The Information system for LHC Distributed Computing

Alexey Anisenkov (BINP)

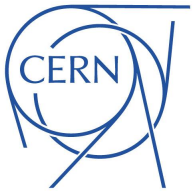

GRID-2018, JINR, Dubna, 11 September 2018

Outline

- The role of Information system in Distributed Computing Environment
- AGIS → CRIC framework
- CRIC main features
- CRIC plugins for WLCG Experiments/other Collaborations

Distributed Computing Environment: Worldwide LHC Computing Grid (WLCG)

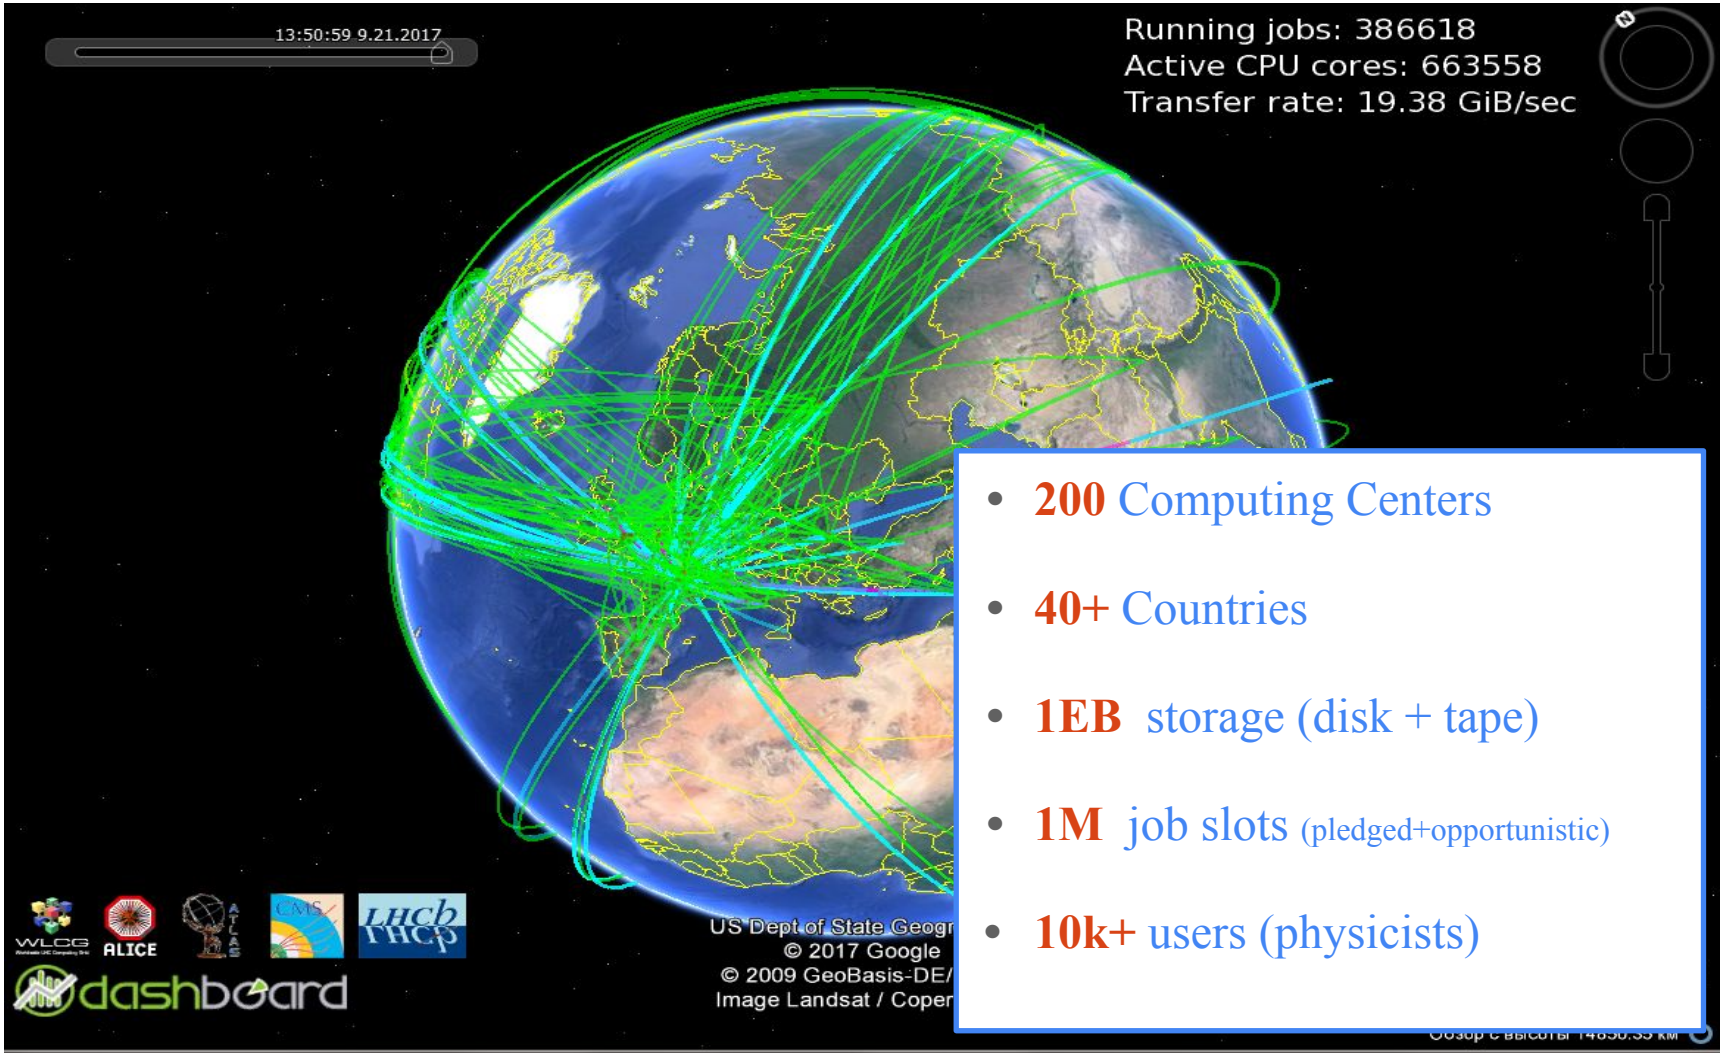

- **200** Computing Centers
- **40+** Countries
- **1EB** storage (disk + tape)
- **1M** job slots (pledged+opportunistic)
- **10k+** users (physicists)

International Collaboration of Computing centers located across the world to distribute and analyse LHC data

Distributed Computing Environment (Resources)

LHC Experiments rely on **heterogeneous** distributed computing

- variety of computing resources involved

Pledged resources

Rented, on demand

Opportunistic

Opportunistic backfilling

- variety of infrastructures and middleware providers

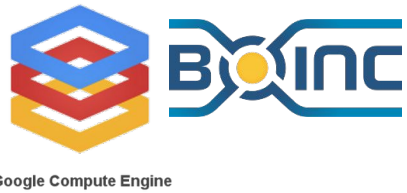

others ..

