

Table of Contents

Tier2ProductionSiteConfiguration.....	1
Introduction.....	1
CE and software installation.....	1
Step 1: Publish CE information to BDII.....	1
Step 2: Get automatic software installations.....	1
Storage and squid issues.....	1
Step 3: Setup SRMv2.2 and configure the required space tokens.....	1
Step 4: Publish space token information to BDII.....	2
Step 5: Set up a Squid.....	2
Step 6: Provide the following information to squad support (Space Tokens and Squid configuration).....	2
Step 7: Ask for FTS channel setup for your site.....	2
Step 8: Ask for site information update for data management.....	2
Step 9: Perform data transfer functional test.....	3
Infrastructure validation.....	3
Addition of New Queue in PanDA.....	3
Run Grid jobs.....	3
Step 10: Analysis activity : Ask for site information update for receiving PanDA analysis jobs ().....	3
Step 11: Production activity : Ask for site information update for receiving production jobs.....	4
Shares between production and analysis queue.....	4
Pledge and beyond shares (detailed).....	4
Shares for 2011.....	5
Shares for 2012.....	5
Useful links.....	5

Tier2ProductionSiteConfiguration

Introduction

This is a step-by-step instruction for Tier-2 and Tier-3 managers to bring their site into the ATLAS production activity. All these actions should be done in close contact with your cloud squad. Please make sure you have the answers of the following questions before going through the steps.

- **Which site is your Tier-1?**
- **Who is your site manager?**
- **Who is your Tier-1 site manager?**

A Tier2 should participate to production and Grid analysis activity. As a consequence, it should go through all steps

A Grid Tier3 (classification defined here) can choose to run:

- production and Grid analysis : All steps
- production only : All steps except
- Grid analysis only (?)

A site without CPU accessible on the Grid should go through step 6 and 7. No space token is requested. The site should be tested only through DDM FT (see step 9)

CE and software installation

Step 1: Publish CE information to BDII

Site responsibility : Site has to check if the CE information is properly published to the BDII. One could check it by using the `ldapsearch` command from a LCG UI.

Step 2: Get automatic software installations

To be done after step 1. is completed.

Site or squad resp. : Once the CE is in the BDII and the ATLAS VO is supported, automatic software installation should start. You can check the status from the installation portal.

Storage and squid issues

Step 3: Setup SRMv2.2 and configure the required space tokens

Independent of step 1 and 2

Site responsibility : Site has the responsibility to setup SRMv2.2 services and configure the space tokens needed by ATLAS production. For the up-to-date space token requirement, please refer to the ATLAS space token specifications and requirements. Please follow the request for acl setup which depends on storage technology.

Step 4: Publish space token information to BDII

To be done after step 3. is completed.

Site responsibility : Site has to check if the space token information is properly published to the BDII. One could check it by using the `ldapsearch` command from a LCG UI. The following commands are the examples for checking ATLASDATADISK space token.

```
ldapsearch -x -h topbdii.grif.fr -p2170 -b o=grid "(&(objectClass=GlueVOInfo)(GlueVOInfoLocalID=a
```

Step 5: Set up a Squid

Independent of previous steps

Tier-2s and large or distant Tier-3s should set up a local Squid; smaller Tier-3s can agree with the nearest Tier-2 to use theirs. Also the fail-over Squid and the default and fail-over Frontier servers must be defined.

Site resp.: The Squid is used for access to the database resident information. The load is small so one machine is sufficient in most of the cases (up to 500 cores dedicated to analysis jobs). This is an action also on the sites hosting the fail-over Squid and the Frontier servers, as most likely they will have to open access to the IP range of the sites WNs and UIs. Installation instructions are here.

Step 6: Provide the following information to squad support (Space Tokens and Squid configuration)

To be done after all previous steps 3-5 have been completed

Site responsibility: The list of of requested informations is available at this address

Remark: Please keep your cloud squad informed in case of site service update (e.g. endpoint or contact person change).

Step 7: Ask for FTS channel setup for your site

To be done after step 6 but within few days to ensure that declared DDM endpoints will not be idle.

squad support manager resp.. The FTS service used for serving data transfer to/from your site is provided by your Tier-1. The FTS administrator at Tier-1 needs to configure (or add) corresponding FTS channels according to the information provided in Step 6. The squad support will provide the information to the Tier-1.

Step 8: Ask for site information update for data management

squad support and DDM Ops team resp.. : Your site will be configured (or added) into the TiersOfAtlasCache according to the information you provided in Step 6. The cloud squad should open a DDM Ops Savannah ticket and fill all informations to update TiersOfAtlasCache.

After update of TiersOfAtlasCache:

- **DDM Ops resp.**
 - ◆ Include DDM endpoint in SiteService and DeletionService
 - ◆ Validate the transfer and deletion steps with one dataset
- DDM endpoints will appear in Datri after 24 hours.

As soon as the site is running production, the DBRelease and condition files will be transferred to HOTDISK.

Step 9: Perform data transfer functional test

To be done after step 8

- **DDM Ops resp.**
 - ◆ Include DDM endpoint in SiteService and DeleteService
 - ◆ Validate the transfer and deletion steps with one dataset
 - ◆ Request the site to be included in DDM FT (T1->site) and Sonar FT (To measure the fastest path for data transfer)
- DDM endpoints will appear in Datri after 24 hours.

