



ARC overview

*Oxana Smirnova
HEPiX Autumn 2019, Amsterdam*

What is ARC

- Middleware to enable computing grids
 - Motivated by the needs of LHC experiments
 - Main goal: common interface to disparate computing facilities
 - First public preview in 2002, first formal release in 2004
 - Developed by NorduGrid since 2001
 - Designed with a distributed Nordic Tier1 in mind
 - Optimised for HPC deployment
 - Built-in data caching
 - Open Source, mostly volunteer contributors
 - Coordinated by the NorduGrid Collaboration
 - Supported by EU in past, NeIC now (partially)

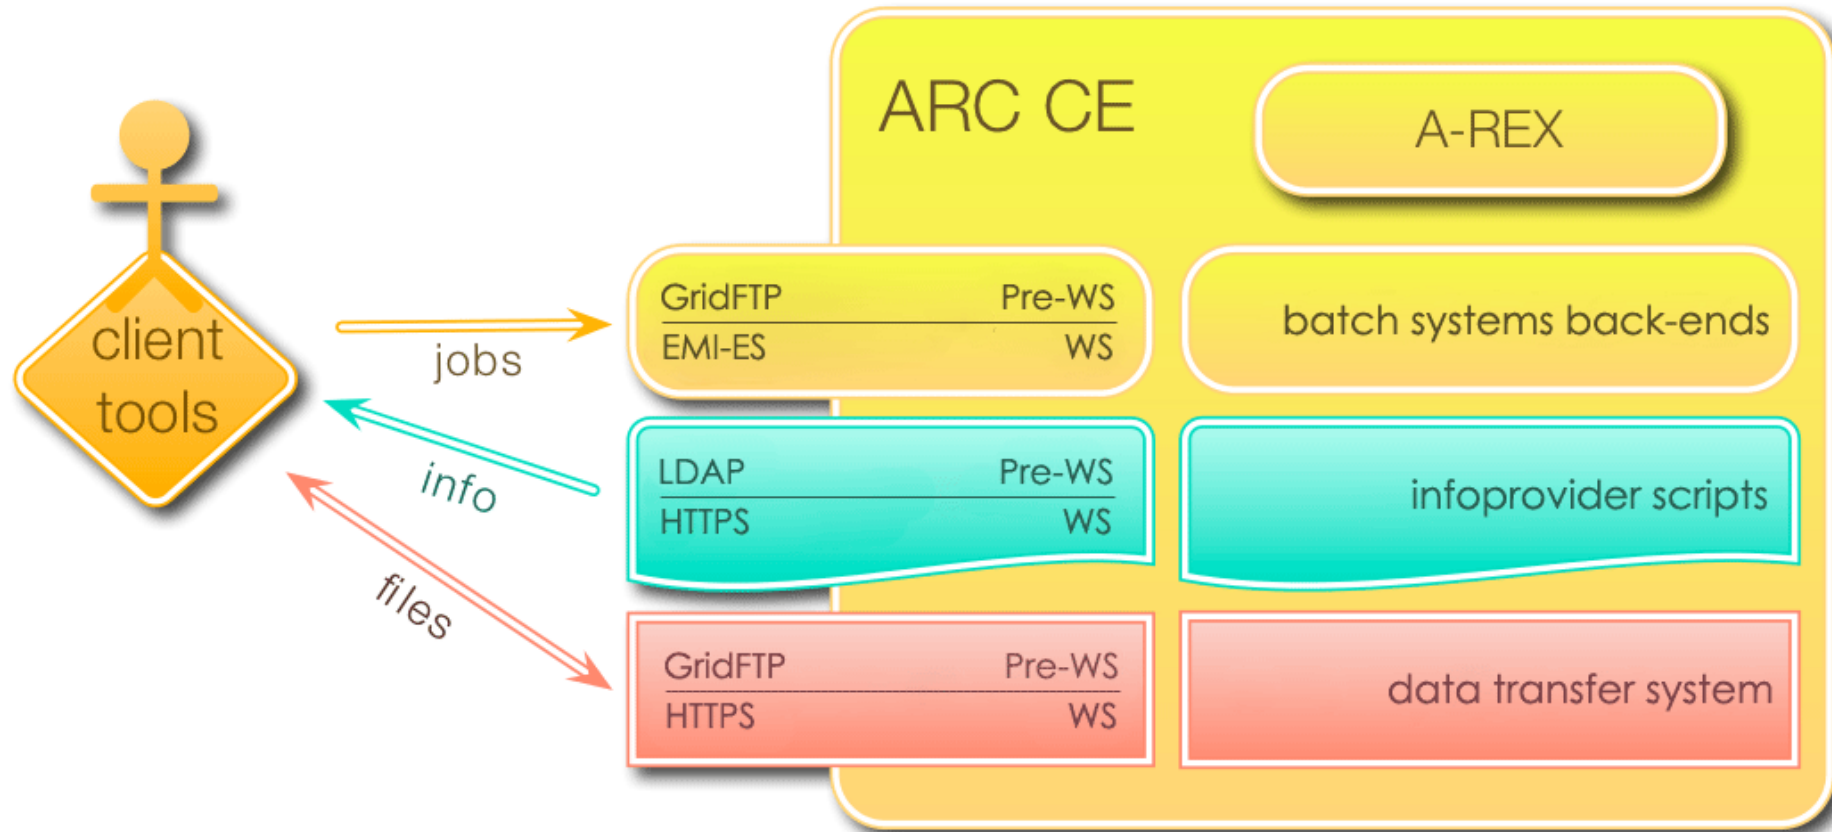


Key ARC components

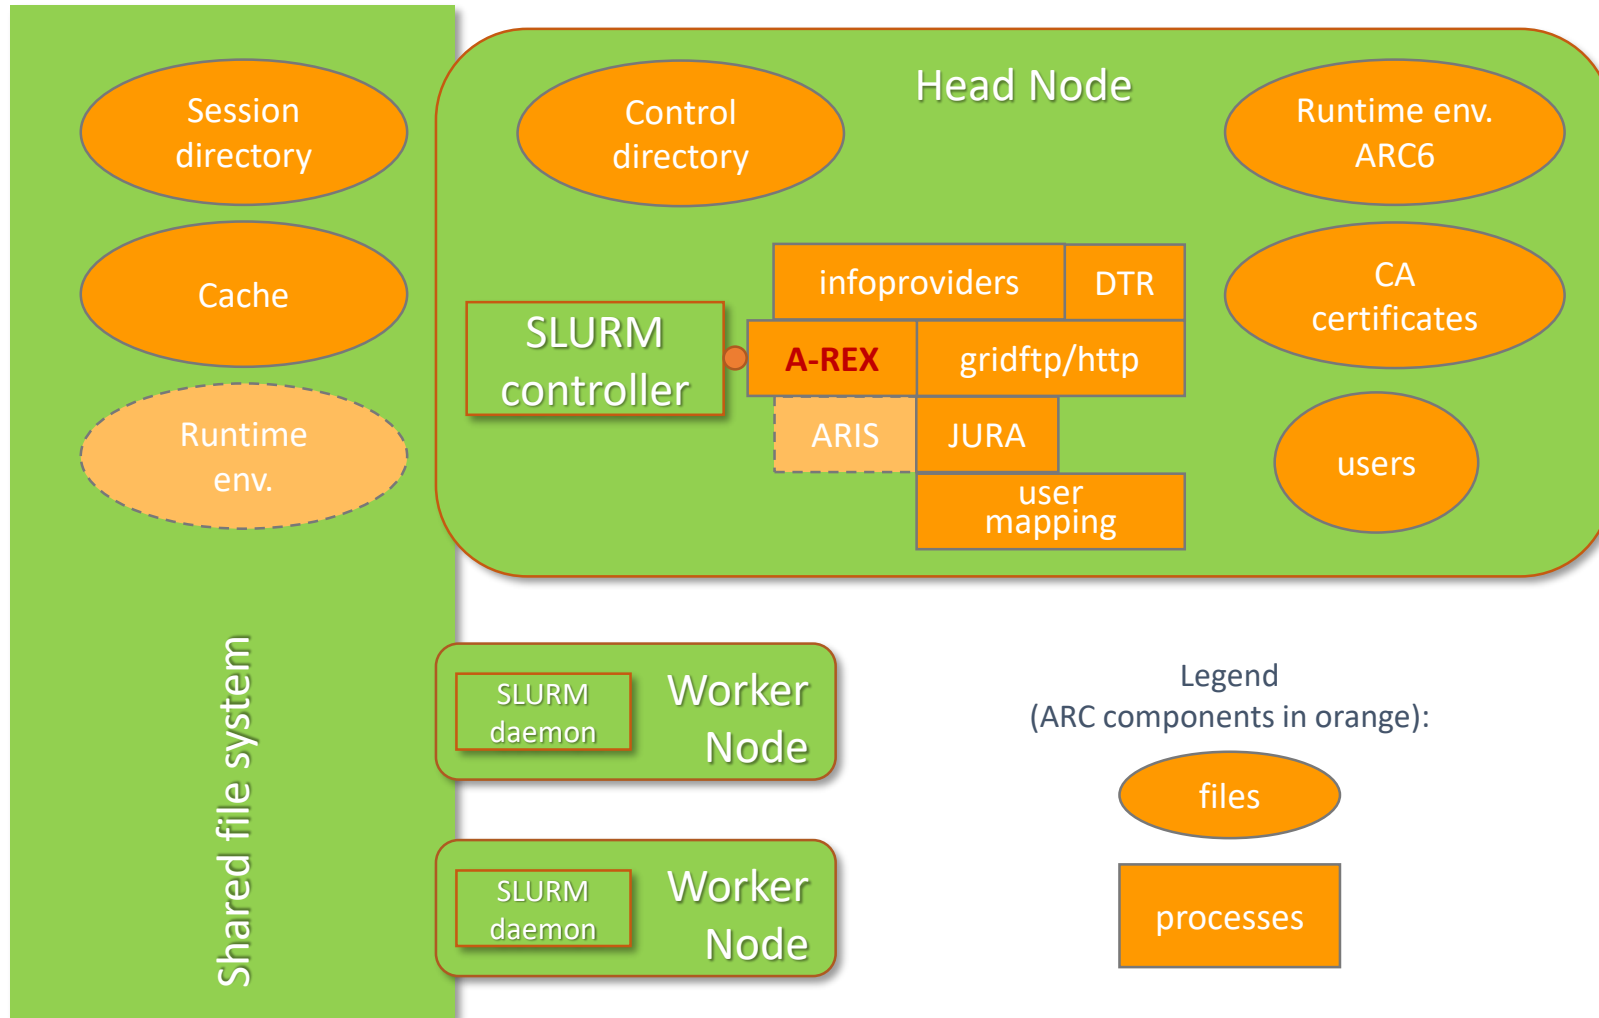
- Key components:
 - **ARC CE** – a Compute Element, providing interfaces to computing resources
 - Modular, consists of several sub-components (services and utilities)
 - Interface for job control
 - Interface for exposing resource and job status info
 - Data staging and shared cache management utilities
 - Jobs do not need to stage data in or out
 - **CLI** client tools to interact with ARC CE and relevant third-party services
 - CLI for jobs management
 - CLI for X509 proxy management (client to VOMS)
 - CLI for file transfer (a wide range of protocols)
 - NEW in ARC6: sysadmin CLI
 - **API**: C++ and Python, for interfacing to full software stack
 - Enables custom services and clients, including arcControlTower (aCT)



ARC CE components and interfaces



ARC-CE components on a (SLURM) cluster



- Supported batch systems and back-ends:
 - SLURM
 - HTCondor
 - SGE
 - PBS/Torque
 - LoadLeveler
 - LSF
 - BOINC
 - Job forking
 - Local submission
 - Remote submission via ssh
- NB: only SLURM is currently available to testers, community help with other systems is much welcomed!

