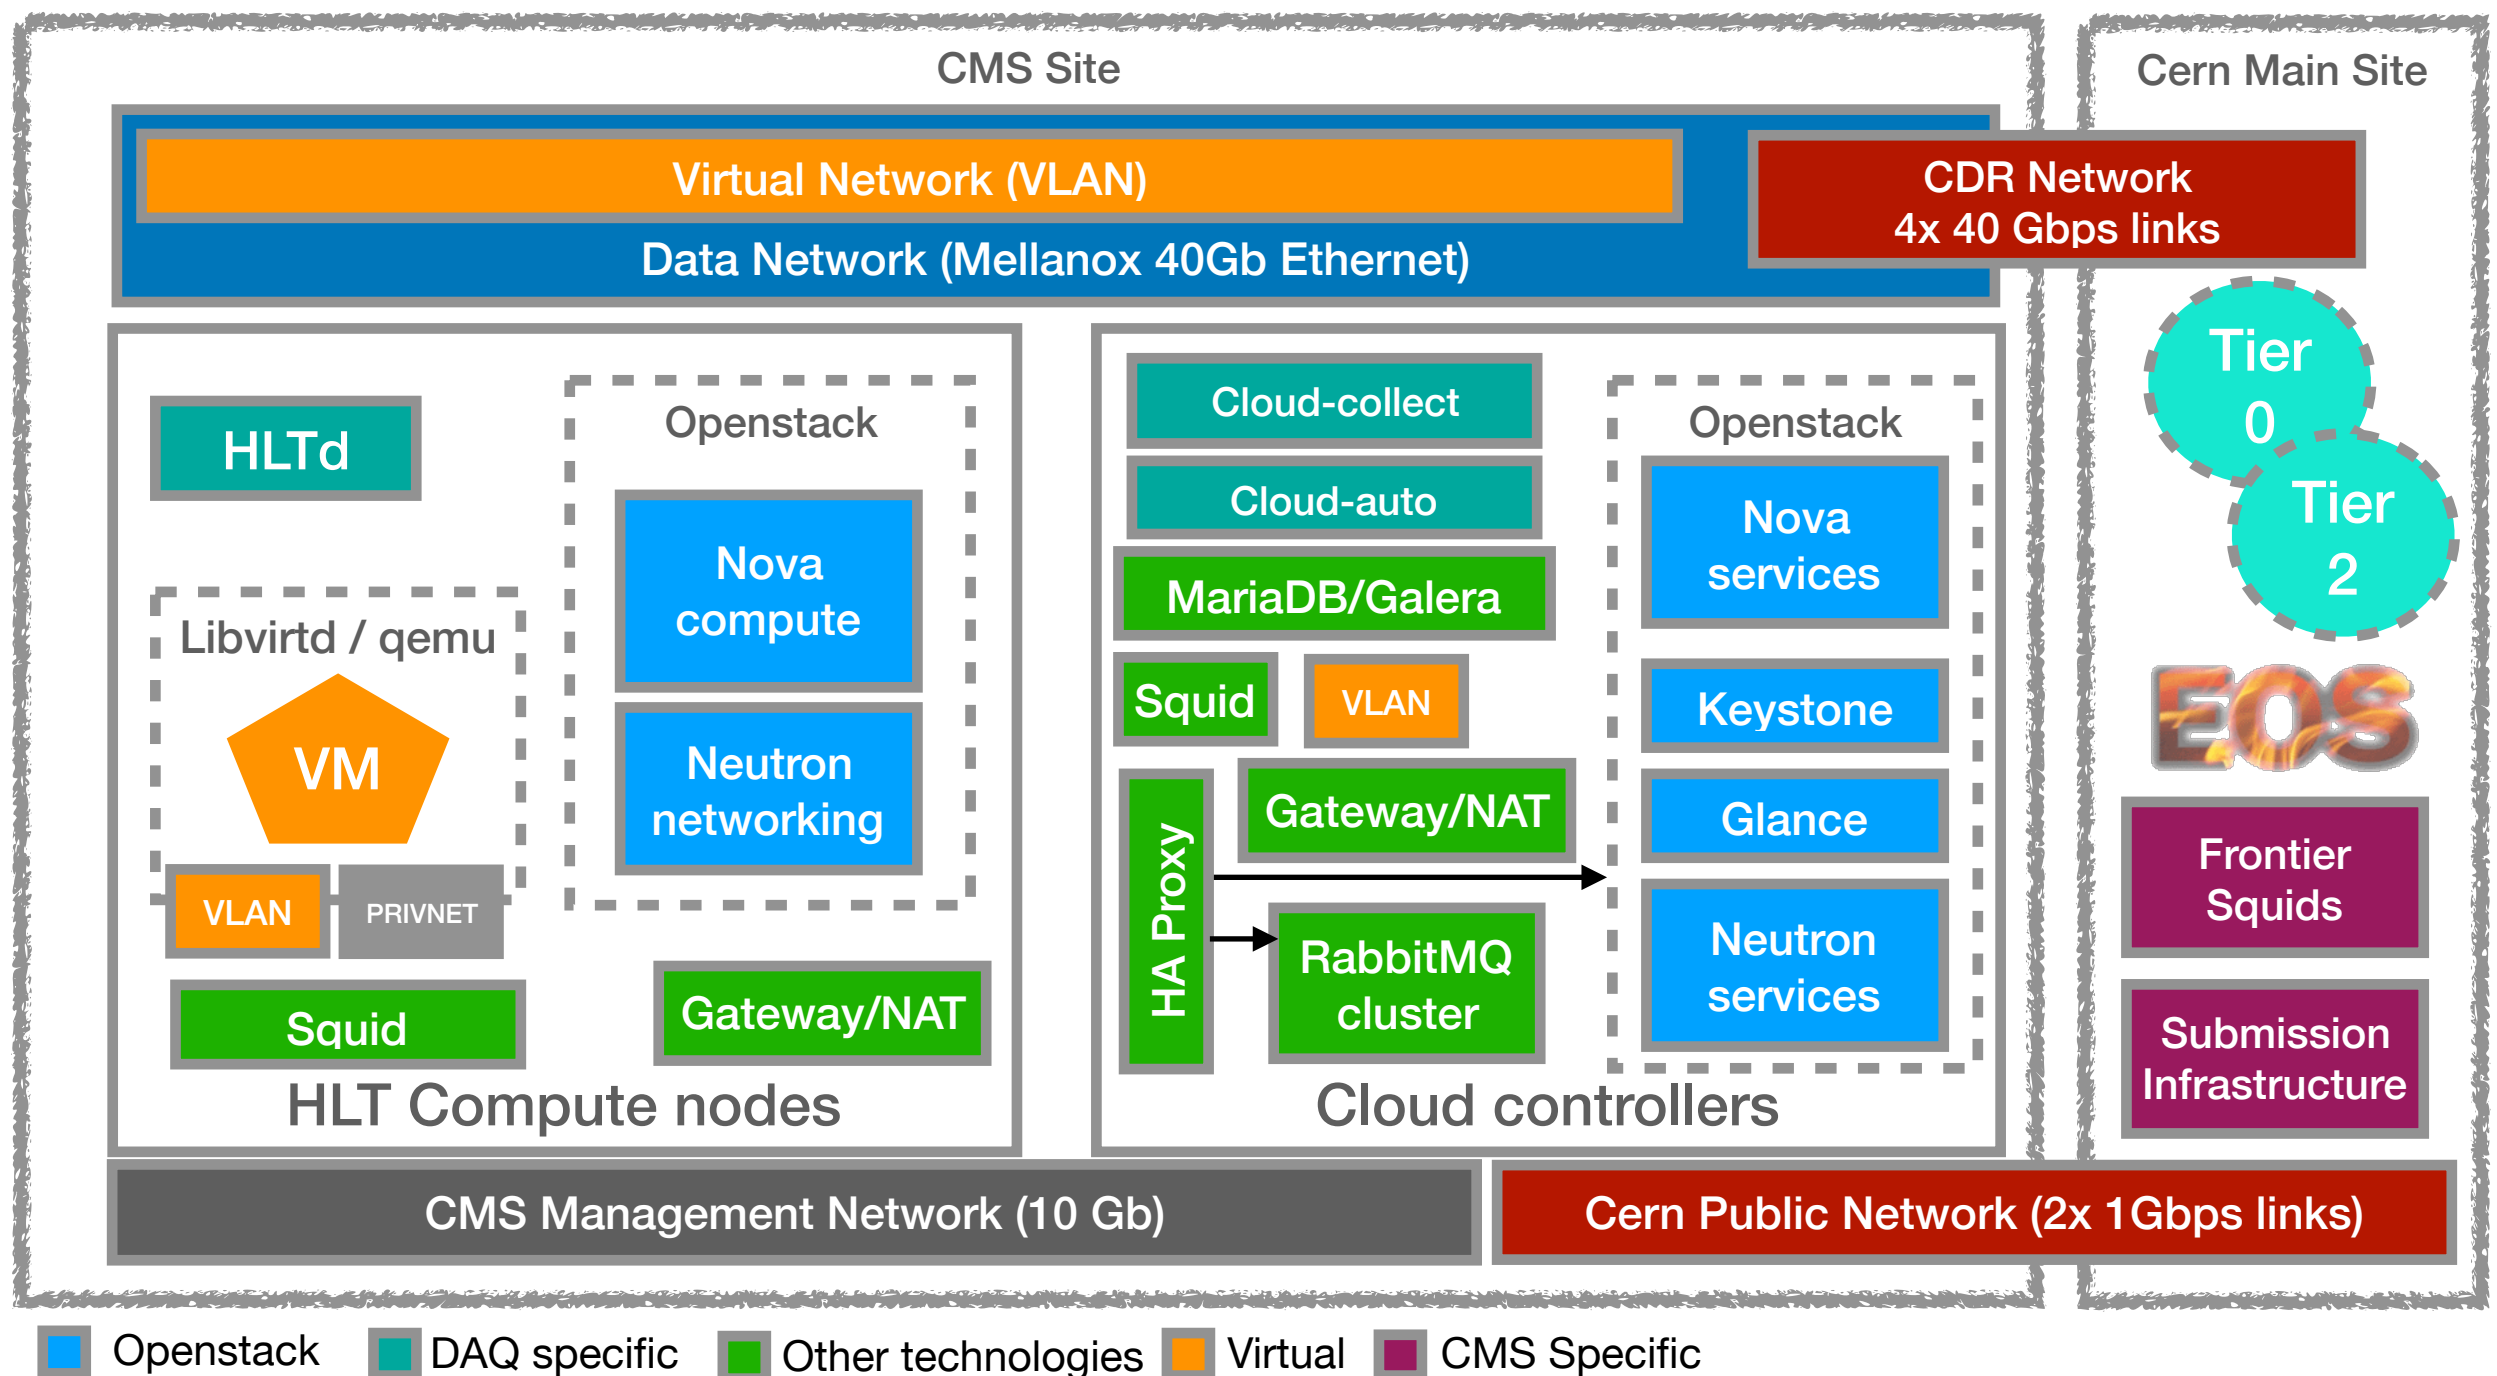


I.e. during periods of LHC Technical Stops, Machine Development weeks but also **during inter-fills** of physics

