

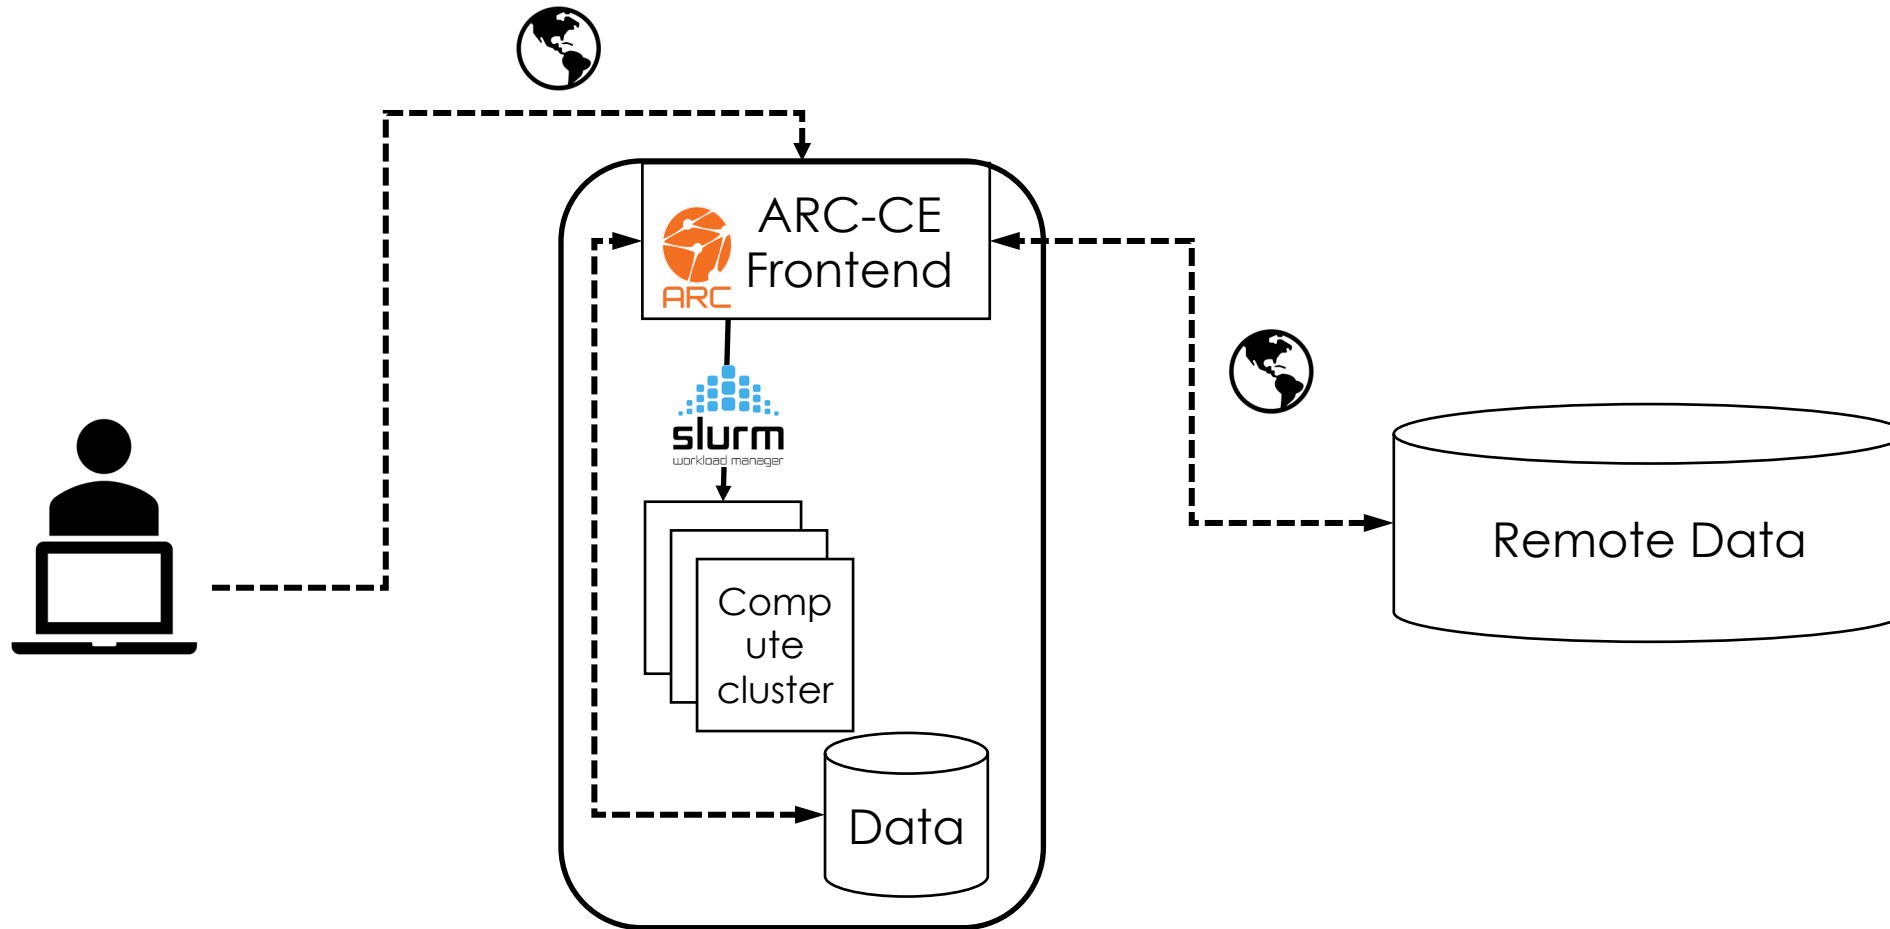
ARC 7 overview



Topics

- Reminder - what is ARC
- Overview of ARC7 components and services
 - Authentication and Authorization
 - Datastaging and cache
 - ARC modes

