

ARC and ARC-CE


Andrej Filipčič

About ARC and Nordugrid

- ARC – advanced resource connector
 - ➔ grid middleware, 10 year in production, most notably in ATLAS ND cloud, CMS and ALICE
 - ➔ Part of EMI, client installed at CERN by default
- Developed and supported by Nordugrid Collaboration (NO,SE,FI,CH,SI,DK,UA,SK,HU)
- Running on ~80 clusters, providing access to ~100k cores (250k since recently)

ATLAS production on ARC

2013-10-24 CEST 15:11:11



Processes: ■ Grid ■ Local



Country	Site	CPUs	Load (processes: Grid+local)	Queueing
Denmark	Steno Tier 1 (DCSC/KU)	5888	832+3886	589+0
	LRZ-C2PAP	4032	2354+896	437+0
Germany	LRZ-LMU	2312	505+1586	198+485
	LRZ-LMU	2312	1410+678	493+190
	RZG ATLAS HYDRA	157072	0+134438	0+0
	wuppertalprod	3268	3088+665	355+18980
Norway	Abel C1(UiO/USIT)	10864	1098+8895	200+2080
	Abel C3(UiO/USIT)	10864	27+9985	2073+204
Slovenia	Arnes	1924	1658+0	288+1
	SiGNET	2200	1982+18	241+11
Sweden	Abisko (HPC2N)	15456	397+13134	174+0
	Alarik (SweGrid, Luna>	3776	400+2736	80+0
	Ritsem (SweGrid, HPC2>	544	437+0	95+0
	Siri (SweGrid, Lunarc)	448	297+94	78+50
	Smokerings (NSC)	520	422+60	151+1
	Tintin (SweGrid, Uppm>	2624	57+2195	0+64
Switzerland	Bern ATLAS WLCG T2	1352	1135+0	726+0
	Bern UBELIX T3	2272	486+1588	95+1411
	Geneva ATLAS T3	418	141+88	68+151
	Manno PHOENIX T2	2490	413+908	328+70
	Manno PHOENIX T2	192	67+52	329+383
UK	arc-ce01 (RAL-LCG2)	5466	1493+3488	320+0
	arc-ce02 (RAL-LCG2)	5458	1261+3711	385+1
	arc-ce03 (RAL-LCG2)	5466	1473+3507	295+0
	cetest01 (UKI-LT2-IC->	4	192+2655	196+1200
TOTAL	25 sites	247222	21617 + 195143	8194 + 25282

ARC-CE approach

- Clear separation of grid and local job execution:
 - All grid-aware job operations done on ARC-CE (including transfers)
 - Worker node executes purely local jobs on local inputs and write outputs locally
- Suitable as an add-on to a pure cluster → well adapted as an extra service to a supercomputer (HPC,...)
- Grid middleware NOT needed on the nodes

2 modes of operation in ATLAS

- **Standard pilot mode:**

- APF submits pilots to ARC-CE
- Job on the WN uses glite middleware to execute the payload
- ARC-CE → a simple replacement for CREAM-CE
- Fully integrated in the ATLAS infrastructure

- **Payload submission mode:**

- arcControlTower talks to PanDA (getjob, heartbeats)
- Fully defined jobs are submitted to ARC-CE
- Downloads/uploads performed by ARC-CE
- Jobs run in network-less mode (apart from Frontier, cvmfs)
- Custom jobs (memory, multi-core, walltime...) are trivial to execute without any additional configuration (whatever batch system supports)
- This mode is used by all the ND cloud clusters

Scalability

- A single ARC-CE can easily sustain 10k running jobs (hard to find larger clusters)
- Pilot jobs are extremely light in terms of job processing on ARC-CE
- Direct submission:
 - ➔ Processing power of heavy I/O jobs (large inputs, outputs) depends on the shared filesystem and network throughput
- Several ARC-CEs can be added for redundancy to the same cluster, but the CE workspace needs to be distinct
- datadelivery servers can extend the transfer capability of a single (or many) ARC-CE

Accounting

- JURA: service on ARC-CE, part of A-REX
- Direct publishing to central APEL, no need to have a local APEL instance

