

ARC pre-GDB summary

NDGF Manager
Mattias Wadenstein
<maswan@ndgf.org>

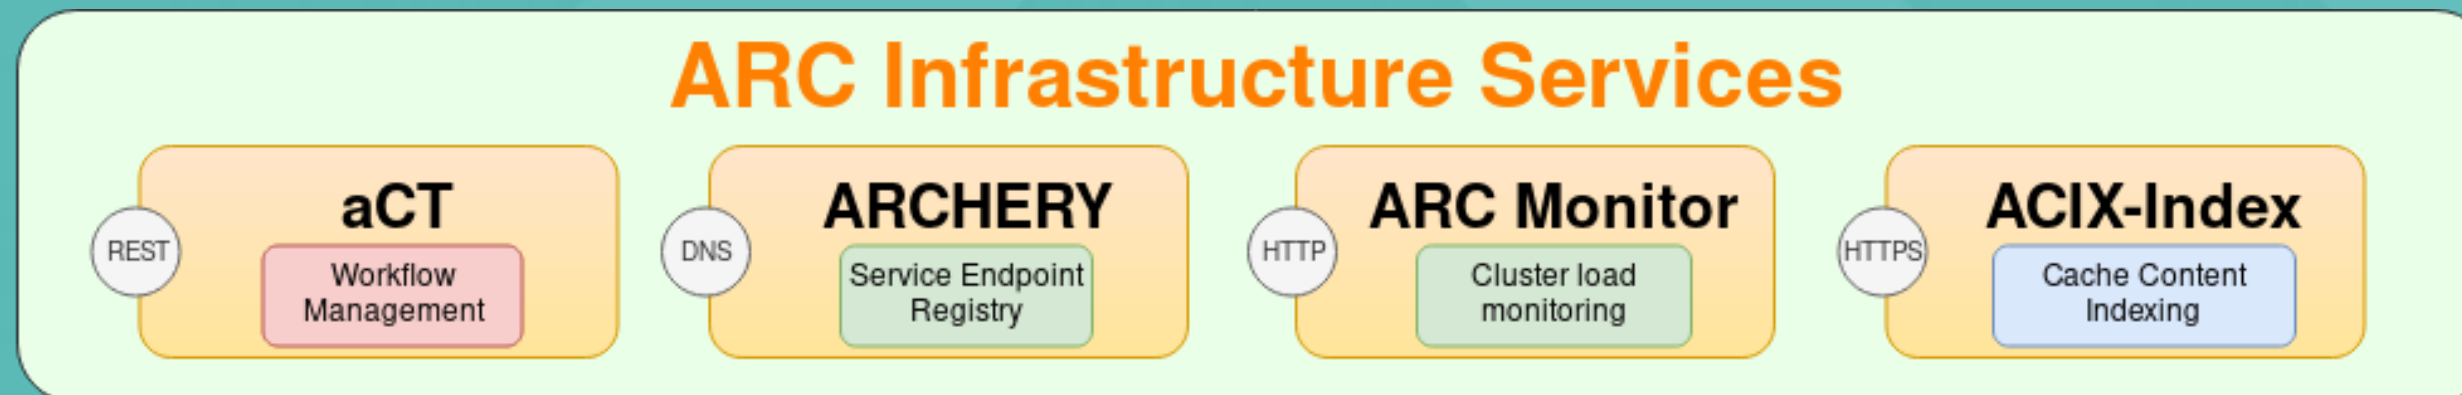
2019-11-20
GDB
CERN

Overview

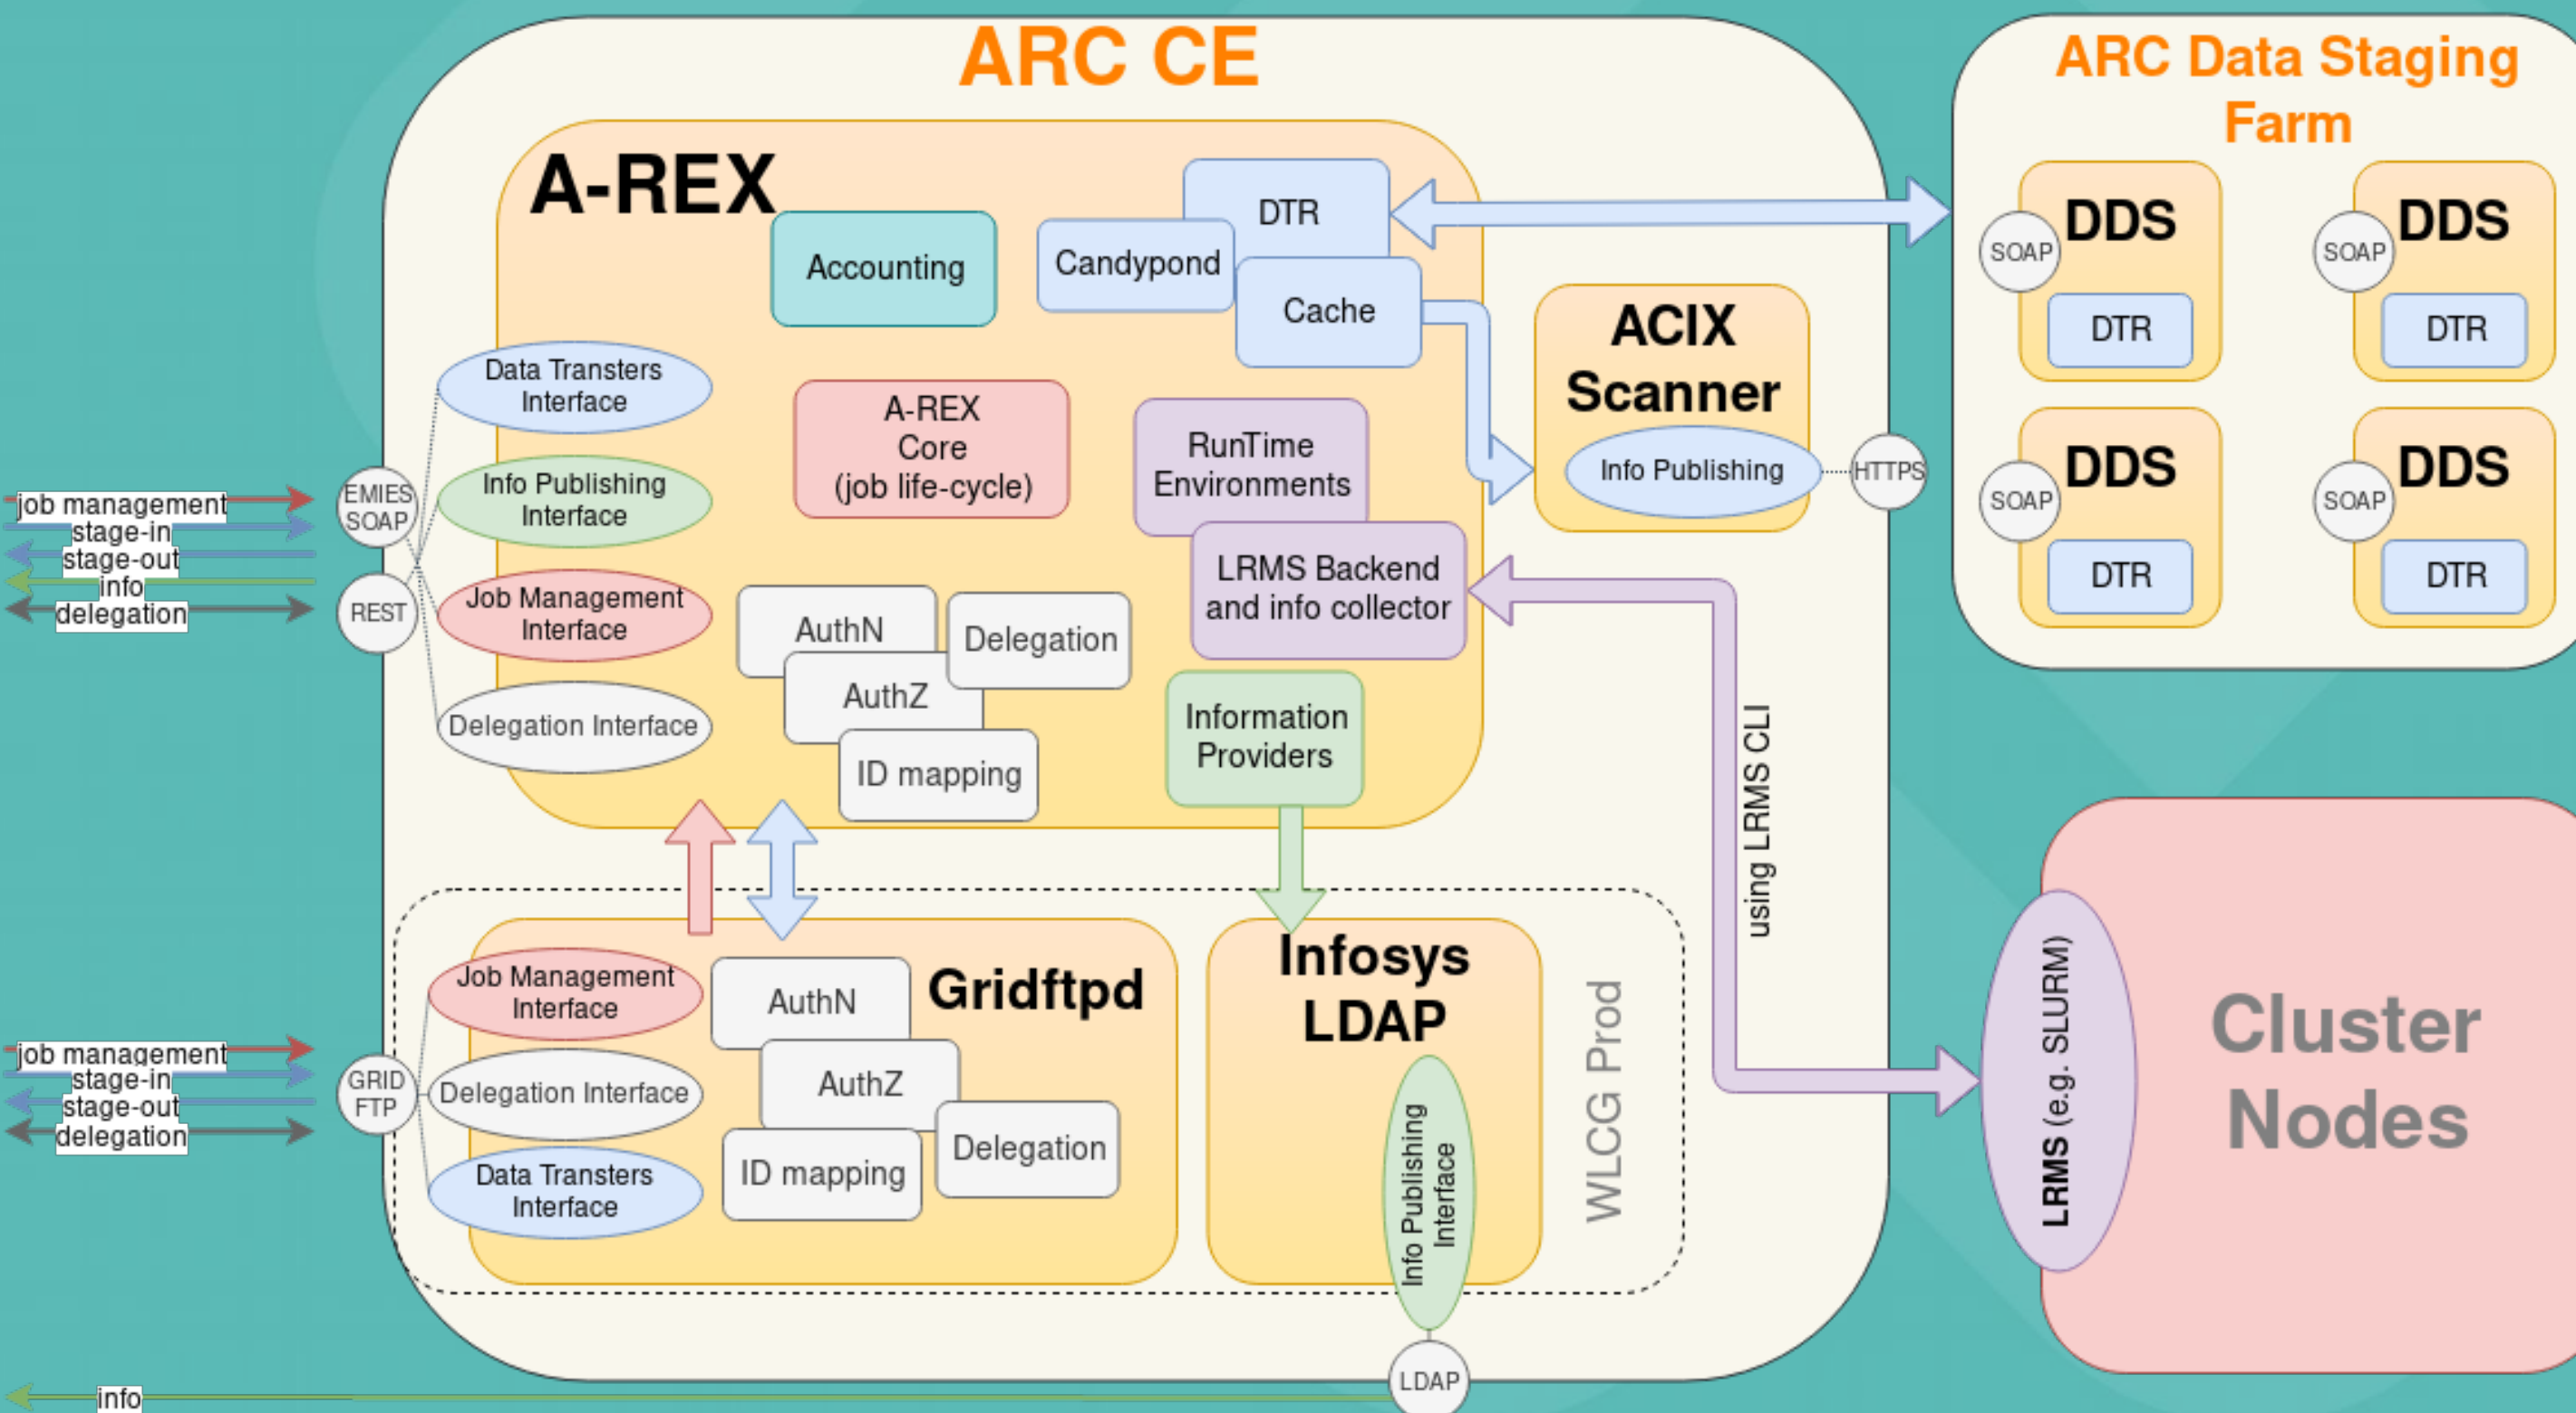
- Full day of talks (10-17)
- Talks
 - Birds-eye overview of ARC-CE
 - Release and development plan
 - ARC for WLCG pilots
 - ARC datastaging and cache, integration with Rucio
 - Deployment and operation experience at the Nordic Tier1
 - Experience with CREAM to ARC-CE migration
 - Accounting and sysadmin tools in ARC6
 - ARCHERY
 - Token-based authorisation plans



