

Information System Topics

Florido Paganelli, Lund University

Nordugrid 2016, Kosice, 31/5-3/6

Information system

Program Settings Help

Browser Choose plugin

Entries

- GLUE2DomainID=NCG-INGRID-P1
- GLUE2DomainID=NCP-LCG2
- GLUE2DomainID=NDGF-T1
- GLUE2DomainID=NIHAM
- GLUE2DomainID=NIKHEF-ELPROD
- GLUE2DomainID=NO-NORGRID-T2
- GLUE2DomainID=OBSPM
- GLUE2DomainID=pic
- GLUE2DomainID=PK-CIT
- GLUE2DomainID=prague_cesnet_lcg2
- GLUE2DomainID=praguelcg2
- GLUE2DomainID=PRISMA-INFN-BARI
- GLUE2DomainID=PSNC
- GLUE2DomainID=RAL-LCG2
 - GLUE2ContactID=general.contact.RAL-LCG2
 - GLUE2ContactID=security.contact.RAL-LCG2
 - GLUE2ContactID=sysadmin.contact.RAL-LCG2
 - GLUE2ContactID=usersupport.contact.RAL-LCG2
 - GLUE2GroupID=resource
 - GLUE2GroupID=services
 - GLUE2ServiceID=urn:ogf:ComputingService:arc-ce01.gridpp.rl.ac.uk:arex
 - GLUE2ContactID=urn:ogf:Contact:arc-ce01.gridpp.rl.ac.uk:ComputingService:arex:con0
 - GLUE2EndpointID=urn:ogf:ComputingEndpoint:arc-ce01.gridpp.rl.ac.uk:gridftpjob:gsiftp://arc-ce01....
 - GLUE2EndpointID=urn:ogf:Endpoint:arc-ce01.gridpp.rl.ac.uk:ldapglue1:2135
 - GLUE2EndpointID=urn:ogf:Endpoint:arc-ce01.gridpp.rl.ac.uk:ldapglue2:2135
 - GLUE2EndpointID=urn:ogf:Endpoint:arc-ce01.gridpp.rl.ac.uk:ldapng:2135
 - GLUE2LocationID=urn:ogf:Location:arc-ce01.gridpp.rl.ac.uk:ComputingService:arex
 - GLUE2ManagerID=urn:ogf:ComputingManager:arc-ce01.gridpp.rl.ac.uk:condor
 - GLUE2ShareID=urn:ogf:ComputingShare:arc-ce01.gridpp.rl.ac.uk:grid3000M
 - GLUE2PolicyID=urn:ogf:MappingPolicy:arc-ce01.gridpp.rl.ac.uk:basic
 - GLUE2ServiceID=urn:ogf:ComputingService:arc-ce02.gridpp.rl.ac.uk:arex
 - GLUE2ServiceID=urn:ogf:ComputingService:arc-ce03.gridpp.rl.ac.uk:arex
 - GLUE2ServiceID=urn:ogf:ComputingService:arc-ce04.gridpp.rl.ac.uk:arex
 - GLUE2LocationID=location.RAL-LCG2