Distributed Computing Environment (Experiments)

- Each **Community** uses and describes **Resources** in its own way

...

  - Computing Models are similar but still have different implementation
  - Various high level VO-specific frameworks & middleware services (e.g. for Data and Workflow management)
  - Cross experiments applications (monitoring, accounting, testing frameworks, resource usage descriptors, etc)

- Apart from resources description, high level VO-oriented middleware services and applications also **require** the diversity of common configurations to be centrally stored and shared

Resource Configurations & VO applications

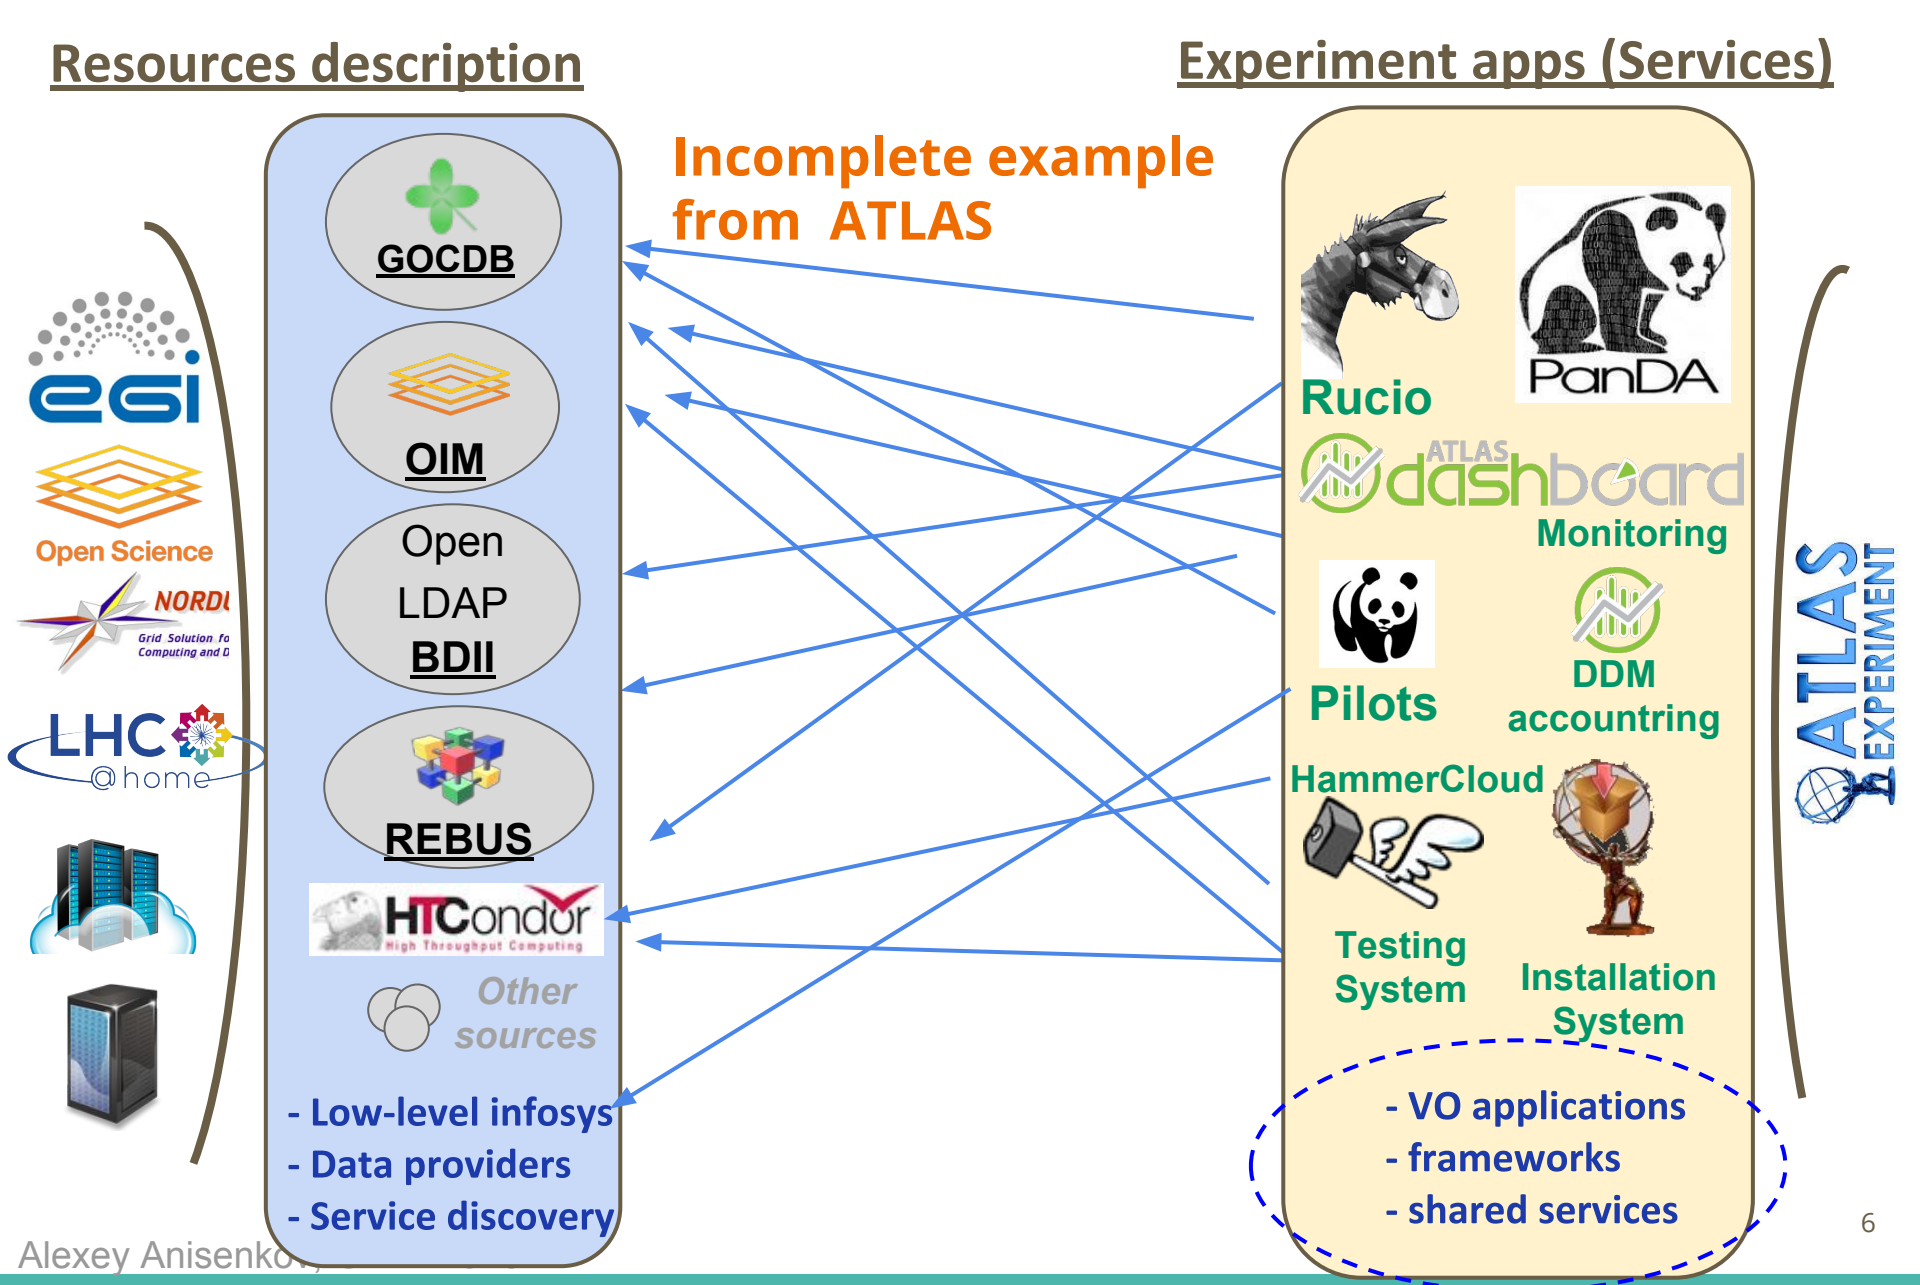

Resource & Services: Gluing them together via high level Information component for WLCG VOs

