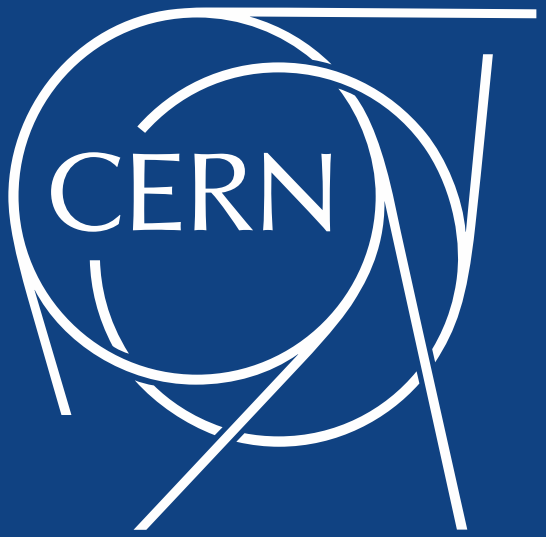
CERN

CERN IT Systems Management

Gavin McCance
CERN IT-CM

Outline

- General strategy
- Puppet for configuration
- Automation
- Data-centre and service monitoring
- Data-centre capacity plans
- Cloud and containers

LHC compute scaling challenges

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

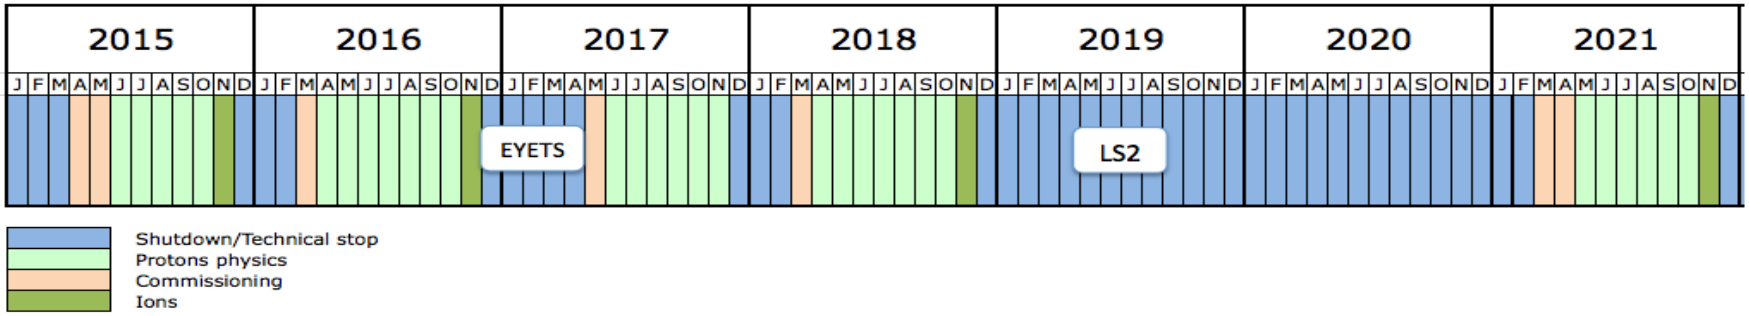

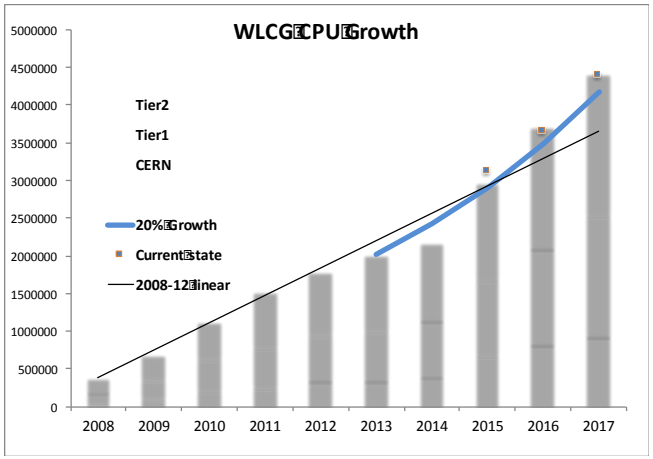

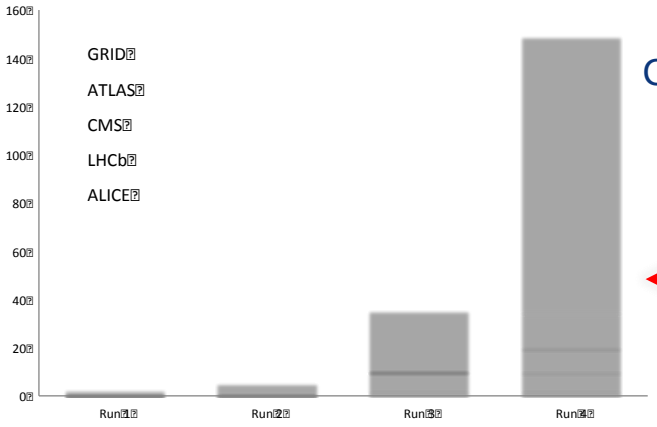

Compute: Growth > x50

What we think is affordable unless we do something differently

Good News, Bad News

- Additional data centre in Budapest now online
- Increasing use of new facility as data rates increase

But…

- Materials budget decreasing, no more money
- Staff numbers are fixed, no more people
- Legacy tools are high maintenance and brittle
