



HammerCloud: benchmarking storages

CRISP Annual Meeting

Valentina Mancinelli
IT/SDC

02/06/2016



What is HammerCloud

HammerCloud is a framework for testing distributed systems.

Two types of tests:

- **Functional (automated) testing:** Are autoscheduled, which means configure once, run forever. Used usually to frequently submit few “ping” jobs for basic continuous validation.
- **Stress (on-demand) testing:** Are on-demand, which means configure once, run once. Used to submit benchmarks, evaluate changes to the infrastructure or configuration, evaluate SW changes, compare performances. Usually used to submits a large number of jobs concurrently during a slot of time.

HammerCloud

Composed by:

- Service to schedule, run, track status, gather results/statistics of tests
- Web interface to:
 - Display test status and test results
 - Configure new tests

Test monitoring

Overview of all the running jobs

Running and Scheduled AFT/PFT Tests

State	Id	Host	Template	Start (CET)	End (CET)	Sites	subm jobs	run jobs	comp jobs	fail jobs	tot jobs
running	20028970	voatlas65	506: AFT Reco_trf 17.2.7.6 SMWZ_NTUP Panda	24/Nov, 3:16	25/Nov, 3:59	ANALY_TRIUMF, ANALY_AUSTRALIA, ANALY_VICTORIA-WG1, 144 more...	47	103	9297	686	10246
running	20028975	voatlas167	487: PFT mc12 AtlasG4_trf 17.2.6.2	24/Nov, 7:18	25/Nov, 5:20	CA-ALBERTA-WESTGRID-T2, CA-VICTORIA-WESTGRID-T2, SFU-LCG2, 193 more...	43	206	6497	268	7096
running	20028983	voatlas49	508: AFT Reco_trf 17.2.7.6 data11 SMWZ_NTUP Panda	24/Nov, 15:56	25/Nov, 15:28	ANALY_TRIUMF, ANALY_AUSTRALIA, ANALY_VICTORIA-WG1, 144 more...	48	88	2436	169	2827
running	20028994	voatlas49	505: AFT UA 17.2.7 Panda	24/Nov, 18:34	25/Nov, 20:39	ANALY_TRIUMF, ANALY_AUSTRALIA, ANALY_VICTORIA-WG1, 144 more...	28	46	2013	78	2249
running	20028995	voatlas65	489: PFT mc12 AtlasG4_trf 16.6.7.34	24/Nov, 19:10	25/Nov, 21:24	CA-ALBE					
running	20028999	voatlas167	485: PFT mc12 AtlasG4_trf 17.2.2.2	24/Nov, 21:14	25/Nov, 19:45	CA-ALBE					