Major requirements

- **Isolation**: should not interfere/impact with Data Taking
- **Fast turnaround**: inter-fills as short as a couple of hours
- **Job resuming**: ability to pause grid jobs and resume them hours later

The Online Cloud overlay



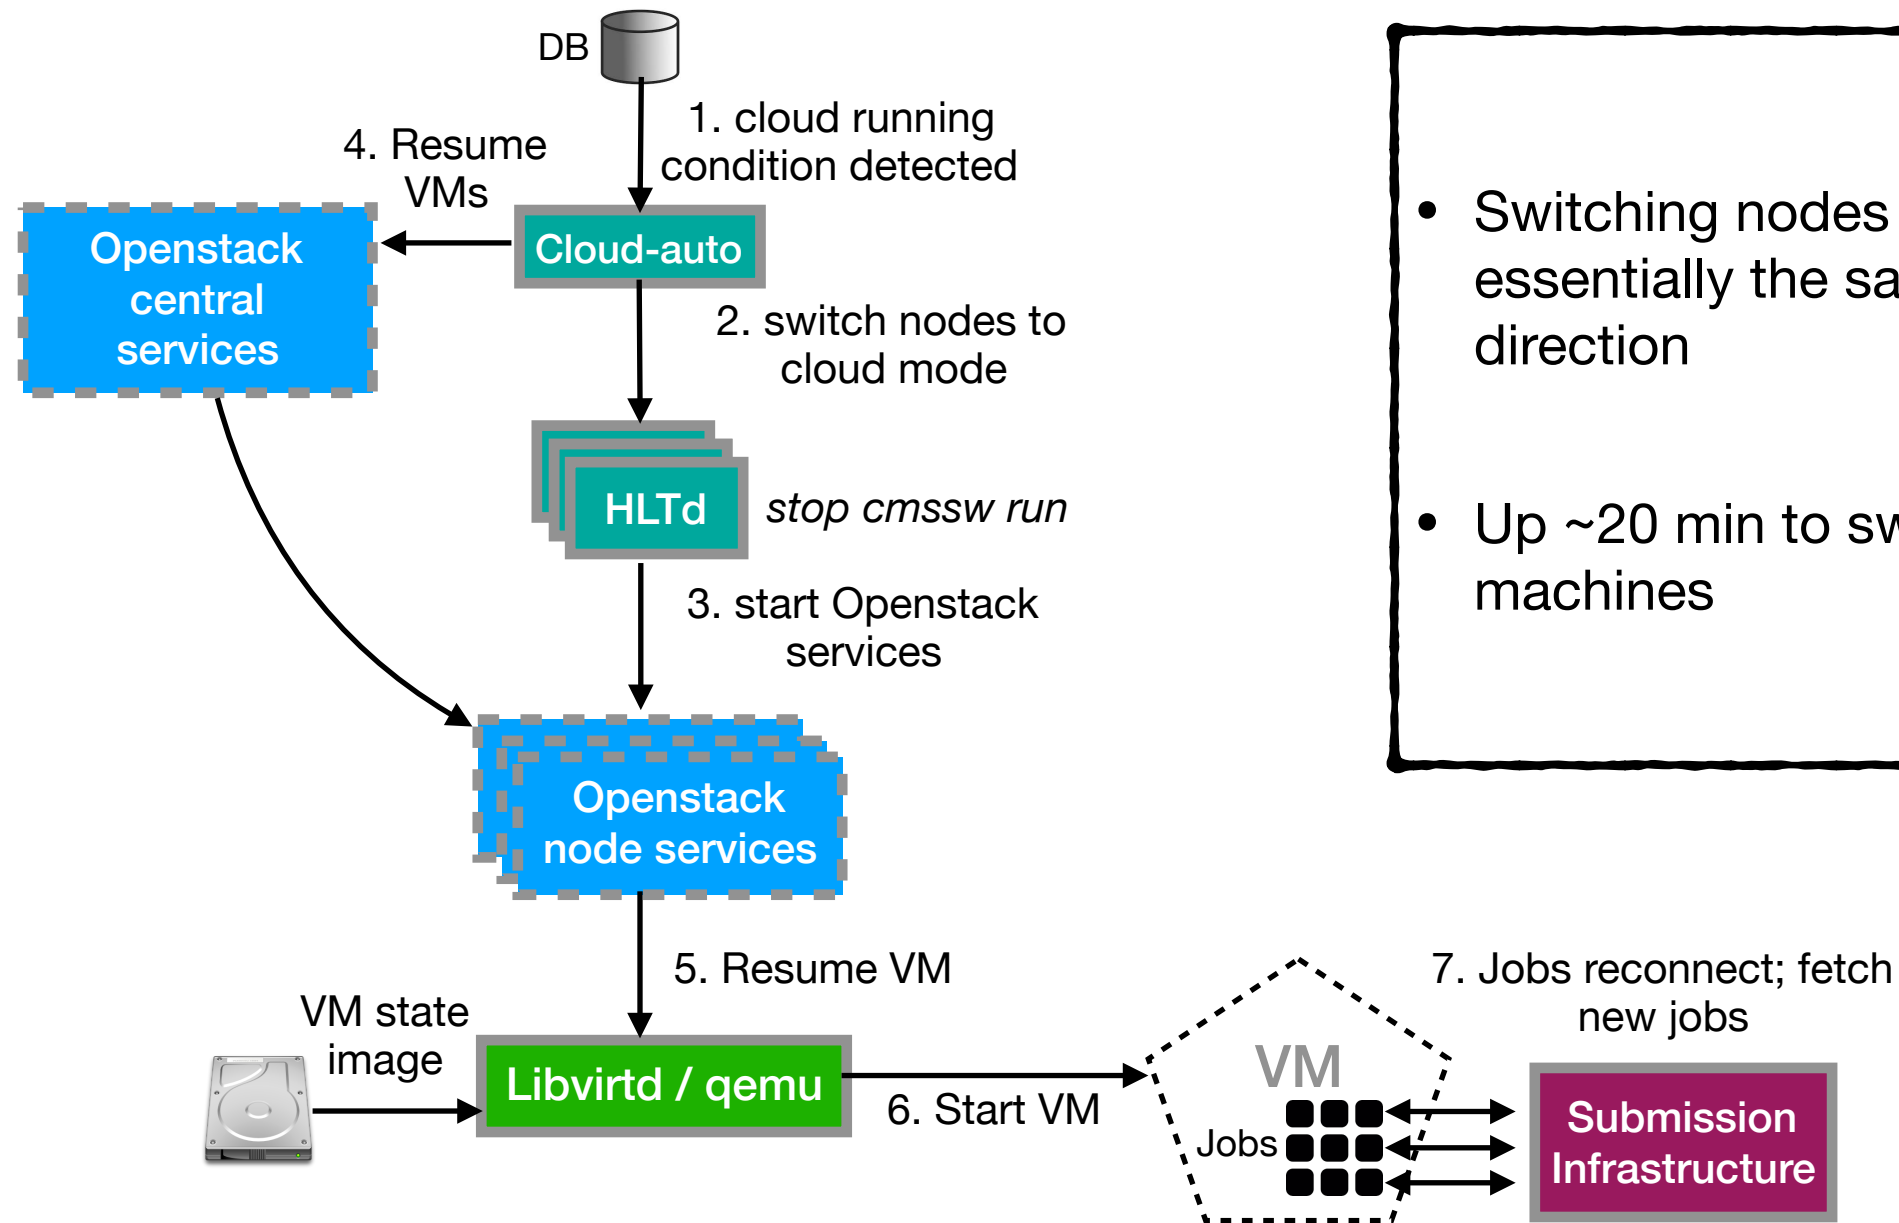
Openstack to manage Virtual Machines running Grid-enabled software