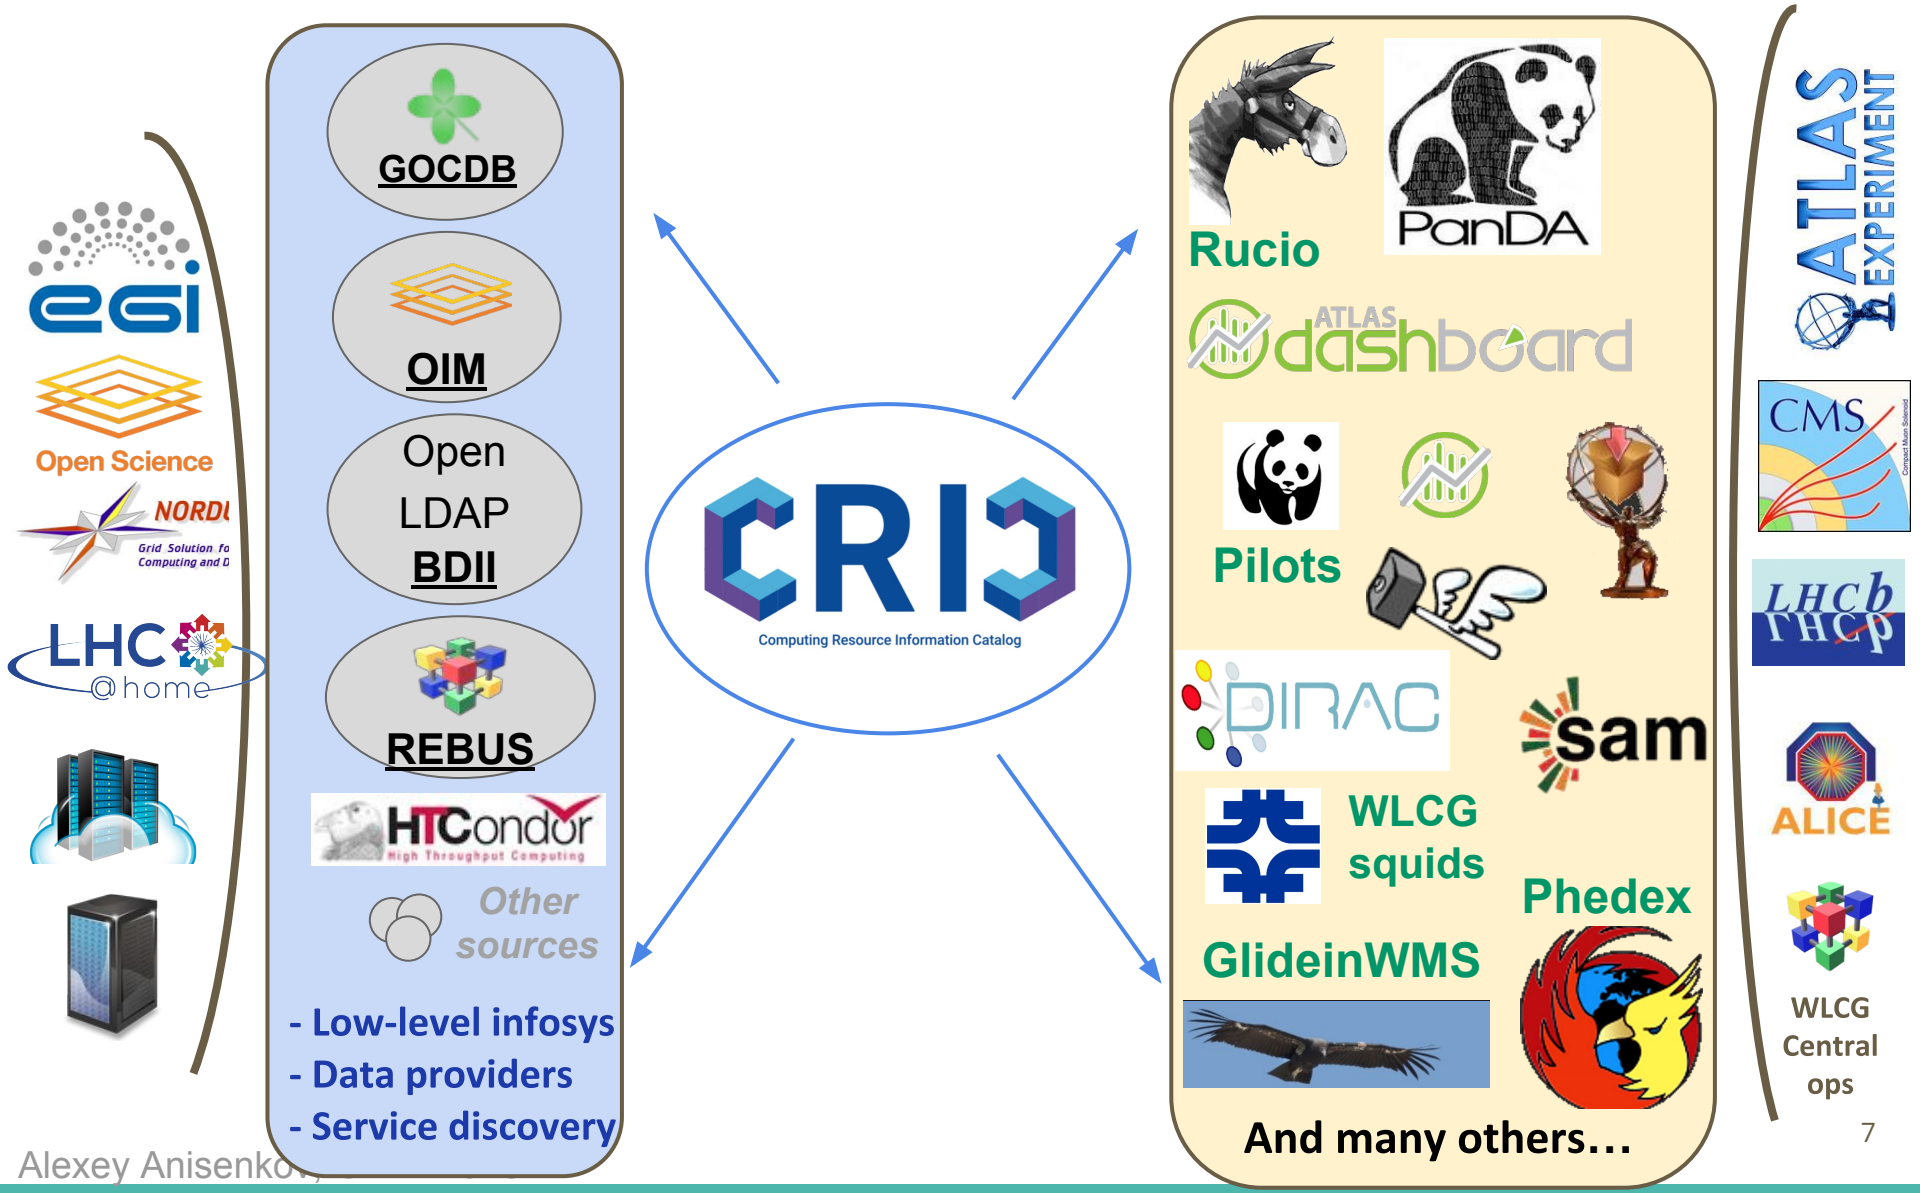

Resource & Services: Gluing them together via high level Information middleware

The concept of central information Experiment-oriented system has been originally implemented for ATLAS in the ATLAS Grid Information System (AGIS)

- Designed in 2009/2010, in full production since LHC Run-1
- ATLAS oriented. Connects **Resources** and **Experiment frameworks (services)** together for the ATLAS Collaboration
- Integrates configuration and status info about resources, services and topology of the whole Computing infrastructure used by ATLAS Distributed Computing

Computing Resource Information Catalog (CRIC):

- 2016/17. Next-generation system, the evolution of AGIS framework
- Not experiment specific (but still experiment oriented)
- Generic solution for the LHC experiments sharing same resources
- Can serves single VO with dedicated resources (e.g. COMPASS)

Why CRIC? Current limitations of the WLCG Information infrastructure

- Multiple sources of information. Data is sometimes contradictory or incomplete. Debugging is complicated.
- No central place where data can be validated
- Integration of new type of resources is not straightforward
- Complex objects like Storage services with variety of access protocols and storage shares are not properly described

No high-level information middleware which completely covers Experiments use-cases and describes resource as VOs need

- **Currently every experiment has to solve all those problems on its own, and resource providers has to deal in different ways with each experiment**

CRIC: a unified Information system

Computing Resource Information Catalog

Focused to describe:

- **the topology of all the WLCG infrastructure (resources provided by the WLCG sites)**
- **experiment-specific configurations required to exploit this infrastructure according to the VOs Computing models.**

CRIC is a framework providing a centralized (and flexible) way to describe which resources LHC experiments are using and also how they use them:

- Built-in information Model for Resource descriptions and Experiment specific objects
- **Experiment independent, but still experiment-oriented**
- Plugin based approach allows customization to address various experiment requirements and implementation of the dedicated experiment instances
- Shared building blocks to optimize development process and to ensure common look and feel. *Think about it in terms of lego bricks*
- Flexibility to address technology evolution and changes in the experiment computing models and applications. *Lego bricks again!*

Core feature of Information model: “provided by” vs “used by” resources

- Clear distinction between resources ***provided by*** (Sites) and resources ***used by*** (Experiments)