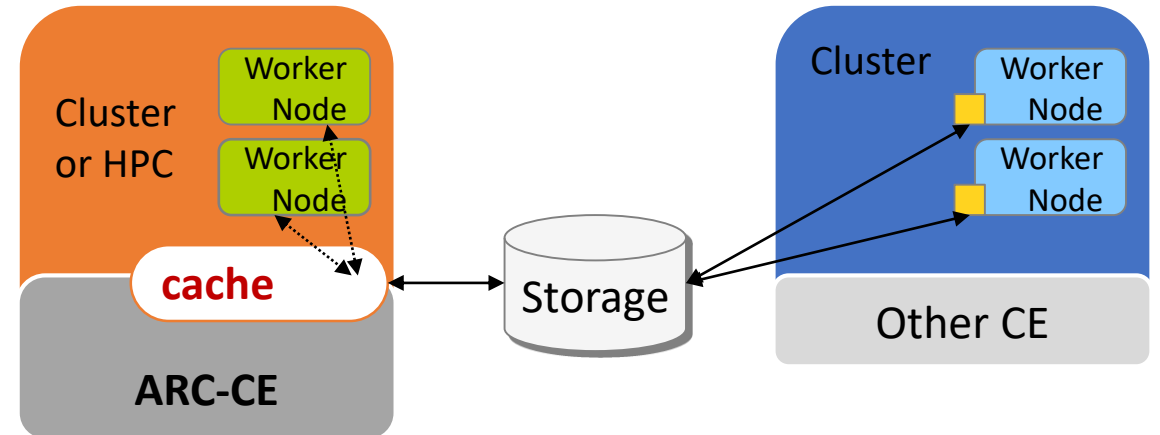
Customisation, optimisation for data-intensive jobs

ARC-CE is a complex service

Highly **customisable**, works on any system

Many sub-components => careful **configuration**

Works best with **high-end storage** for cache



ARC-CE can do all the data transfers

Allows to cache **frequently used** files

Minimizes **bandwidth**

Maximizes worker nodes usage **efficiency**



Supported workflows and deployment modes

- Out of the box: ARC client discovers matching resources and pushes jobs
 - Classical HPC workflow
 - Suitable for smaller communities and testing
- ARC-CE + ARC Control Tower (aCT): aCT acts as a centralised client, still a push model
 - Used for the Nordic Tier1 and some other HPC sites
- “Dry” ARC-CE deployment with pilot jobs
 - No use of ARC-CE data management and cache
 - Most ARC-CE deployments at WLCG are like this
 - Closest analogy to CREAM-CE



