



# ARC-CE: updates and plans

*Oxana Smirnova, NeIC/Lund University*

*1 July 2014*


*Grid 2014, Dubna*

*using input from: D. Cameron, A. Filipčič,  
J. Kerr Nilsen, B. Kónya*

## ARC-CE: brief summary

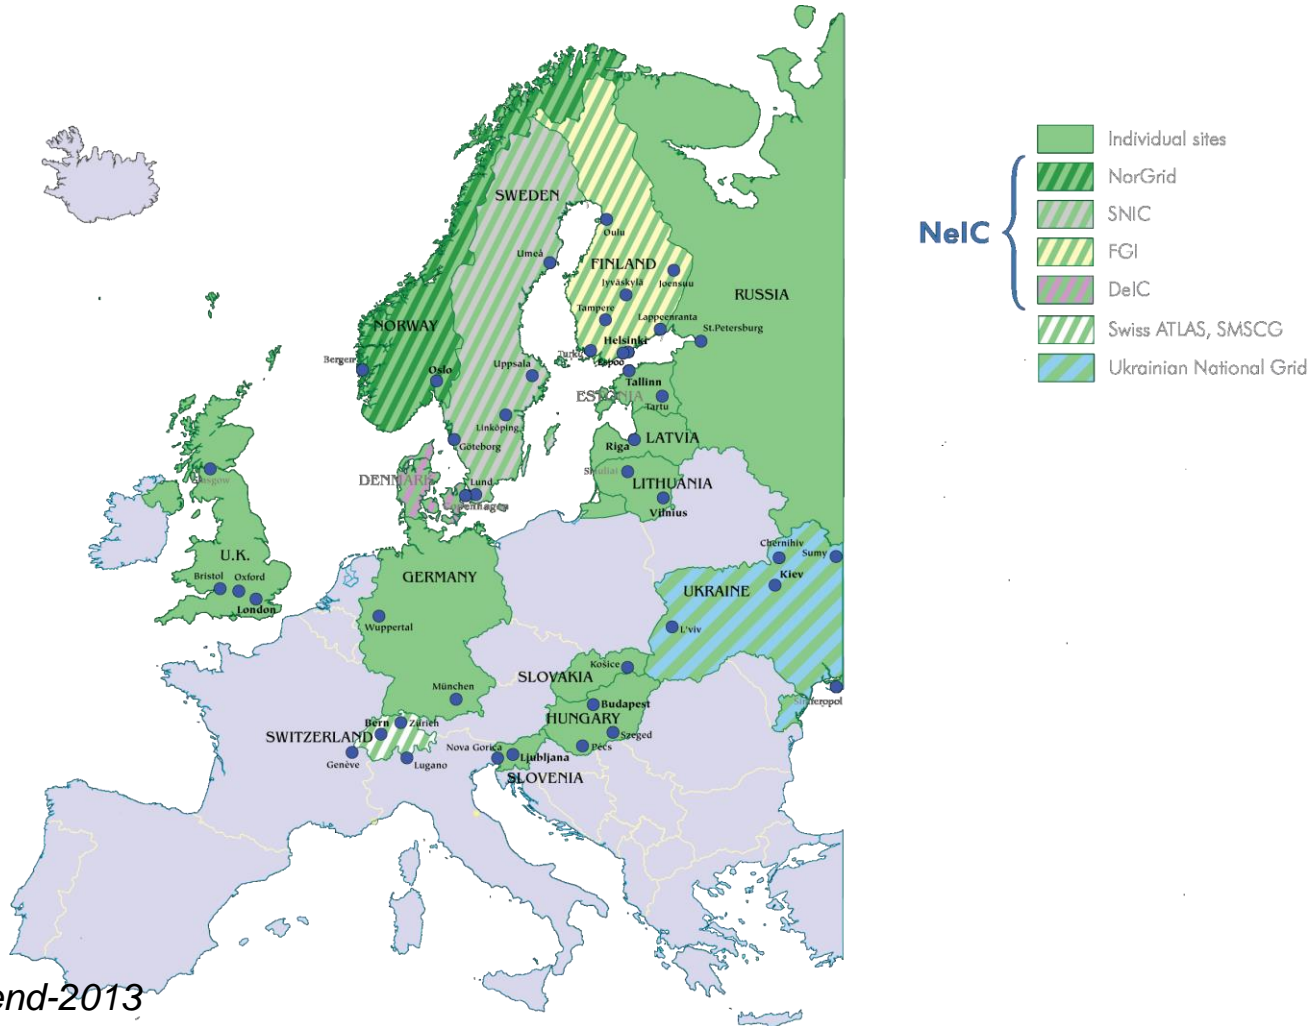
- Compute element by NorduGrid
- A key component of ARC middleware
  - Other components: clients, information services
- A key enabler of the Nordic Tier-1
  - Along with the distributed dCache storage
- Used by all LHC experiments
  - Mostly by ATLAS (so far)
    - Increasing use by CMS
  - The only CE used in Nordic and Baltic countries
  - Main CE in Slovenia and Ukraine
  - Increasing use in the UK and Germany

