

ANALYSIS FUNCTIONAL AND STRESS TESTING

Dan van der Ster, CERN IT-ES-DAS

for the HC team: Johannes Elmsheuser,
Federica Legger, Mario Úbeda García

WLCG Workshop, 8 July 2010

- Overview – what should we test in Distributed Analysis?
- HammerCloud and GangaRobot – Tools for stress and functional testing
- Recent Developments – HammerCloud v3 deployment
- Next Steps – the AFT, integration with SSB

- **Functional Testing:**

- Test the basic infrastructure
 - SAM...Nagios. Not covered in this talk.
- Basic test of the complete analysis workflow
 - Client...Workload Mgmt...Site...Worker Node...Storage
- “Special” workflows:
 - Complete chain test with Frontier/Squid access
 - Tier 3 analysis

- **Stress Testing:**

- On-demand test to help commission/tune/benchmark the analysis sites.
- Standardized tests – end-to-end test with real analyses of “real” data

- **HammerCloud** (HC) is a Distributed Analysis testing system serving these two use-cases:
 - **Robot-like Functional Testing**: frequent “ping” jobs to all sites to perform end-to-end DA testing
 - **DA Stress Testing**: on-demand (large-scale) stress tests using real analysis jobs to test one or many sites simultaneously to:
 - Help commission new sites
 - Evaluate changes to site infrastructure
 - Evaluate SW changes
 - Compare site performances...
- ATLAS has already made a big investment in HC stress testing:
 - ~210,000 CPU-wallclock days (that’s 576 CPU-years)
 - But this is only a few percent of the global DA resources



Hammercloud | ATLAS

You are connected as gangarbt, [click here to Logout!](#)

Home	Clouds	Tests	HC Stats	Ganga Robot	Panda Dashb.	Administration
------	--------	-------	----------	-------------	--------------	----------------

Welcome to HammerCloud-ATLAS. Click "HC Stats" above to see the currently running jobs.

Running and Scheduled Stress Tests

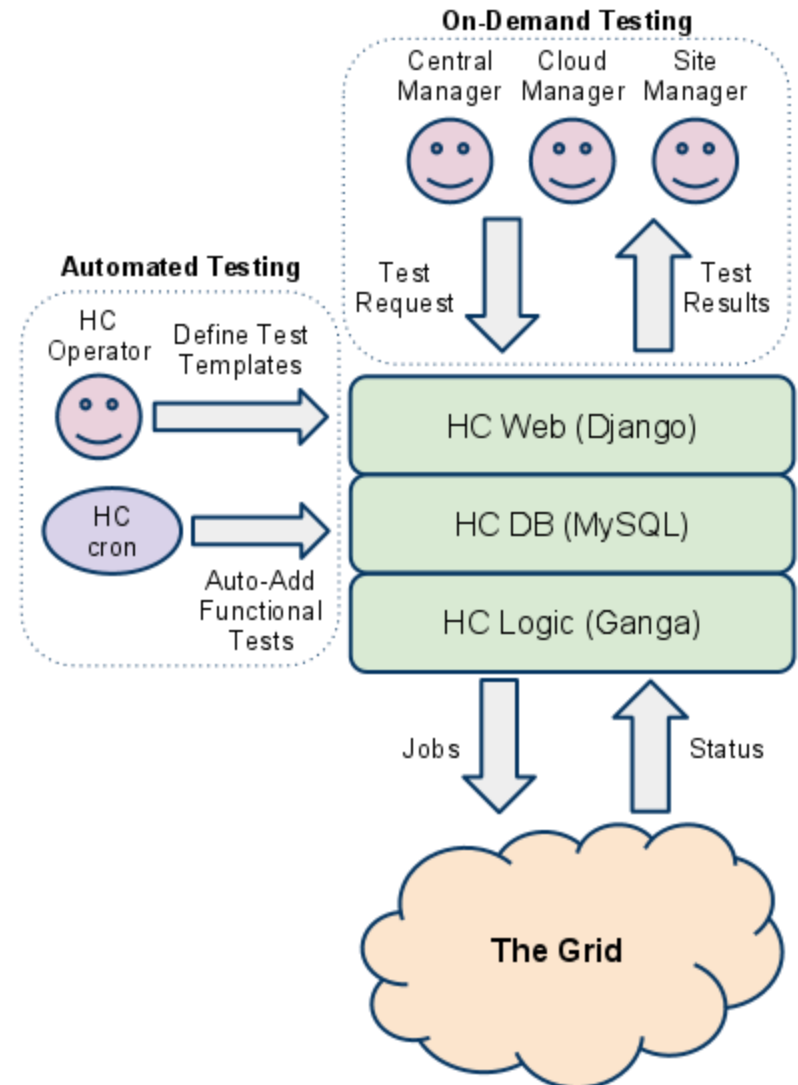
state	id	host	template	start time (CET)	end time (CET)	clouds	sites	subm jobs	run jobs	comp jobs	fail jobs	tot jobs
running	10000340	voatlas73	Muon 15.6.6 PANDA default data-access	2010-07-07 09:03:00	2010-07-07 13:03:00	US	ANALY_MWT2, ANALY_MWT2_X, ANALY_SLAC, 2 more...	99	744	977	94	1914