Birds-eye overview of ARC-CE



Birds-eye overview of ARC-CE
Balazs Konya
Overview of architecture and power features of ARC



Ongoing developments - ARC 6.x



Current focus: completing the last parts of ARC 6 roadmap

New roadmap to be prepared in the following months

Ongoing development

- Improvements on the CLI subsystem (e.g. review of the cli options)
- Packaging improvements (localizing dependencies into separate packages)
- Standalone arcctl
- Extending the flexible injection and managing capability of Archery-manage
 - Increase sys-admin friendliness

• RESTful interface

- Documented
 - <http://www.nordugrid.org/arc/arc6/tech/rest/rest.html>
- Simplification of WS interfaces - skips all SOAP related things , HTTP based REST
- Can be used with simple clients based on something like curl for example
- Easy to integrate into third-party clients which do not want to adopt ARC library
- Technical preview available as of 6.0.0 - work ongoing to finalize for regular release
- ARC metrics
 - Implement missing DTR metrics in integrated arex-ganglia
 - Provide metrics in json format not depending on ganglia, to be consumed by e.g. Prometheus
- Running jobs in containers via RTE's
 - Docker for HTCondor
 - Singularity containers

Longer-term development for ARC 7

ARC development
Maiken Pedersen
Release strategy and
planned development

Following the technological development

- Support of token-based security layer
- Going towards container orchestrators
 - So far via RTE's – jobs can be run in docker or singularity containers
 - We support docker for htcondor
 - Singularity
- NEW ARC deployment mechanisms: ARC installation via containers
- DATA subsystem scalability improvements
- Controldir scalability improvement
- ARC DATA Service, a sort of DDS2 that can perform data transfers independently of jobs.
 - Modularize the already existing Data Delivery Component as a generic ARC Data service
 - Authorization security layer - token based

- LRMS-less ARC: an ARC CE that can be deployed directly on top of set of dedicated resources
- Autonomous ARC?



Volvo 360c - Autonomous

Always ongoing: improving code-base for scalability and stability

- ▶ **The ARC CE has been operated with the WLCG pilot model of ALICE, ATLAS, CMS and LHCb for years**
 - ▶ Also operated in parallel with the CREAM CE at sites like CSCS, and in a mix of modes
- ▶ **For the experiments, it has been proved to be a drop-in replacement for CREAM**
 - ▶ It integrates seamlessly with the experiment WMSs
 - ▶ The SDK allows the support of the experiment specific frameworks evolving needs
- ▶ **Site admins find it easy to deploy and operate**
 - ▶ Minor adaptations at sites might be needed, mostly depending on the local job accounting scheme for EGI, or for the infosys
 - ▶ Helps managing efficiently resources at sites, specially useful at shared sites