# ARC-CE instances in GOCDDB

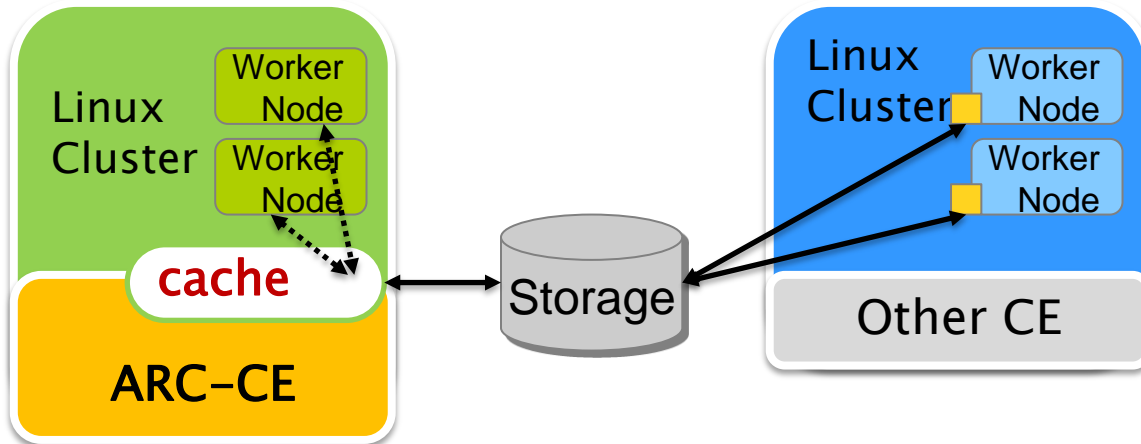


*cf. CREAM-CE: 566 instances as of today*

# ARC-CE geography



# Optimisation for data-intensive jobs



ARC-CE can do all the data transfers

Allows to cache frequently used files