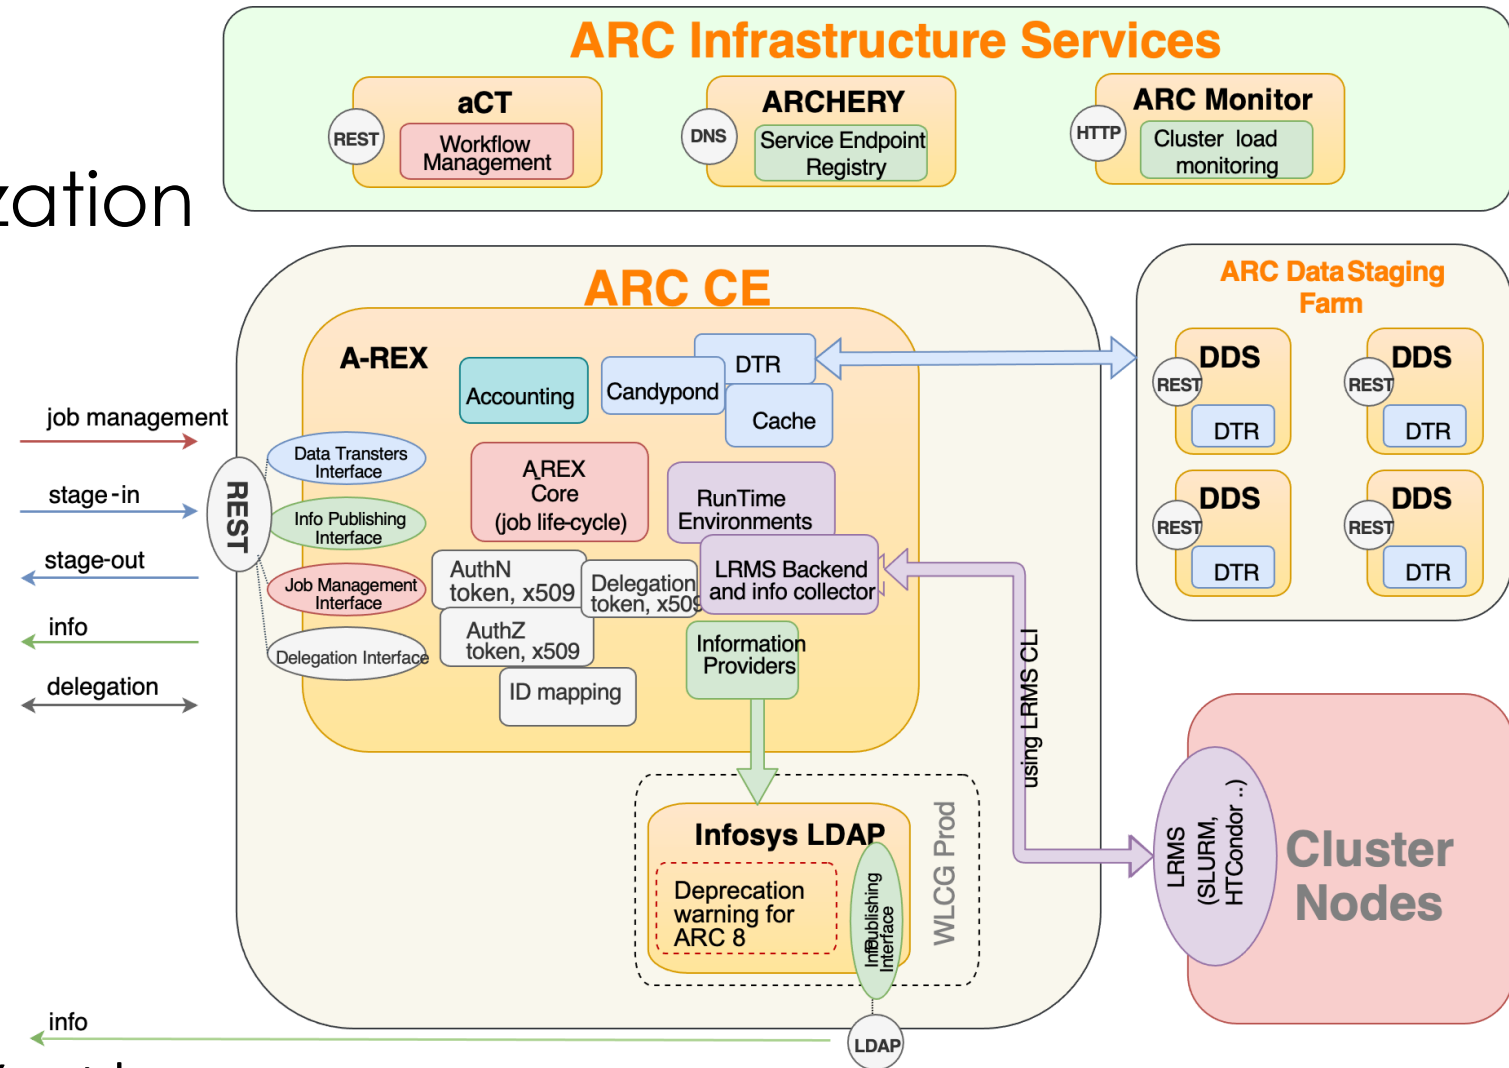
What is ARC



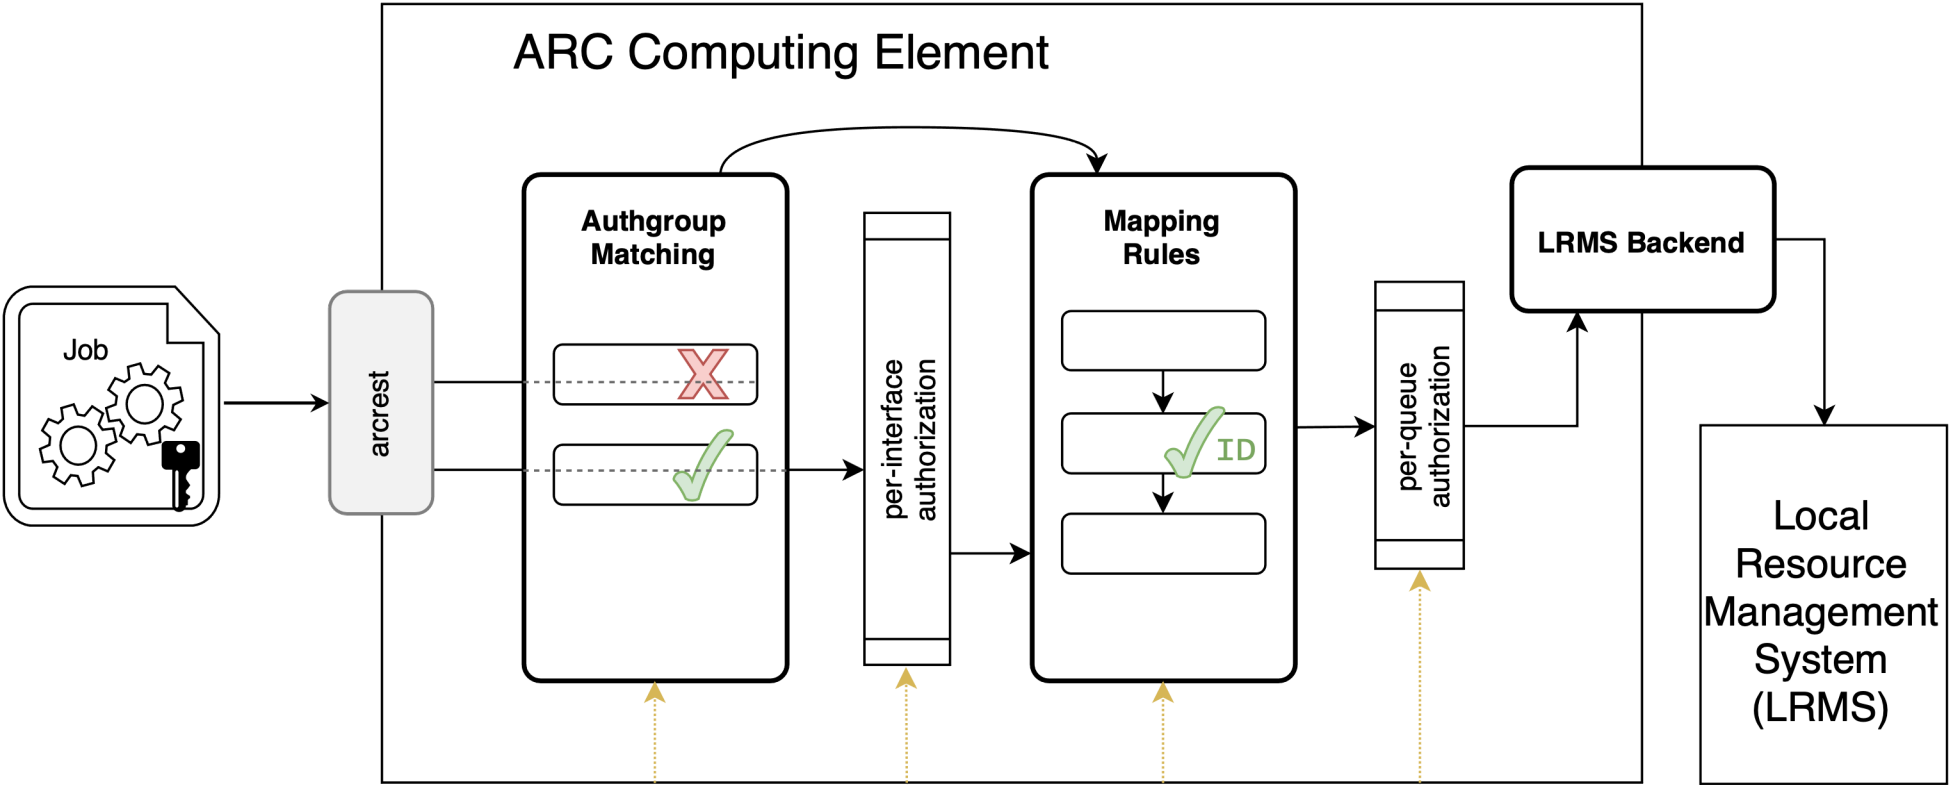
ARC specs

- Authentication, authorization
 - tokens
 - x509
- Job submission
 - ARC REST interface
- Information exchange
 - REST interface
 - Ldap interface (will be deprecated in ARC 8)
- File access
 - Protocols:
 - file, HTTP(s/g), GridFTP, SRM, Xrootd, LDAP, Rucio, S3, RFIO/DCAP/LFC (through GFAL2 plugins)