ARC for WLCG pilots
Gianfranco Sciacca
Swiss and experiment
experiences with ARC



Advantages

ARC with datastaging
David Cameron
ARC in datastaging mode
Advantages and usage

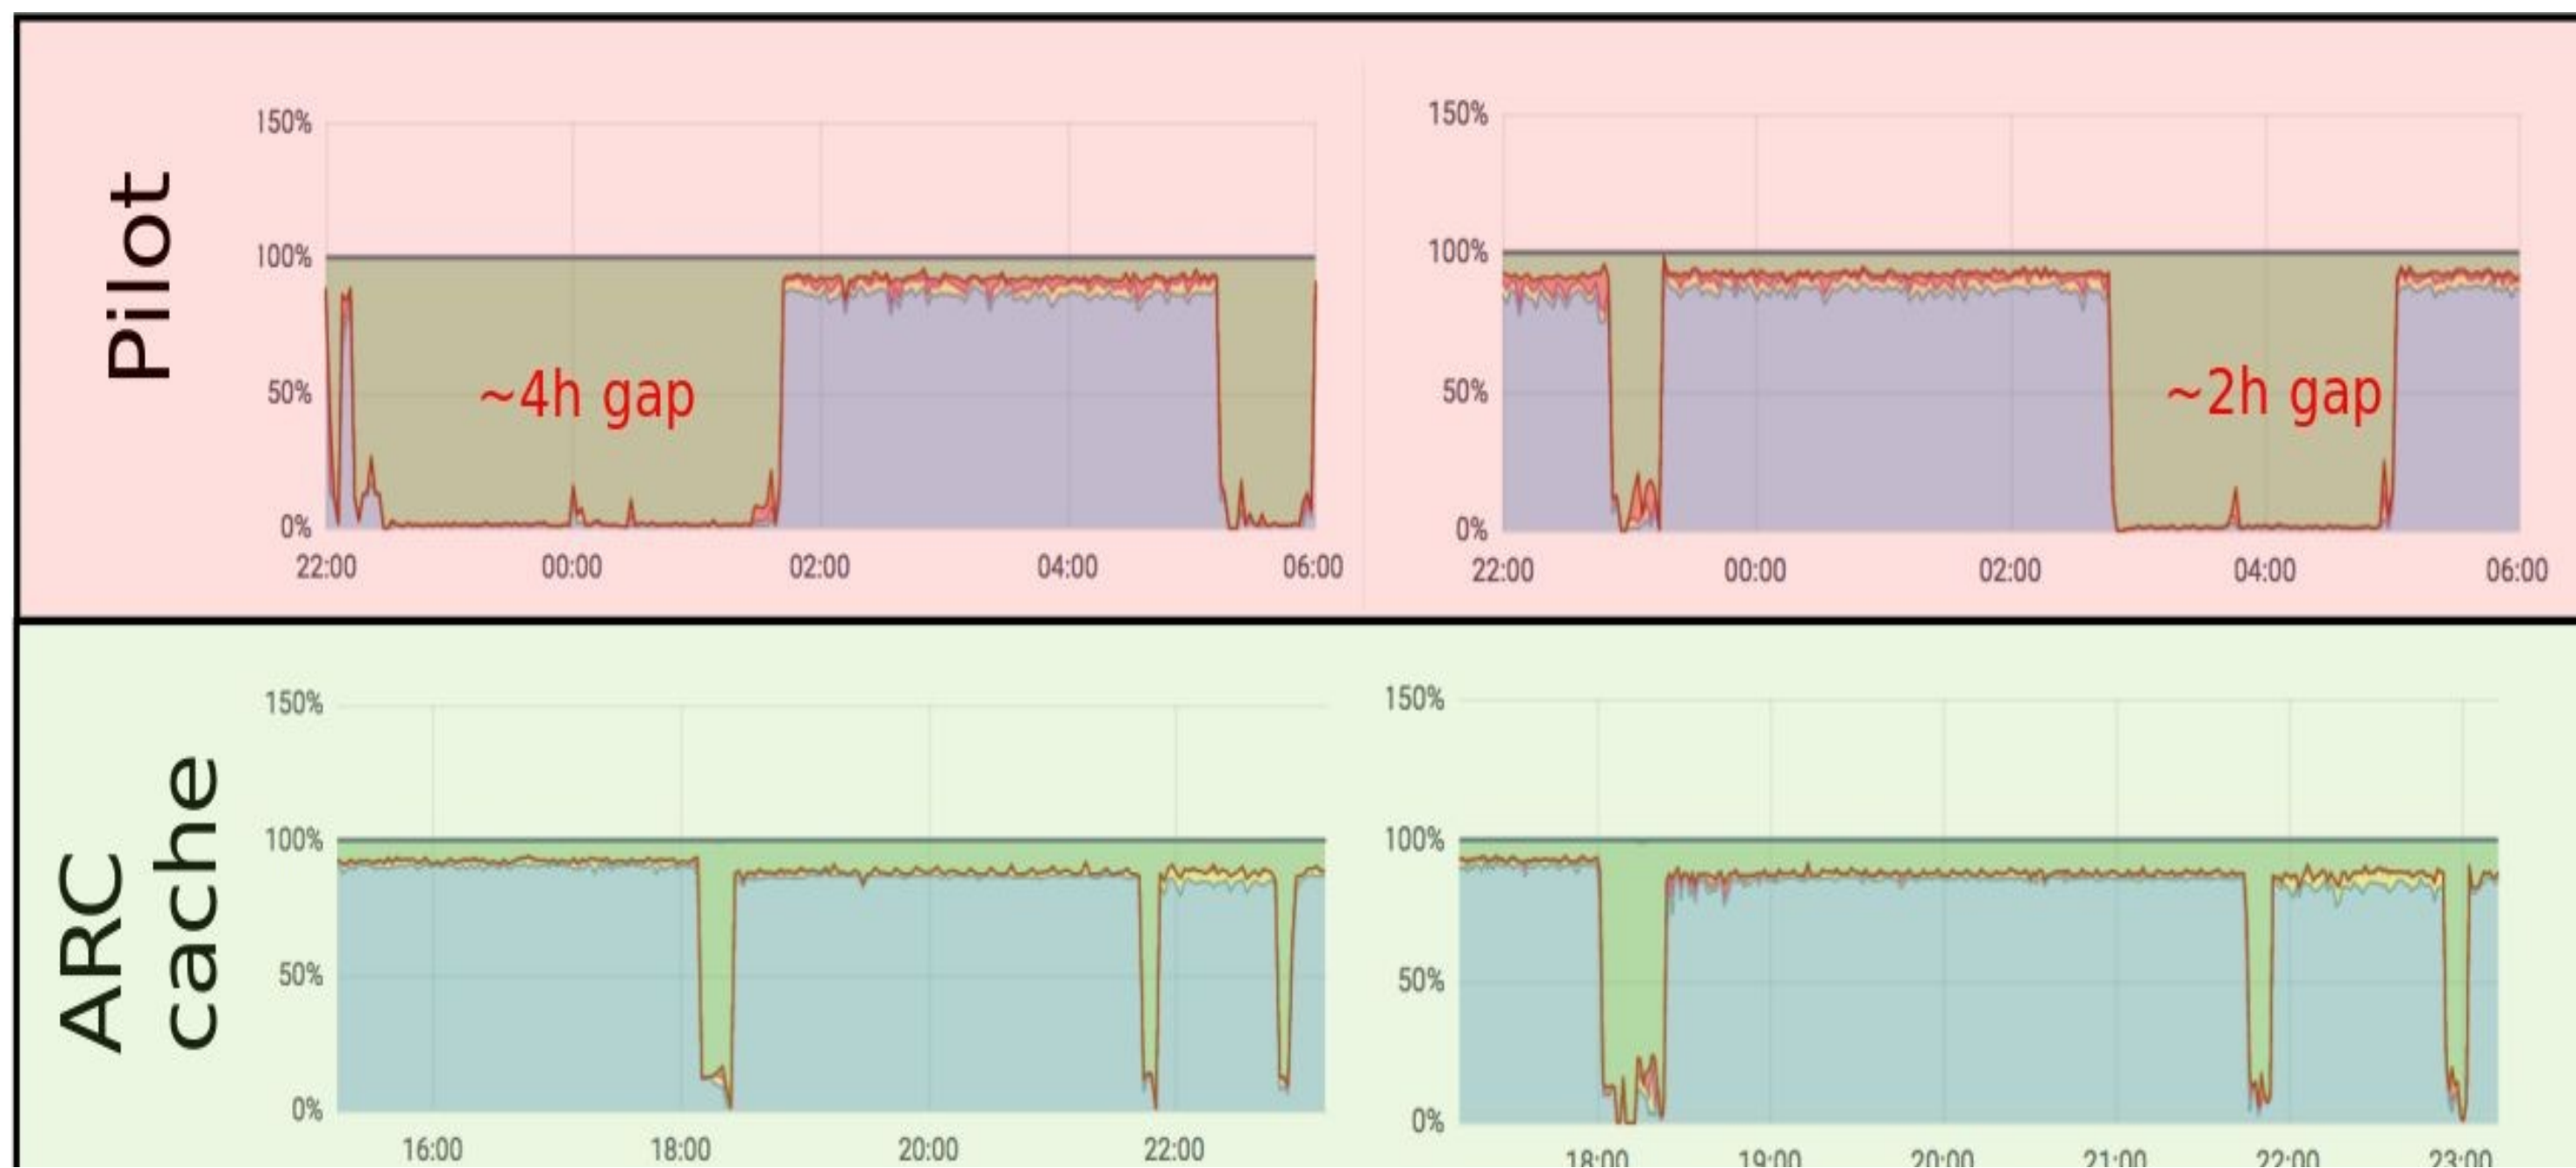
- Overall efficiency
 - Data access is on low-latency local filesystems
 - Download before submission to batch system - better CPU efficiency
- Non-local storage
 - Like NDGF with distributed storage
 - Or a “compute only” site
- Limited external connectivity
 - Like HPC sites where external connectivity might be blocked or only available through a slow NAT
- No need for grid-aware computational software