GLUE2DomainID=RAL-LCG2, GLUE2
GroupID=grid,o=glue

ObjectClasses

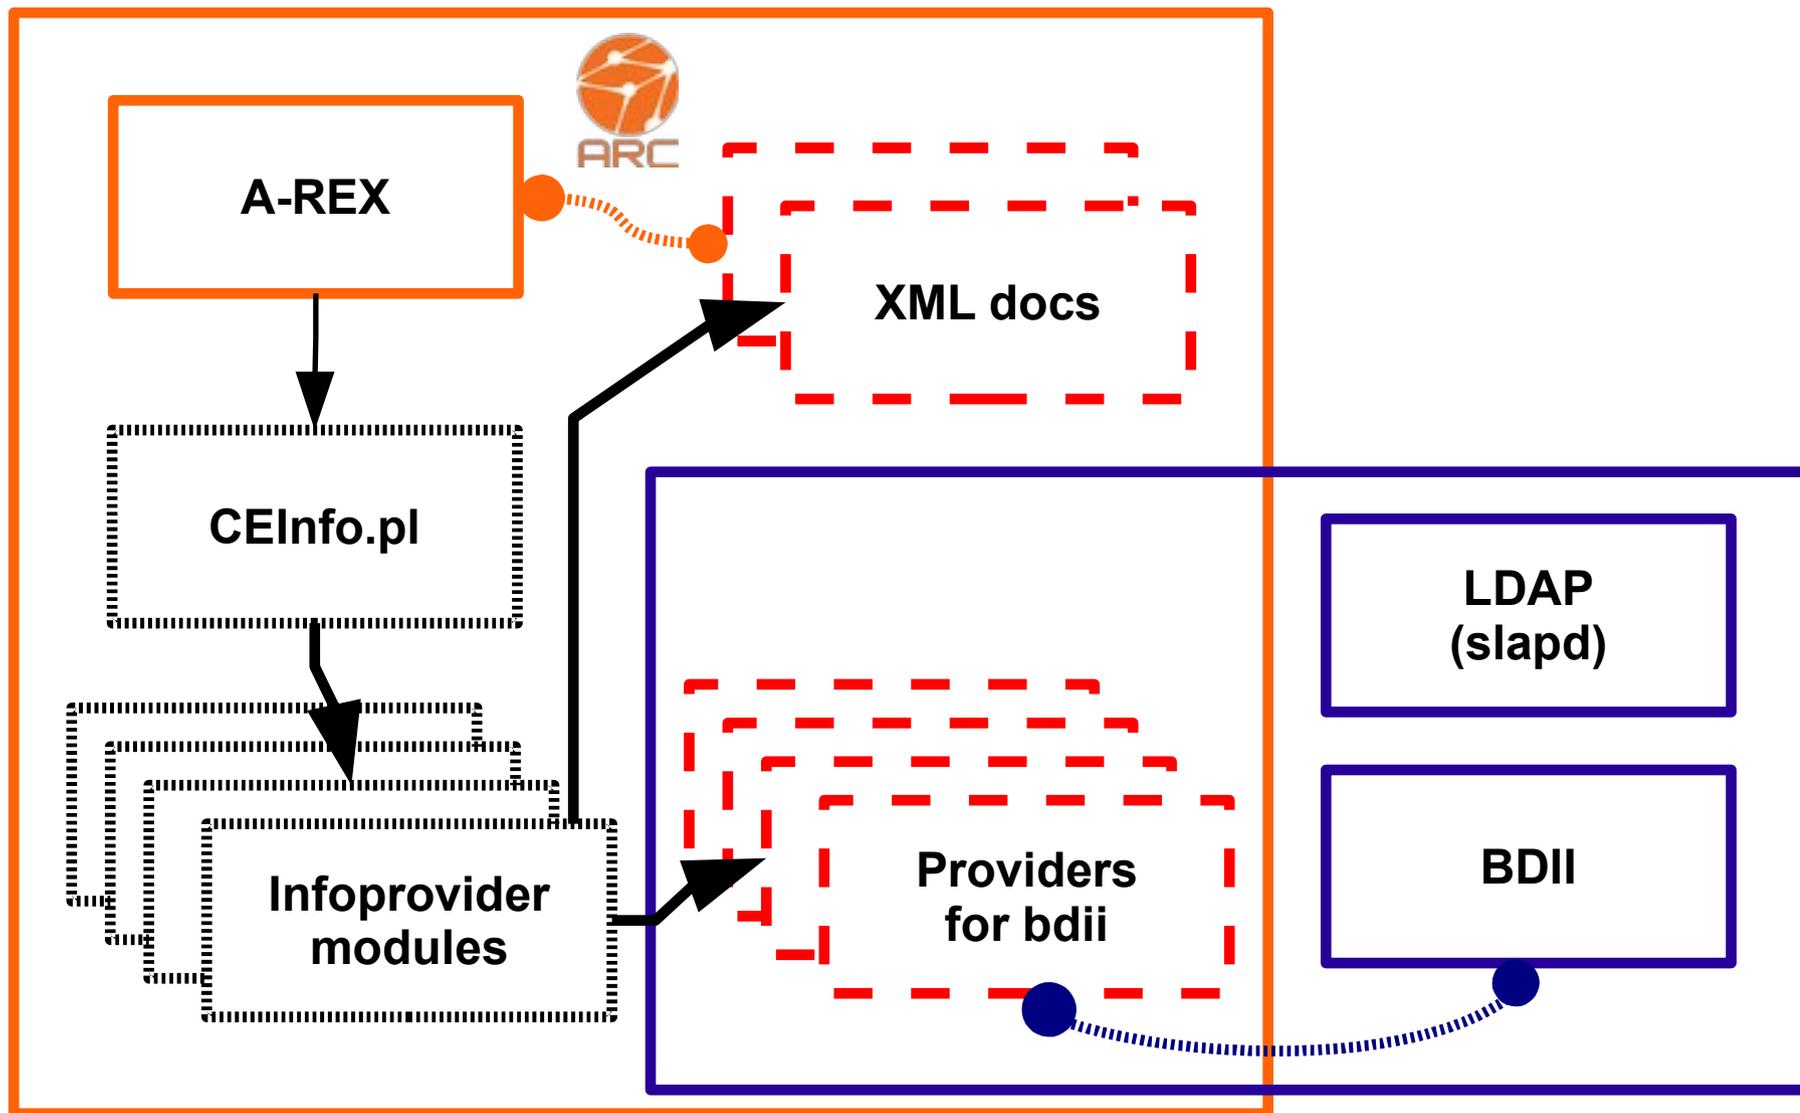
GLUE2Share

GLUE2ComputingShare 

Attributes

GLUE2ComputingShareComputingEndpointForeignKey	urn:ogf:ComputingEndpoint:arc-ce01.gridpp.rl.ac.uk:gridftpjob:gsiftp://arc-ce01.gridpp.rl.ac.uk:2811/jobs	 
GLUE2ComputingShareComputingServiceForeignKey	urn:ogf:ComputingService:arc-ce01.gridpp.rl.ac.uk:arex	 
GLUE2ComputingShareExecutionEnvironmentForeignKey	urn:ogf:ExecutionEnvironment:arc-ce01.gridpp.rl.ac.uk:execenv0	 
GLUE2ComputingShareFreeSlots	460	 
GLUE2ComputingShareFreeSlotsWithDuration	460	 
GLUE2ComputingShareLocalRunningJobs	13326	 
GLUE2ComputingShareLocalSuspendedJobs	0	 
GLUE2ComputingShareLocalWaitingJobs	0	 
GLUE2ComputingShareMappingQueue	grid3000M	 
GLUE2ComputingShareMaxRunningJobs	16968	 
GLUE2ComputingShareMaxSlotsPerJob	16968	 
GLUE2ComputingShare...	16968	 

Infosys ingredients



```
<ComputingShare BaseType="Share" CreationTime="2016-06-01T08:22:30Z" Validity="1200">
  <ID>urn:ogf:ComputingShare:mowgli.hep.lu.se:fork_atlas</ID>
  <Name>fork_atlas</Name>