ARC 7



Powerful and transparent authentication and authorization functionality



[authgroup]

Who

allowaccess
denyaccess

What

[mapping]

Mapping
external to
local user

allowaccess
denyaccess

What

Powerful and transparent authentication and authorization functionality

...

```
[authgroup:testers]
file = /etc/grid-security/testCA.allowed-subjects

[authgroup:wlcg_iam]
authtokens = * https://wlcg.cloud.cnaf.infn.it/ * compute.* *

[authgroup:atlas]
subject=/DC=ch/DC=cern/OU=Organic Units/OU=Users/CN=ddmadmin/CN=531497/CN=Robot: ATLAS Data Management
voms=atlas * * *

[authgroup:wlcg]
authgroup = atlas
authgroup = wlcg_iam
```

Who

```
[mapping]
map_to_user = testers gridadm:gridadm
map_to_user = wlcg grid:grid
```

Mapping external to local user

...

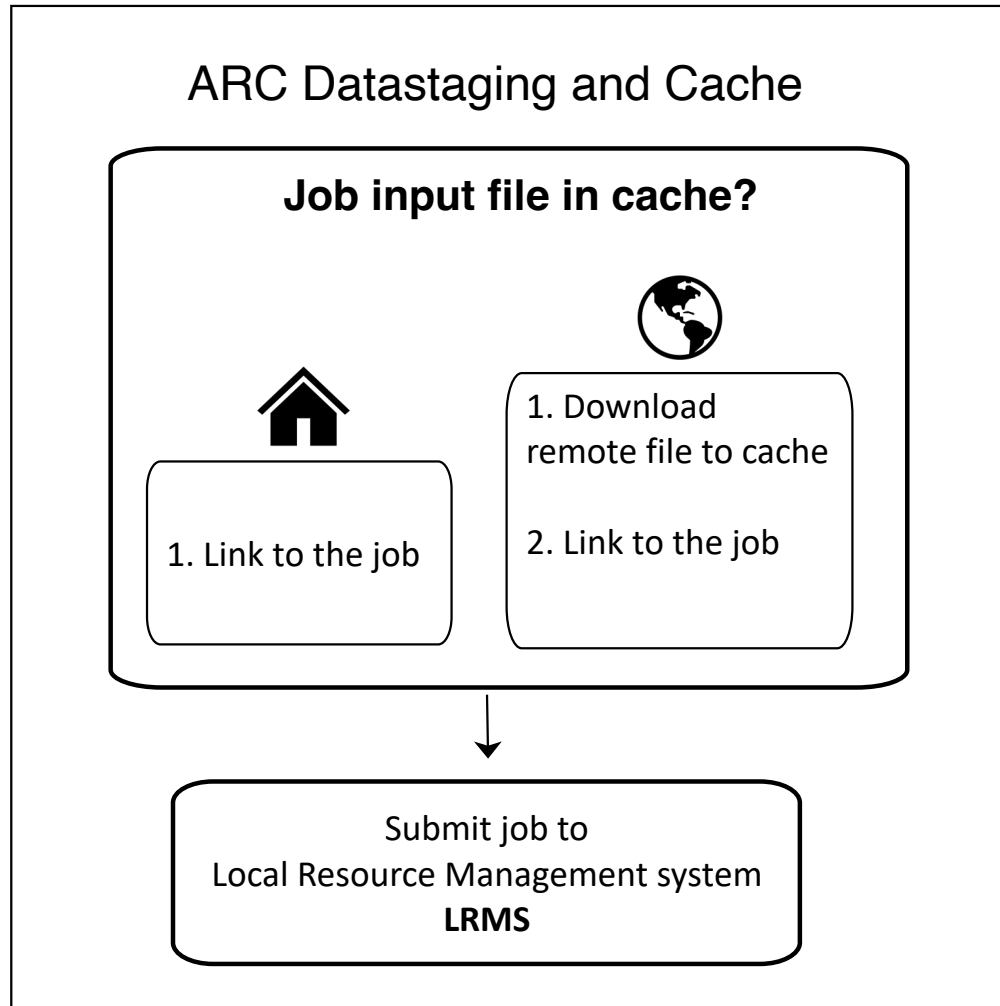
```
[arex/ws/jobs]
allowaccess = testers
allowaccess = wlcg
```

What

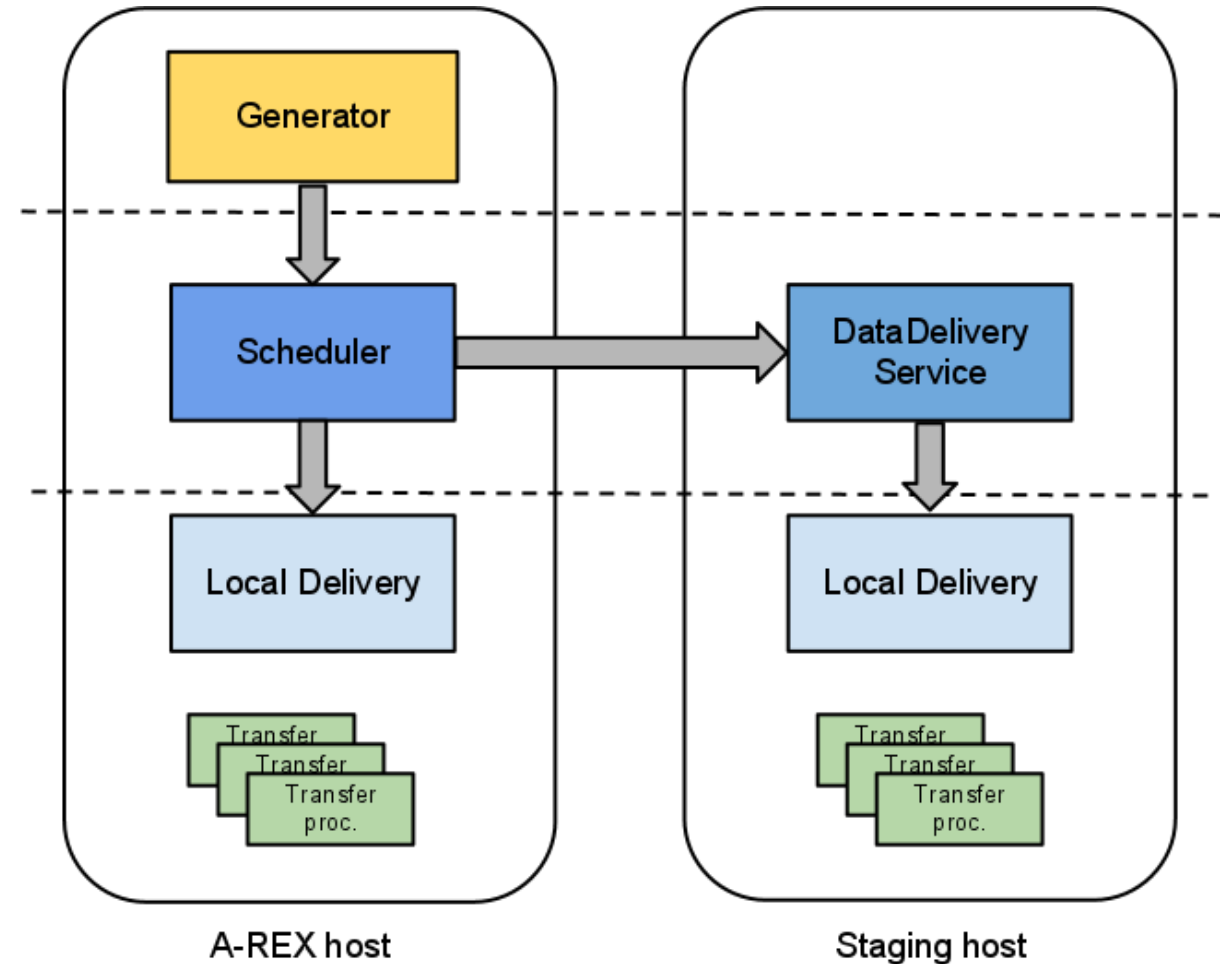
...

```
[queue:main]
allowaccess = wlcg
```