Efficiency study

ARC with datastaging
David Cameron
ARC in datastaging mode
Advantages and usage

- Comparison of pilot jobs and ARC caching presented by Maiken at CHEP
- Pilot staging from remote storage lowers CPU efficiency
- Asynchronous staging by ARC CE eliminates the holes in CPU usage





Efficiency study

ARC with datastaging
David Cameron
ARC in datastaging mode
Advantages and usage

Comparison of NDGF sites CPU efficiency



Operations

- The site view is a fairly reliable service
 - More issues than apache or openssh
 - But on par with SLURM, Lustre/GPFS, etc
- Responsive development team
 - Via the nordugrid-discuss mailinglist
 - Often easy to get a test build with a fix for your issue fast
 - If you have identifiable symptoms, “sometimes it is slow” might be more problematic
- The old and deprecated parts are worst
 - Like Glue1 that keep recurring as something needed by sites

NDGF experience
Mattias Wadenstein
Deployment and operation
experience at the Nordic Tier1



Operations

NDGF experience
Mattias Wadenstein
Deployment and operation
experience at the Nordic Tier1

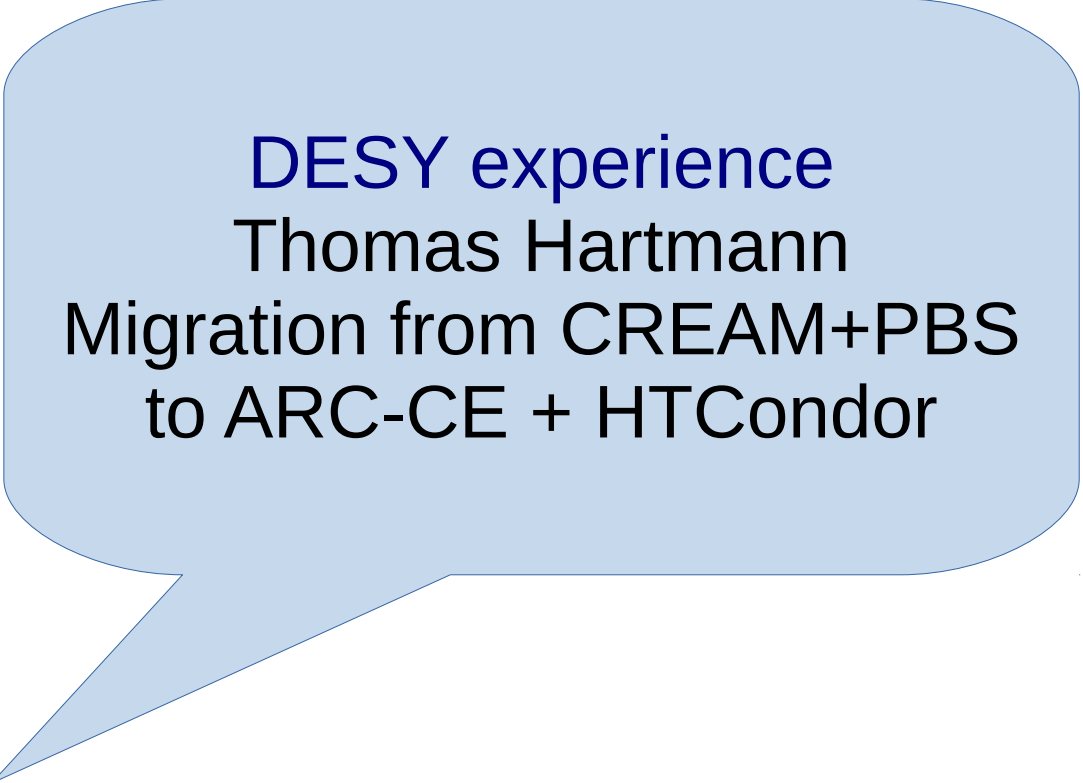
- Caching mode also hides storage downtimes
 - No running jobs affected
 - New jobs will just stage the files currently available from whatever sources registered in rucio
 - A couple of hours of full srm.ndgf.org downtime or a day or two of a few pools down before preparing jobs starts to queue up
- Setup for a new site not quite trivial
 - Biggest complexity in auth, due to the many choices involved
 - Map to local unix user based on DN, VOMS, callout to site service, combination?
- Once set up, tends to just keep working
 - Random breakage is rare



Experiences etc.