Running and Scheduled Functional Tests

state	id	host	template	start time (CET)	end time (CET)	clouds	sites	subm jobs	run jobs	comp jobs	fail jobs	tot jobs
running	10000342	voatlas49	UA 15.6.9 Panda Filestager	2010-07-07 10:38:07	2010-07-08 10:38:07	CA_PANDA, DE_PANDA, ES_PANDA, 9 more...	ANALY_AGLT2, ANALY_ALBERTA, ANALY_ANLASC, 89 more...	58	60	258	11	405
running	10000338	voatlas73	D3PDMaker 15.6.10.6 PANDA default data-access Frontier/Squid test	2010-07-06 15:47:03	2010-07-07 15:47:03	CA_PANDA, DE_PANDA, ES_PANDA, 9 more...	ANALY_AGLT2, ANALY_ALBERTA, ANALY_ANLASC, 89 more...	156	32	567	167	941
running	10000337	voatlas73	D3PDMaker 15.6.10.6 LCG Frontier/Squid test	2010-07-06 15:41:03	2010-07-07 15:41:03	CA, DE, ES, 7 more...	ALBERTA-LCG2_MCDISK, AUSTRALIA-ATLAS_MCDISK, BEIJING-LCG2_MCDISK, 98 more...	8	26	445	183	662
running	10000336	voatlas49	UA 15.6.9 Panda	2010-07-06 13:12:04	2010-07-07 13:12:04	CA_PANDA, DE_PANDA, ES_PANDA, 9 more...	ANALY_AGLT2, ANALY_ALBERTA, ANALY_ANLASC, 89 more...	64	54	5366	150	5649
running	10000335	voatlas73	UA 15.6.9 LCG DQ2 Local	2010-07-06 13:12:04	2010-07-07 13:12:04	CA, DE, ES, 7 more...	ALBERTA-LCG2_MCDISK, AUSTRALIA-ATLAS_MCDISK, BEIJING-LCG2_MCDISK, 98 more...	19	44	2001	174	2240
running	10000334	voatlas73	UA 15.6.9 LCG DQ2 Filestager	2010-07-06 13:11:04	2010-07-07 13:11:04	CA, DE, ES, 7 more...	ALBERTA-LCG2_MCDISK, AUSTRALIA-ATLAS_MCDISK, BEIJING-LCG2_MCDISK, 98 more...	22	44	2073	116	2257

<http://hammercloud.cern.ch/atlas/>

- The HC UI is implemented as a Django web app:
 - View test results
 - View cloud/site evolution
 - DB Admin
- State and results are maintained in MySQL
- HC Logic (job submission, monitoring, resubmission) implemented on top of the *Ganga Grid Programming Interface (GPI)*



- HammerCloud v3 was recently deployed
- What's new?
 - **Test Templates** – standardized tests are templated. Templates are instantiated as a Test
 - **Functional Testing** – automatic instantiation of “functional” Templates at a defined frequency (these are the GangaRobot tests)
 - **Robot Report** – graphical display of site efficiencies for the functional tests
 - **Behind the scenes refactoring** – move to SL5, RPMs, generalizing the code for non-Athena/non-ATLAS tests
 - Plus many small interface changes...

