




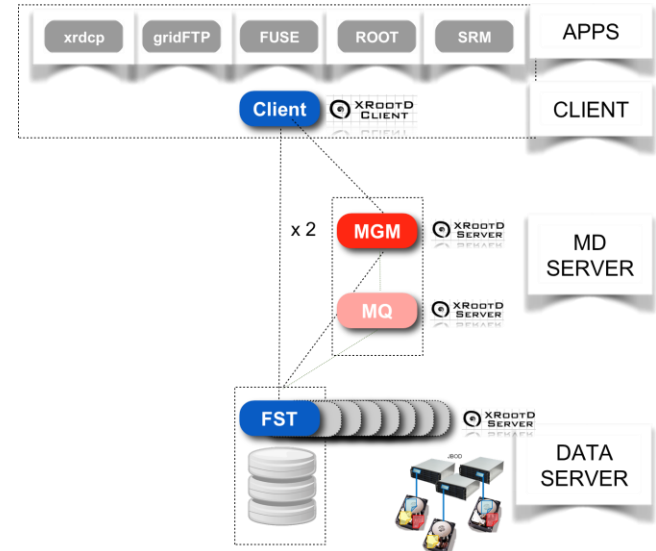


Testing of complex, large-scale distributed storage systems: a CERN disk storage case study

J. Makai, A. J. Peters, G. Bitzes, E. A. Sindrilaru, M. K. Simon, A. Manzi (speaker)
CERN

EOS

- **Disk** file storage
- Designed for **Massive storage** with high performances (> 50PB, 500 M files and Khz metadata operations)
- **In memory namespace**
- **New distributed namespace backend (QuarkDB)**
- <https://indico.cern.ch/event/587955/contributions/2936873/>



EOS CI till beginning of 2017

- Jenkins CI
 - Not state of the art anymore
 - Lacks some functionalities
 - Harder to maintain (lot of plugins from different sources)
- No real CI present for EOS
 - Only building RPMs for couple of platforms (CC7, SLC6 maintained mostly)
 - Using mock for build
 - No regression testing, regular static analysis in place
 - No integration with other build facilities (e.g. Koji)
 - Lot of manual steps, error-prone

Long journey behind us

- We started again almost from scratch
- Migration to Gitlab CI for EOS, XrootD and all related projects
- Completely new technologies
 - Docker containers instead of mock
 - Docker images/containers for system testing and distribution
 - Compiler caching
 - Static analysis
 - Google test for unit testing
- New build cluster and new repos
- Support for brand new platforms
- Agile mindset, easy and quick to respond to new issues, needs




























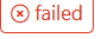























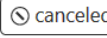











Migration to Gitlab CI

dss > eos > Pipelines

All 2211 Pending 0 Running 0 Finished 2181 Branches Tags

Run Pipeline

CI Lint

Status	Pipeline	Commit	Stages		
	#295063 by  latest	 master  13e02395 MGM: FsCmd clone and compilation fix	    	 01:24:12  about 14 hours ago	   
	#294843 by  latest	 master  13e02395 MGM: FsCmd clone and compilation fix	   	 00:28:17  a day ago	 
	#294827 by  latest	 master  75715cf4 MGM: FsCmd dropfiles improvement	   	 00:10:22  a day ago	 
	#294673 by  latest	 master  7c78f808 FUSEX: update README	   	 00:25:46  a day ago	 
	#294665 by  latest	 master  c0d4507a FUSEX: correct README	   	 00:20:38  a day ago	 



Docker containers for builds

- Isolated build environment
- Faster to support new platforms
 - mock moves slowly
 - Images are immediately present
- Support other distributions with different packaging
- Still one homogeneous build environment to manage

New platforms supported

- Mac OS
 - Not completely new but revived from its ashes
 - Updated build to Sierra
- Fedora
 - We can update quickly to latest versions thanks to the new ecosystem
 - Fedora 27, 28 and rawhide currently in CI
- Debian
 - Ubuntu artful & bionic in CI
- Docker image for each release available via GitLab Registry
 - gitlab-registry.cern.ch/dss/eos/eos-prod:4.3.2