Infrastructure validation

The site should demonstrate its stability over 10 days for the following tests:

- DDM FT : Test storage and connectivity stability
- SAM test : Test CE and storage stability

Addition of New Queue in PanDA

Information about PanDA queues are available in schedconfig. Schedconfig contains subdirectories corresponding to existing clouds. Each of these clouds contain a series of sites, and each site contains a number of python files. Each file contains the specifications for one queue, and is named with the queue's name. Changes made to this file will reflect in the database after they have been checked in and harvested by the update code. 'All.py' file gets automatically created when one commit these codes while creation of new site in schedconfig.

Ask your squad to update schedconfig (more information about addition of new queue in PanDA can be found at SchedConfigNewController) or escalate to Schedconfig support.

In addition, the site should be included in pilot factory (documentation to be found)

Run Grid jobs

Step 10: Analysis activity : Ask for site information update for receiving PanDA analysis jobs ()

cloud squad and the ADC operation resp: Your site will be added into the PanDA database in order to make your site available for pathena/prun analysis jobs. Once the configuration is done, your site should appear in the PanDA Cloud Monitor page within your cloud. The cloud squad is also responsible to contact the responsible of the pilot factory to include the site. The analysis queue should be set 'brokeroff' (only jobs forced to go at the site will be run at the site).

For 10 days, the site will have to run GangaRobot jobs with a success rate > 95% . Cloud squad should request to preplace input datasets used as input for GangaRobot.

When the site is validated, the analysis queue is set to 'online'

When the panda analysis queue is setup and in production, the site will receive data through PD2P

Step 11: Production activity : Ask for site information update for receiving production jobs

cloud squad and the ADC operation team: Your site will be added into the PanDA database in order to load your site with production jobs. Once the configuration is done, your site should appear in the PanDA Cloud Monitor page within your cloud. The cloud squad is also responsible to contact the responsible of the pilot factory to include the site.

To validate the site, the production queue should be setup in 'test' mode. Few jobs are run. If the jobs are successful, the production queue is set online.

Shares between production and analysis queue

Pledge and beyond shares (detailed)

The next chapters explain the values to be implemented for batch queues associated to ATLAS. It includes details about implementation of the 'beyond pledge' mechanism in Panda. The target CPU resources means that, if the site is fully busy with all possible activities, the associated activity will get this amount of resources. It is assumed that

Let's define (all values in HSPEC06) as:

- H_{tot} : total CPU resources potentially managed by the site batch cluster
- H_{na} : Target resources for non ATLAS Panda jobs (local ATLAS jobs or local users or other VOs)
- H_{ap} : Target resources for Panda analysis within pledge (atlas/Role=pilot)
- H_{ab} : Target resources for Panda analysis beyond pledge (atlas/Role=pilot)
- H_{pp} : Target resources for Panda production within pledge (atlas/Role=production)
- H_{pb} : Target resources for Panda production beyond pledge (atlas/Role=production)

Most of the sites who did not implement 'beyond pledge' Panda mechanism have pilots associated to the country.

- H_{np} : Target resources for ATLAS private panda pilots (usually atlas/<country>/Role=pilot)

One should get : $H_{tot} = H_{na} + H_{np} + H_{ap} + H_{ab} + H_{pp} + H_{pb}$

The shares requested by ATLAS apply to the relative ratio between H_{ap} and H_{pp} . When the site has defined the 'beyond pledge resources' (H_{ab} and H_{pb}), it can define the shares for the 3 pilots as:

- atlas/<country>/Role=pilot : $(H_{np}) / H_{tot}$
- atlas/Role=pilot : $(H_{ap} + H_{ab}) / H_{tot}$
- atlas/Role=production : $(H_{pp} + H_{pb}) / H_{tot}$

Within Panda schedconfig, the relative ratio between pledge and beyond pledge are defined by:

- availableCPU : $H_{ap} + H_{ab}$
- pledgeCPU : H_{ap}

The weight will be multiplied by available/pledge in the brokerage when countryGroup is matched.

For the moment, the beyond pledge resources is monitored in Panda through *Special Handling*='localpool' . It is not yet separated in Historical View.

For T3 sites without pledge resources but accepting general ATLAS analysis/production jobs when idle , `pledgeCPU` should be set to 0. If `pledgeCPU = 0`, all the mtached countryGroup jobs will start first.

Shares for 2011

ATLAS requests a share between the production and analysis queues within pledge ressources:

- T2 : 50/50 %
- T1 : Was 20 for analysis and 80% for production

For the non-pledge CPUs, it is up to the site to define the share between production and analysis

Shares for 2012

ATLAS requests a share between the production and analysis queues within pledge ressources:

- T2 : 50/50 %
- T1 : 20% for analysis and 80% for production

For the non-pledge CPUs, it is up to the site to define the share between production and analysis

Useful links

- ATLAS dashboard to monitor the data transfer status to/from your site.
- PanDA Monitor to monitor the production jobs on your site.
- ADC Operation Mailing List

Major updates:

-- HurngChunLee - 30 Apr 2008 -- DarioBarberis - 31-Mar-2011

Responsible: StephaneJezequel

Last reviewed by: **Never reviewed**

This topic: Atlas > Tier2ProductionSiteConfiguration

Topic revision: r19 - 15-Aug-2012 - UedaI



Copyright &© by the contributing authors. All material on this collaboration platform is the property of the contributing authors.

Ideas, requests, problems regarding TWiki? Send feedback