ARC Datastaging and cache

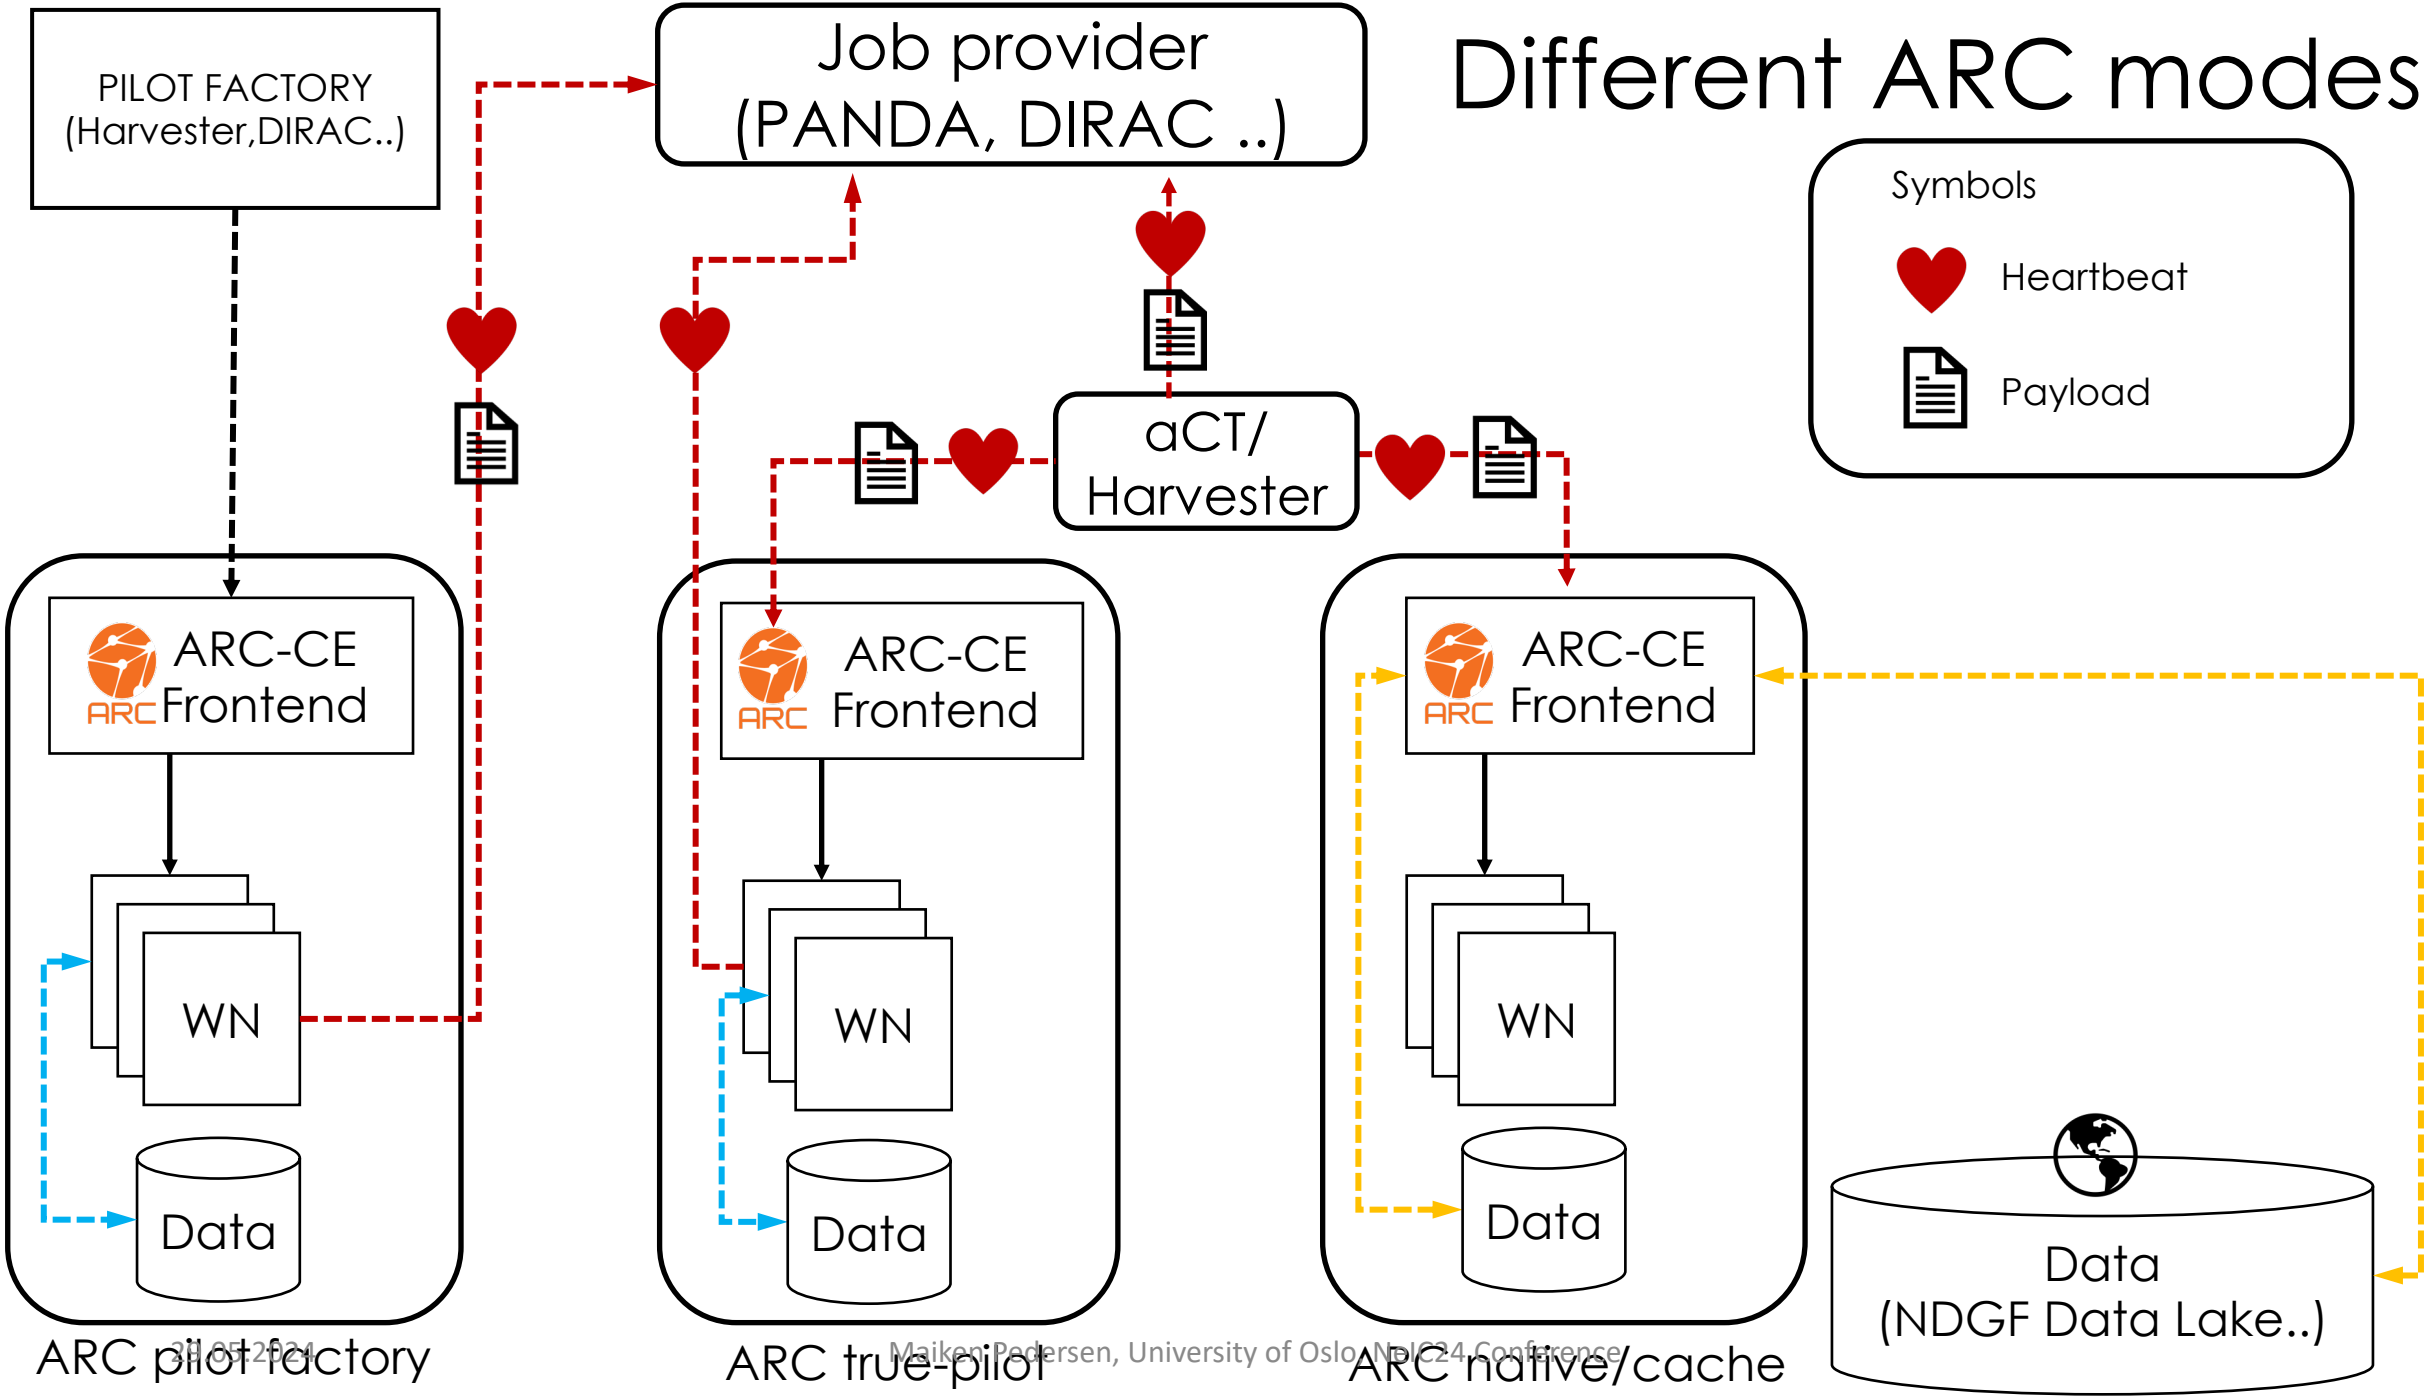


Optional remote multihost datastaging