- Currently active Functional Tests:
- UserAnalysis, Athena 15.6.9 with mc0*.merge.AOD.e*_r* on Panda and LCG
 - Data access: Panda “schedconfig”, local direct, FileStager
 - 4 tests in total
- D3PDMaker, Athena 15.6.10.6 (Frontier/Squid test) on Panda and LCG
 - Data access: Panda “schedconfig”, local direct
 - 2 tests in total
- Each tests is set to keep 1 job running at all sites continuously

<< Back

Ganga Robot

All - Panda - LCG - Nordugrid

Panda

CA_PANDA	ANALY_ALBERTA ANALY_SFU ANALY_TORONTO ANALY_TRIUMF ANALY_VICTORIA ANALY_AUSTRALIA ANALY_SFU_bugaboo
DE_PANDA	ANALY_CYF ANALY_FREIBURG ANALY_FZK ANALY_LRZ ANALY_PSNC ANALY_CSCS ANALY_DESY-HH ANALY_GOEGRID ANALY_HEPHY-UIBK ANALY_wuppertalprod ANALY_DESY-ZN ANALY_MPPMU ANALY_UNI-DORTMUND ANALY_FZU
ES_PANDA	ANALY_IFIC ANALY_PIC ANALY_IFAE ANALY_UAM ANALY_LIP-Lisbon ANALY_LIP-Coimbra ANALY_NCG-INGRID-PT
FR_PANDA	ANALY_BEIJING ANALY_CPPM ANALY_GRIF-IRFU ANALY_GRIF-LAL ANALY_GRIF-LPNHE ANALY_LAPP ANALY_LPC ANALY_LPSC ANALY_LYON ANALY_ROMANIA07 ANALY_TOKYO_RFIO ANALY_TOKYO ANALY_ROMANIA02 ANALY_LYON_DCACHE
IT_PANDA	ANALY_INFNO-NAPOLI ANALY_INFNO-MILANO ANALY_INFNO-ROMA1 ANALY_INFNO-FRASCATI ANALY_CNAF ANALY_INFNO-MILANO-ATLASC
NG_PANDA	ANALY_ARC
NL_PANDA	ANALY_Itep ANALY_JINR ANALY_IHEP ANALY_SARA ANALY_TR-10-ULAKBIM ANALY_PNPI ANALY_NIKHEF-ELPROD ANALY_WEIZMANN ANALY_RRC-KI ANALY_CSTCDIE ANALY_IL-TAU-HEP ANALY_TECHNION-HEP
TO_PANDA	ANALY_CERN ANALY_CERN_XROOTD
TW_PANDA	ANALY_TAIWAN ANALY_TW-FTT
UK_PANDA	ANALY_GLASGOW ANALY_LANCS ANALY_QMUL ANALY_RHUL ANALY_BHAM ANALY_CAM ANALY_MANC1 ANALY_UCLCENT ANALY_OX ANALY_RALPP ANALY_SHEF ANALY_MANC2 ANALY_LIV ANALY_MANC ANALY_ECDF ANALY_UCL
US	ANALY_BNL_ATLAS_1 ANALY_NET2 ANALY_AGLT2 ANALY_MWT2 ANALY_SLAC ANALY_SWT2_CPB ANALY_LONG_BNL_ATLAS ANALY_OU_OCHEP_SWT2 ANALY_ANLASC ANALY_MWT2_X

Gangarobot test efficiency
in the last 24 hours


No test <80% 80%<eff<95% >95%

Click to see Ganga blacklist for WMS
(efficiency <80% in the last 6 hours)

Tests used for the efficiency calculation:

LCG: UA 15.6.9 DQ2Local and Filestager - Panda: UA 15.6.9 Default data access and Filestager

- <http://dashb-sam-atlas.cern.ch>



- Email robot report (currently sent to DAST)

Daily GangaRobot report: Wed Jul 7 08:27:53 2010 Inbox | X Lists/ATLAS/DAST Internal | X

