Sitemon migration to CRIC

Panos Paparrigopoulos on behalf of the CRIC team
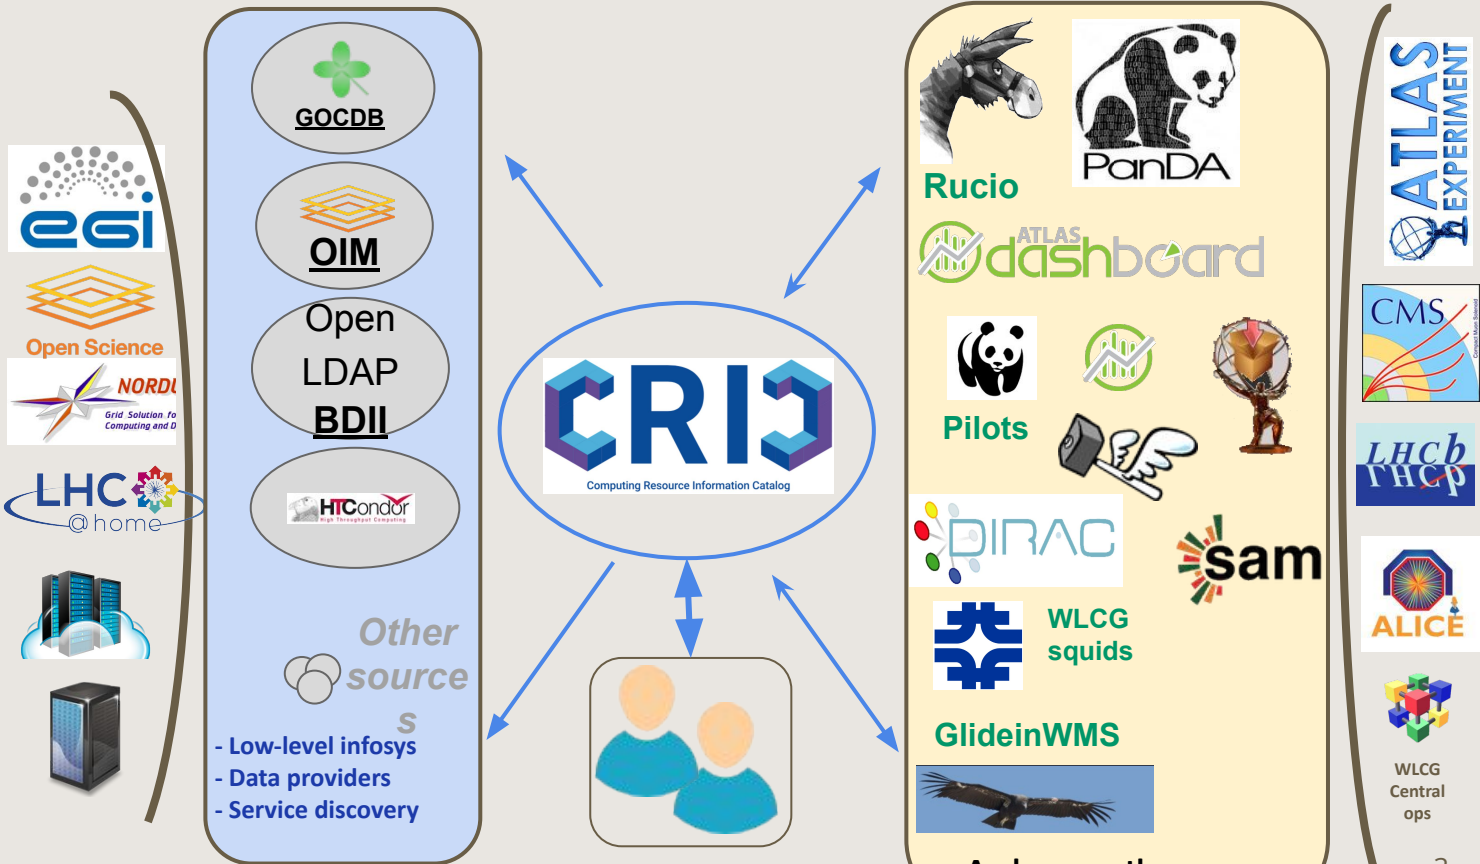
Migration Challenges

- Modeling the concept of service usage

- Incorporating the logic which decides which services are tested/appear in AR and moving this calculation away from the MONIT side