Minimizes bandwidth

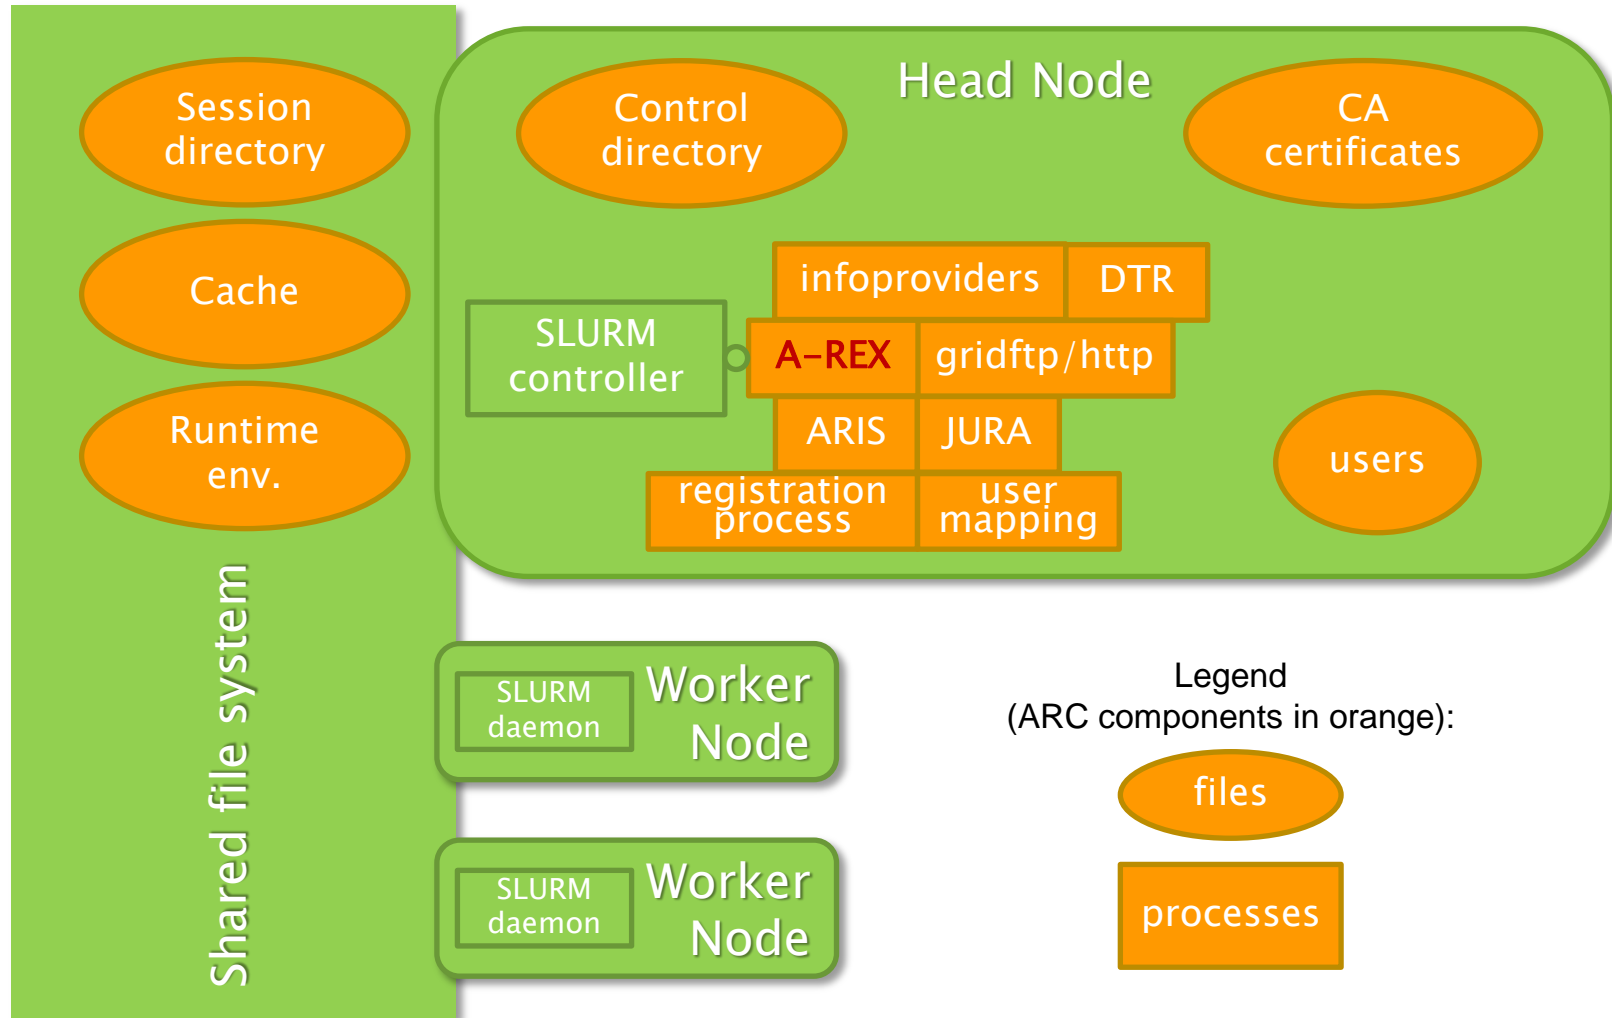
Maximizes worker nodes usage efficiency

ARC-CE is a rather complex service

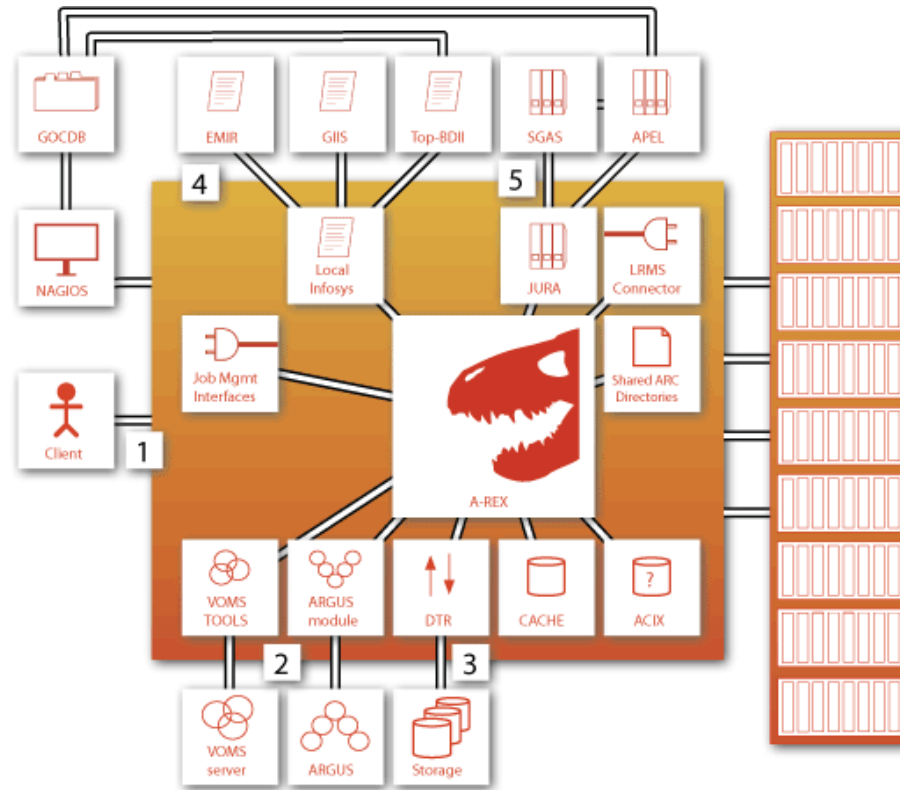
Is built of many individual services and tools