  <MappingPolicy BaseType="Policy" CreationTime="2016-06-01T08:22:30Z" Validity="1200">
    <ID>urn:ogf:MappingPolicy:mowgli.hep.lu.se:basic:atlas</ID>
    <Scheme>basic</Scheme>
    <Rule>vo:atlas</Rule>
  </MappingPolicy>
  <MappingQueue>fork</MappingQueue>
  <MaxVirtualMemory>600</MaxVirtualMemory>
  <ServingState>production</ServingState>
  <TotalJobs>0</TotalJobs>
  <RunningJobs>0</RunningJobs>
  <LocalRunningJobs>0</LocalRunningJobs>
  <WaitingJobs>0</WaitingJobs>
  <LocalWaitingJobs>0</LocalWaitingJobs>
  <SuspendedJobs>0</SuspendedJobs>
  <LocalSuspendedJobs>0</LocalSuspendedJobs>
  <StagingJobs>0</StagingJobs>
  <PreLRMSWaitingJobs>0</PreLRMSWaitingJobs>
  <FreeSlots>0</FreeSlots>
  <FreeSlotsWithDuration>0</FreeSlotsWithDuration>
  <RequestedSlots>0</RequestedSlots>
  <Associations>
```

```
ComputingEndpointID>urn:ogf:ComputingEndpoint:mowgli.hep.lu.se:gridftpjob:gsiftp://mowgli.hep.lu.se:50000
ComputingEndpointID>
```

```
ComputingEndpointID>urn:ogf:ComputingEndpoint:mowgli.hep.lu.se:xbes:https://mowgli.hep.lu.se:50000
ComputingEndpointID>
```