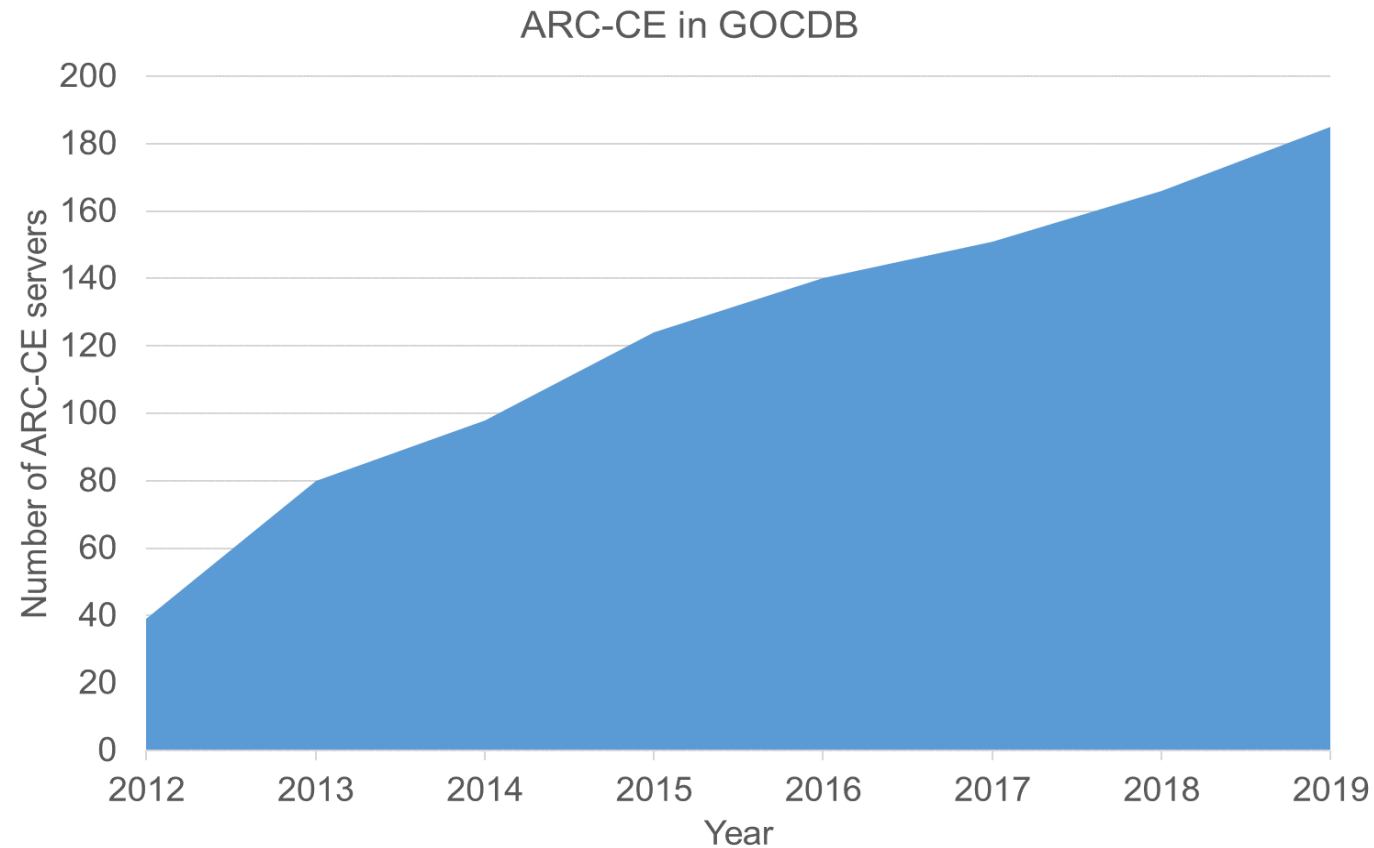
Integration with WLCG and EGI

- ARC-CE is fully integrated into WLCG and EGI operations
 - Registered service in GOCDB
 - Accounting reports sent to APEL by ARC's JURA module
 - GLUE2 info
 - Glue1 is possible but strongly discouraged and not supported
 - Part of UMD releases
 - Undergoes standard EGI certification
 - User support via GGUS
 - ARC unit
- Widely used by ATLAS sites, also by LHCb, CMS, ALICE and smaller VOs that are supported by respective WLCG sites