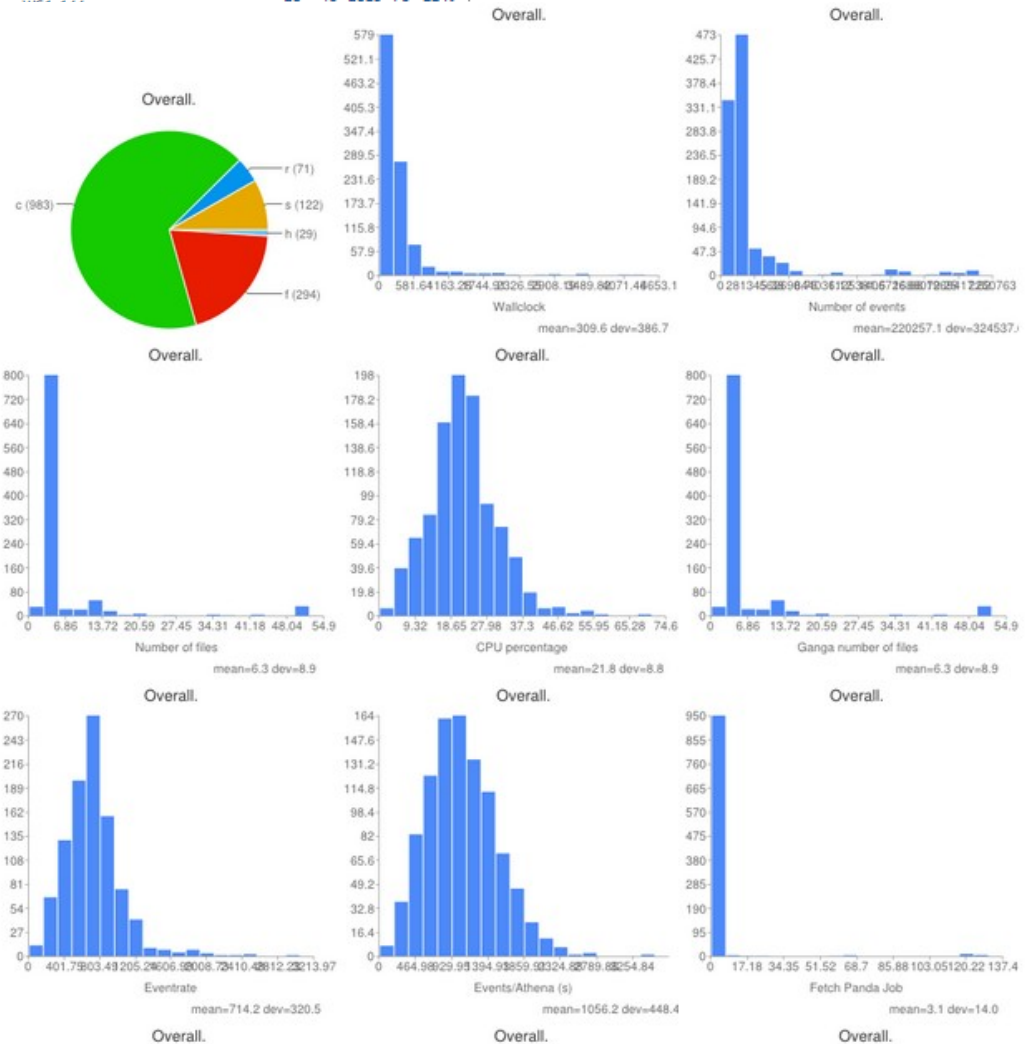


Running and Scheduled Stress

State	Id	Host	Template	Start (CET)	End (CET)	Cloud
running	20028982	voatlas65	555: UA 17.2.7 Panda (default Panda job splitting, special DS)	24/Nov, 18:00	25/Nov, 3:00	TW_PA
running	20028987	voatlas167	512: CERN-P1 StressTest - mc12 AtlasG4_trf 17.2.2.2	24/Nov, 21:50	30/Nov, 21:12	FR_PR
running	20029000	voatlas49	533: ROOT HWWNtupleCode-00-02-07 DATA p1067 17.2.7 Panda SL6	24/Nov, 22:20	28/Nov, 21:20	FR_PA

Running and Scheduled Function

State	Id	Host	Template	Start (CET)	End (CET)	Cloud
running	20028968	voatlas285	544: FT FAX direct rc_test : Simple histogram making 17.5.0	24/Nov, 0:52	25/Nov, 3:13	DE_PANDA, US, UK_PANDA, 6 mor
running	20028973	voatlas285	547: FT FAX direct cloud: Simple histogram making 17.2.7	24/Nov, 4:50	25/Nov, 6:19	DE_PANDA, US, UK_PANDA, 6 mor
running	20028974	voatlas285	532: ProdTrans ptest mc12 AtlasG4_trf 17.2.6.2	24/Nov, 7:02	25/Nov, 5:18	CA_PROD, DE_PROD, ES_PROD, 9 more
running	20028977	voatlas285	546: FT FAX copy: Simple histogram making 17.5.0	24/Nov, 7:02	25/Nov, 5:18	DE_PANDA, US, UK_PANDA, 6 mor



View of the results of a test



Test configuration

Change template

✖ Delete

Type information

Category:

Description:

Lifetime:

Period:

Active

Is golden

Obsolete

Files

Jobtemplate:

Usercode:

Optionfile:

Metricperm:

Inputtype:

Testoption:

Gangabin: Please, use the HEAD version.