Published info

- LDAP (slapd + bdii):
 - Query through Idapsearch
 - No authentication
 - Produces Ldif:
 - One big doc with
 - Nordugrid schema
 - GLUE2 schema
 - Glue 1.3 document is translated from Nordugrid schema (very inaccurate, bad code)
- Web Services (A-REX native):
 - Query through EMIES interface
 - Authenticated (X.509)
 - Produces XML documents in GLUE2 format:
 - For services (Resouce document)
 - For each job

Performance

Current performance data

- Uses Perl NYTProf
- Currently incomplete wrt ARC plans: creates one NYTprof file per Ceinfo.pl run (lots of files!)
- External script aggregates info all together into one single file (avg 30MB per day on low end machine, doesn't change depending on load)
- Plans to generate a file according to ARC specs using NYTProf libraries

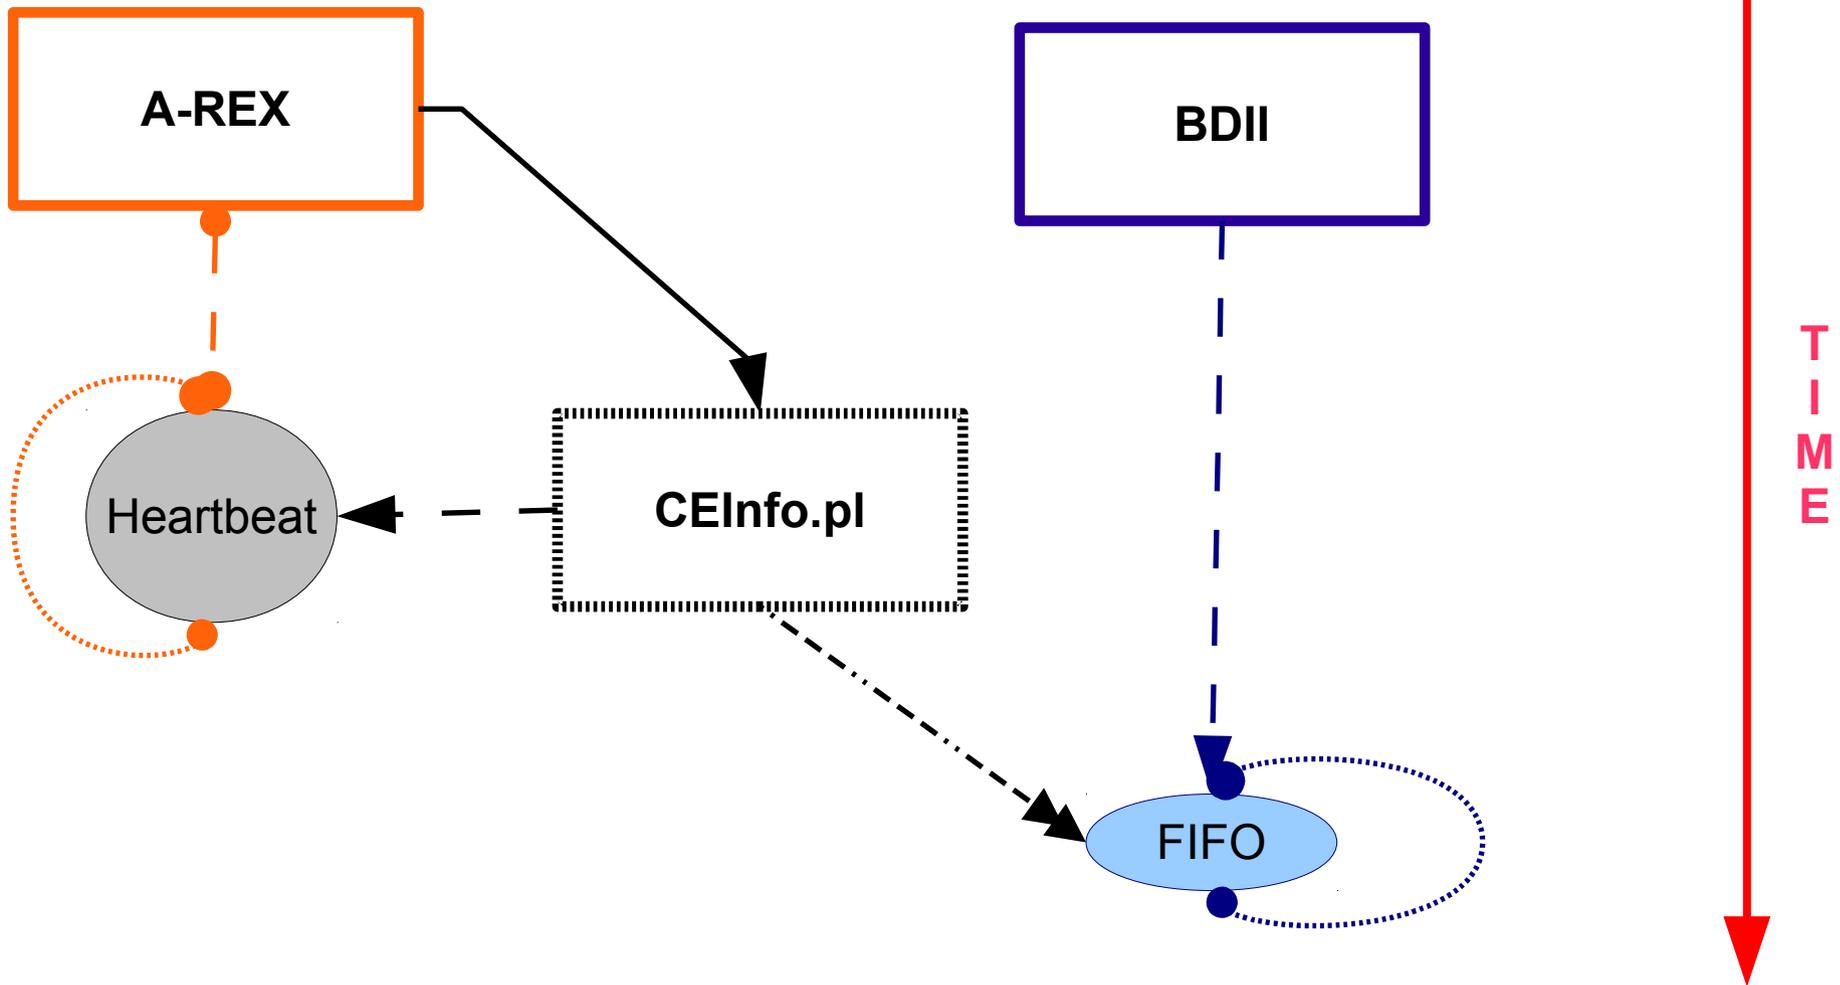
Known performance issues

- Huge delays in controldir scanning
 - Improved a bit with an heartbeat system (i.e. are infoproviders actually doing something?)
 - Currently no good solution. In big clusters scan can take hours, which causes outdated data or no data for some jobs. Equivalent to no timeout or no heartbeat
- Way out:
 - use database and store job info
 - have some scheduling strategy to parse job info and generate well formed documents in case of missing information