New yum and apt repositories

storage-ci.web.cern.ch/storage-ci/


Index of /storage-ci

<u>Name</u>	<u>Last modified</u>	<u>Size</u>	<u>Description</u>
 Parent Directory		-	
 cernbox/	23-Nov-2017 11:00	-	
 debian/	11-Aug-2017 10:22	-	
 eos-cms-redirect-prefix/	20-Jul-2017 10:33	-	
 eos/	06-Oct-2017 13:46	-	
 eosrucioplugin/	20-Jul-2017 13:15	-	
 grid-hammer/	13-Jun-2017 17:38	-	
 quarkdb/	25-Oct-2017 20:05	-	
 storageci.key	11-Aug-2017 10:15	929	
 xrootd-dsi/	25-Jul-2017 09:45	-	
 xrootd/	23-Nov-2017 16:30	-	

- On EOSUSER (backup for release packages on CEPH volume)
- Non-release packages kept for 10 days, automatic maintenance
- Apt repo with Ubuntu packages under debian
- Release RPMs and all DEB packages are automatically signed with GPG key
- `rpm --import http://storage-ci.web.cern.ch/storage-ci/storageci.key`
- `curl -sL http://storage-ci.web.cern.ch/storage-ci/storageci.key | apt-key add -`

Docker based testing

- Docker image created for each build with fresh packages of the changes
- Simulates a fully functional system with all its features
- <https://gitlab.cern.ch/eos/eos-docker> project
 - Independent from the CI environment
 - Used also for development and easy functionality showcase



Docker based testing – System test

- Both with old, in-memory and new, quarkdb namespace
- New fully functional instance of EOS is created: tests booting, authentication, etc.
- Runs eos functional tests: testing EOS commands with eos and xrootd clients
- FUSE(X) testing: fuse microtests for both, benchmarks for FUSEX and POSIX compliance tests for FUSEX daily
- <https://indico.cern.ch/event/587955/contributions/3012718/>

Docker based testing – Stress test

- To put some load on the system
- Write, stat, read, delete files from multiple threads simultaneously using xrootd protocol

```
10000 files :: 100 threads :: Run 2/3

PASS xroot :: write (13.60 sec)
Rate: 735.42 Hz
Latency span: from 0.31 to 8.91 sec

PASS xroot :: stat (0.61 sec)
Rate: 16489.76 Hz
Latency span: from 0.00 to 0.56 sec

PASS xroot :: read (8.54 sec)
Rate: 1171.22 Hz
Latency span: from 0.21 to 7.06 sec

PASS xroot :: delete (3.82 sec)
Rate: 2619.79 Hz
Latency span: from 0.00 to 3.58 sec
```

```
10000 files :: 100 threads :: Run 3/3

PASS xroot :: write (16.61 sec)
Rate: 602.00 Hz
Latency span: from 0.12 to 8.39 sec

PASS xroot :: stat (1.46 sec)
Rate: 6864.65 Hz
Latency span: from 0.01 to 1.44 sec

PASS xroot :: read (16.73 sec)
Rate: 597.70 Hz
Latency span: from 0.53 to 5.21 sec

PASS xroot :: delete (5.06 sec)
Rate: 1976.40 Hz
Latency span: from 0.00 to 4.62 sec
```