- User expectations are for fast, dynamic self-service

General strategy

- As we scaled-up and made things more dynamic, our home-baked tools started to break and we struggled to find the effort to fix them

- How can we avoid the sustainability trap?
  - Really try hard to not develop our own stuff
  - Avoid accumulating technical debt
  - The major driver for the Agile Infrastructure was to control the technical debt

- How can we learn from others and share?
  - Find compatible open source communities
  - Contribute back where there is missing functionality
  - Stay mainstream

Are CERN computing needs really special ?

O'Reilly Consideration

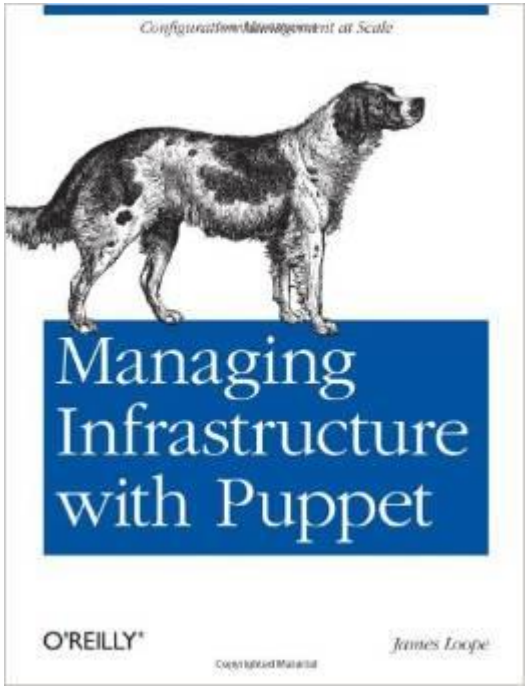

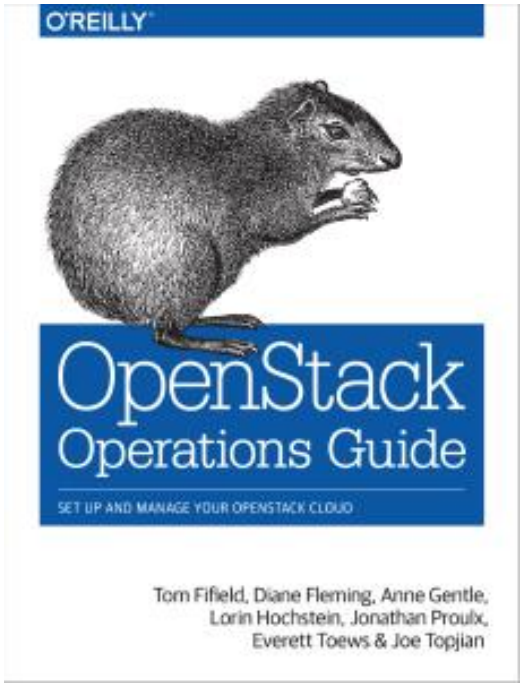

CERN Tool Chain

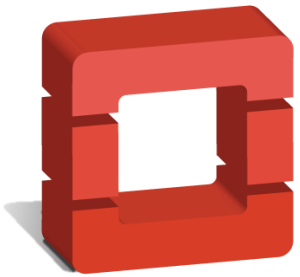

elasticsearch.