VLAN on top of the DAQ Data network to deal with the network segmentation and isolation

Cloud daemons monitor LHC/Beam and DAQ status and switches nodes to/from cloud accordingly

The Online Cloud overlay

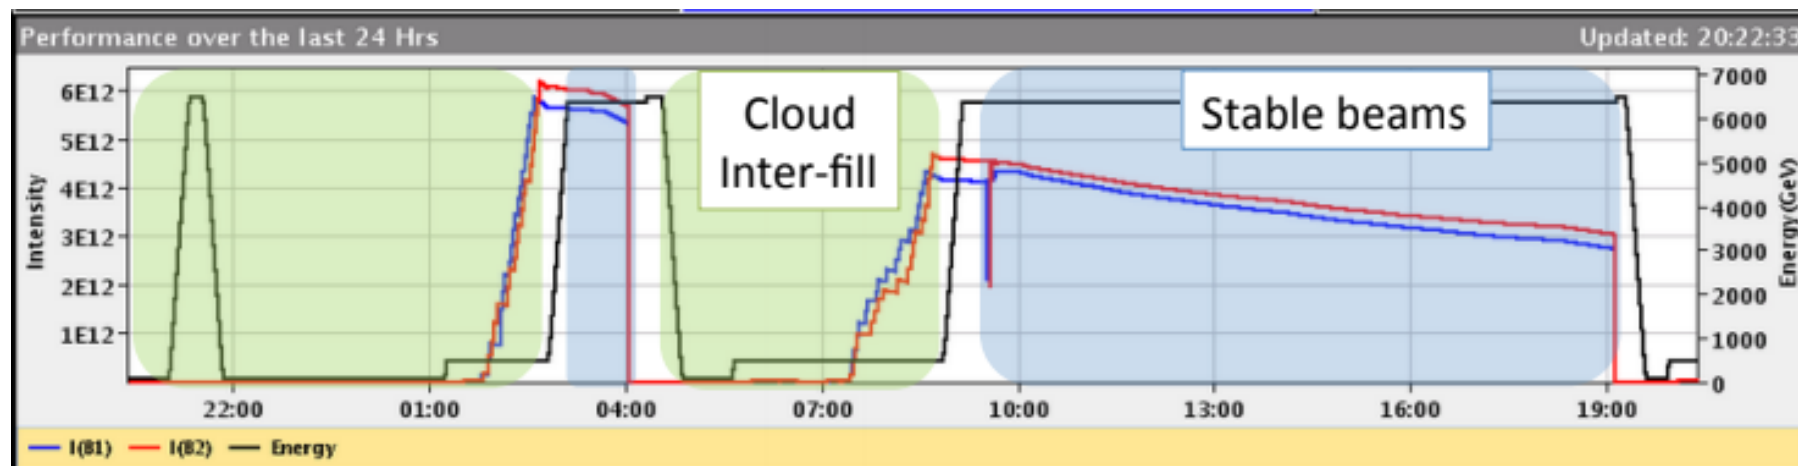
Switching HLT nodes from DAQ to cloud mode




