

HammerCloud / Switcher

Jose Caballero <jcaballero@bnl.gov>

What are they? (reminder from the past WBS meeting at BNL last June)

They communicate with AGIS, suggesting to switch a panda queue ON / OFF based on several policies:

- Site admins need to understand the different policies being involved ([a link to the documentation is included in this presentation](#))
- The final decision made by AGIS takes into account:
 - the suggestion from HC
 - the suggestion from Switcher
 - plus any manual setup

- In the case of HammerCloud, policies are based on the performance (failed/succeeded) of a constant stream of testing jobs.
- In the case of Switcher, based on coming scheduled downtimes.

ATLAS Grid Information System

RC Site ATLASite DDMEndpoint PANDA Queue Service Central Services DDM Groups PandaQueue combined resources Docs TWiki OLD

Show 200 entries

First Previous 1 Next Last

give me url of this page *hold shift + click column for Multi-column ordering* VO ATLAS Site PanDA Site Template PanDA Resource PanDA Queue state Final Status Manual HC Switcher type capability rtype CLOUD TIER use newmover DAL deprecate oldmover

atlas																			
VO	ATLAS Site	PanDA Site	Template	PanDA Resource	PanDA Queue	state	Final Status	Manual	HC	Switcher	type	capability	CLOUD	TIER	use newmover	DAL	deprecate oldmover		
atlas	AGLT2	GreatLakesT2	GreatLakesT2_VIRTUAL	ANALY_AGLT2_TEST_SL6-condor	Clone ANALY_AGLT2_TEST_SL6-condor	ACTIVE	TEST	AUTO	TEST	ONLINE	analysis	score	US	T2D	True		tru		

Showing 1 to 1 of 1 entries

ATLAS Grid Information System

RC Site ATLASite DDMEndpoint PANDA Queue Service Central Services DDM Groups PandaQueue combined resources Docs TWiki OLD JSON

Show 200 entries

First Previous 1 Next Last

give me url of this page *hold shift + click column for Multi-column ordering* VO ATLAS Site PanDA Site Template PanDA Resource PanDA Queue state Final Status Manual HC Switcher type capability rtype CLOUD TIER use newmover DAL deprecate oldmover

atlas																			
VO	ATLAS Site	PanDA Site	Template	PanDA Resource	PanDA Queue	state	Final Status	Manual	HC	Switcher	type	capability	CLOUD	TIER	use newmover	DAL	deprecate oldmover		
atlas	CA-SFU-T2	CA-SFU-T2	CA-SFU-T2_VIRTUAL	ANALY_SFU_T2	Clone ANALY_SFU_T2	ACTIVE	BROKEROFF	AUTO	ONLINE	BROKEROFF	analysis	score	CA	T2D	True		false		

Showing 1 to 1 of 1 entries

ATLAS Grid Information System

RC Site ATLASite DDMEndpoint PANDA Queue Service Central Services DDM Groups PandaQueue combined resources Docs TWiki OLD JSON

Show 200 entries

First Previous 1 Next Last

give me url of this page *hold shift + click column for Multi-column ordering* VO ATLAS Site PanDA Site Template PanDA Resource PanDA Queue state Final Status Manual HC Switcher type capability rtype CLOUD TIER use newmover DAL deprecate oldmover

atlas																			
VO	ATLAS Site	PanDA Site	Template	PanDA Resource	PanDA Queue	state	Final Status	Manual	HC	Switcher	type	capability	CLOUD	TIER	use newmover	DAL	deprecate oldmover		
atlas	BNL-ATLAS-OPP	BNL-ATLAS-T3	BNL_VIRTUAL	ANALY_BNL_LOCAL	Clone ANALY_BNL_LOCAL	ACTIVE	OFFLINE	OFFLINE	ONLINE	ONLINE	analysis	score	US	T3	True		true		

Showing 1 to 1 of 1 entries

Useful documentation


Some links to documentation you should always have close at hand when your site is blacklisted

- HammerCloud:
 - Main twiki <https://twiki.cern.ch/twiki/bin/viewauth/IT/HammerCloud>
 - **Twiki for site admins**
<https://twiki.cern.ch/twiki/bin/view/IT/HammerCloudTutorialATLASsiteAdmins>

- Switcher:
 - <https://twiki.cern.ch/twiki/bin/viewauth/AtlasComputing/SwitcherBlacklisting>