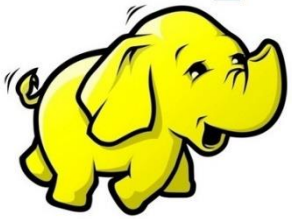

FOREMAN

Jenkins

Puppet for configuration

- Puppet roll-out has worked well for us
  - ~300 service manager with ~200 distinct hostgroups, both IT and experiment services
- Positive user experience
- Excellent open community support
  - Lots of standard modules
  - We've complemented this with a library of standard CERN configurations
- Solid APIs hooks for automation

Why Puppet?

- Several other obvious alternatives (Ansible, Chef, cfEngine,…)
- Puppet's declarative language model fitted our users' expectations (similar to previous tool)
- Centralised service model works well in our environment
- Solid community and good upstream support
  - This is the most critical aspect for anything you might chose
- We've invested in Puppet and we're still very happy
  - We've no plans to move to anything else anytime soon :)

Configuration service plans

- Faster versions of Puppet in the pipeline
- Potentially replace some of our own stuff (e.g. secrets management) with products that community have subsequently released

- Focus now on automation

Ops automated testing

- Focus is now on automation to control ops cost as we expand capacity

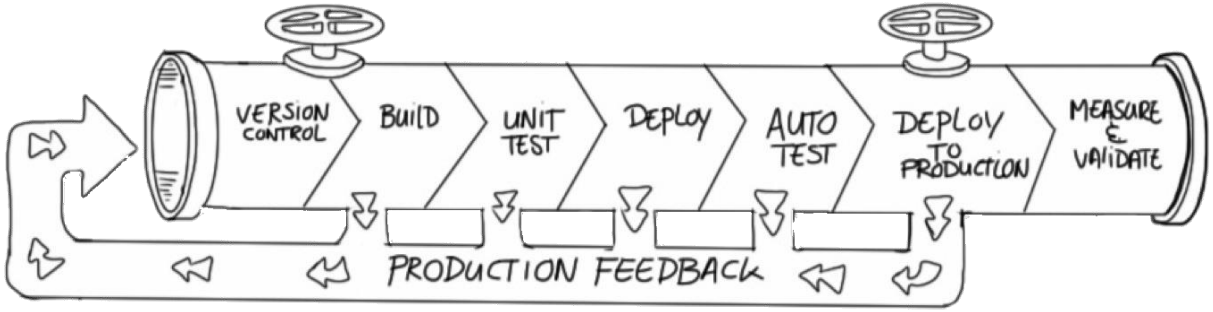

- "Continuous Deployment" combining gitlab-ci, puppet, Openstack orchestration
  - Aim is to automatically test as much as possible before rolling out a change to production
  - Working on making it easy to do this for all service changes with easy-to-add validation tests for automated "qa" testing

Ops automation tools

- Lots on the open and open++ market
  - Stackstorm, Cloudify, Saltstack, Rundeck, …
  - Several teams in IT using Rundeck just now
- Aiming at event-based automated recovery
- "Monitoring has not heard from this service for a while / has seen high-load"
  - → kick off a canned investigation / fix job
  - Everything in our infrastructure is now an API
- Rule #42: Only raise to a human when you run out of things to try

Infrastructure monitoring

- ‘Lemon monitoring service” backend now replaced by standard, open technologies
- Node metrics collected by Flume
  - Funneled to ElasticSearch for dashboards, Spark for online analysis and Hadoop for subsequent analytics
- Node exceptions (i.e. metrics out of bound) transported by ActiveMQ messaging
  - Can be subscribed to for automation events
  - Main “subscriber” is our ServiceNOW

New infrastructure monitoring (meter.cern.ch, based on Kibana/ES)