- Consuming various experiment sources and unifying all the information

- Adapting to changes in the experiment sources

- Develop a mechanism to compare CRIC to external sources and get alerts for any inconsistencies
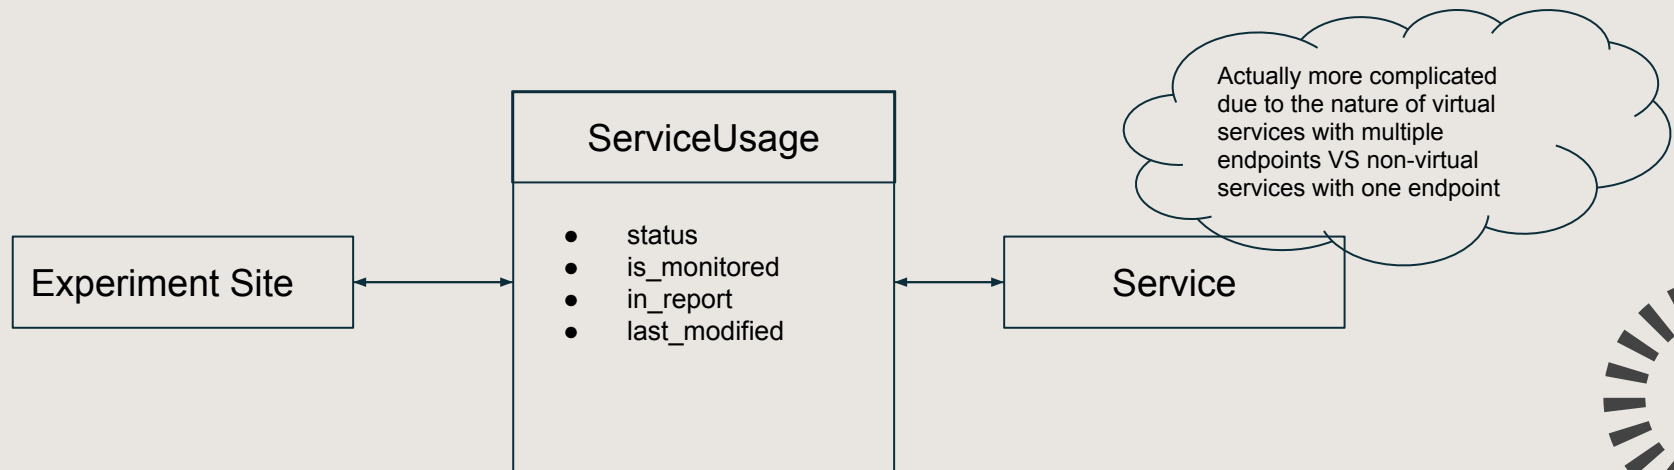
"Provided by" VS "Used by"

GOCDB

OIM

Open LDAP BDII

HTCondor

Other sources
- **Low-level infosys**
- **Data providers**
- **Service discovery**