Extraargs:

Template hosts

Host	Delete?
559 (functional) - PFT_MCORE mc12 AtlasG4_trf 17.2.2.2 - voatlas284.cern.ch <input type="text" value="voatlas284.cern.ch"/>	<input type="checkbox"/>
559 (functional) - PFT_MCORE mc12 AtlasG4_trf 17.2.2.2 - voatlas285.cern.ch <input type="text" value="voatlas285.cern.ch"/>	<input type="checkbox"/>

+ Add another Template Host

Template backends

Resubmit enabled	Resubmit force	Num datasets per bulk	Min queue depth	Max running jobs	Backend
<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="text" value="1"/>	<input type="text" value="0"/>	<input type="text" value="1"/>	<input type="text" value="-----"/>
<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="text" value="1"/>	<input type="text" value="0"/>	<input type="text" value="1"/>	<input type="text" value="-----"/>
<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="text" value="1"/>	<input type="text" value="0"/>	<input type="text" value="1"/>	<input type="text" value="-----"/>

+ Add another Template Backend

Template clouds

Resubmit enabled	Resubmit force	Num datasets per bulk	Min queue depth	Max running jobs	Cloud
<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="text" value="1"/>	<input type="text" value="0"/>	<input type="text" value="1"/>	<input type="text" value="-----"/>
<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="text" value="1"/>	<input type="text" value="0"/>	<input type="text" value="1"/>	<input type="text" value="-----"/>
<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="text" value="1"/>	<input type="text" value="0"/>	<input type="text" value="1"/>	<input type="text" value="-----"/>

+ Add another Template Cloud

Template sites


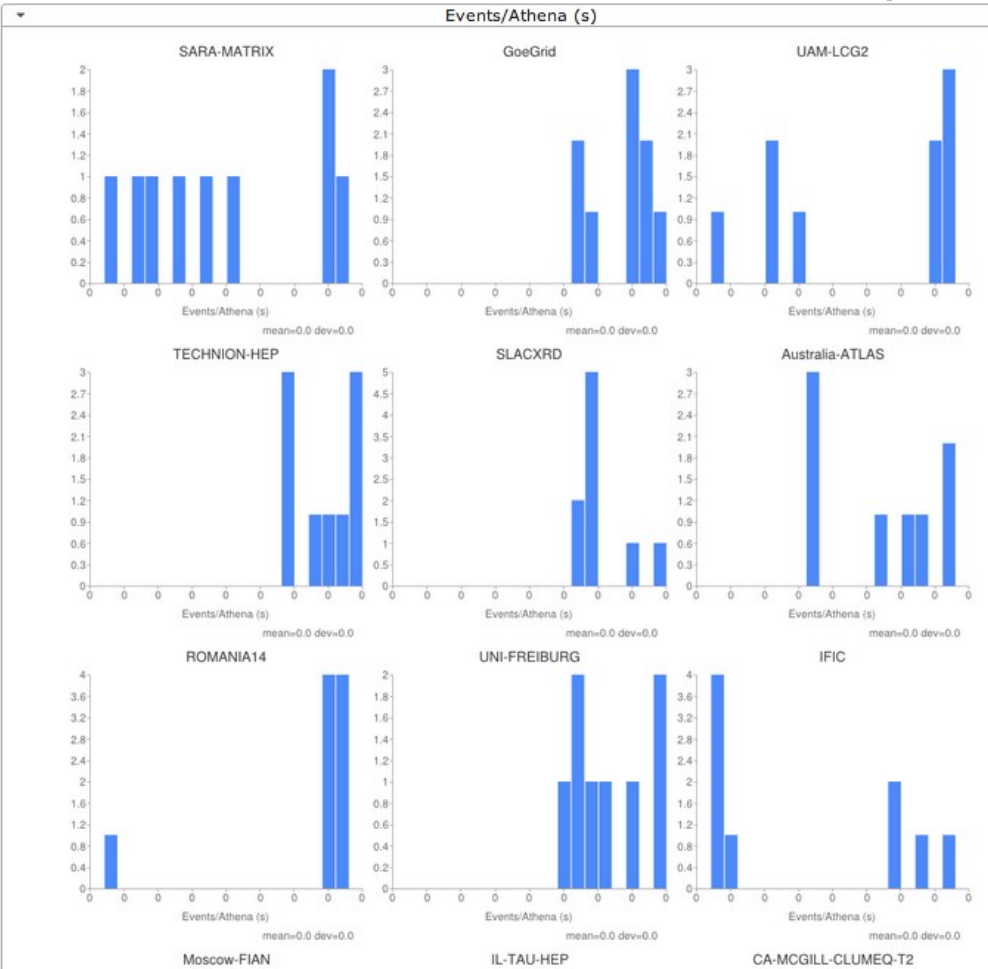
Resubmit enabled	Resubmit force	Num datasets per bulk	Min queue depth	Max running jobs	Site
<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="text" value="1"/>	<input type="text" value="0"/>	<input type="text" value="0"/>	<input type="text" value="FZK-LCG2_MCORE"/>
<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="text" value="1"/>	<input type="text" value="0"/>	<input type="text" value="0"/>	<input type="text" value="wuppertalprod_MCORE"/>
<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="text" value="1"/>	<input type="text" value="0"/>	<input type="text" value="0"/>	<input type="text" value="BEIJING_MCORE"/>

