

# Sites compute and storage description

---

Some clean-up is needed...

# **This is not a sexy talk**

This is about the legacy stuff we have around and makes our life difficult

# Many site attributes “leftover”

- in the next slides: list few key places where we have a lot of things that are “maybe used - maybe not used anymore but we have still that small bit that need fixing”

# Why?

- Why? Everybody is talking about simplification, consolidation, automation, etc...
  - But how can we do this if it's taking several experts days to explain how our system works?
- Lightweight sites??
  - No way! With the examples above bootstrapping a site is complicated
  - Experts always needed, not clear what happens if someone touch something.
- Whenever we make a change is supercomplex to understand the impact

# DDM endpoints

## DDM endpoint info

Name:	AGLT2_DATADISK
Type:	DATADISK
SRM:	token:ATLASDATADISK:srm://head01.aglt2.org:8443/srm/manager2?SFN=/pnfs/aglt2.org/atlasdatadisk/
Token:	ATLASDATADISK
Phys Group:	
Is Rucio enabled:	No
Domain:	.*aglt2.org.*@atlasdatadisk/.*
mkdir:	No
Is cache:	No
Is Deterministic:	Yes
Is Volatile:	No
Space method:	other
Space Usage:	gsiftp://dcdum02.aglt2.org:2811/pnfs/aglt2.org/atlasscratchdisk/space-usage.json
Tape:	No
Pledged:	No
Tool Assigner::	lcg
LFC:	CERN-PROD_RUCIO_Catalog
Site:	AGLT2
ATLAS Site:	AGLT2
SE info:	
<b>Resource:</b>	<b>New Storage Relation: NULL (NULL)</b>
Storage element:	AGLT2-SRM-head01.aglt2.org (srm://head01.aglt2.org:8443/srm/manager2?SFN=)
Default FTS:	BNLFTS3/FTS (BNL-ATLAS)
Tier1s FTS:	None
CERN FTS:	None
REST FTS master1:	BNLFTS3_REST/FTS (BNL-ATLAS)
REST FTS master2:	CERNFTS3_REST/FTS (CERN-PROD)
REST FTS master3:	RALFTS3PILOT_REST/FTS (RAL-LCG2)
REST FTS test:	CERNFTS3PILOT_REST/FTS (CERN-PROD)
Permission Group:	defaultPermission
Permissions	<b>Role:</b> /atlas/Role=production <b>Permissions</b> frwpuda <b>Role:</b> /atlas/Role=NULL <b>Permissions</b> fr---- <a href="#">Change permissions</a> <a href="#">Create new Permission group for this endpoint</a>
DDM Groups:	ALL TIER2DS UMICH USASITES USTIER2DS USTIER2S
Data Policy:	TODisk, Nucleus
State:	ACTIVE
State comment:	
State update:	None

Operations:



[Clone DDM Endpoint](#)



[Update DDM Endpoint information](#)



[Show Changes log](#)

# PandaQueue 1

atlas-agis.cern.ch/agis/pandaqueue/detail/ANALY\_CERN\_TO\_SHORT/full/ 80% Search ☆ 📁 ✓ ↓ 🏠 ABP