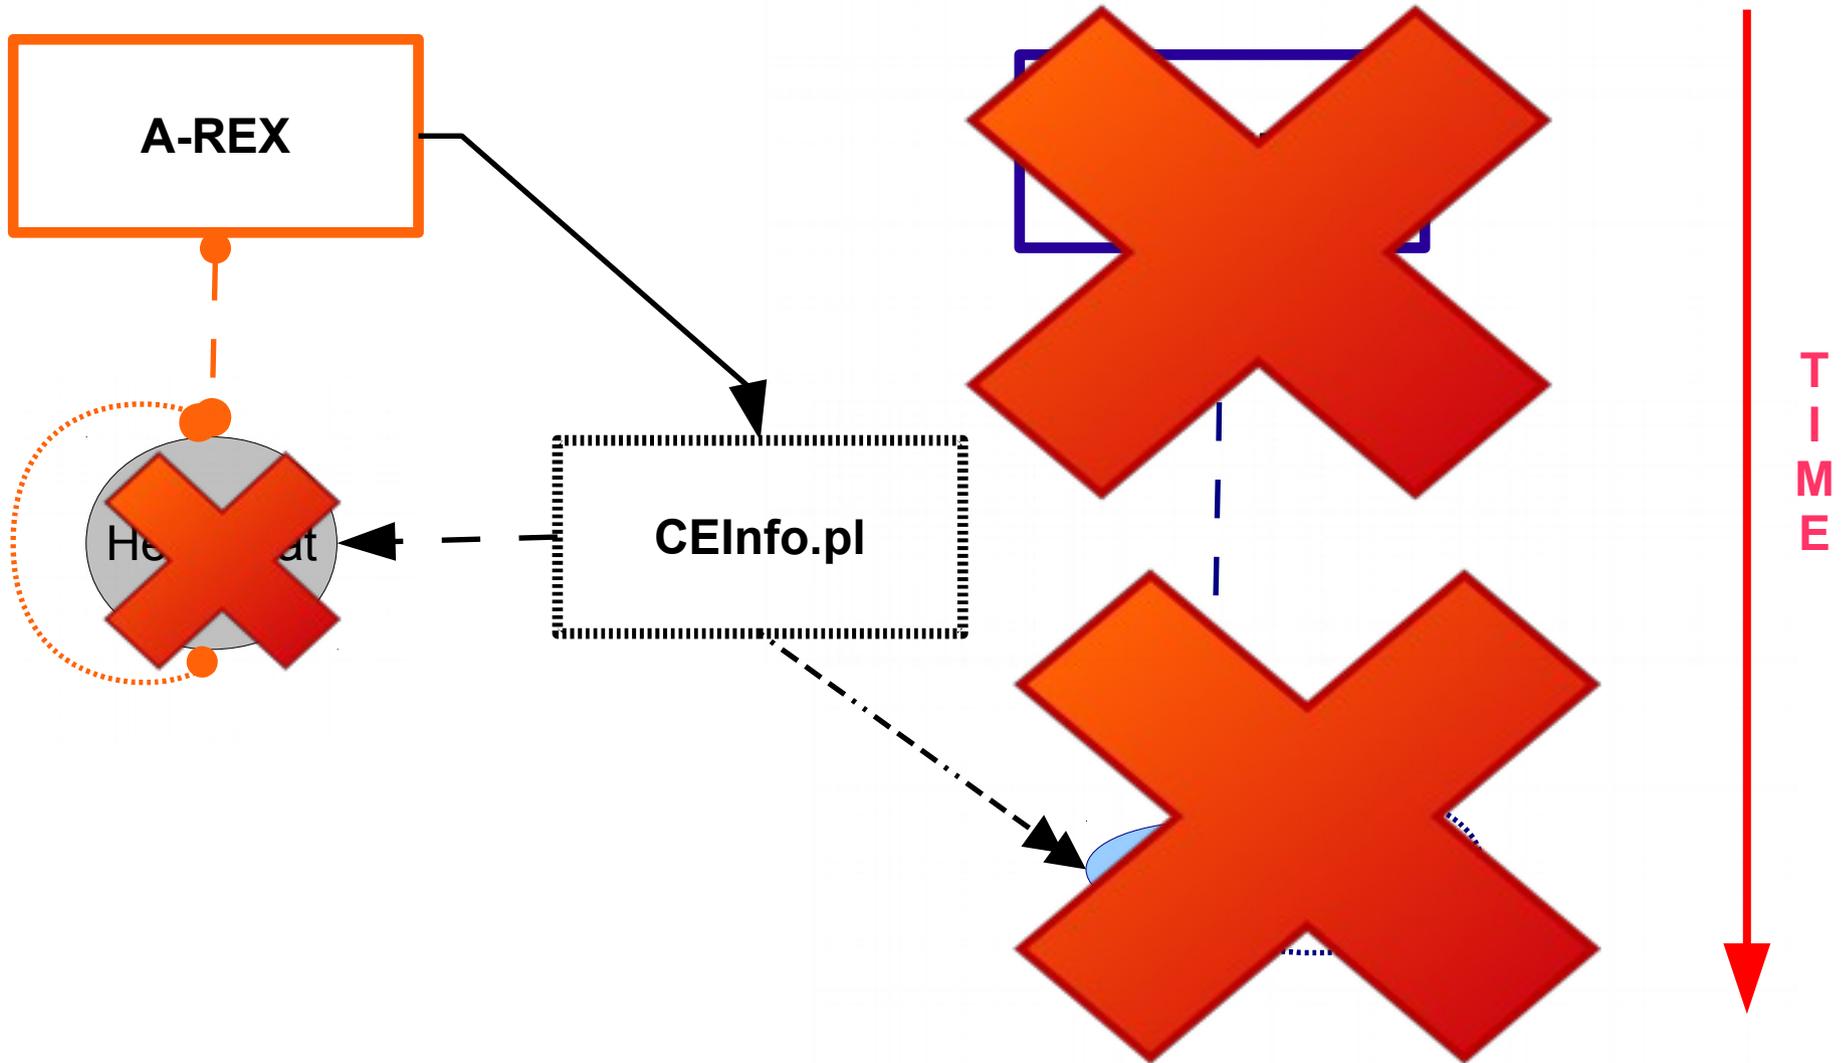
Known performance issues

- May hang due to batch system hanging
 - Former solution was to kill Ceinfo.pl after a timeout. Assumption: restarting infoproviders would help.
This caused information never to be produced.
 - Not sufficient: restarting infoproviders rarely helps and data is still incomplete
- Possible way out:
 - be more clever in information production: generate intermediate documents that can be still published
 - handle hangs in perl instead than in C++.
 - Make only one component (i.e. scan-jobs) do all the LRMS scanning

Infosys execution cycle



Wishes: infosys execution cycle



VO support status

VO requirements

- ALICE and DIRAC query the information system (LDAP BDII, Glue 1.3) for job statistics
- Glite/CREAM used to provide that info, ARC did not
- It seems that condor used to push this info inside the LRMS (that's why RAL crafted a patch to calculate this info with ARC)
- ARC will collect this info from the attributes in the user proxy for each job

Reference bugs

- **#3495** A-REX should write which VO a job "belongs to", in controldir .local file
- **#3464** ComputingShare and MappingPolicy should be created on a VO basis for WLCG/CERN use case, and job data calculated per VO
- **#3433** Publish authorised VOs per queue
- **#3384** Support for per-queue authorisation configuration and publishing
- **#3538** ARC Client libraries should take the batch system queue name from GLUE2 ComputingShare.MappingQueue and NOT from ComputingShare.Name

Reference bugs

- **#3495** A-REX should write which VO a job "belongs to", in controldir .local file
 - Done by Aleksandr, we take the first VO in the certificates
- **#3433** Publish authorised VOs per queue
 - Done for all renderings