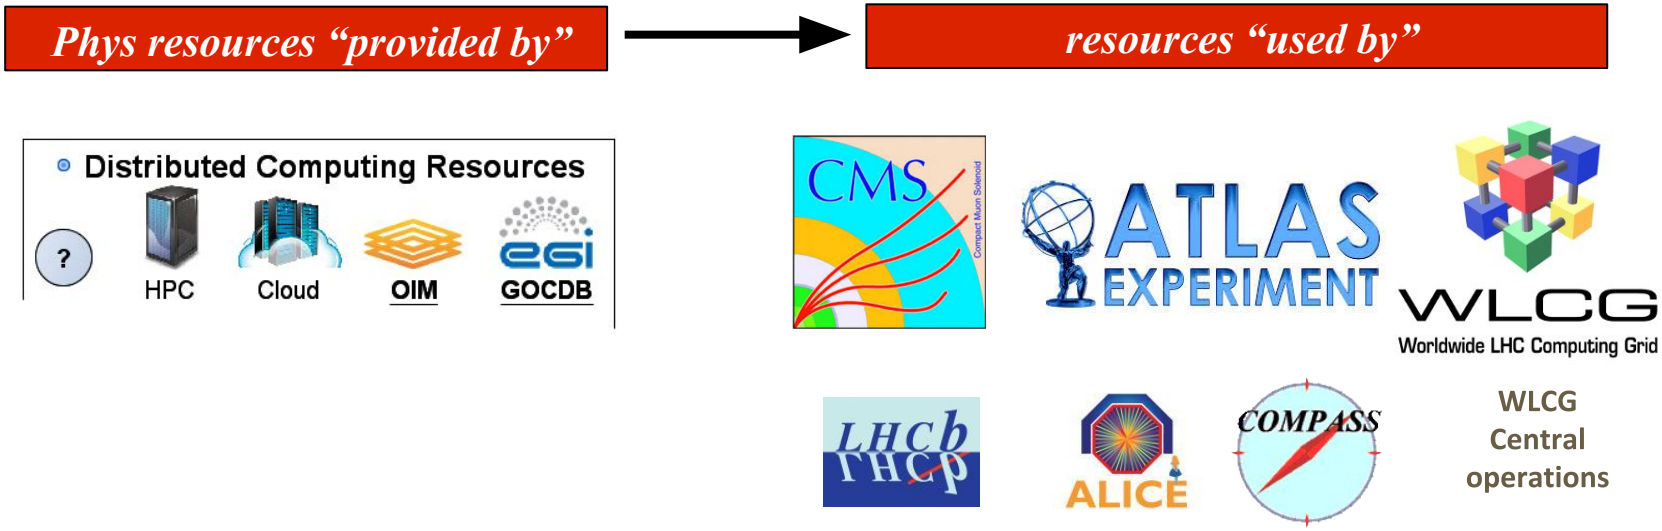

By Providing such abstraction layer from the physical Resources CRIC allows Experiments to define their own real organization of resources and required experiment specific structures.

CRIC Architecture: plugin based

- Modular architecture is based on the Django framework, implementing different apps for CORE (provided by) and experiment (used by) parts
- Data are exposed via REST API which is configurable by filters and various presets (views).
- Bootstrap, jQuery, Web services and many other modern tools and technologies involved

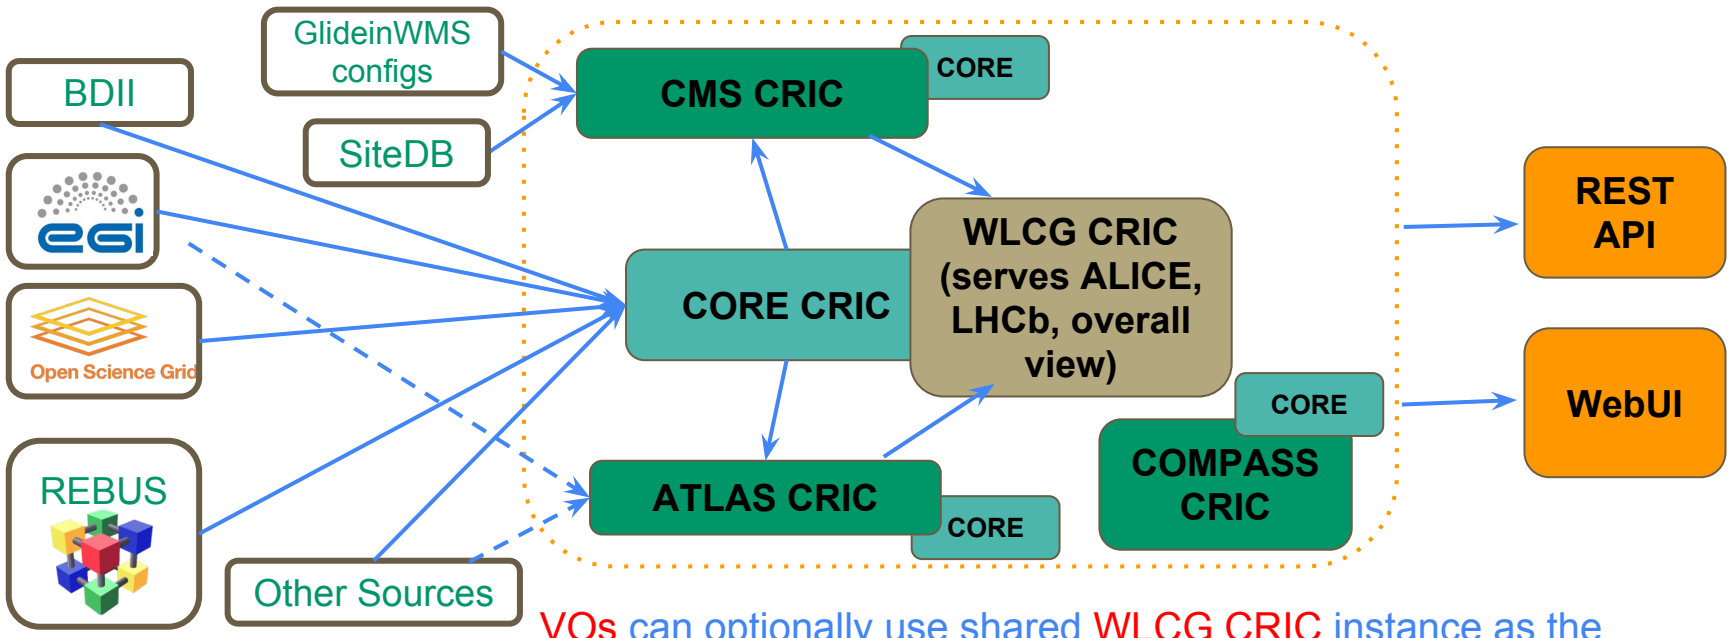

VOs can optionally use shared WLCG CRIC instance as the source or fetch data directly from low-level info providers

Authorization and Authentication (A&A)

- CRIC supports enhanced **Access controls** and user Group management
- Several Authentication methods enabled (SSO, SSL, VOMS, local)
- Flexible utilisation of Permissions, Roles and Groups at various levels
- Fine grain A&A on the level of object (class, instance, global permissions)
- Ability to bootstrap User info/DB from whatever external source (CERN DB, Experiment DBs, config files, e-groups, etc)