Current uses

- Used at CERN for Frequent validation and stress testing for different resources.
- Enabling VO- and site-administrators to run tests in an automated or on-demand manner.
- Used by 3 experiments (Atlas, CMS, LHCb)
- Many use cases

Benchmarking

- HammerCloud tests to submit MC simulation to crosscheck HEPSPROC (measure CPU performances)

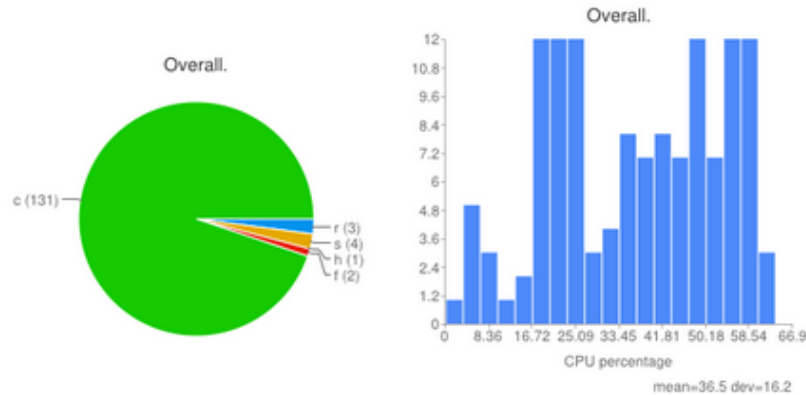


Storage testing

- Functional test to constantly monitor DPM performances
- DPM (Disk Pool Manager) is a grid enabled disk storage management system operated at more than 200 sites in the WLCG.
- Use HammerCloud to submit jobs to test data access performances:
 - Jobs use root (analysis application) to access storage, different protocols
 - 2 different read pattern: 1% and 100% of input file
 - Upload metrics to a database for analysis and visualization

state	id	host	clouds	start time (CET)	end time (CET)	total jobs
running	20036284	it-hammercloud-submit-atlas-03	TO_PANDA, UK_PANDA, FR_PANDA, 3 more...	21/5/2014 18:28	22/5/2014 20:14	141

Input type: PANDA
 Output DS: user.gangarbt.hc20036284.*
 Input DS Patterns: ROOTIOTests/datasetsToUse.txt
 Ganga Job Template: ROOTIOTests/ROOTIOTests1667_panda.tpl
 Athena User Area: ROOTIOTests/ROOTIOTests1667_svn.tar.gz
 Athena Option file: ROOTIOTests/run-DPM.sh
 Template: SDC Tests
[View Test Directory \(for debugging\)](#)