Useful documentation

← → ↻ 🏠 🔒 https://twiki.cern.ch/twiki/bin/view/IT/HammerCloudTutorialATLASiteAdmins ☆ 👤

 Jump Search IT All webs

JoseCaballero
Log Out
IT

IT Web
Create New Topic
Index
Search
Changes
Notifications
Statistics
Preferences

Public webs

TWiki > IT Web > HammerCloud > HammerCloudATLASOperations > HammerCloudTutorialATLAS > HammerCloudTutorialATLASiteAdmins (2018-11-26, JoseCaballero) Edit Attach PDF

HammerCloud Tutorial for Site Administrators

- 1 Introduction
 - 1.1 Useful links/Info
 - 2 ATLAS HammerCloud tests in a nutshell
 - 2.1 ATLAS Automatic Site Exclusion
 - 2.1.1 ATLAS Automatic Site Exclusion Policies Overview
 - 2.1.2 Change status and comment of the queue
 - 2.1.3 Monitoring
 - 2.1.4 Integration with AGIS
 - 2.1.4.1 Updating database from AGIS
 - 2.1.4.2 Disable site from AutoExclusion
 - 2.1.5 MCORE queues
 - 2.1.6 HighMem queues
 - 2.1.7 Setting a contact email for Sites
 - 2.1.8 Data Replication for AFTs/PFTs
 - 2.2 Tests other than AFT/PFTs
 - 3 Setting up HC stress tests in a nutshell
 - 3.1 HC admin page
 - 4 HC for ATLAS Shifters
 - 4.1 HC for ATLAS Shifters
 - 5 Q&As
 - 5.1 How do I configure my queue in AGIS?
 - 5.2 Why did my queue not change status even though I got an email notification from HammerCloud?
 - 5.3 Why is my queue not getting jobs from HammerCloud?
 - 5.4 Why is my newly created Unified panda queue (UCORE) not getting jobs from HammerCloud?
 - 5.5 Why has my queue been blacklisted when all HammerCloud jobs ran OK?
 - 5.6 Why my stress test does not run?
 - 5.7 Why my site is not blacklisted if it is failing jobs (either AFTs or PFTs)?
 - 5.8 Why is my site not whitelisted if I see successful gangarbt jobs in the last hour?
 - 5.9 Why didn't I get any email notification from HammerCloud when my PanDA queue was set back to ONLINE?
 - 5.10 How can I check the status of HC tests/jobs/AFT/PFT at my site?
 - 5.11 Why has been my PanDA queue blacklisted by HammerCloud upon migration to CentOS7?
 - 5.12 How can I edit a template?
 - 5.13 Why am I getting "reset online" email notifications every half an hour, but the queue does not change status?
 - 5.14 How do I configure Template files?
 - 5.15 How do I disable a PandaQueue so it does not get more HC jobs?
 - 5.16 How do I kill a running test?

1 Introduction

Troubleshooting pages

There are a couple of pages you should visit when your site has been blacklisted:

- AGIS blacklisting setup page <http://atlas-agis.cern.ch/agis/pandablacklisting/list/>
- AGIS blacklisting history page http://atlas-agis.cern.ch/agis/pandablacklisting/table_view/

ATLAS Grid Information System

RC Site ATLASite DDMEndpoint PANDA Queue Service Central Services DDM Groups

Prioritized list of probes: manual (A) switcher (A) hammercloud (A) hctest1 (D)

State values: OFFLINE TEST BROKEROFF PAUSED ONLINE

Experiment Site	PandaQueue	Status default
AGLT2	ANALY_AGLT2_TEST_SL6-condor	hammercloud
ANLASC	ALCF_Cooley	manual
	ALCF_Theta	manual
	ALCF_Theta_ES	manual
ARNES	ANALY_ARNES_DIRECT	hammercloud
	ARNES	manual
Australia-ATLAS	ANALY_AUSTRALIA_TEST	manual
	Australia-ATLAS_MCORE	manual
	Australia-ATLAS_MCORE_TEST	manual
	Australia-ATLAS_TEST	manual
BEIJING-LCG2	BEIJING-CS-TH-1A_MCORE	manual
	BEIJING-ERAIL_MCORE	manual
	BEIJING-TIANJIN-TH-1A_MCORE	manual
BNL-ATLAS	ANALY_BNL_Test_2_CE_1-condor	hammercloud
	BNL_PROD_CONTR_TEST-condor	hammercloud
	BNL_Test_2_CE_1-condor	hammercloud
	ES_ORNL_Titan	manual
	HARVESTER_AFF_BNL_TEST	hammercloud