## SchedConfig parameters

Final value of parameter marked with green color

Parameter	Default Value from parent PQ (CERN-PROD_VIRTUAL)	Ovewritten Value (ANALY_CERN_TO_SHORT)
accesscontrol:		
<b>allowdirectaccess:</b> (deprecated)	True	
allowednode:		
allow FAX:	False	True
allow JEM:	False	False
appdir:	/cvmfs/atlas.cern.ch/repo/sw/nightlies^/cvmfs/atlas-nightlies.cern.ch/repo/sw/nightlies	/cvmfs/atlas.cern.ch/repo/sw/nightlies^/cvmfs/atlas-nightlies.cern.ch/repo/sw/nightlies
autosetup post:		
autosetup pre:		
availablecpu:	None	None
cachedse:	None	None
capability:	score	
catchall:		
cmd: (deprecated)		
cmdconfig: (deprecated)		
copyprefix: (deprecated)	srm://srm-eosatlas.cern.ch/^root:/eosatlas.cern.ch/	srm://srm-eosatlas.cern.ch/^root:/eosatlas.cern.ch/
copyprefixin: (deprecated)	srm://srm-eosatlas.cern.ch,srm://srm-atlas.cern.ch/^root:/eosatlas.cern.ch,/root://castor.atlas.cern.ch/	
copyprefixin fax direct: (deprecated)		srm://srm-eosatlas.cern.ch/^root:/atlas-xrd-eos-rucio.cern.ch:1094
copyprefixin fax xrdcp: (deprecated)		
copysetup: (deprecated)	/cvmfs/atlas.cern.ch/repo/sw/local/xrootdsetup.sh	/cvmfs/atlas.cern.ch/repo/sw/local/xrootdsetup.sh^False^False
copysetup fax direct: (deprecated)		/cvmfs/atlas.cern.ch/repo/sw/local/xrootdsetup.sh^False^True
copysetup fax xrdcp: (deprecated)		
copysetupin: (deprecated)	/cvmfs/atlas.cern.ch/repo/sw/local/xrootdsetup.sh^srm://srm-eosatlas.cern.ch/^root:/eosatlas.cern.ch/^False^False	/cvmfs/atlas.cern.ch/repo/sw/local/xrootdsetup.sh^False^True
copysetupin fax direct: (deprecated)		
copysetupin fax xrdcp: (deprecated)		/cvmfs/atlas.cern.ch/repo/sw/local/xrootdsetup.sh^False^False
copytool: (deprecated)	xrdcp	xrdcp
copytoolin: (deprecated)	xrdcp^dummy	xrdcp^dummy
corecount:	None	None

# PandaQueue 2

corepower:	9.86	0-0
countrygroup:		
<b>datadir:</b> (deprecated)		
ddm:	CERN-PROD_DATADISK,CERN-PROD_DATATAPE,CERN-PROD_MCTAPE	CERN-PROD_SCRATCHDISK,CERN-PROD_DATADISK
deprecate_oldmover:	False	
depthboost:	None	None
description:	Short analysis queue for ATLAS using Tier0 resources	
direct access lan:	False	True
direct access wan:	True	True
dq2url:		
email:		
environ:		DATAPATH=/afs/cern.ch/atlas/conditions/poolcond/catalogue
envsetup:		
envsetuppin:	export XrdSecSSLUSERCERT=\$X509_USER_PROXY;export XrdSecSSLUSERKEY=\$X509_USER_PROXY	export XrdSecSSLUSERCERT=\$X509_USER_PROXY;export XrdSecSSLUSERKEY=\$X509_USER_PROXY
fairsharepolicy:		priority>6000:100
faxredirector:	None	atlas-xrd-eu.cern.ch:1094   XROOTD XROOTD-atlas-xrd-us/Redirector (BNL-ATLAS)
<b>gatekeeper:</b> (deprecated)	to.be.set	
glxexec:		
globusadd:		
hc_param:	OnlyTest	AutoExclusion
ignore swreleases:	False	False
CVMFS:	True	True
is default:	False	
is virtual:	False	
<b>jdk:</b> (deprecated)		
<b>jdkadd:</b> (deprecated)		
jobmanager:	lcglsf	lcglsf
last modified:	Sept. 21, 2017, 3:32 p.m.	
lfc host:	prod-lfc-atlas.cern.ch	prod-lfc-atlas.cern.ch
lfc path:	/grid/atlas/users/pathena	/grid/atlas/users/pathena
lfcprodpath:	/grid/atlas/dq2	/grid/atlas/dq2
lfcregister:	server	server
localqueue:	grid_atlas	grid_atlas
maxinputsize:	14336	14336
maxmemory:	6000	0