[more plots »](#)

Site	S	R	C	F	Eff	T	Datasets	Queue	Max R	Resubmit	R. Force	Link
ANALY_TAIWAN_XROOTD_SL6	0	0	0	0	0.00	0	1	0	1	yes	no	»
ANALY_LAPP	1	0	5	0	1.00	6	1	0	1	yes	no	»
ANALY_INFN-FRASCATI	1	0	29	0	1.00	31	1	0	1	yes	no	»
ANALY_GLASGOW_SL6	0	2	27	1	0.96	30	1	0	1	yes	no	»
ANALY_ECDF_SL6	1	0	11	0	1.00	12	1	0	1	yes	no	»
ANALY_CERN_XROOTD	1	0	28	0	1.00	29	1	0	1	yes	no	»
ANALY_BNL_LONG	0	1	31	1	0.97	33	1	0	1	yes	no	»

DPM Root Read 1% TTC



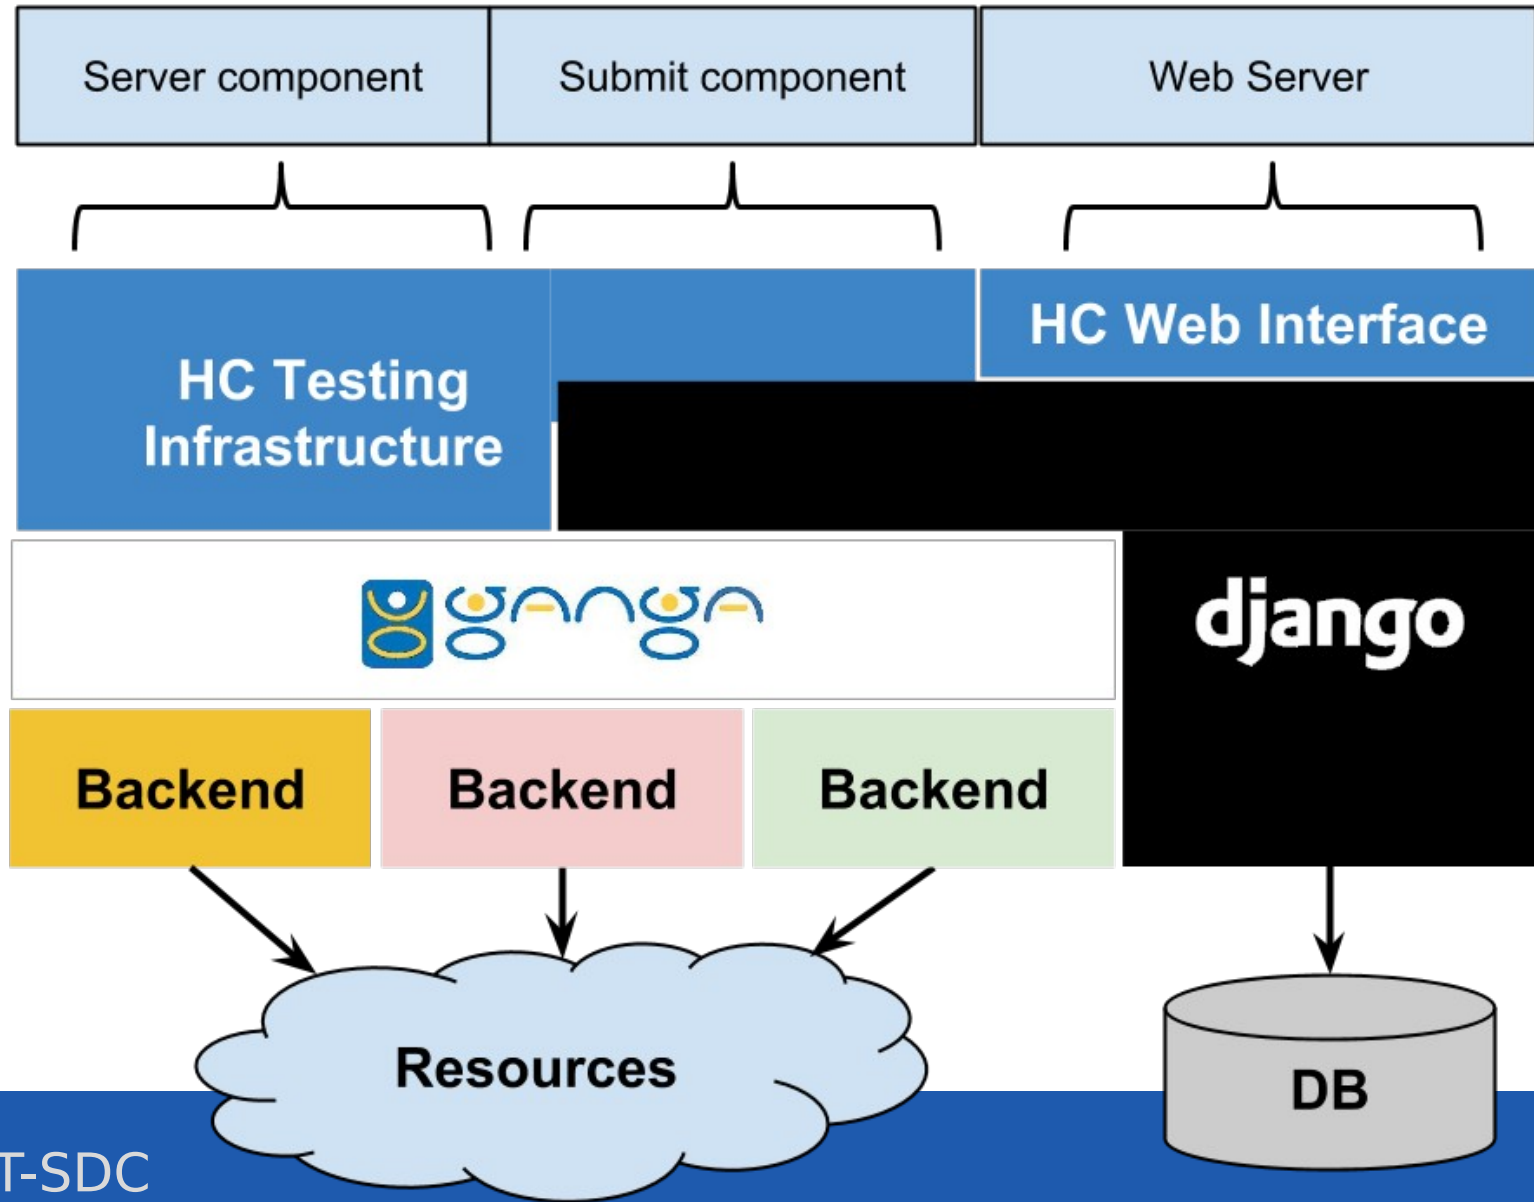
WALL time [s] for DPM Root Read 1% TTC



HammerCloud for CRISP

- Use HammerCloud as a tool to submit benchmarks for storage solution analysis
- Develop and run a suite for quantitative performance **evaluation of data access interfaces**
- synthetically reproduce arbitrary access patterns through appropriate parametrisation

HammerCloud Components



Modular implementation

- Core that implements basic components
- Composed by applications that extend core model. Easily extensible and customizable
- Ganga (submission component) provides:
 - various backend to submit jobs
 - plugin system that allows to develop specific backend

Start using HammerCloud

Basic deployment:

- 1 host with web server, server component, submit component
- DB (hosted local or its own server)
- Run executable locally to access local/remote storage systems.

Then...

- HC is a distributed system, easy to add new submission nodes to scale load
- Use different backend provided by Ganga to submit to various batch systems