ATLAS Grid Information System

RC Site ATLASite DDMEndpoint PANDA Queue Service Central Services DDM Groups

Show 200 entries First Previous 1 2 3 4 Next Last

give me url of this page **hold shift + click column for Multi-column ordering**

Panda Queue	Probe	Value	Reason	Expiration	DN
BNL_LOCAL-condor	hammercloud	ONLINE	HC.Esblacklist.set.online	Never	/DC=ch/DC=
BNL_PROD-condor	hammercloud	ONLINE	HC.Blacklist.set.online.to.mitigate.massive.blacklisting.event	Never	/DC=ch/DC=
BNL_ATLAS_2-condor	hammercloud	ONLINE	HC.Blacklist.set.online.to.mitigate.massive.blacklisting.event	Never	/DC=ch/DC=
ANALY_BNL_MCORE-condor	hammercloud	ONLINE	HC.Blacklist.set.online.to.mitigate.massive.blacklisting.event	Never	/DC=ch/DC=
ANALY_BNL_SHORT-condor	hammercloud	ONLINE	HC.Blacklist.set.online.to.mitigate.massive.blacklisting.event	Never	/DC=ch/DC=
ANALY_BNL_LONG-condor	hammercloud	ONLINE	HC.Blacklist.set.online.to.mitigate.massive.blacklisting.event	Never	/DC=ch/DC=
BNL_LOCAL-condor	hammercloud	TEST	HC.Esblacklist.set.test	Never	/DC=ch/DC=
BNL_PROD-condor	hammercloud	TEST	HC.Blacklist.set.test	Never	/DC=ch/DC=
BNL_ATLAS_2-condor	hammercloud	TEST	HC.Blacklist.set.test	Never	/DC=ch/DC=
ANALY_BNL_MCORE-condor	hammercloud	TEST	HC.Blacklist.set.test	Never	/DC=ch/DC=
ANALY_BNL_SHORT-condor	hammercloud	TEST	HC.Blacklist.set.test	Never	/DC=ch/DC=
ANALY_BNL_LONG-condor	hammercloud	TEST	HC.Blacklist.set.test	Never	/DC=ch/DC=
BNL_LOCAL-condor	hammercloud	ONLINE	HC.Esblacklist.set.online	Never	/DC=ch/DC=
BNL_LOCAL-condor	hammercloud	TEST	HC.Esblacklist.set.test	Never	/DC=ch/DC=
ANALY_BNL_SHORT-condor	hammercloud	ONLINE	HC.Blacklist.set.online	Never	/DC=ch/DC=
ANALY_BNL_SHORT-condor	hammercloud	TEST	HC.Blacklist.set.test	Never	/DC=ch/DC=
ANALY_BNL_SHORT-condor	hammercloud	ONLINE	HC.Blacklist.set.online	Never	/DC=ch/DC=
ANALY_BNL_MCORE-condor	hammercloud	ONLINE	HC.Blacklist.set.online	Never	/DC=ch/DC=
ANALY_BNL_LONG-condor	hammercloud	ONLINE	HC.Blacklist.set.online	Never	/DC=ch/DC=
ANALY_BNL_SHORT-condor	hammercloud	TEST	HC.Blacklist.set.test	Never	/DC=ch/DC=
ANALY_BNL_MCORE-condor	hammercloud	TEST	HC.Blacklist.set.test	Never	/DC=ch/DC=

To fix broken queues in AGIS

Sometimes PanDA Queues in AGIS don't have the write combination of parameters to allow HammerCloud to behave correctly. Not always site's fault, it could be a change in the Workload Management Workflow that suddenly requires a change in the setup.