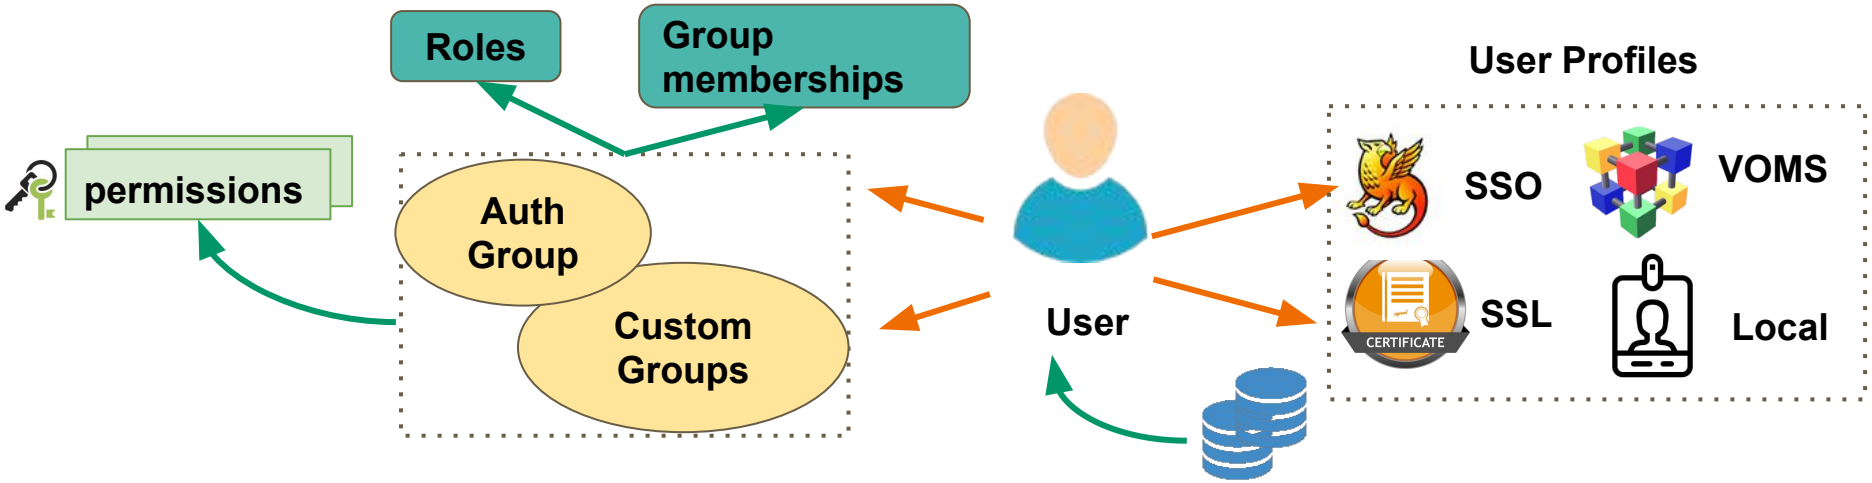

Each Experiment could configure own Data access policies!

Example of A&A use-cases for different VOs

- **CMS** considers CRIC not only to define access rights within the system, but also to control user privileges for **CMS applications** (CRAB, WMAgent, Phedex, etc…). Relies on CERN SSO and local authentication.
- **ATLAS** uses a simpler Auth concept based on user's DNs coming from VOMS

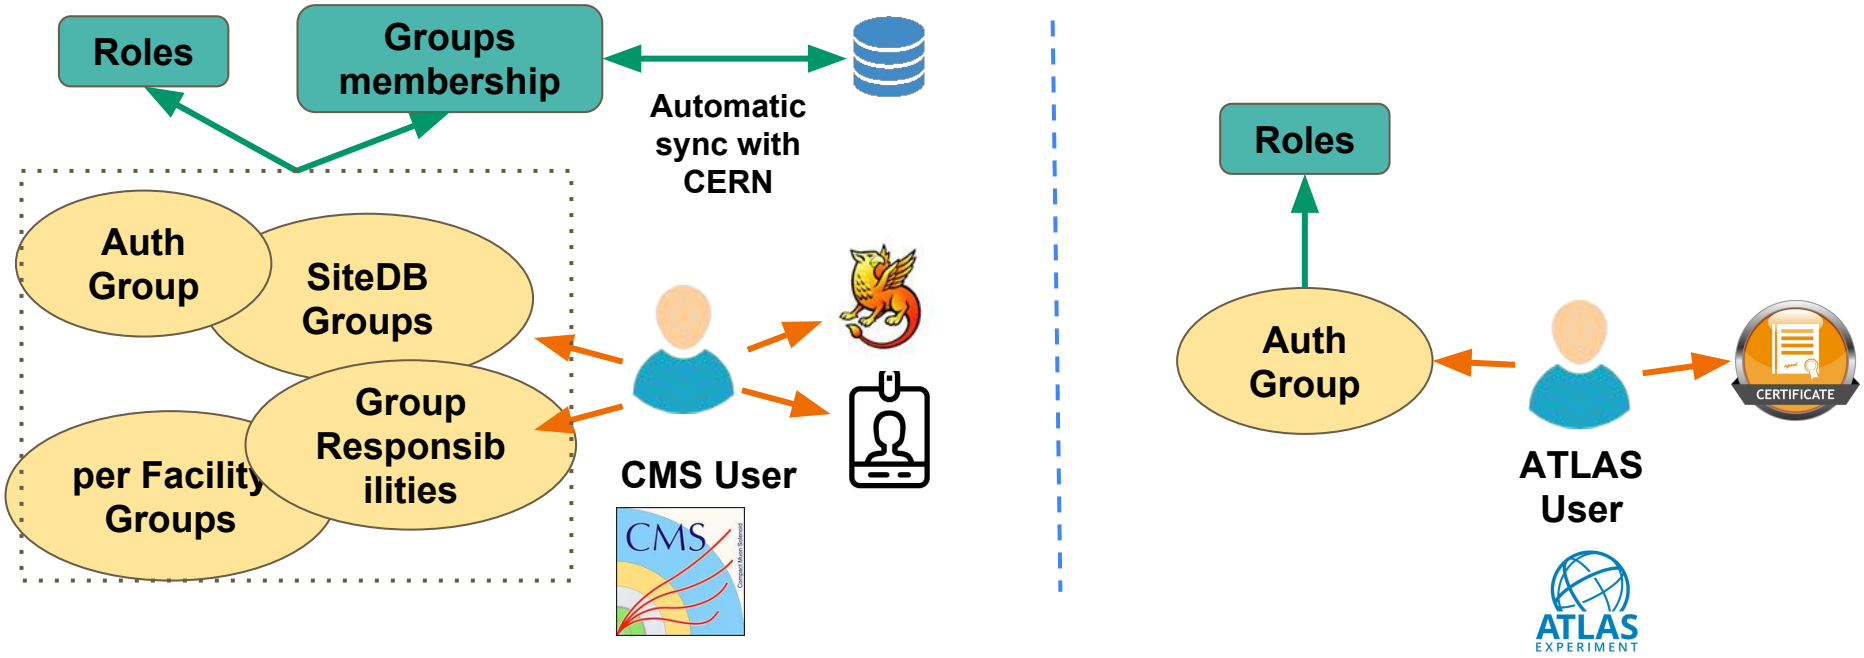

Experiment decides what elements should be used out of the CRIC box to implement own policies and follow own workflow.