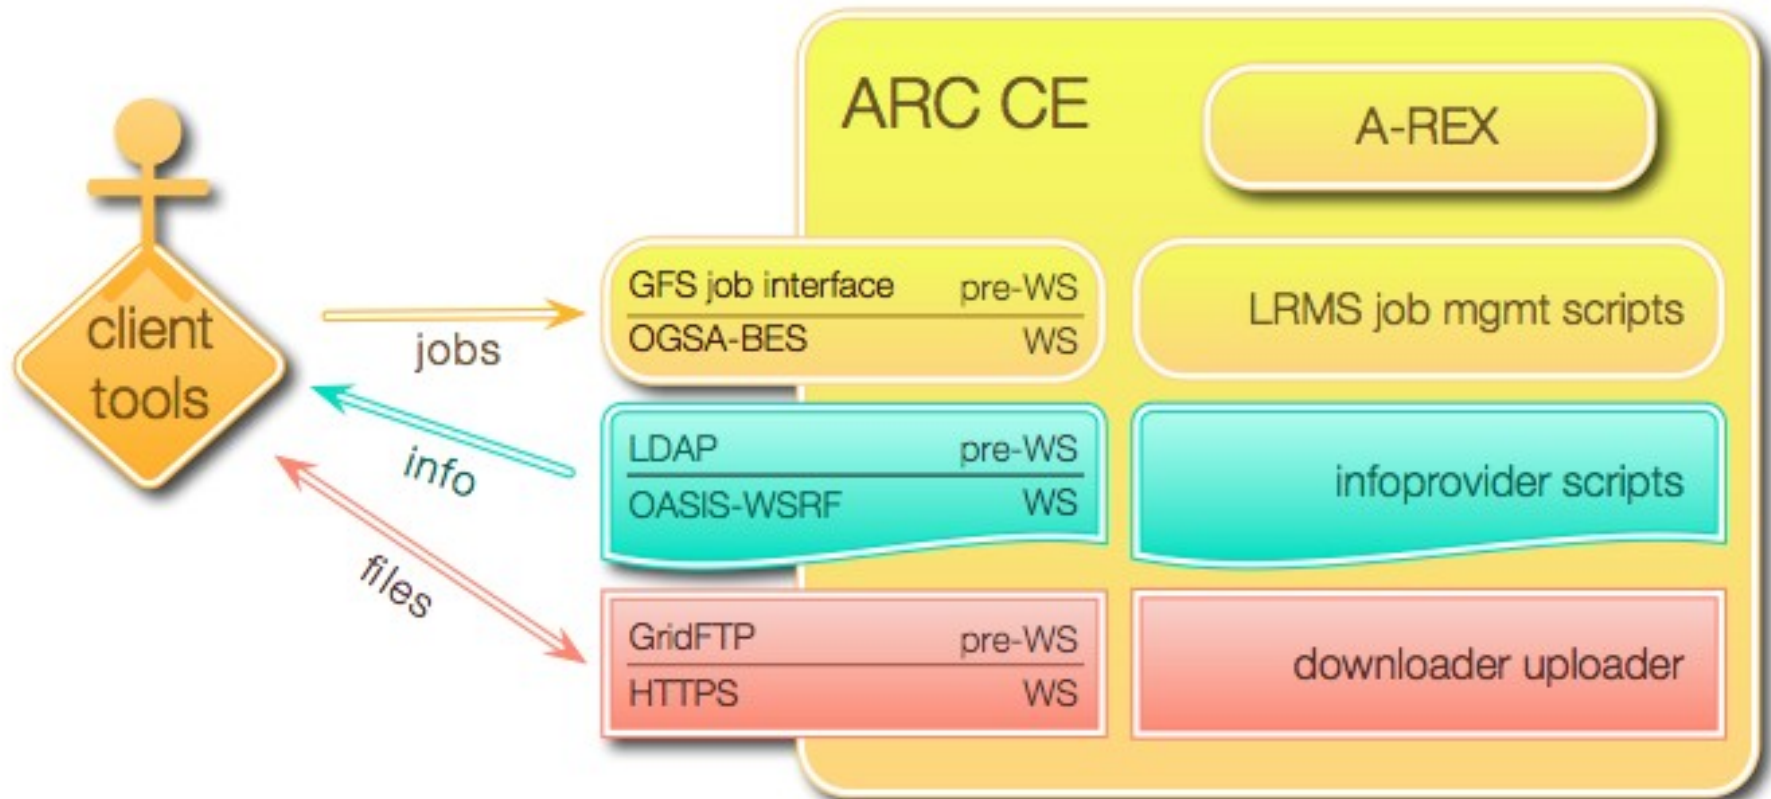
Release Cycle

- Major release once or twice per year
 - brings new functionality
 - Upgrade can be done on-the-fly within minutes, no draining is needed (not even downtime)
- Minor releases
 - Bugfixes
 - As needed
- Express releases
 - For urgent bug fixes

Support

- Technical discussion list
 - nordugrid-discuss@nordugrid.org
- Bugzilla
 - <http://bugzilla.nordugrid.org/>