Requires high-end storage for cache

# ARC-CE components on a (SLURM) cluster



# Integration with EGI services

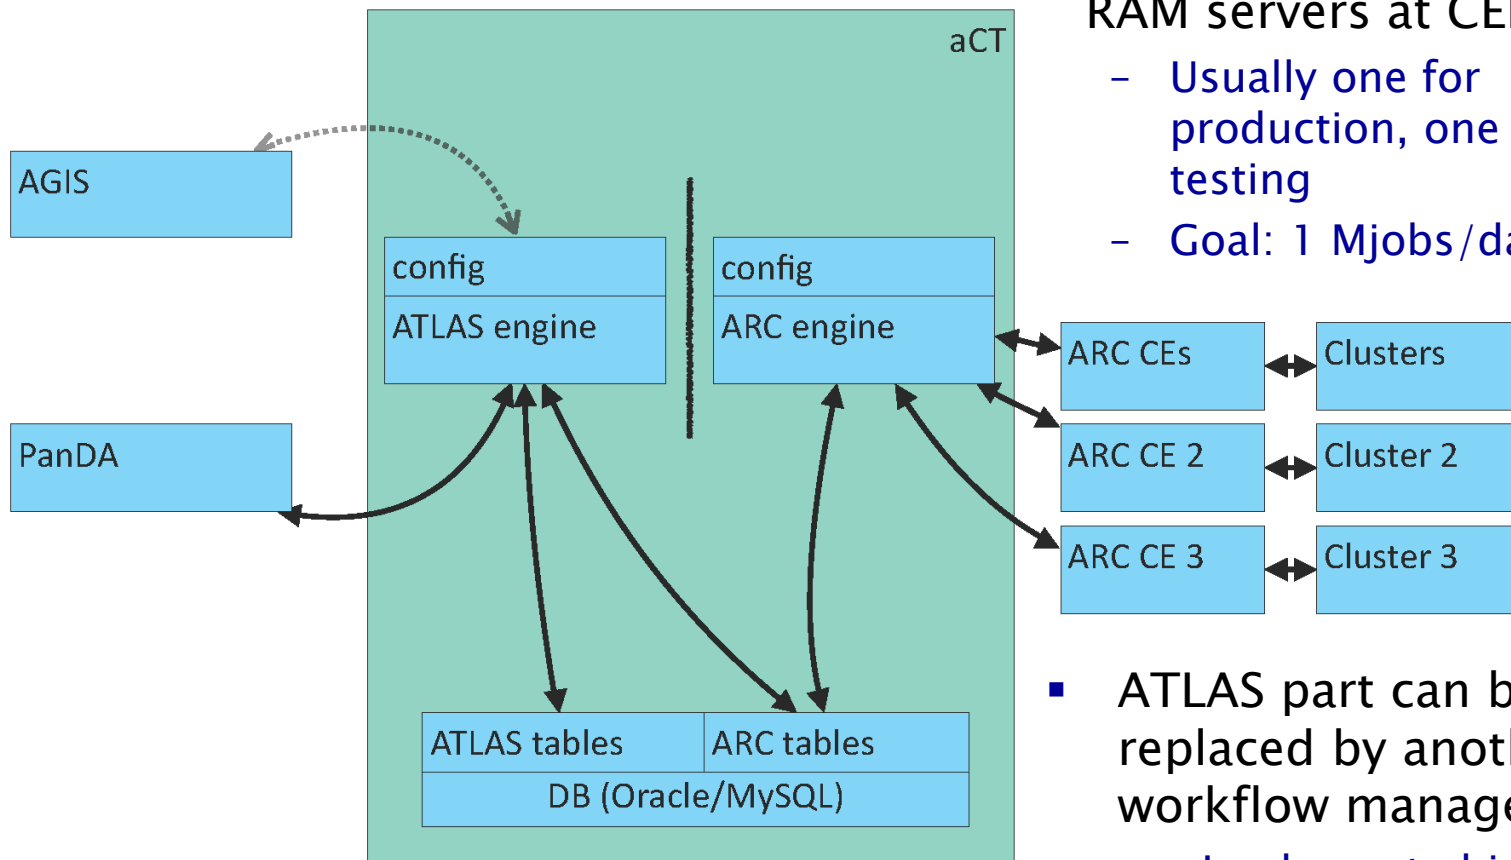


1. Job submission (brokering based on info from GIIS, EMIR, Local Infosys and ACIX)
2. Check credentials (VOMS, ARGUS, etc.)
3. Data staging from/to external storage
4. Registration to information indices (EGIS, EMIR); serving information requests of global aggregators (Top-BDII)
5. JURA parses job logs, prepares and sends job usage records to either SGAS or APEL accounting databases

# Control tower for pilots: aCT

- ARC Control Tower is the layer between ATLAS and ARC
  - Picks up job descriptions from Panda
  - Converts them to XRSL job description
  - Submits and manages jobs on ARC-CEs
  - Fetches output, handles common failures and updates Panda
- aCT2: based on ARC API
- Modular
  - ARC actors: submitter, status checker, fetcher, cleaner
  - ATLAS actors: autopilot, panda2arc, atlas status checker, validator






- Runs on two 24 CPU, 50GB RAM servers at CERN
  - Usually one for production, one for testing
  - Goal: 1 Mjobs/day
  
- ATLAS part can be replaced by another workflow manager
  - Implemented in Python

# Gateways to supercomputers

- ATLAS is pursuing usage of HPC resources
  - Long-term funding guaranteed
  - Not suitable for all kinds of jobs still
- In Europe, ARC is often used as a “gateway”
  - C2PAP/SuperMUC, Hydra: via aCT
  - Piz Daint: via ARC-CE ssh back-end (testing)
- Still, a lot remains to be done on both HEP and HPC sides to make it useable



# HPC challenges to overcome

- WAN access on nodes is limited or not available
- Whole-node, whole-socket, whole-partition scheduling
- Shared filesystem might suffer with heavy I/O jobs
- Limited access to login/edge nodes
- HPC policies and procedures
  - Sites are tuned for few classes of massively parallel applications with relatively low I/O
  - Limited time slots allocated
  - Access via a SSH-login front node