Configure with shared or local cache storage

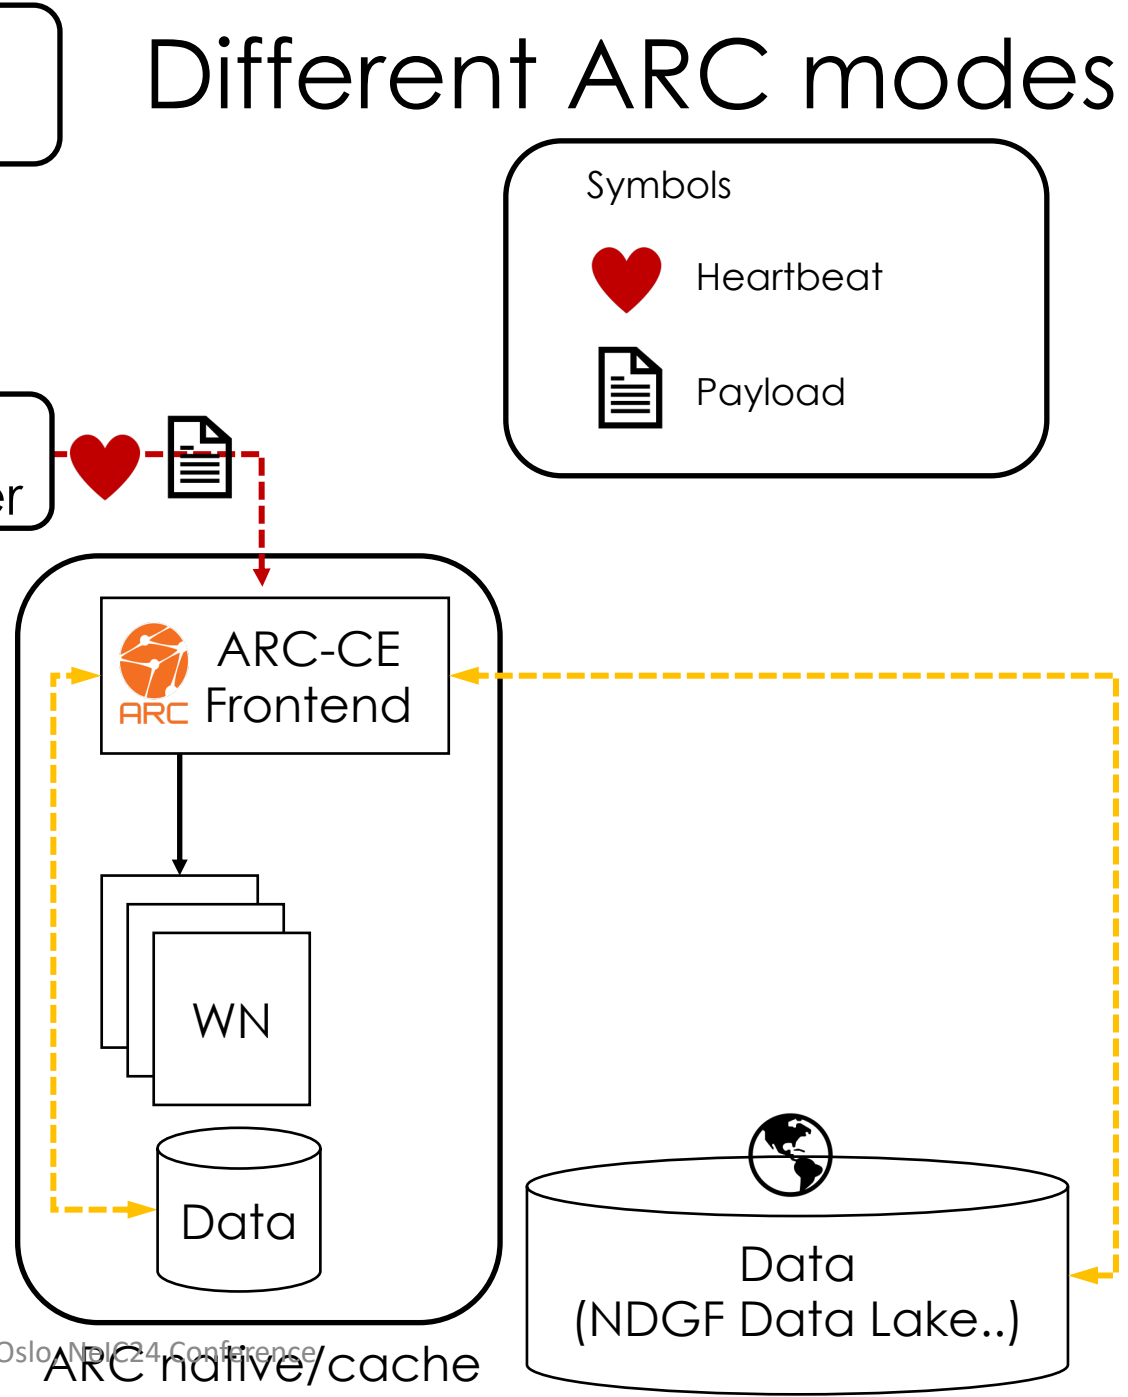