Subtitle here

- migrated from 2015 both CE/LRMS
 - from CREAM CE + PBS/MySched to ARC CE + HTCondor
- two main ARC CE nodes + pre-prod/testing/... ARCs
 - no shared fs
 - config in Puppet/Hiera
- HTCondor
 - ~20k cores
 - supporting ATLAS, Belle II, CMS, ILC, LHCb,...
- generally: when ARC CE are running, they are running :)
 - not much intervention needed



DESY experience
Thomas Hartmann
Migration from CREAM+PBS
to ARC-CE + HTCondor

ARCCTL: CE under your control

Deployment Automation

- CA Certificates
- VOMS
- Firewall
- Package install

Service Control

manage
ARC services based
on configuration

Configuration Management

- overview
- validation
- in-line help
- runtime config

Jobs Control

- get info, logs, script
- kill and clean
- ownership
- statistics

arcctl

Accounting, Usage Analysis

- who is using CE
- how they using CE
- how much CPU, memory, disc, network

Worker Node Environment

- features available during job runtime
- environment setup
- containers support

Data Staging

- cache control
- transfers stats

ARC Control Tool
designed to automate
ARC CE operations
for sysadmins

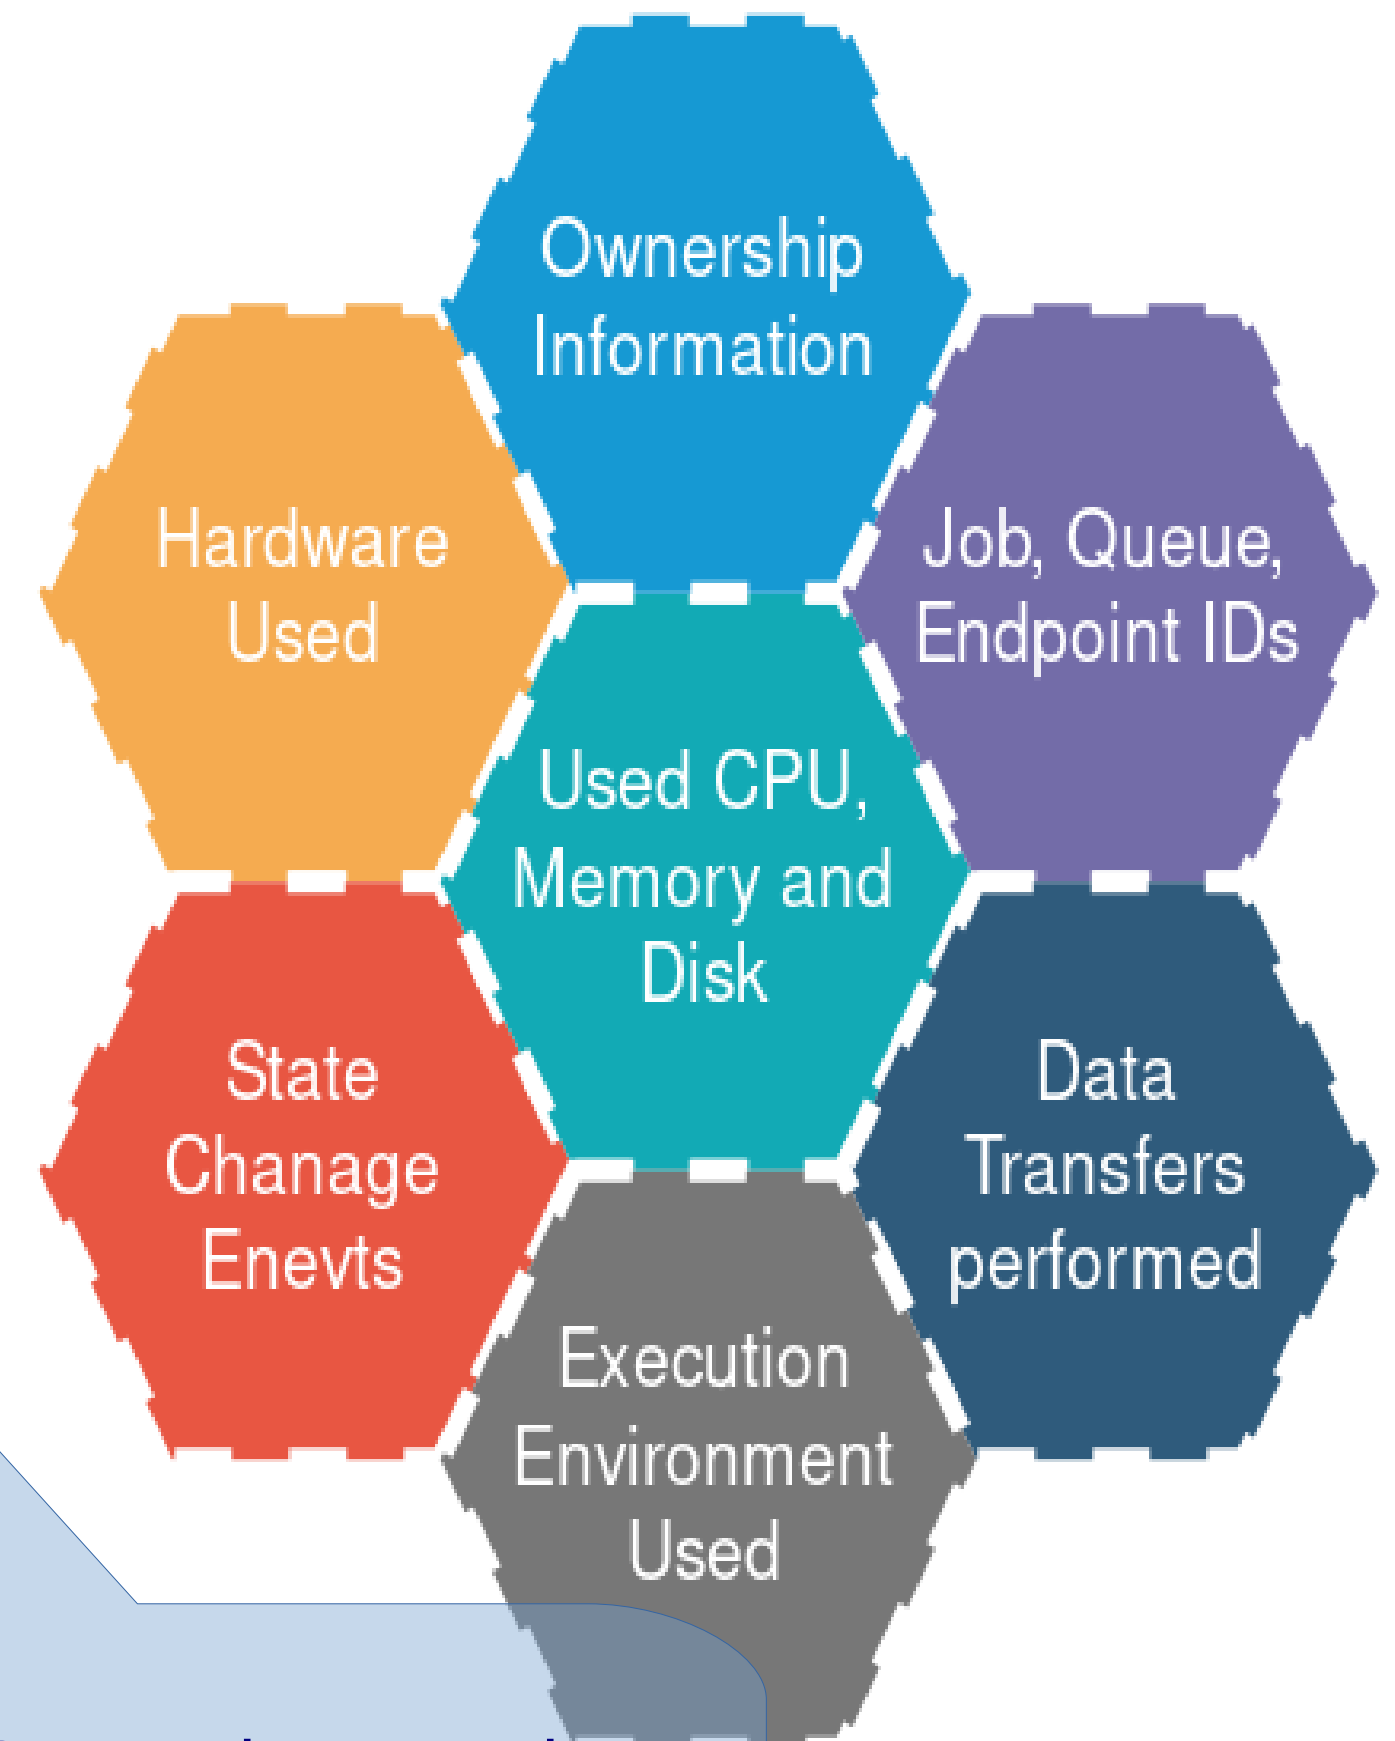


ARCCTL and Accounting
Andrii Salnikov
Showcase of arcctl and
the new accounting DB

Accounting: keep track on resource usage

- ARC CE accounting subsystem collects data about **per-job resource usage** in the local **SQLite database** for **on-site CE auditing** and **publishing**
- A-REX Accounting Record* (**AAR**) defines data stored about the job

**Complete AAR Definition can be found in the documentation.*



ARCCTL and Accounting
Andrii Salnikov
Showcase of arcctl and
the new accounting DB

The Idea

- Understand what is really needed for **service discovery**
- Define the **minimalistic data model** of e-Infrastructure services
- Store the whole thing **embedded in DNS records!**



ARCHERY
Balazs Konya
The new service registry
and service discovery system

Summary

- **Good discussions**
 - Clarified both current status and future plans
 - Participation by most LHC experiments and several sites
- **Valuable feedback to the development team**
 - Where current issues are
 - That Glue1 can finally be abandoned
 - How to proceed on token authentication
- **Hopefully useful for sites too**
 - Experience sharing
 - Developer intent
 - Half the pain points listed already fixed in ARC6





Questions?