 - Simply adding *authorizedvo* config option to the [cluster] or [queue/queuenam]
- **#3384** Support for per-queue authorisation configuration and publishing
 - Enforcement must be done manually using [group] and *groupcfg=groupname* in [gridftp]

Reference bugs

- **#3464** ComputingShare and MappingPolicy should be created on a VO basis for WLCG/CERN use case, and job data calculated per VO
 - Only done for GLUE2
 - One GLUE2 Share for each VO
 - Job stats calculated per VO, accuracy must be checked
 - Not tested for LDAP/LDIF, only XML

```
<ComputingShare BaseType="Share" CreationTime="2016-06-01T08:22:30Z" Validity="1200">
  <ID>urn:ogf:ComputingShare:mowgli.hep.lu.se:fork_atlas</ID>
  <Name>fork_atlas</Name>
  <MappingPolicy BaseType="Policy" CreationTime="2016-06-01T08:22:30Z" Validity="1200">
    <ID>urn:ogf:MappingPolicy:mowgli.hep.lu.se:basic:atlas</ID>
    <Scheme>basic</Scheme>
    <Rule>vo:atlas</Rule>
  </MappingPolicy>
  <MappingQueue>fork</MappingQueue>
  <MaxVirtualMemory>600</MaxVirtualMemory>
  <ServingState>production</ServingState>
  <TotalJobs>0</TotalJobs>
  <RunningJobs>0</RunningJobs>
  <LocalRunningJobs>0</LocalRunningJobs>
  <WaitingJobs>0</WaitingJobs>
  <LocalWaitingJobs>0</LocalWaitingJobs>
  <SuspendedJobs>0</SuspendedJobs>
  <LocalSuspendedJobs>0</LocalSuspendedJobs>
  <StagingJobs>0</StagingJobs>
  <PreLRMSWaitingJobs>0</PreLRMSWaitingJobs>
  <FreeSlots>0</FreeSlots>
  <FreeSlotsWithDuration>0</FreeSlotsWithDuration>
  <RequestedSlots>0</RequestedSlots>
  <Associations>
```

```
ComputingEndpointID>urn:ogf:ComputingEndpoint:mowgli.hep.lu.se:gridftpjob:gsiftp://mowgli.hep.lu.se:50000
ComputingEndpointID>
```

```
ComputingEndpointID>urn:ogf:ComputingEndpoint:mowgli.hep.lu.se:xbes:https://mowgli.hep.lu.se:50000
ComputingEndpointID>
```