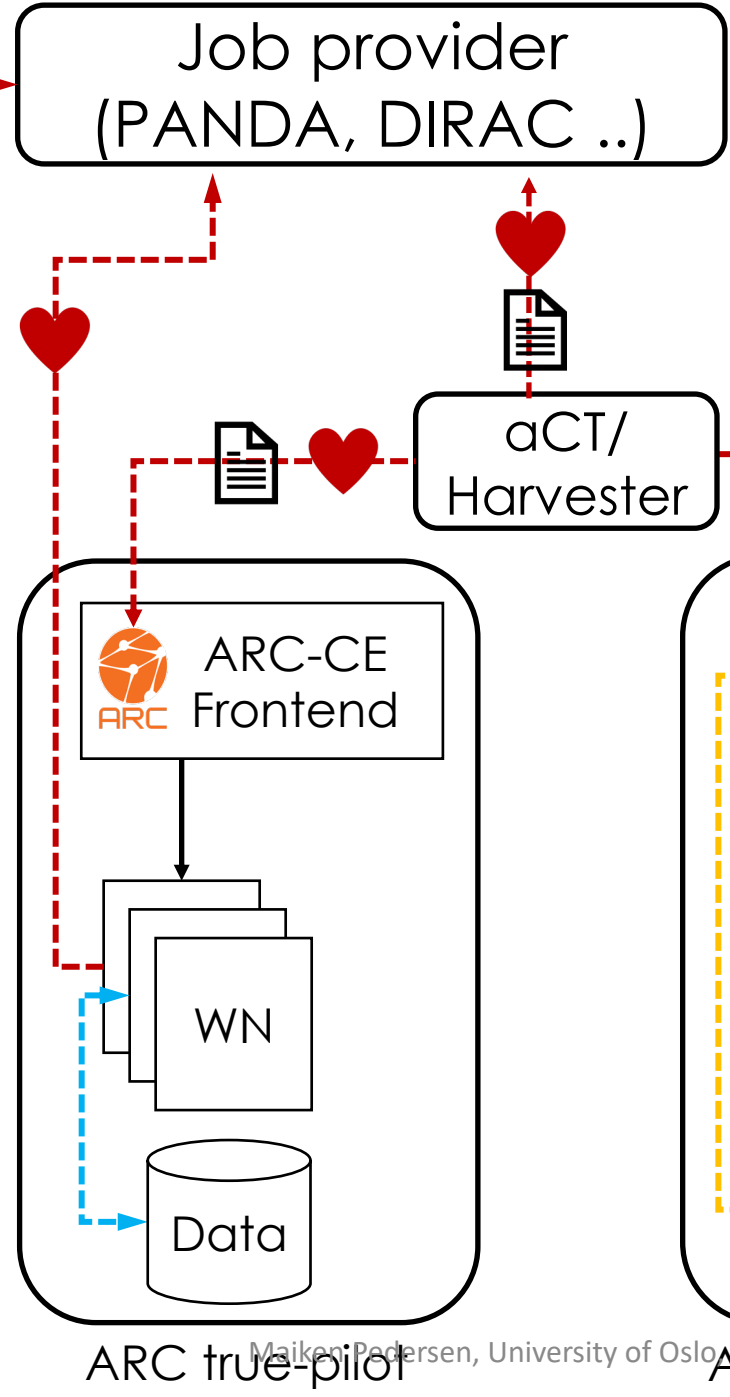
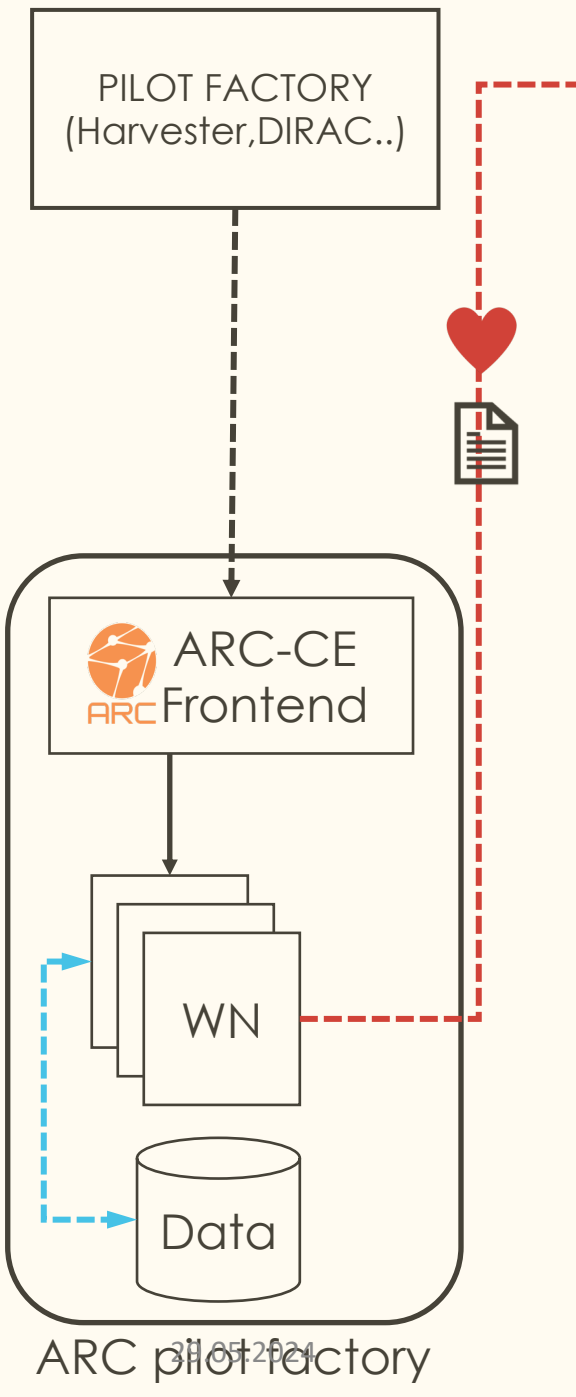
Different ARC modes