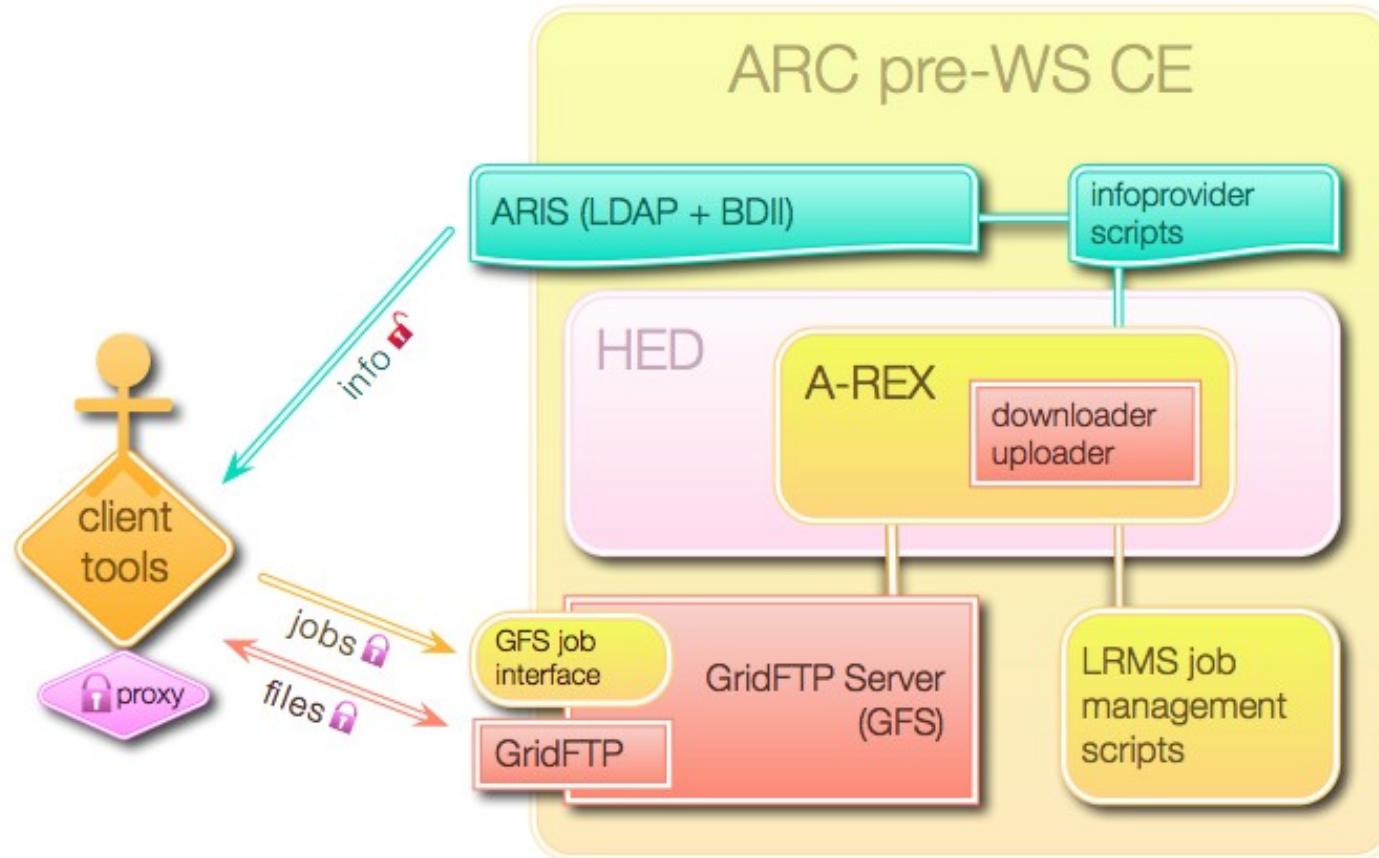
ARC Computing Element (CE)



ARC-CE services

- **gridftpd service**
 - ➔ jobs submission, output retrieval through gsiftp protocol (https EMI-ES interface will supersede it)
- **ldap information system + information index:**
 - ➔ Provides detailed information on the clusters
 - ➔ Provides information on all the grid jobs in the system (monitor)
 - ➔ Index service (EGIIS) aggregates several sites and provides a generic submission entry point
- **A-REX**
 - ➔ Job processor
 - ➔ Converts a grid job to a local batch job and supervises job execution
 - ➔ Transfers the requested inputs from remote SEs (https,gsiftp,xrootd,lfc,srm,...)
 - ➔ Uploads the outputs to remote SEs
- **JURA**
 - ➔ the accounting publishing interface (APEL,...)
- **/etc/arc.conf** → the single almighty configuration file

Pre-WS ARC-CE



WS ARC-CE