- Entries
- GLUE2DomainID=Australia-ATLAS
 - GLUE2DomainID=AUVERGRID
 - GLUE2DomainID=BEgrid-BELNET
 - GLUE2DomainID=BEgrid-ULB-VUB
 - GLUE2DomainID=BEIJING-LCG2
 - GLUE2DomainID=BelGrid-UCL
 - GLUE2DomainID=BG01-IPP
 - GLUE2DomainID=BG05-SUGrid
 - GLUE2DomainID=BIFI
 - GLUE2DomainID=BIOCOMP
 - GLUE2DomainID=BMRZ-FRANKFURT
 - GLUE2DomainID=BUDAPEST
 - GLUE2ContactID=general.contact.BUDAPEST
 - GLUE2ContactID=security.contact.BUDAPEST
 - GLUE2ContactID=sysadmin.contact.BUDAPEST
 - GLUE2ContactID=usersupport.contact.BUDAPEST
 - GLUE2GroupID=resource
 - GLUE2ServiceID=816
 - GLUE2EndpointID=816_org.glite.ce.ApplicationPublisher_2053530043
 - GLUE2EndpointID=grid106.kfki.hu_org.glite.ce.CREAM
 - GLUE2PolicyID=grid106.kfki.hu_org.glite.ce.CREAM_Policy
 - GLUE2EndpointID=grid107.kfki.hu_org.glite.ce.CREAM
 - GLUE2EndpointID=grid109.kfki.hu_org.glite.ce.CREAM
 - GLUE2ManagerID=816_Manager
 - GLUE2ResourceID=kfki.hu
 - GLUE2ShareID=alice_alice_816
 - GLUE2PolicyID=alice_alice_816_policy
 - GLUE2ShareID=cms_cms_816
 - GLUE2ShareID=dteam_dteam_816
 - GLUE2ShareID=hungrid_hungrid_816
 - GLUE2ShareID=ops_ops_816
 - GLUE2ToStorageServiceID=816_grid143.kfki.hu
 - GLUE2ServiceID=grid11.kfki.hu_VOMS_2715298739
 - GLUE2ServiceID=grid143.kfki.hu

teedVirtualMemory			
GLUE2ComputingShareMappingQueue	alice		 
GLUE2ComputingShareMaxCPU Time	259200		 
GLUE2ComputingShareMaxMainMemory	9223372036854775807		 
GLUE2ComputingShareMaxRunningJobs	2147483647		 
GLUE2ComputingShareMaxSlotsPerJob	999		 
GLUE2ComputingShareMaxTotalJobs	2147483647		 
GLUE2ComputingShareMaxVirtualMemory	9223372036854775807		 
GLUE2ComputingShareMaxWaitingJobs	0		 
GLUE2ComputingShareMaxWallTime	259200		 
GLUE2ComputingShareRunningJobs	675		 
GLUE2ComputingShareServiceState	Production		
GLUE2ComputingShareTotalJobs	707		 
GLUE2ComputingShareUsedSlots	0		 
GLUE2ComputingShareWaitingJobs	32		 
GLUE2EntityCreationTime	2015-12-09T17:07:54Z		 
GLUE2EntityOtherInfo	CREAMCEId=grid106.kfki.hu:8443/cream-sge-alice		 
	ServiceType=org.glite.ce.CREAM		 
	InfoProviderName=glite-ce-glue2-share-static		 
	InfoProviderVersion=1.1		 

Reference Bugs and GGUS tickets

- https://ggus.eu/index.php?mode=ticket_info&ticket_id=112543
- https://bugzilla.nordugrid.org/show_bug.cgi?id=3495
- https://bugzilla.nordugrid.org/show_bug.cgi?id=3433
- https://bugzilla.nordugrid.org/show_bug.cgi?id=3464
-