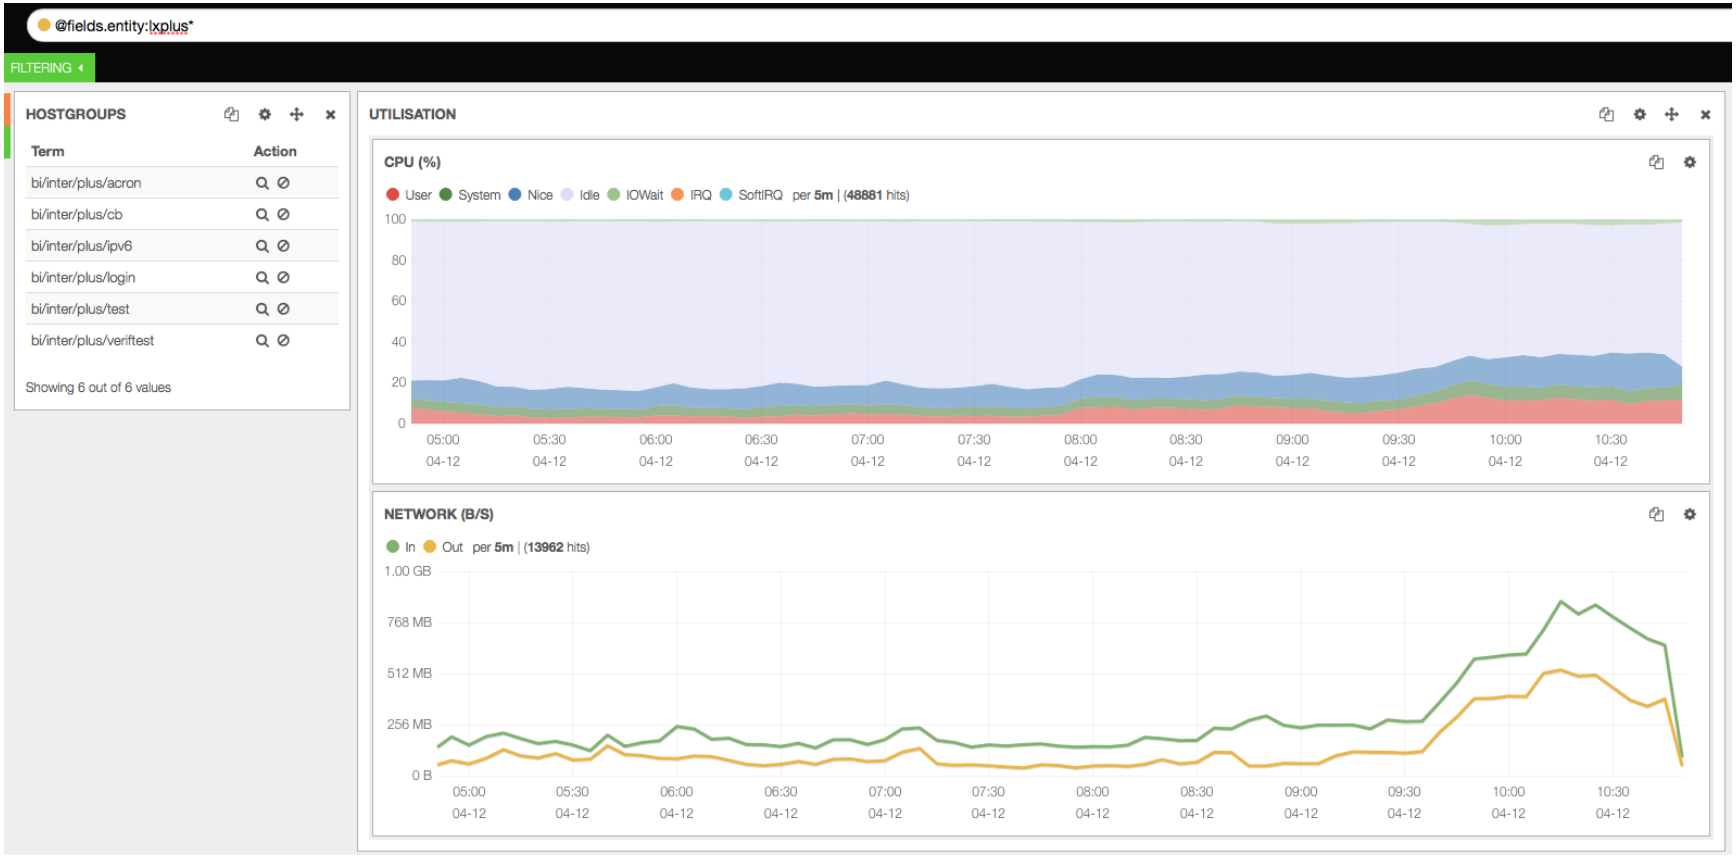

Infrastructure Monitoring futures

- Investigating further dashboard technologies and backend data stores
  - Currently ElasticSearch and Kibana
  - InfluxDB as possible backend
  - Graphana as possible dashboard

- Lemon agent replacement -> collectd?
  - Lower maintenance: replace all standard metrics
  - Make it cheap to create new metrics
  - Translation layer for residual Lemon sensors

Service monitoring

- ETF framework (from the WLCG SAM) based around checkMK and Nagios
  - Huge library of service check available
- Currently used for monitoring services across the WLCG for availability reporting