# PandaQueue 3

maxrss:	2000	3200
maxswap:	0	0
maxtime:	0	259200
maxwdir:	16336	19000
minmemory:	0	0
minrss:	0	0
mintime:	0	0
name2: (deprecated)	default	
nodes:	0	0
objectstore:		
panda site:	CERN-PROD	
pilot manager:	APF	APF
pilot version:	current	
pilotlimit:	None	None
pledgedcpu:	None	None
proxy: (deprecated)		
Pilot test:	False	False
python path:		
queue: (deprecated)		
queuehours:	0	0
recoverdir:		
region:	CERN	CERN
releases:		
resource type:	GRID	GRID
retry:	False	False
se: (deprecated)	token:ATLASSCRATCHDISKsrm://srm-eosatlas.cern.ch:8443/srm/v2/server?SFN=	token:ATLASSCRATCHDISKsrm://srm-eosatlas.cern.ch:8443/srm/v2/server?SFN=
sein: (deprecated)		
seinopt: (deprecated)		
seopt: (deprecated)	token:ATLASDATADISKsrm://srm-eosatlas.cern.ch:8443/srm/v2/server?SFN=,token:ATLASDATATAPEsrm://srm-atlas.cern.ch:8443/srm/manager?SFN=,token:ATLASMCTAPEsrm://srm-atlas.cern.ch:8443/srm/manager?SFN=	token:ATLASSCRATCHDISKsrm://srm-eosatlas.cern.ch:8443/srm/v2/server?SFN=,token:ATLASDATADISKsrm://srm-eosatlas.cern.ch:8443/srm/v2/server?SFN=,token:ATLASDATATAPEsrm://srm-atlas.cern.ch:8443/srm/manager?SFN=,token:ATLASMCTAPEsrm://srm-atlas.cern.ch:8443/srm/manager?SFN=
sepath: (deprecated)	/eos/atlas/atlasscratchdisk/ruccio	/eos/atlas/atlasscratchdisk/ruccio
seproddir: (deprecated)	/eos/atlas/atlasscratchdisk/ruccio	/eos/atlas/atlasscratchdisk/ruccio
setokens: (deprecated)	ATLASSCRATCHDISK,ATLASDATADISK,ATLASDATATAPE,ATLASMCTAPE	ATLASSCRATCHDISK,ATLASDATADISK,ATLASDATATAPE,ATLASMCTAPE
special par:		
stageinretry:	2	2
stageoutretry:	2	2
AGIS internal state:	ACTIVE	
State comment:	Object was cloned from ANALY_CERN_SHORT via WebUI	
State modification time:	March 29, 2017, 3:51 p.m.	
system:	lcg-cg	lcg-cg
tags:		
timefloor:	None	60
tmpdir:		



# PandaQueue 4

transferringlimit:	None	None
type:	production	analysis
use_newmover:	True	
validatedreleases:	14.5.0 14.5.2 15.2.0 15.3.1	14.5.0 14.5.2 15.2.0 15.3.1
wansinklimit:	None	None
wansourcelimit:	None	None
wnconnectivity:	full	
wntmpdir:		

# Attributes to the Pilot

- SchedConfig is under evolution
  - The “configurator”
- Still the pilot needs several other attributes which are not in SchedConfig
  - E.g. the storage protocols
- For now using CVMFS
  - Plus retries directly on AGIS if CVMFS not available
  - Couple of hours minimum refresh time, i.e. change a value you’ve to wait to see if the change is ok.
- If we need faster turnaround we should think:
  - Several options, just listing 2:
    - having these infos in Panda which can serve them to the pilot (as today panda does for the “queuedata”)
    - Query directly AGIS