To validate, and correct when needed, your queues, this table can be extremely handy:

https://hc-ai-core.cern.ch/testdirs/atlas/check_agis_config.html

SUMMARY as of 2023-11-29T15:00:19.002801

SUMMARY of PanDA queues possibly mis-configured for blacklisting

Cloud	PanDA queue	Change this or a different PQ?	Which PanDA queue to change?	What change?	URL
-------	-------------	--------------------------------	------------------------------	--------------	-----

SUMMARY of PanDA queues possibly mis-configured for testing

Cloud	PanDA queue	Perform this change	Current hc_suite	Desired hc_suite	Capability in AGIS	URL
CERN	ANALY_CERN-PROD_DATALAKES_TESTA	add AFT		AFT	score	https://atlas-agis-api.cern.ch/request/pandaqueue/update/set_attributes/?name=ANALY_CERN-PROD_DATALAKES_TESTA&hc_suites=1
CERN	ANALY_CERN-PROD_DATALAKES_TESTA_PNPI	add AFT		AFT	score	https://atlas-agis-api.cern.ch/request/pandaqueue/update/set_attributes/?name=ANALY_CERN-PROD_DATALAKES_TESTA_PNPI&hc_suites=1
CERN	ANALY_CERN-PROD_DATALAKES_TESTA_RU	add AFT		AFT	score	https://atlas-agis-api.cern.ch/request/pandaqueue/update/set_attributes/?name=ANALY_CERN-PROD_DATALAKES_TESTA_RU&hc_suites=1
CERN	ANALY_CERN-PROD_DATALAKES_TESTA_RU_RAIN6	add AFT		AFT	score	https://atlas-agis-api.cern.ch/request/pandaqueue/update/set_attributes/?name=ANALY_CERN-PROD_DATALAKES_TESTA_RU_RAIN6&hc_suites=1
CERN	ANALY_CERN-PROD_DATALAKES_TESTB	add AFT		AFT	score	https://atlas-agis-api.cern.ch/request/pandaqueue/update/set_attributes/?name=ANALY_CERN-PROD_DATALAKES_TESTB&hc_suites=1
CERN	ANALY_CERN-PROD_DATALAKES_TESTB_PNPI	add AFT		AFT	score	https://atlas-agis-api.cern.ch/request/pandaqueue/update/set_attributes/?name=ANALY_CERN-PROD_DATALAKES_TESTB_PNPI&hc_suites=1
CERN	ANALY_CERN-PROD_DATALAKES_TESTB_RU	add AFT		AFT	score	https://atlas-agis-api.cern.ch/request/pandaqueue/update/set_attributes/?name=ANALY_CERN-PROD_DATALAKES_TESTB_RU&hc_suites=1
CERN	ANALY_CERN-PROD_DATALAKES_TESTB_RU_RAIN6	add AFT		AFT	score	https://atlas-agis-api.cern.ch/request/pandaqueue/update/set_attributes/?name=ANALY_CERN-PROD_DATALAKES_TESTB_RU_RAIN6&hc_suites=1
CERN	ANALY_CERN-PROD_DATALAKES_TESTC	add AFT		AFT	score	https://atlas-agis-api.cern.ch/request/pandaqueue/update/set_attributes/?name=ANALY_CERN-PROD_DATALAKES_TESTC&hc_suites=1
CERN	ANALY_CERN-PROD_DATALAKES_TESTC_PNPI	add AFT		AFT	score	https://atlas-agis-api.cern.ch/request/pandaqueue/update/set_attributes/?name=ANALY_CERN-PROD_DATALAKES_TESTC_PNPI&hc_suites=1
CERN	ANALY_CERN-PROD_DATALAKES_TESTC_RU	add AFT		AFT	score	https://atlas-agis-api.cern.ch/request/pandaqueue/update/set_attributes/?name=ANALY_CERN-PROD_DATALAKES_TESTC_RU&hc_suites=1
CERN	ANALY_CERN-PROD_DATALAKES_TESTC_RU_RAIN6	add AFT		AFT	score	https://atlas-agis-api.cern.ch/request/pandaqueue/update/set_attributes/?name=ANALY_CERN-PROD_DATALAKES_TESTC_RU_RAIN6&hc_suites=1
CERN	ANALY_CERN-PROD_DATALAKES_TESTD	add AFT		AFT	score	https://atlas-agis-api.cern.ch/request/pandaqueue/update/set_attributes/?name=ANALY_CERN-PROD_DATALAKES_TESTD&hc_suites=1
CERN	ANALY_CERN-PROD_DATALAKES_TESTD_PNPI	add AFT		AFT	score	https://atlas-agis-api.cern.ch/request/pandaqueue/update/set_attributes/?name=ANALY_CERN-PROD_DATALAKES_TESTD_PNPI&hc_suites=1
CERN	ANALY_CERN-PROD_DATALAKES_TESTD_RU	add AFT		AFT	score	https://atlas-agis-api.cern.ch/request/pandaqueue/update/set_attributes/?name=ANALY_CERN-PROD_DATALAKES_TESTD_RU&hc_suites=1
CERN	ANALY_CERN-PROD_DATALAKES_TESTD_RU_RAIN6	add AFT		AFT	score	https://atlas-agis-api.cern.ch/request/pandaqueue/update/set_attributes/?name=ANALY_CERN-PROD_DATALAKES_TESTD_RU_RAIN6&hc_suites=1
CERN	CERN-PROD_EOS	add PFT		PFT	ucore	https://atlas-agis-api.cern.ch/request/pandaqueue/update/set_attributes/?name=CERN-PROD_EOS&hc_suites=2