☆ from Daniel Colin Van Der Ster <gangarbt@mail.cern.ch>
 sender-time Sent at 7:27 AM (GMT+02:00). Current time there: 1:45 PM.
 to atlas-project-adc-operations-analysis-shifts@cern.ch
 date Wed, Jul 7, 2010 at 7:27 AM
 subject Daily GangaRobot report: Wed Jul 7 08:27:53 2010
 mailing list <atlas-project-adc-operations-analysis-shifts.cern.ch> [Filter messages from this mailing list](#)
 mailed-by cern.ch

Daily GangaRobot report: Wed Jul 7 08:27:53 2010

Sites currently failing jobs by the GangaRobot for Panda in 15.6.9:

ANALY_BEIJING
 ANALY_QMUL

Sites with job submission failures jobs by the GangaRobot for Panda in 15.6.9:

ANALY_NCG-INGRID-PT brokerage gave wrong PandaSiteID:ERROR : No candidate. Release:Atlas-15.6.9 was not found in
 ANALY_LIP-Lisbon brokerage gave wrong PandaSiteID:ERROR : No candidate. Release:Atlas-15.6.9 was not found in [ANALY
 ANALY_UAM brokerage gave wrong PandaSiteID:ERROR : No candidate. [ANALY_UAM] are inactive (no pilots for last 3 ho
 ANALY_TW-FTT brokerage gave wrong PandaSiteID:ERROR : No candidate. [ANALY_TW-FTT] are inactive (no pilots for las
 ANALY_RAL brokerage gave wrong PandaSiteID:ERROR : No candidate. [ANALY_RAL] are inactive (no pilots for last 3 hou
 ANALY_INFN-MILANO brokerage gave wrong PandaSiteID:ERROR : No candidate. Release:Atlas-15.6.9 was not found in [ANALY
 ANALY_CNAF brokerage gave wrong PandaSiteID:ERROR : No candidate. [ANALY_CNAF] are inactive (no pilots for last 3

Sites currently failing the GangaRobot jobs for Panda in [15.6.3.6](#):

Sites with job submission failures jobs by the GangaRobot for Panda in [15.6.3.6](#):

Sites currently failing the GangaRobot jobs for Panda in 15.6.7:

- And available on the web:
http://gangarobot.cern.ch/blac_klist_hammercloud.html

- Panda processingTypes:
 - Functional tests use processingType=gangarobot
 - Stress tests use processingType=hammercloud
- Browse the HC results in the Panda monitor:
 - <http://tiny.cc/panda-gangarobot>
 - <http://tiny.cc/panda-hammercloud>

Analysis job summary, last 24 hours

Analysis Sites	Job Nodes	Jobs	Latest	defined	assigned	waiting	activated	sent	running	holding	transferring	finished	failed tot trf other
ANALY_IHEP	34	2	07-07 11:23	<u>1</u>	<u>0</u>	<u>0</u>	<u>0</u>	<u>0</u>	<u>0</u>	<u>2</u>	<u>0</u>	<u>69</u>	<u>2</u> 3% 0% 3%
ANALY_IL-TAU-HEP	16	0	07-07 11:23	<u>1</u>	<u>0</u>	<u>0</u>	<u>3</u>	<u>0</u>	<u>0</u>	<u>0</u>	<u>0</u>	<u>92</u>	<u>0</u>
ANALY_AUSTRALIA	19	2	07-07 11:12	<u>0</u>	<u>0</u>	<u>0</u>	<u>1</u>	<u>0</u>	<u>0</u>	<u>0</u>	<u>0</u>	<u>113</u>	<u>2</u> 2% 0% 2%
ANALY_MPPMU	64	0	07-07 11:21	<u>0</u>	<u>0</u>	<u>0</u>	<u>0</u>	<u>0</u>	<u>0</u>	<u>1</u>	<u>0</u>	<u>141</u>	<u>0</u>
ANALY_SWT2_CPB	78	1	07-07 11:21	<u>0</u>	<u>0</u>	<u>0</u>	<u>0</u>	<u>0</u>	<u>0</u>	<u>0</u>	<u>0</u>	<u>128</u>	<u>1</u> 1% 0% 1%
ANALY_PIC	58	0	07-07 11:23	<u>1</u>	<u>0</u>	<u>0</u>	<u>3</u>	<u>0</u>	<u>0</u>	<u>0</u>	<u>0</u>	<u>77</u>	<u>0</u>
ANALY_GLASGOW	101	9	07-07 11:14	<u>0</u>	<u>0</u>	<u>0</u>	<u>0</u>	<u>0</u>	<u>0</u>	<u>1</u>	<u>0</u>	<u>133</u>	<u>9</u> 6% 0% 6%

