



ALICE Offline services and operations

L. Betev, G. Eulisse, T. Wilken, M. Litmaath

ALICE Offline operations

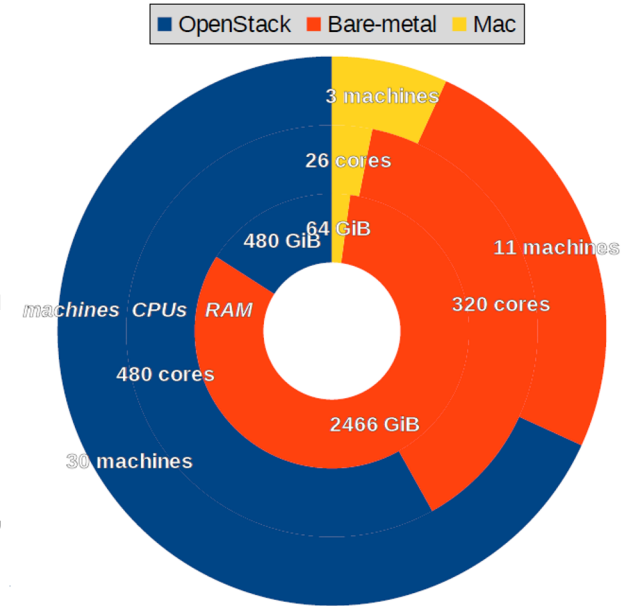
- A coin with two sides, each handled by its own team
 - What to run → build and CI operations (next pages)
 - Where to run → grid operations
- Grid operations
 - Mostly managed through central services in the offline cluster
 - LDAP, CA, file catalog, task queue, job broker, job manager, MonALISA, ...
 - MonALISA is crucial for many aspects
 - Monitoring of central services
 - Orchestration of productions, analysis trains and bulk data transfers
 - GUI for user job and data management
 - SE tests, to inform automatic and operator data management decisions
 - Tracking of grid site network metrics, ditto
 - Monitoring of site VOboxes, jobs, WN properties, SE disk servers
 - Accounting: VO perspective
 - And more...

WLCG connections

- VOMS, MyProxy, IAM
 - For VO management and pilot job submissions to CEs
- SAM
 - ETF only used for CE tests
 - SE and VObox test results are forwarded by MonALISA
- CRIC
 - SAM VO feed contents are determined from ALICE LDAP service
- Accounting: site perspectives
- Ops Coordination
 - Handling of matters not specific to ALICE, e.g. MW campaigns

ALICE Offline build and CI system

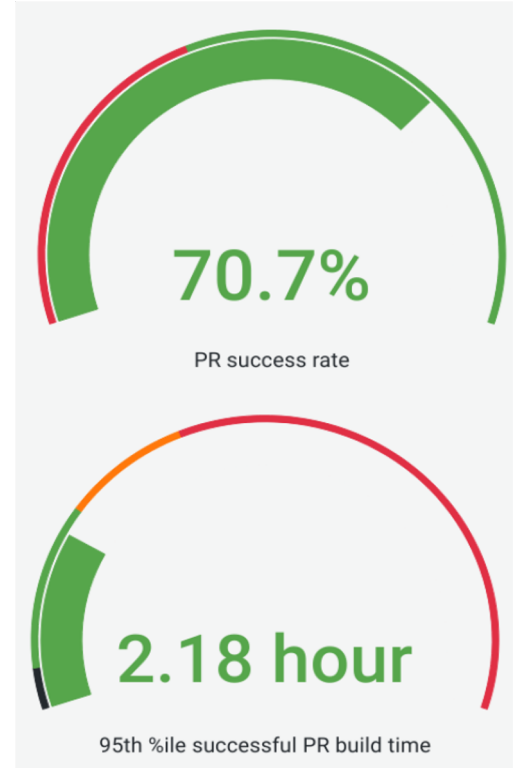
- 42 build machines with 826 CPU cores in total build pull requests for ALICE Offline repositories
 - across different platforms (CentOS, *Ubuntu*, *OSX*) and architectures (x86, *arm64*) on OpenStack and the *offline cluster*
 - using bespoke system of CI scripts + aliBuild as the build system
 - for simpler tasks and code formatting: GitHub Actions
 - 37 different checks across 13 different repositories
 - 20–30 (non-trivial) builds + ~200 fast re-checks* per hour
- Builds software every night, e.g. for use on the Grid, from CVMFS, as RPMs, etc
- Build processes scheduled using Mesos and Aurora on Linux, builds are containerised using Docker