# ToACache

- Is anyone still using ToACache????



```
-bash-4.1$ ls -la /afs/cern.ch/atlas/www/GROUPS/DATABASE/project/ddm/releases/Tier*
-rw-----, 1 15861 zp 56427 Jul 27 2006 /afs/cern.ch/atlas/www/GROUPS/DATABASE/project/ddm/releases/TiersOfATLASCache_new.py
-rw-----, 1 15861 zp 21662 Jul 27 2006 /afs/cern.ch/atlas/www/GROUPS/DATABASE/project/ddm/releases/TiersOfATLASCache_new.pyc
-rw-----, 1 15861 zp 54805 Jul 13 2006 /afs/cern.ch/atlas/www/GROUPS/DATABASE/project/ddm/releases/TiersOfATLASCache_new.py_with_toolAssigners
-rw-----, 1 1702 zp 29714 Jun 23 2006 /afs/cern.ch/atlas/www/GROUPS/DATABASE/project/ddm/releases/TiersOfATLASCache_old.py
lrwxr-xr-x, 1 iueda zp 25 Sep 24 2012 /afs/cern.ch/atlas/www/GROUPS/DATABASE/project/ddm/releases/TiersOfATLASCache.py -> TiersOfATLASCache.py.AGIS
lrwxr-xr-x, 1 iueda zp 49 Sep 21 2012 /afs/cern.ch/atlas/www/GROUPS/DATABASE/project/ddm/releases/TiersOfATLASCache.py.AGIS -> /afs/cern.ch/atlas/GRID/AGIS/TiersOfATLASCache.py
-rw-----, 1 15861 zp 21649 Jul 26 2006 /afs/cern.ch/atlas/www/GROUPS/DATABASE/project/ddm/releases/TiersOfATLASCache.pyc
lrwxr-xr-x, 1 iueda zp 24 Sep 12 2012 /afs/cern.ch/atlas/www/GROUPS/DATABASE/project/ddm/releases/TiersOfATLASCache.py.DDM -> ToA/TiersOfATLASCache.py
-rw-----, 1 15861 zp 41686 Jul 1 2006 /afs/cern.ch/atlas/www/GROUPS/DATABASE/project/ddm/releases/TiersOfATLASCache.pySat_Jul_1_10:30:00_2006
-rw-----, 1 15861 zp 45509 Jul 9 2006 /afs/cern.ch/atlas/www/GROUPS/DATABASE/project/ddm/releases/TiersOfATLASCache.py_Sun_Jul_9_10:25:10_2006
-rw-----, 1 15861 zp 42888 Jul 6 2006 /afs/cern.ch/atlas/www/GROUPS/DATABASE/project/ddm/releases/TiersOfATLASCache.py_Thu_Jul_6_09:35:27_2006
-rw-----, 1 15861 zp 34211 Jun 29 2006 /afs/cern.ch/atlas/www/GROUPS/DATABASE/project/ddm/releases/TiersOfATLASCache.py_Thu_Jun_29_15:57:28_2006
-rw-----, 1 dcameron zp 54124 Jul 11 2006 /afs/cern.ch/atlas/www/GROUPS/DATABASE/project/ddm/releases/TiersOfATLASCache_test.py
-bash-4.1$ ls -l /afs/cern.ch/atlas/GRID/AGIS/TiersOfATLASCache.py
lrwxr-xr-x, 1 atagadm root 76 Sep 20 11:00 /afs/cern.ch/atlas/GRID/AGIS/TiersOfATLASCache.py -> /afs/cern.ch/atlas/GRID/AGIS/ToACache/TiersOfATLASCache_2017_09_20_110006.py
-bash-4.1$
```

# SW releases (publication)

- We are (since a while) automatically tagging the CVMFS sites
- Do we still need the whole granularity for all the sites?
- nightlies?

# Why? -- again

- Why? Everybody is talking about simplification, consolidation, automation, etc...
  - But how can we do this if it's taking several experts days to explain how our system works?
- Lightweight sites??
  - No way! With the examples above bootstrapping a site is complicated
  - Experts always needed, not clear what happens if someone touch something.
- Whenever we make a change is supercomplex to understand the impact



- What do we clean?
- ... and when?