Cloud	PanDA queue	Perform this change	Current hc_suite	Desired hc_suite	Capability in AGIS	URL
ES	UAM-LCG2_CREAM_HTCONDOR_MCORE	add PET	PET_MCORE	PET_MCORE	ucore	https://atlas-agis-api.cern.ch/request/pandaqueue/update/set_attributes/?name=UAM-LCG2_CREAM_HTCONDOR_MCORE&hc_suites=32

New features recently implemented in HammerCloud

I finished the implementation of the Blacklisting features for EventService:

- ES queue blacklisted based on number of failed jobs
- ES queue blacklisted based on number of processed events (many jobs finished but very few events processed)
- ES queue blacklisted if it is already blacklisted for Production Functional Tests (PFTs)

Queues recovery:

We have seen that, sometimes, when a person modifies manually the value of some AGIS parameters while the queue was ES Blacklisted, the queue may end up in an undefined status and HC does not know what to do anymore.

I implemented a recovery for the EventService queues that are left in those "undefined" statuses.

Custom tests for US ATLAS?

The people who has access to the HammerCloud admin page (I do) can create custom templates and trigger the submission to test jobs on demand. These jobs will run in parallel to the centrally scheduled jobs.

If we think this infrastructure can be useful for some US ATLAS specific testing campaign, in principle it is technically possible.

New (coming soon) features in Switcher

After a whole refactoring, the new version of the service will include:

- Disabling individual CEs in AGIS, instead of the whole Queue, when there are other CEs that can keep the queue running.
 - Only when all CEs are disabled at the same time the queue will also blacklisted.
 - Email notifications will be more verbose in those cases, but will include all information.
- New policies to decided when to blacklist a queue based on a DDM downtime (next slide)
- Only source of information will be AGIS. Currently is AGIS, SiteStatusBoard (SSB), GOCDB, ...
- [Possibly] the DDM Blacklisting, currently done inside AGIS, will be integrated into Switcher

New (coming soon) features in Switcher

pandaqueue get assigned DDM endpoint

AGIS pandaqueues (a.k.a. schedconfig)

AGIS ddmendpoints

```
{
  "AGLT2_LMEM_SL7-condor": {
    ...
    "astorages": {
      ...
      "write_lan": [
        "AGLT2_DATADISK",
        "AGLT2_CALIBDISK",
        "AGLT2_LOCALGROUPDISK",
        "AGLT2_SCRATCHDISK"
      ]
    }
  },
  ...
}
```

```
"AGLT2_DATADISK": {
  ...
  "arprotocols": {
    ...
    "write_wan": [
      {
        "endpoint": "srm://head01.aglt2.org:8443/srm/managerv2?SFN=",
        "flavour": "SRMv2",
        "id": 256,
        "path": "/pnfs/aglt2.org/atlasdatadisk/rucio/"
      },
      {
        "endpoint": "davs://head01.aglt2.org:2880",
        "flavour": "WEBDAV",
        "id": 257,
        "path": "/atlasdatadisk/rucio/"
      }
    ]
  }
  ...
}
```

1. check the "write_lan" activity in the "astorages" block
2. pick the first spacetoken (the one with higher priority)
3. search for that spacetoken in the DDM endpoints view, SRC2
4. in block "arprotocols", search for protocol "write_lan" activity if it exists, otherwise use "write_wan"
5. first endpoint in that protocol (higher priority), irrespective its flavor, will be associated to that panda queue
6. turn the queue OFF if that endpoint is in downtime

Communication

For questions, when the documentation does not help:

HammerCloud: atlas-adc-hammercloud-support@cern.ch

Switcher: atlas-adc-autoexclusion@cern.ch

For feedback / feature requests / complaints ...