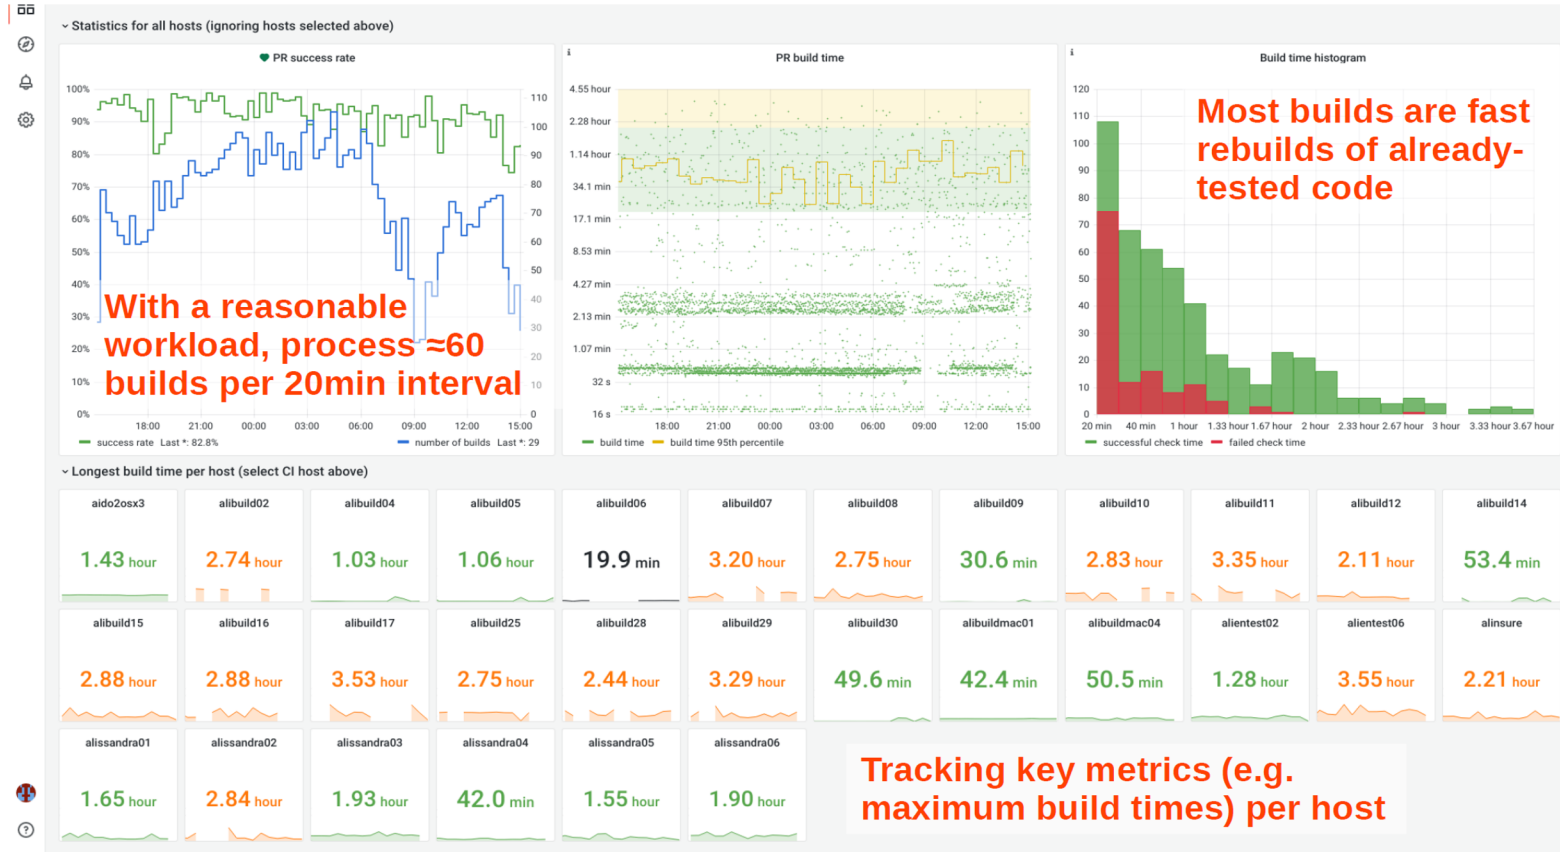
* re-builds of already tested code, to make sure it still builds against any code updated since its submission

CI system operation

- Scripts and configuration tracked using git
- GitLab used for some of the repositories
- Changes to CI scripts can be deployed on a rolling basis for testing
- OpenStack VMs managed through Puppet
- S3 for storing the build artefacts to use as a cache
- Bare-metal machines managed manually, as maintenance burden is fairly low
- Monitoring and alerting through Grafana



Monitoring



CI monitoring: incident example

- CI system monitored using CERN IT administered Grafana instance
- Metrics pushed to CERN IT administered InfluxDB instance
- Alerting is easy to set up



Possible improvements

- On the ALICE side
 - Room for performance improvements
 - CI performance eaten up by stringent isolation of different pull requests from each other – there are likely ways to reduce this without compromising sandboxing
 - Unit testing coverage of some scripts can be improved, in order to make testing of simple changes to existing scripts simpler
- On the CERN IT side
 - Making the backend services **rock solid**, instead of providing more features
 - Cf. incidents affecting GitLab etc.
 - In other respects the CI system is rather low maintenance
 - Having macOS resources provided by CERN IT would be desirable