Logging functionality - track the changes

- CRIC provides an advanced logging functionality to monitor, administer and troubleshoot the system
- Logging is performed at the object level (a given object or any other objects related to it)
- Full list of history changes (including old values) is provided through build-in WebUI views.
- You can check **who**, **when** and **how** interacted with an object.

Common concepts. Storage example

CRIC introduces unified description of **Service Resources** in particular to target and resolve the complexity of SE definition

- Link together all protocols, activities, closeness metrics, space tokens, other experiment specifics belong to same Storage into unified Resource
- Multiple protocols concept
- Connect associated CE to default SE for given activity
- Integration of new SE technologies (e.g. ObjectStores) within the experiments

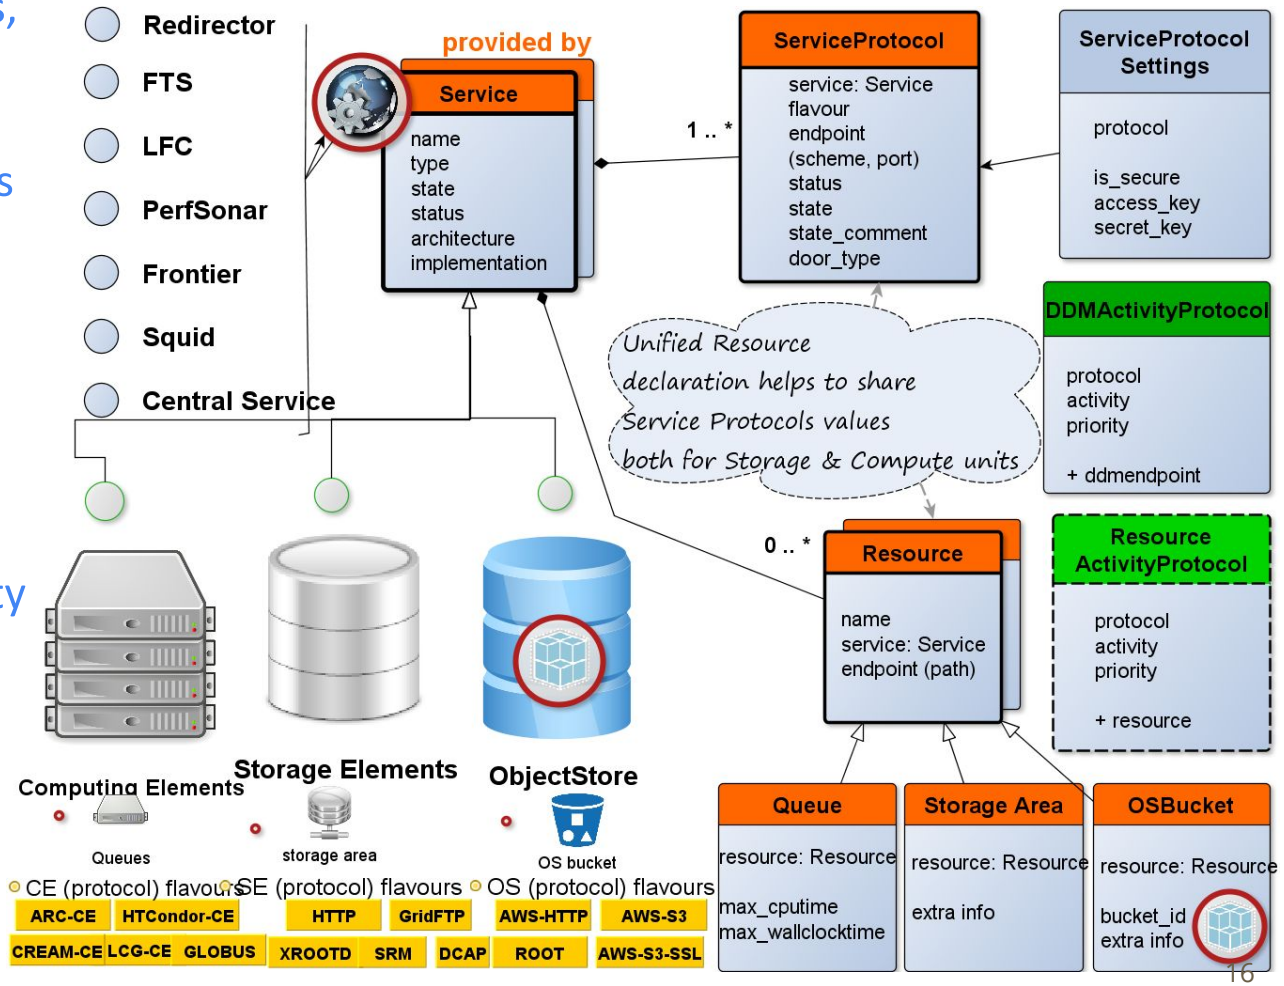

Experiment specific concepts. Computing resource example

- **Experiments use different systems to submit jobs into the grid**
- **Different workflows and data structures**
  - **CMS** submits Pilots through GlideinWMS (Job submission via CRAB3+WMAgent). Currently GlideinWMS configuration is described in the XML files hosted by Github.
    - This configuration has been already imported into CMS CRIC
  - **ATLAS** submits Pilots through AutoPilotFactories (APF), ARC ControlTower (aCT), and Harvester. Job submission via PanDA.
    - These frameworks require conf files which describe the Computing Elements (and batch system settings underlying)
    - Currently the configs are auto-generated from AGIS, exploiting the resource description in "core", and later will come from **ATLAS CRIC**

CRIC manages all these structures individually and delivers them to the Experiments

Experiment specific concepts. CMS Facility example

CRIC as experiment oriented info tool covers VO needs for models descriptions

General details and links to other objects.