Different ARC modes

Symbols

- ♥ Heartbeat
- 📄 Payload



Different ARC modes

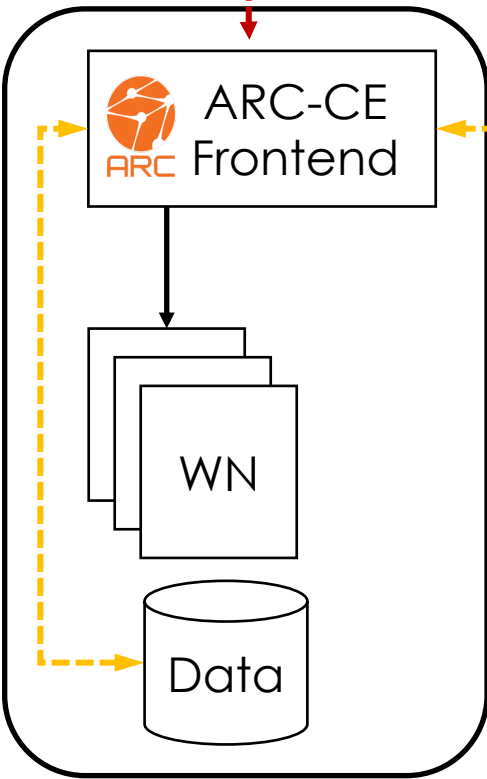
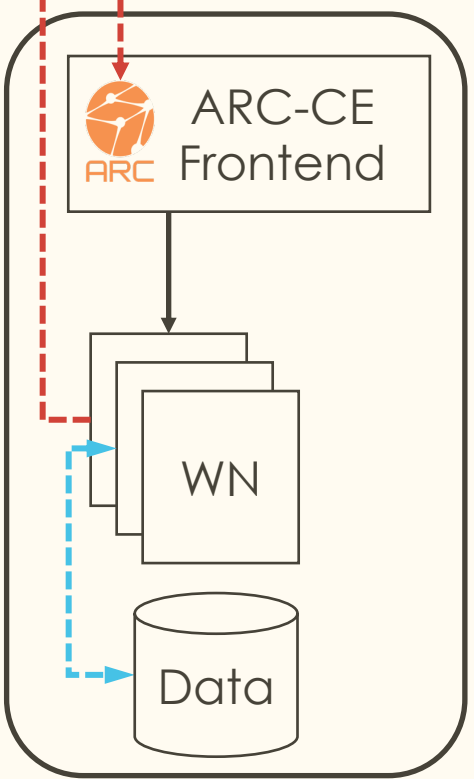
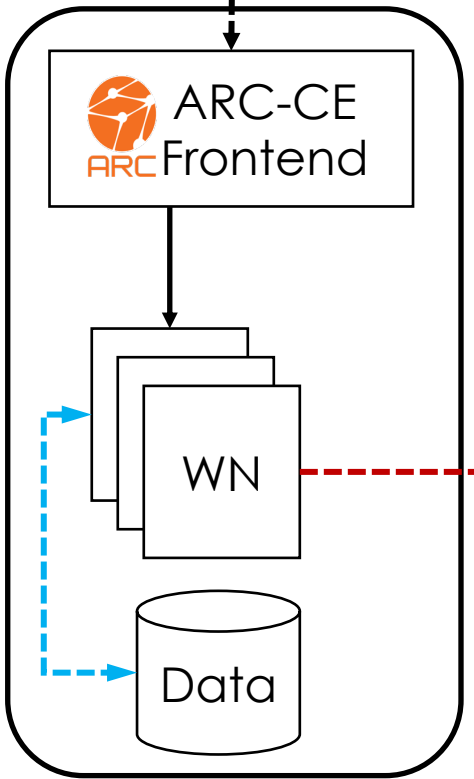
Symbols

- ♥ Heartbeat
- 📄 Payload

PILOT FACTORY
(Harvester, DIRAC..)

Job provider
(PANDA, DIRAC ..)