- Switching nodes back to DAQ: essentially the same but in the reverse direction
- Up ~20 min to switch mode on all machines

Cloud Operation modes

- **Inter-fill:** follow LHC and Beam states and run the cloud during inter-fills and during LHC maintenance periods — *most common operation mode*



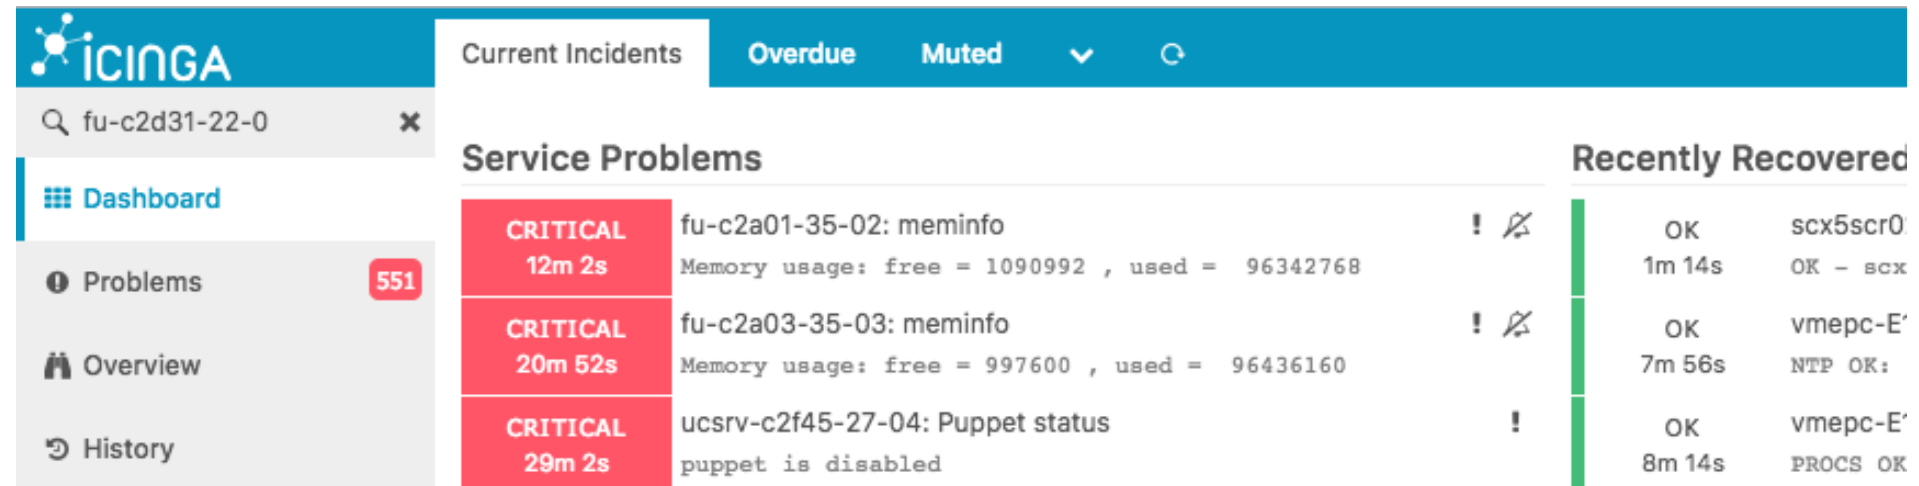
- **Fill:** follows beam states and progressively ramps up the cloud when DAQ conditions allow