Facility is linked to a GocDB/OIM site

(which is VO independent concept defined in CORE part)

* Facility = group of resources (sites, services) belong to same administrative domain

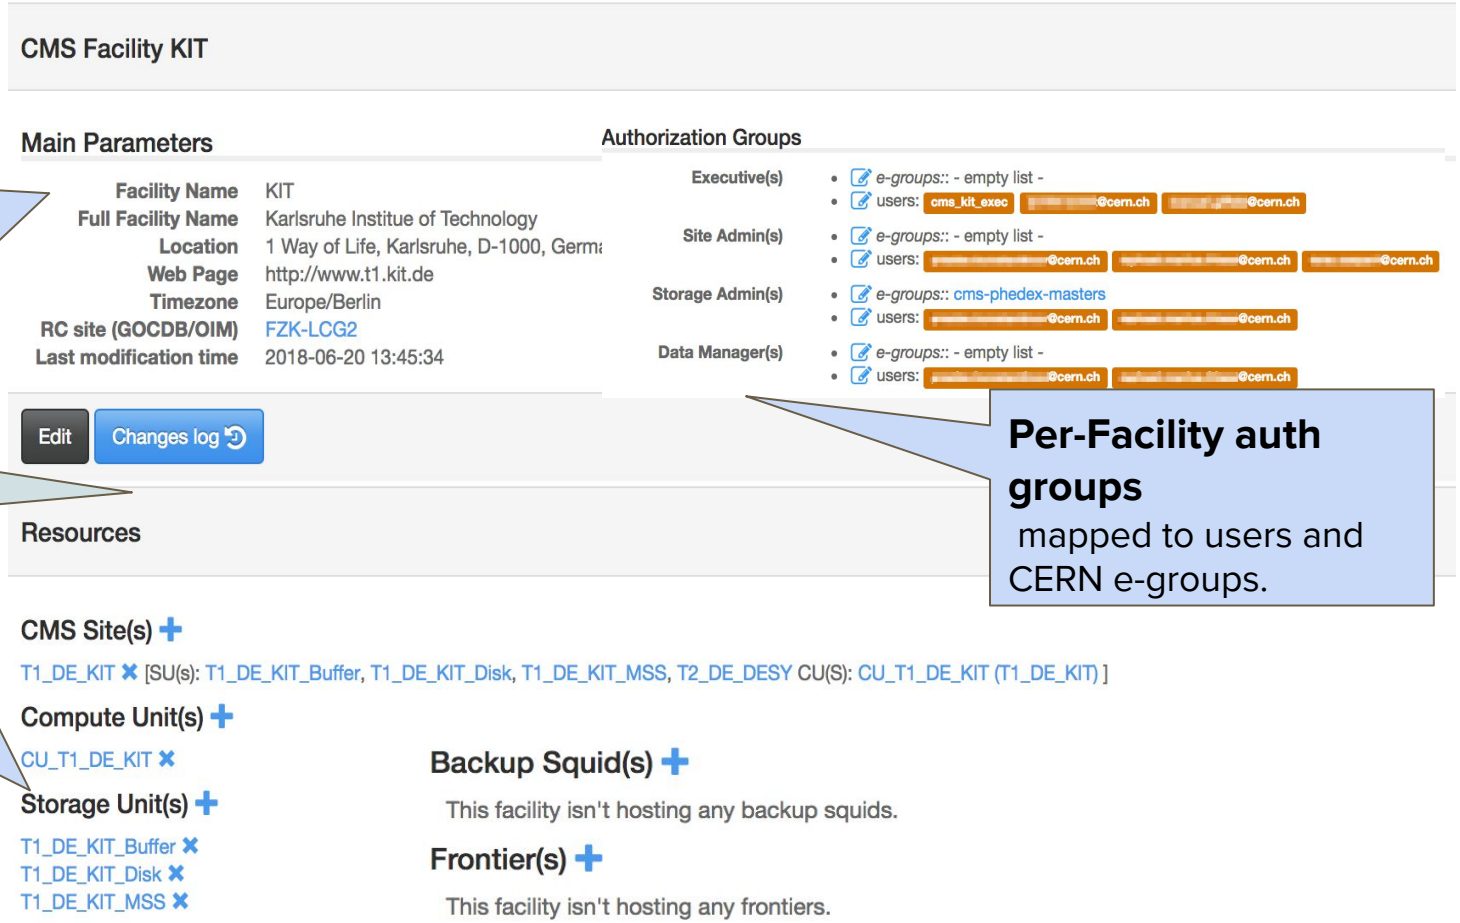

Track the history of changes for the object

Per-Facility auth groups mapped to users and CERN e-groups.

Aggregation of Experiment specific objects: CMS Sites, Compute and Storage units …

CRIC features as the info middleware for VOs

- Helps to easily integrate new Computing technologies which have not yet appeared in WLCG as the services or can not be part of WLCG in general, for example:
  - newer type of SE based on ObjectStore technology
  - Federated Access to storage (FAX redirectors, direct access to remote files from Worker Nodes)
  - Description of opportunistic/volunteer resources

- Helps to minimize side effects for end-user applications of various internal migrations/changes/tests/evolution of Distributed Computing components/infrastructure:
  - Consolidation of protocols description that should be applied only for few sites, unification of resources, migration to HTCondor
  - Keeps data export in several format for backward compatibility

- Masks incompatible updates in external data providers, implement missing functionality/overwrite/fulfill data:
  - e.g. fix wrongly published number of cores, core-power
  - remove direct dependency to ext sources (obsolete data providers)

Status and plans

- **CMS CRIC** instance has been implemented:
  - At the validation and integration step into the CMS production (iterative approach)
  - CMS topology and SiteDB information have been imported into CRIC
  - Implemented backward compatible API for CMS applications (CRAB3, WMAgent, SiteDB export)
- Ongoing implementation of dedicated **WLCG CRIC** instance to represent Computing topology for sites and services used by all 4 LHC VOs
  - Will be used for WLCG central operations and administration
  - As the main info provider for the cross-experiments WLCG tools:
    - central test system (SAM), monitoring (WLCG transfer dashboard), accounting (Storage Space Accounting) ...

Next steps:

- Integration of accounting and resource usage desc into the **WLCG CRIC**
- Extend **CMS CRIC** functionality following CMS feedback
- Implementation of **COMPASS CRIC** (ongoing)
- **AGIS** migration into **ATLAS CRIC**

Conclusions

- All LHC experiments are sharing common computing infrastructure. **CRIC** offers a common framework describing this infrastructure with also an advanced functionality to describe all necessary experiment-specific configuration.

  The way the system is designed each experiment can independently describe it's world and still coexist with the others under the same roof.

- First CRIC version mainly focused on the **CMS-required** functionality is under the validation and going into the production for CMS collaboration.

- Check CRIC:
  - http://cms-cric.cern.ch (CMS-CRIC)
  - http://wlcg-cric.cern.ch (WLCG-CRIC)