- A number of test templates are ready for site or cloud admins to schedule on-demand...
 - An HC account is needed. Contact us if you want one.
- On Panda or LCG using any data access method (including Panda FileStager or direct access):
 - Muon Analysis, Athena 15.6.6, mc09*merge.AOD*.e*r12*
 - D3PDMaker, Athena 15.6.10.6, data10_7TeV*physics_*AOD*, Frontier/Squid
- Panda Tier 3 Test:
 - Muon Analysis, Athena 15.6.6 – you mail us a list of PFNs

Overall Sites Metrics Incidents

Summary

state	id	host	clouds	start time (CET)	end time (CET)	total jobs
running	10000340	voatlas73	US	2010-07-07 09:03:00	2010-07-07 13:03:00	2346

Input type: PANDA

Output DS: user.elmsheus.hc.10000340.*

Input DS Patterns: mc09*merge.AOD*.e*r12*

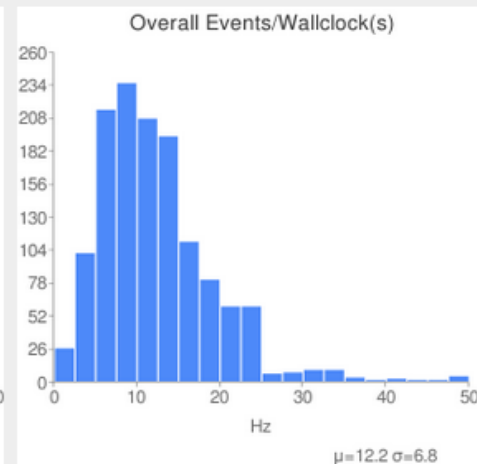

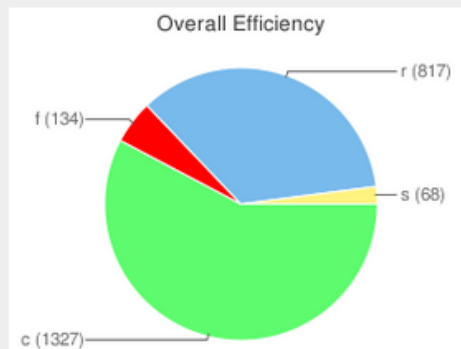
Ganga Job Template: /data/hammercloud/atlas/inputfiles/muon1566/muon1566_panda.tpl

Athena User Area: /data/hammercloud/atlas/inputfiles/muon1566/MuonTriggerAnalysis_1566.tar.gz

Athena Option file: /data/hammercloud/atlas/inputfiles/muon1566/MuonTriggerAnalysis_1566.py

Test Template: 1 (stress) - Muon 15.6.6 PANDA default data-access

[View Test Directory](#) (for debugging)



- **Fix the frontier/squid test... ~25% of the jobs are currently crashing**
 - Is correlated in a non-obvious way to the dESD dataset used – currently replicating a known working dataset to all DE sites (globally later)
- **Integrate with Site Status Board – this is the long-discussed ADC “Analysis Functional Test”**
 - SSB implements the policy for site “exclusion”
 - Will provide better communication to sites in case of exclusion
- **Auto-approval for some test requests**
 - Manual approval isn’t needed if the test is simply an instance of one of the approved Templates
- **Site Ranking Tool**
 - Provides a score (per test template) to make comparisons
- **Robot Web display to present results separated by Test Type (template)**
 - E.g see only the Frontier/Squid test results.

- HammerCloud has taken over responsibilities for all central DA tests: functional and stress testing
- New “template” model makes the usage more user friendly
 - Not much room for error – enables automatic test approval
- Test results are available in a variety of places:
 - HC web, Email reports, SAM Dashboard, Panda Monitor, (SSB coming soon)
- Acknowledgements to the HC team: Johannes Elmsheuser, Federica Legger, Mario Úbeda García