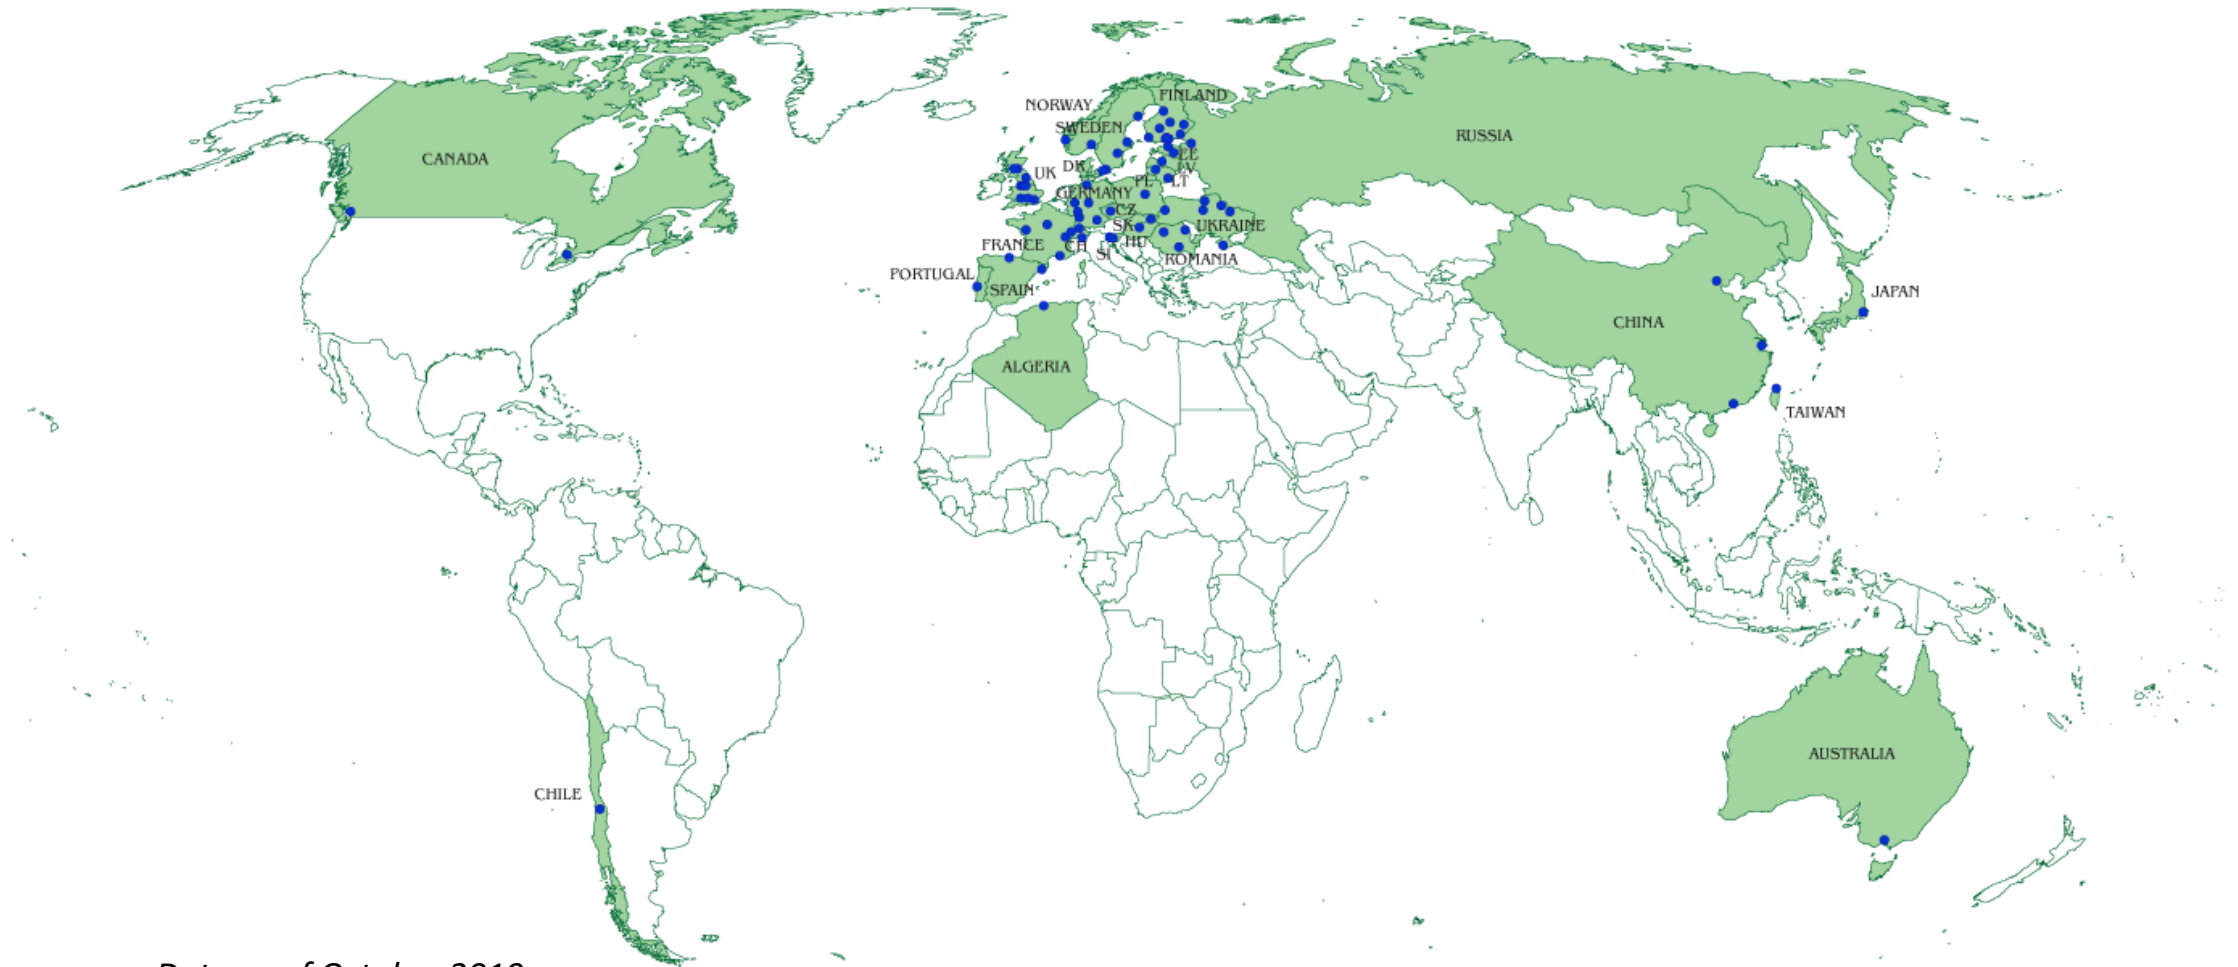


ARC-CE instances in GOCDDB

cf. CREAM-CE: 339 instances, HTCondor-CE: 40 registered instances, as of today



ARC-CE geography



Data as of October 2019



Latest hit: ARC6

- All details in the next presentations!
- Same interfaces, mostly the same design
 - Backwards compatible with clients
 - REST interface added
 - Handy administrator tool added
- Completely different configuration
 - Everything still in a single config file, like in ARC5
 - But a totally different structure, no automatic conversion
 - Default values added
- Some new services and cleaned-up internals
- New code repository (NeIC Coderefinery gitlab) and release procedures



ARC release cycle

- Overview: <http://www.nordugrid.org/arc/releases/>
- Current stable release: ARC 6.3
 - Includes only ARC packages themselves
 - Documentation is online, separate releases for Nagios probes
 - Tagging scheme: 6.m.0 (for backwards compatibility)
 - No distinction between feature and bugfix releases
 - Releases expected to be monthly (or so), but no fixed schedule yet
- Legacy release: ARC 15.03 update 20
 - A.k.a. “ARC5”
 - Contains ARC core packages tagged 5.4.4
 - Includes documents (tag 2.0.21), Nagios probes (tag 1.9.1) etc
 - Since ARC6 is out, will have only critical bug fix releases
 - Apart of that, only the current release is supported, no backporting of fixes to older tags



Documentation, support, availability

- Documentation:
 - ARC6: complete documentation online at <http://www.nordugrid.org/arc/arc6>
 - Continuously improving
 - ARC5: documentation is distributed with the software
 - ARC CE sysadmin guide is the key reference
- Support:
 - For those familiar with GGUS, submit tickets to “ARC” unit
 - For community support, subscribe to either:
 - nordugrid-discuss@nordugrid.org – generic
 - CERN e-group wlcg-arc-ce-discuss@cern.ch – WLCG-specific
 - For bug reports and feature requests, submit tickets to:
 - <https://bugzilla.nordugrid.org>
- Code:
 - <https://source.coderefinery.org/nordugrid/arc>
- Linux packages:
 - Global Linux repositories (CentOS, Debian, EPEL)
 - Upstream: <http://download.nordugrid.org/repos.html>