# Overcoming the difficulties

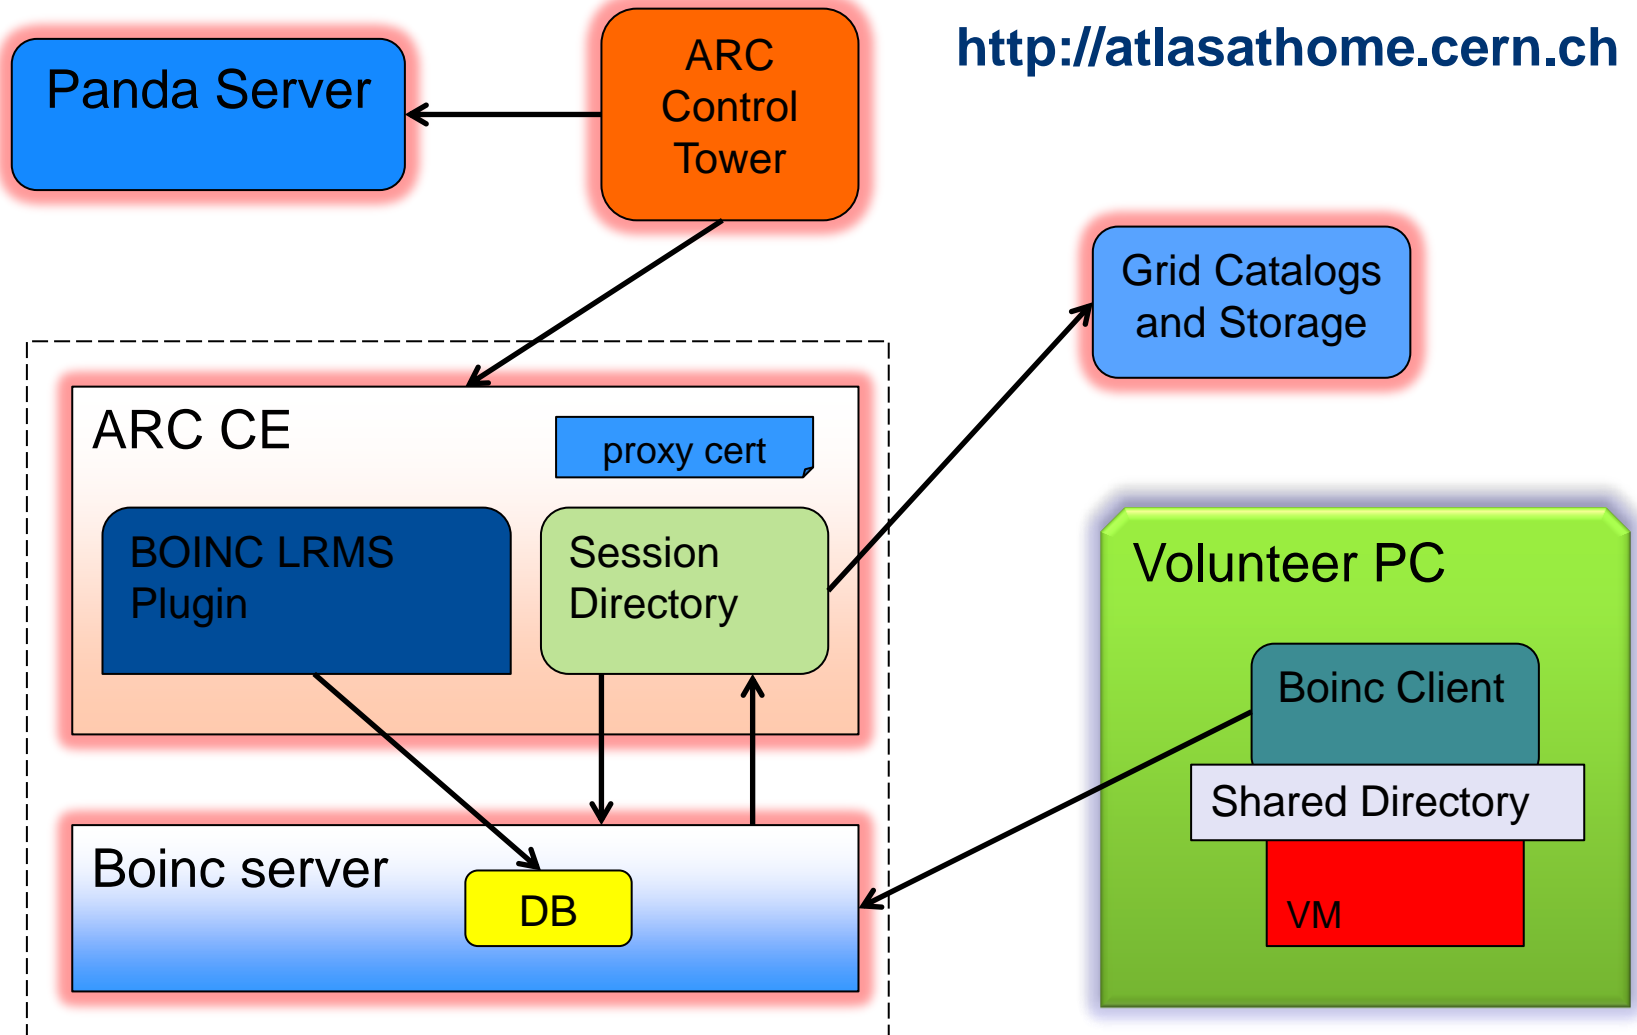
- Custom ARC-CE service machines:
  - Using RHEL6 packages on unsupported OS
  - Porting to unsupported OS (SLES11)
- User mode services:
  - Adapting ARC-CE to run as a non-privileged user
  - Limited usability (uid mapped to one batch account)
- No WAN access:
  - aCT, manual software sync, Parrot, no db access
    - Experienced GPFS lock-up due to heavy software area usage (bug fixed by IBM)
- Limited connectivity to edge nodes:
  - ARC SSH backend in the works