- Extending soon to local service monitoring
  - i.e. is a service doing what it's supposed to be doing
  - Useful directly for service managers
    - e.g. service dashboards
    - …but also an automated service recovery trigger

Streaming and analytics

- Lemon metric / exception / ETF streams can be monitored in real-time by Spark / Kafka or Esper
  - Register jobs hunting in real-time for more complex patterns ... firing a trigger
  - Event based automation
- Hadoop analytics
  - All data poured into Hadoop
  - Longer term problem investigation using analytics
    - e.g. understanding inefficiencies in the batch system

Physical DC Infrastructure

- CERN Main Data Centre
  - 3.8 MW for computing of which ~2.3 MW currently used
- Remote hosting site (Wigner Data Centre in Budapest)
  - 2.7 MW for computing of which ~1 MW currently used
    - Due to low rack power density, majority of available rack space currently used
  - Contract until end of 2019 but could be extended by 1-2 years (if necessary)

Options for additional capacity

- No firm plans yet.
- No plan to upgrade the Meyrin Data Centre further as it is considered that it would not be cost effective

- Investigating commercial clouds or external hosting as an option for providing additional capacity

- There is the option to do another tender for a remote hosting facility such as Wigner either as a sole solution or in combination with commercial cloud resources

- A final option, although not yet investigated in any way yet, would be a modular Data Centre addition, similar to the approach foreseen for LHCb and ALICE for post LS2 needs

How to integrate external capacity?

- Critical for ops cost to manage any external resources with same tools as we do here
  - Puppet / common monitoring
  - Possible restricted use-cases
  - Being tested now on IBM Softlayer / TSystems
  - Further procurements planned (HNSciCloud)

- Aim to expose (compute) resources only via common HTCondor lxbatch interface
  - one place to submit to for any CERN compute work

Private cloud

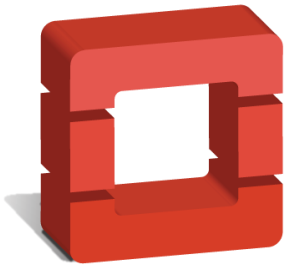

openstack™
CLOUD SOFTWARE

- Happy with Openstack
- Cloud model has allowed us to scale within fixed manpower
  - Emulate legacy environment – most services moved to VMs
  - Enable new ways of working
  - But difficult to avoid divergence / segmentation in fixed DC
- Collaborations with open source and industry covers technical debt
  - Future maintenance and testing
  - External mentoring avoids special solutions
  - Enhance staff job opportunities

Private cloud plans

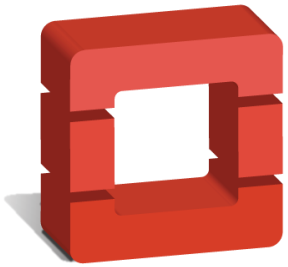

openstack™
CLOUD SOFTWARE

- Keep running the cloud
  - New releases every 6 months
  - Around 2,000 servers / year to renew
  - Integrate bare-metal management with Openstack
- Change the cloud – variety of new things in the pipeline
  - Containers
  - Software Defined Networking
  - Fine grained accounting and quota

Container technology

- Potentially, a very lightweight way to deploy services
  - Focus on application code and scaling behaviour rather than deployment

- Various use cases: R as a Service, Jupyter notebooks for HEP analysis, pre-packaged batch jobs
- We're testing Magnum from Openstack
  - Will deploy and scale for you a Docker Swarm, Kubernetes or Mesos instance

- Technology evolving very quickly in this space
- Aim to track and offer stable service
  - CERN ITTF: https://indico.cern.ch/event/506245/

Summary

- We'll keep doing configuration with Puppet
- Now focusing on the change testing and recovery automation
- Monitoring based on standard technologies; understanding which tools are best
  - Looking how best to integrate the monitoring with our automation efforts
- Various options for extending compute capacity
  - Validating providers, tools and experiment payloads in various providers now
  - No firm decisions yet
- Private cloud with Openstack – very happy with it
  - Manpower scaling well

Summary

- Moving to a set of solid, open-source based community products has helped us a lot
- New technologies, as they come along, are typically integrated quickly
  - Though you have to invest manpower to keep things up to date
  - Plus side -> any development effort you do (should be) integrated upstream

- We’re happy with the Puppet / Openstack stack eco-system and would recommend it
- We believe we have a solid base which will serve us for a long time yet