EGI

Open Science

NORDU Grid Solution for Computing and D

LHC @home

CRIC
Computing Resource Information Catalog

Rucio

PanDA

ATLAS dashboard

Pilots

DIRAC

sam

WLCG squids

GlideinWMS

And many others…

ATLAS EXPERIMENT

CMS

LHCb

ALICE

WLCG Central ops

Used by sources

- Usage information can be very well described in each specific CRIC (ATLAS-CRIC, CMS-CRIC, etc)
- In WLCG-CRIC we needed a way to track the usage of Services by different VOs, so we created the ServiceUsage model:
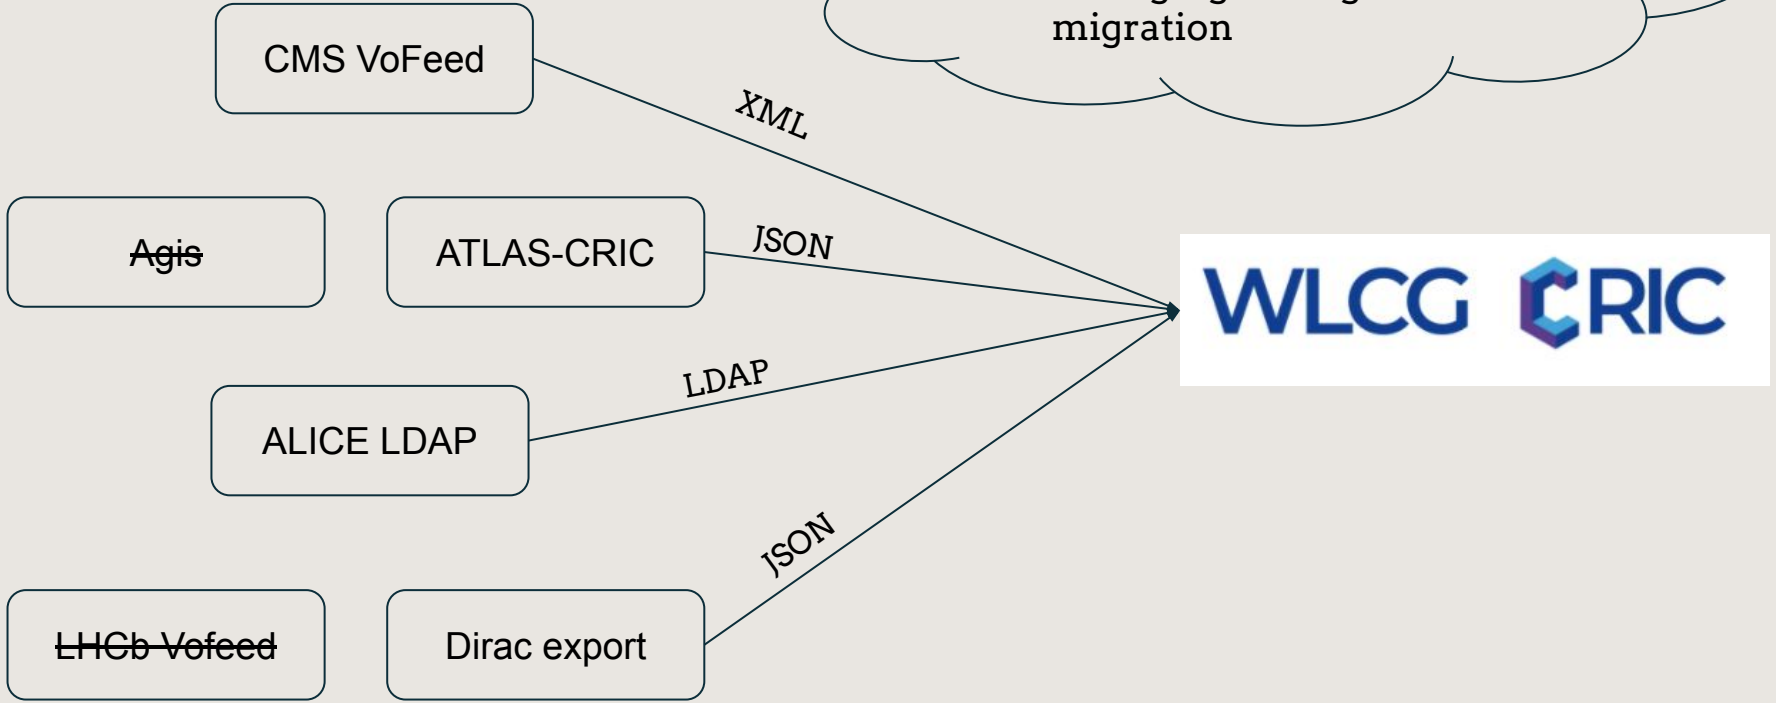
ServiceUsage

- status
- is_monitored
- in_report
- last_modified

Experiment Site

Service

Actually more complicated due to the nature of virtual services with multiple endpoints VS non-virtual services with one endpoint

Experiment Sources

- Many sources
- Different ways to advertise data
- Changing during the migration

CMS VoFeed

XML

Agis

ATLAS-CRIC

JSON

ALICE LDAP

LDAP

LHCb Vofeed

Dirac export

JSON

WLCG CRIC

Comparing info in various sources with CRIC

- Inconsistencies will always be there, a way to spot them fast enough and take action is needed.
- We developed an independent service that can compare the data on the various VO sources with the data in CRIC.
- This comparator is deployed in OpenShift as a REST api that can be queried to compare various sources.
- In case new inconsistencies appear they are reported in our mailing list and we can take action in a timely manner. Most common issues would be:
  - A new site that is not declared in official sources appears.
  - New/unknown flavours appear and endpoints are not imported in CRIC
  - A service is modified by a site-admin directly in CRIC, leading to inconsistencies with the VO source.