- The scale of volunteer computing:
  - At Amazon EC2 price: 1.16USD/hour for a CPU intensive Windows/Linux instance (~1.5GFLOPS)  
All BOINC activities would cost ~  
 $(7.2\text{PetaFLOPS}/1.5\text{GFLOPS}) * 1.16\text{USD} =$   
**=5.56M USD/hour**
  - Successful Individual projects:
    - Einstein@home (280 TeraFLOPS),
    - SETI@home (581 TeraFLOPS)
    - LHC@home (2TeraFLOPS)

- Why use volunteer computing?
  - It's free!
  - Public outreach
- Considerations
  - Low priority jobs with high CPU-I/O ratio
    - Non-urgent Monte Carlo simulation
  - Need virtualisation for ATLAS sw environment
    - CERNVM image and CVMFS
  - No grid credentials or access on volunteer hosts
    - ARC middleware for data staging
  - The resources should look like a regular Panda queue
    - ARC Control Tower

# ATLAS@Home Architecture

<http://atlasathome.cern.ch>



# Participants

gilda117.ihep.ac.cn/ATLAS/top\_users.php?sort\_by=expavg\_credit Google

## Top participants

David Cameron · lc

Rank	Name	Recent average credit	Total credit	Country	Participant since
1	<a href="#">Andrej Filipcic</a>	592	17,237	Slovenia	19 Feb 2014, 12:45:04 UTC
2	<a href="#">Bok</a>	409	5,126	United States	29 Apr 2014, 15:24:12 UTC
3	<a href="#">David Cameron</a>	264	3,843	Switzerland	15 Jan 2014, 13:18:02 UTC
4	<a href="#">lancon</a>	144	2,175	Switzerland	20 Jan 2014, 16:10:41 UTC
5	<a href="#">Tomas Kouba</a>	129	1,372	Czech Republic	12 May 2014, 8:22:51 UTC
6	<a href="#">Simone Campana</a>	86	956	Switzerland	9 May 2014, 9:21:13 UTC
7	<a href="#">kxt</a>	69	765	Czech Republic	12 May 2014, 13:09:16 UTC
8	<a href="#">zombie67 [MM]</a>	43	450	United States	9 May 2014, 13:15:39 UTC
9	<a href="#">Andrej Gorisek</a>	22	340	France	7 May 2014, 19:29:49 UTC
10	<a href="#">fialal</a>	14	145	Czech Republic	12 May 2014, 13:07:16 UTC
11	<a href="#">kanwx</a>				4 Apr 2014, 5:47:34 UTC
12	<a href="#">Laurence</a>				29 Apr 2014, 12:03:12 UTC
13	<a href="#">wuwj</a>	1	880	International	31 Dec 2013, 2:11:20 UTC
14	<a href="#">APOATLAS</a>	0	0	France	29 Apr 2014, 10:08:41 UTC

**No idea who they are!**



- ARC–CE is well established beyond Nordics now
  - More contributors, too
- aCT2 opens up many more possibilities, such as usage of HPC and volunteer computing
  - Very new, much optimisation is still needed
  - ... and also documentation and packaging
- Future directions:
  - Focus on LHC requirements
  - Enhance support of HPC systems, more batch system options
  - Develop more user–friendly task schedulers, using aCT2 experience