- **Fill + Inter-fill**
- **On/off:** regardless of any condition

Monitoring

- **Icinga** for host and service checks

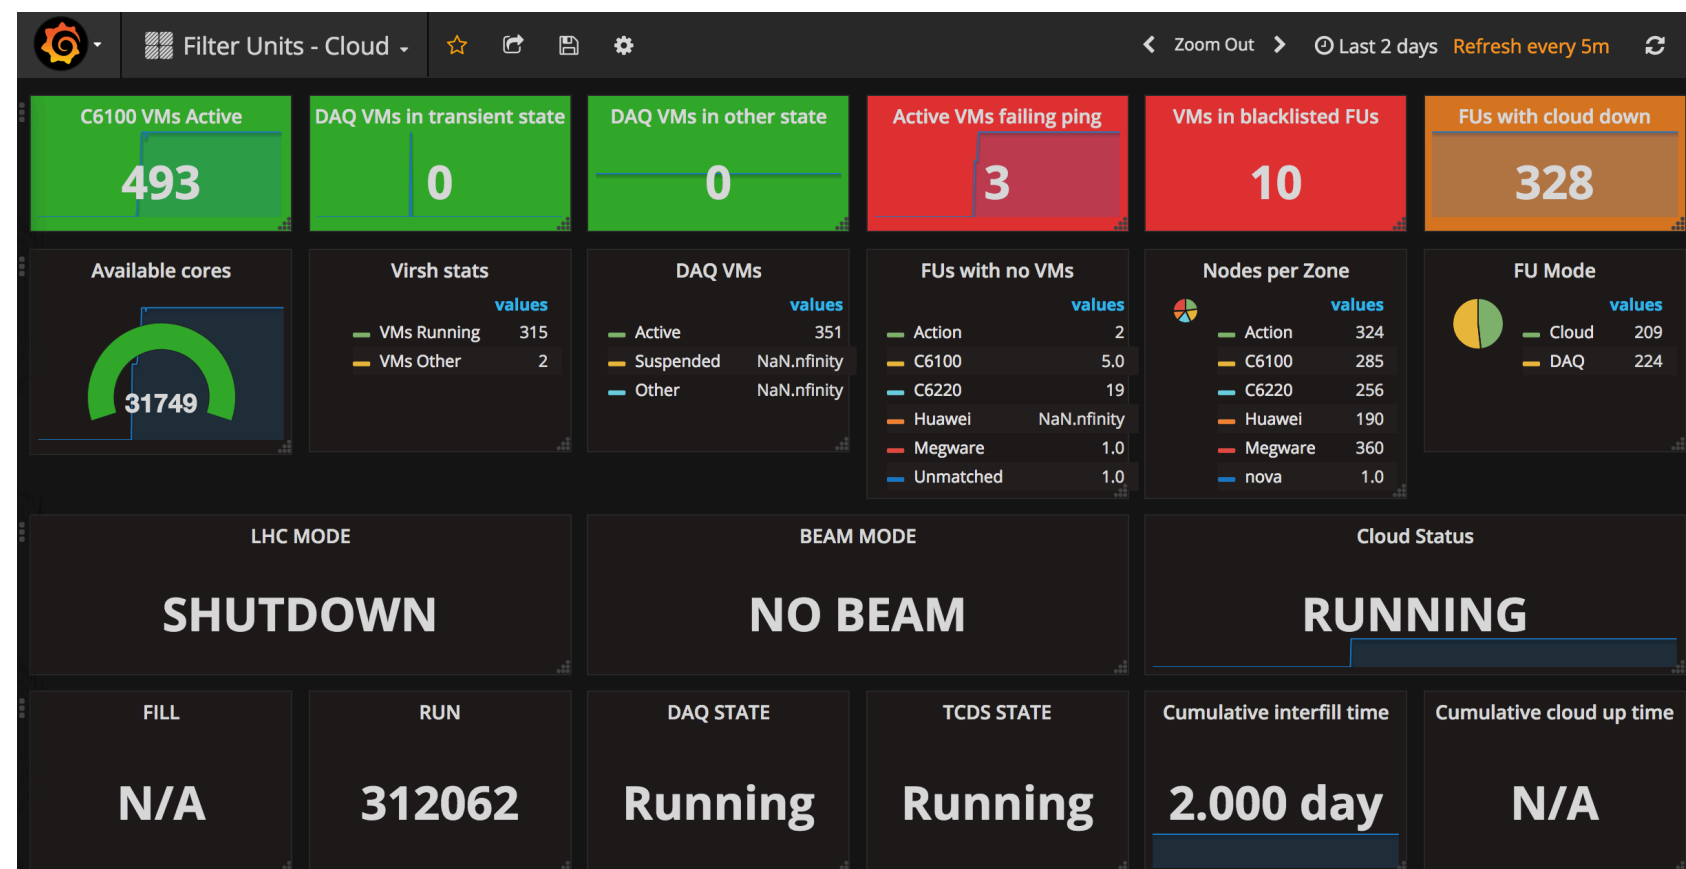


- **SMS** alerts to on call phone

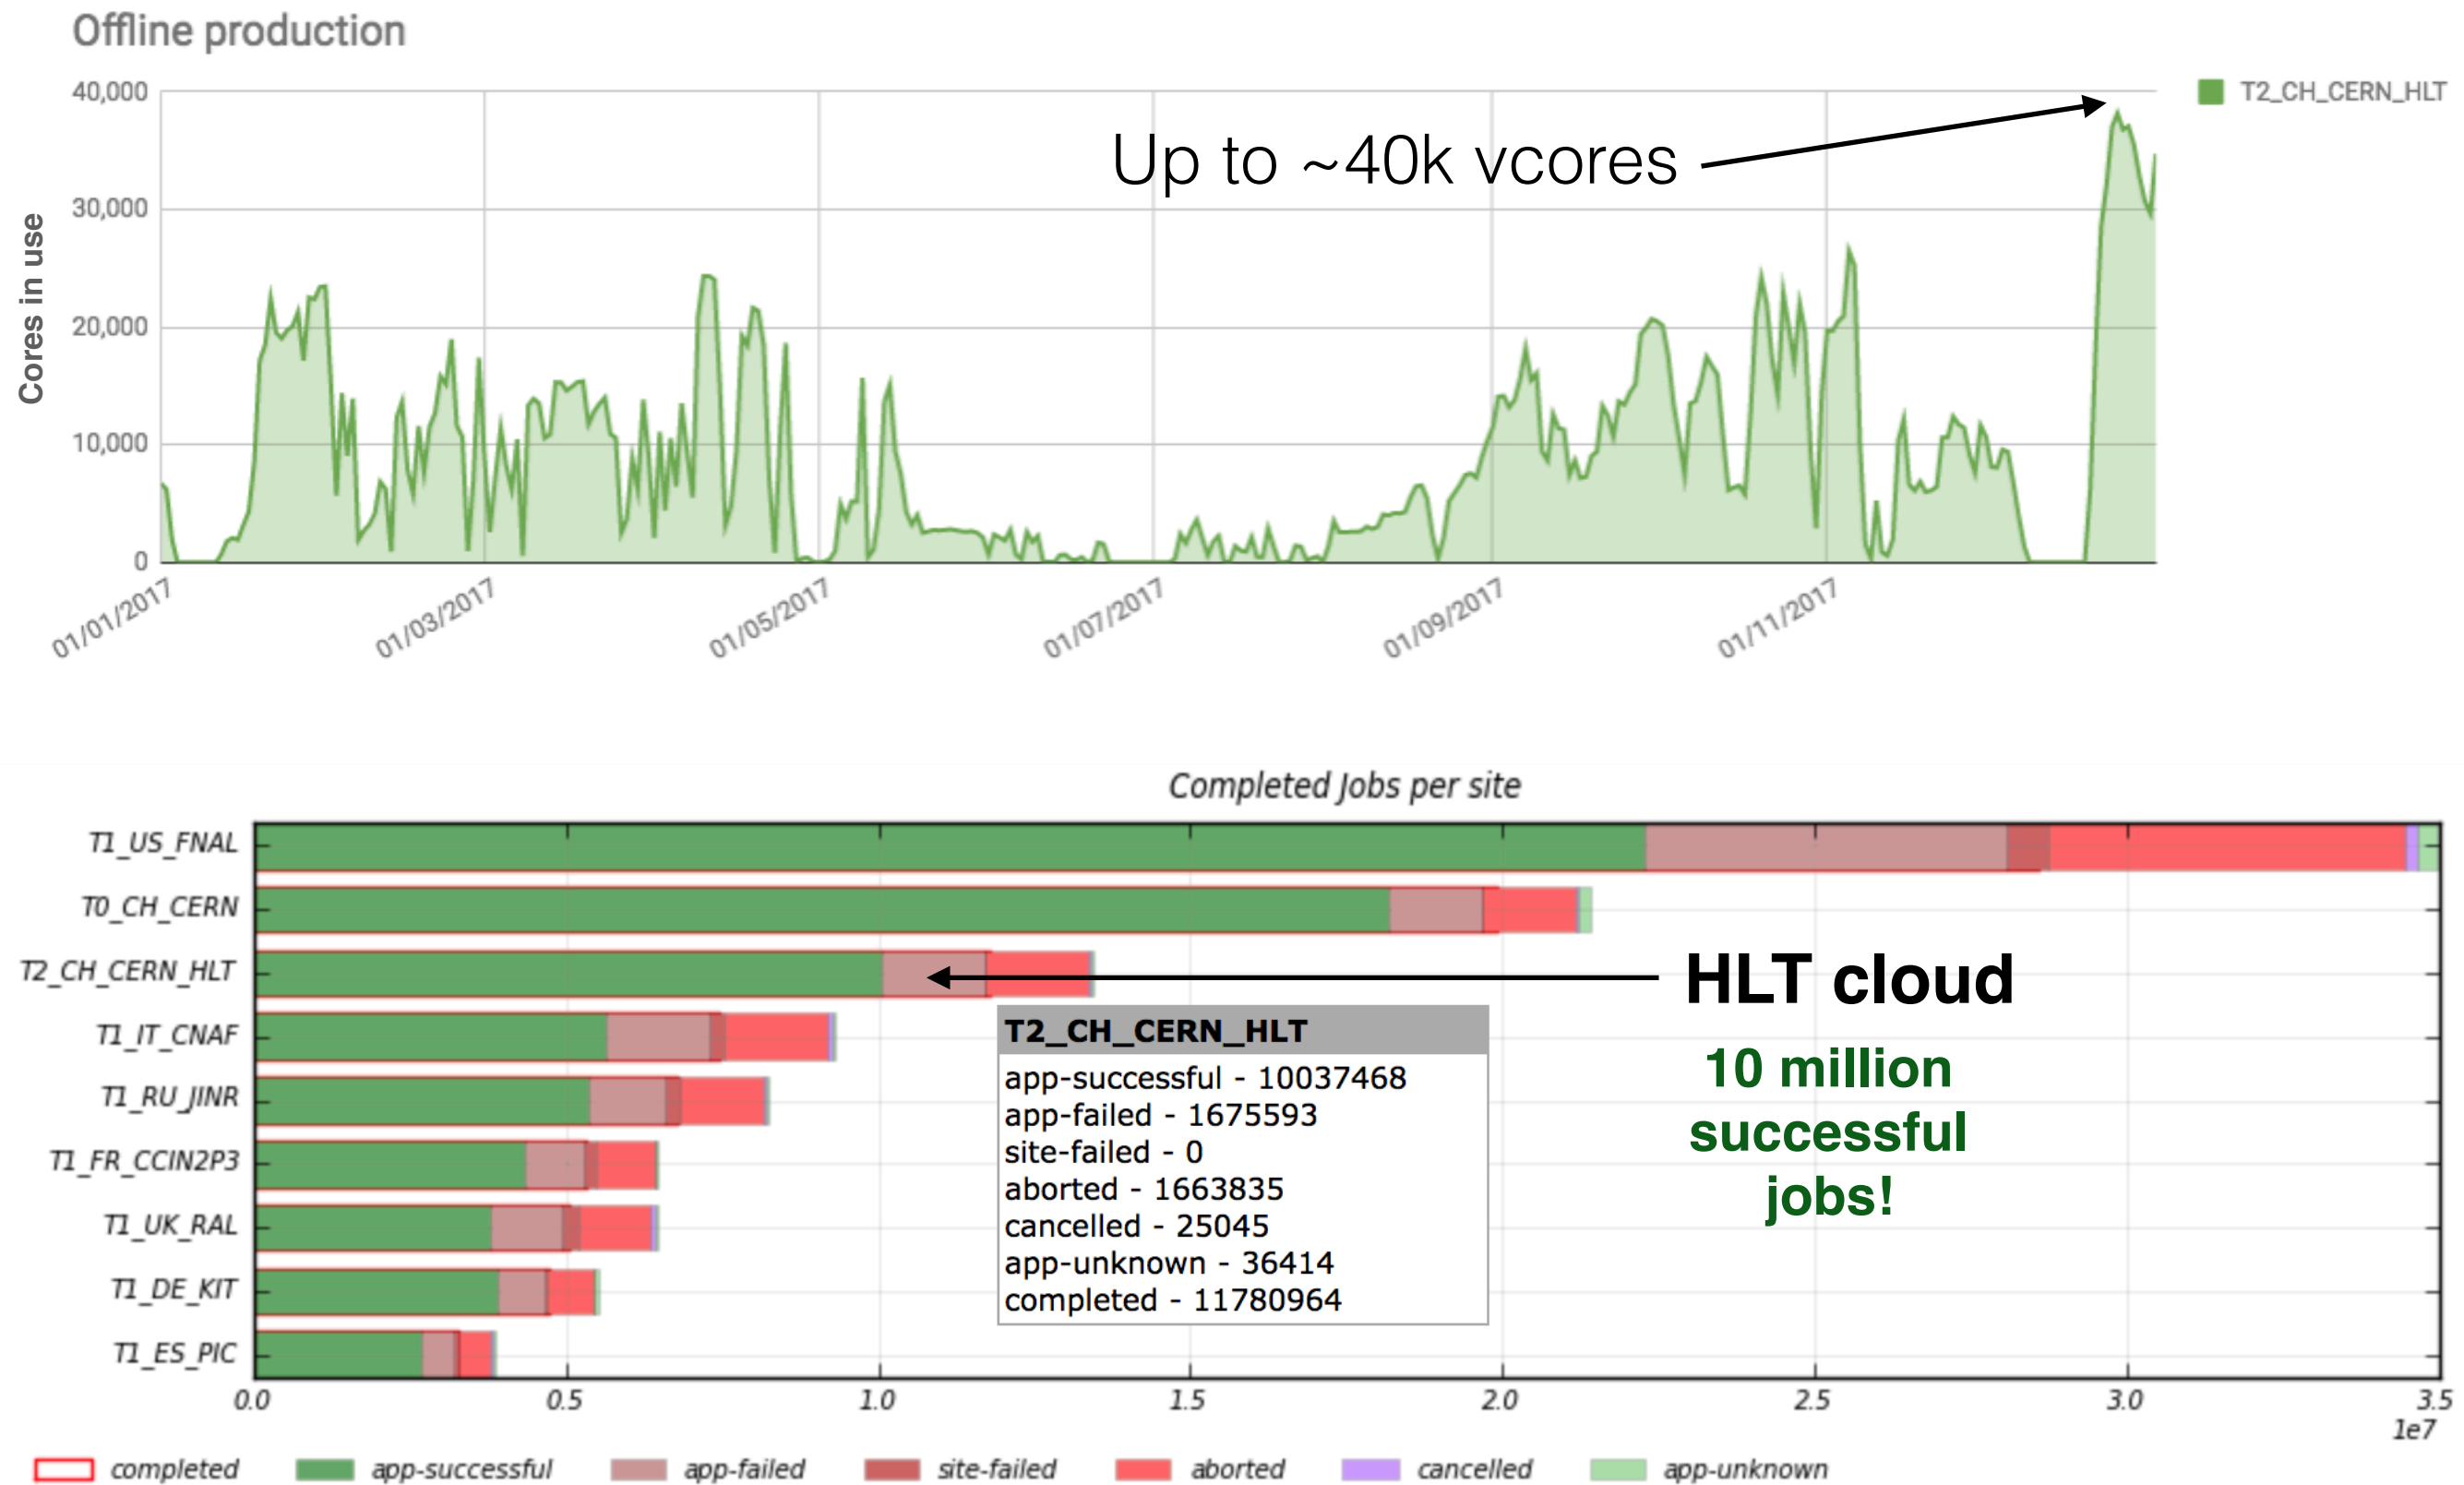
- **Graphite**: for collecting cloud specific metrics

- **Grafana**: for aggregation, correlation and displaying

- **Scripts**: consistency checks, auto-killing protections





The HLT Cloud Contribution: 2017



Contribution in 2018

1st Jan to 22 Jun 2018




Known difficulties from experience

- **Catching up with DAQ changes** — the cloud overlay must be provided on top of what's in use for CMS-DAQ as Operating System Software, Networking and Computing Hardware
- **Making efficient use of resources** — to achieve, we must follow closely the LHC and CMS schedules and plans to set the appropriate cloud operation mode
 - Also, some conditions are hard to anticipate and thus not reacting accordingly may result in massive job failures or waste of resources
- **Troubleshooting job failures** — current monitoring tools help a lot spotting known issues, but unknown issues require log/state analysis.
 - However, cloud VMs available only for a few hours (i.e. during inter-fills) or get vanished due to protections kicking in upon issues.

What's next...

- Exploit more running cloud **during fills** - feature made available late in 2017



- Use of Singularity for containerisation under way — will allow the cloud to match both CC7/SLC6 job workflows
- Update to the latest Openstack release: need to work out best time to do this according to need for the cloud resources
- Monitoring: Elasticsearch/Logstash/Kibana for log analysis and job troubleshooting

Thank you!