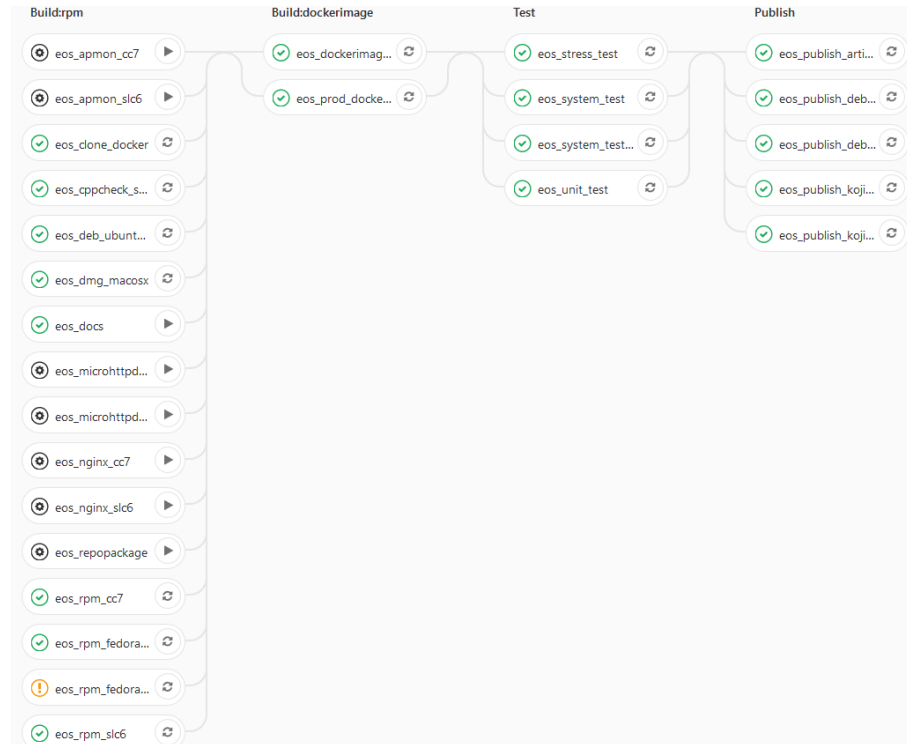

Gitlab CI Pipeline – 30 jobs, 4 stages

- No manual, no error-prone steps
- Integrated with everything possibly
 - SRPM → CERN Koji

Overall statistics

- Total: **3534 pipelines**
- Successful: **2012 pipelines**
- Failed: **977 pipelines**
- Success ratio: **67%**

Pipelines for last year




Static analysis – Coverity & Sonar

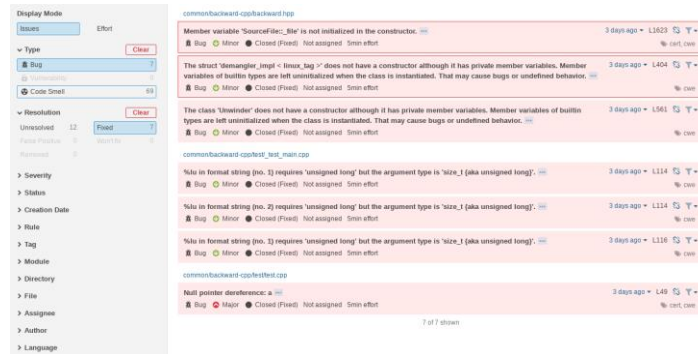
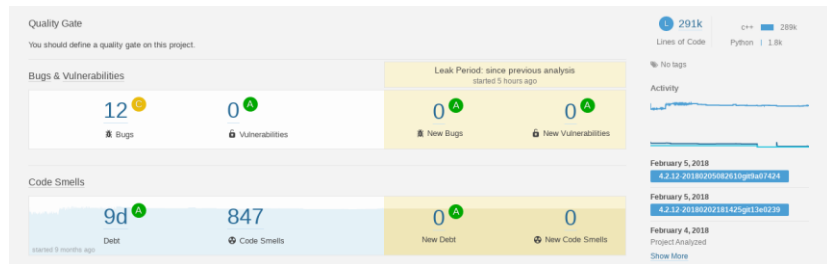

- Runs once a day
- Special instrumented compilation required

- For each push to master
- Cppcheck: analysis on source code level, no compilation

Outstanding vs Fixed defects over period of time



High impact Outstanding Defect per Category



Conclusions & Future steps

- Successfully implemented Continuous Testing for EOS since more than a year
- Next Steps :
 - Implement CD via Kubernetes integration in Gitlab
 - Add EOS HTTP/GridFTP gateways testing
 - Add more regression tests for functionality
 - Implement Robustness testing



dss / eos Master
EOS project.



updated about 5 hours ago