aCT/
Harvester



ARC pilot factory

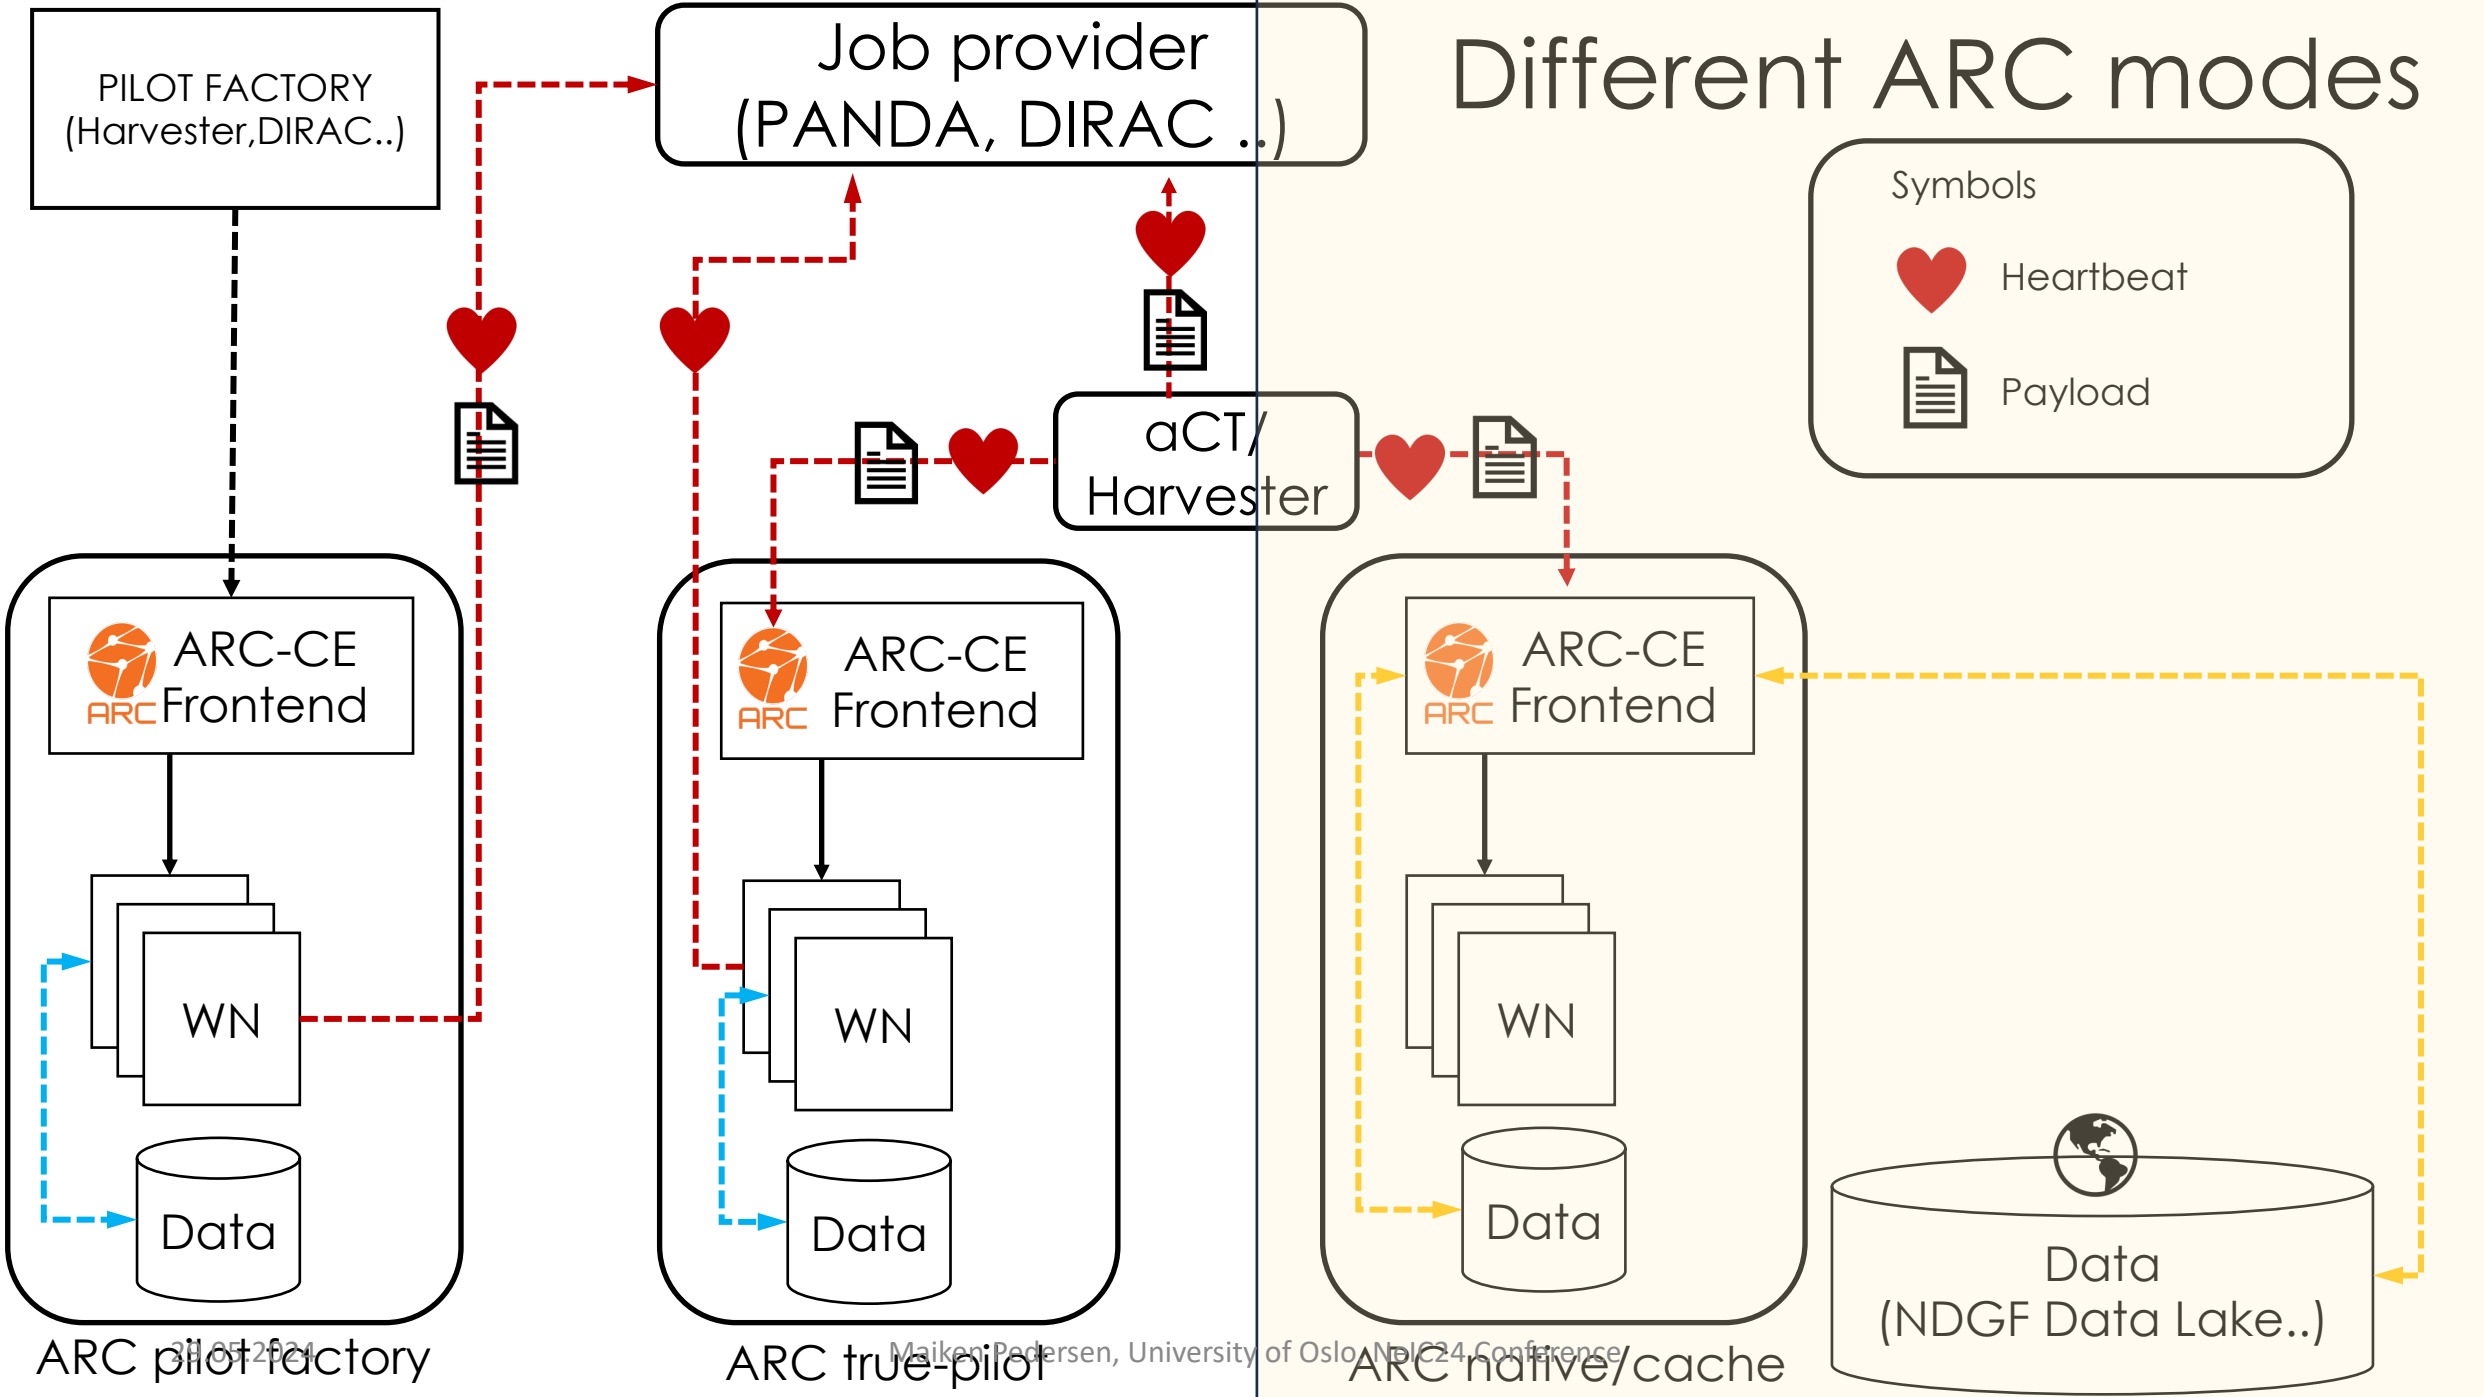
ARC true-pilot

ARC native/cache

Different ARC modes

Symbols

- ♥ Heartbeat
- 📄 Payload



ARCHERY – for brokering a set of sites



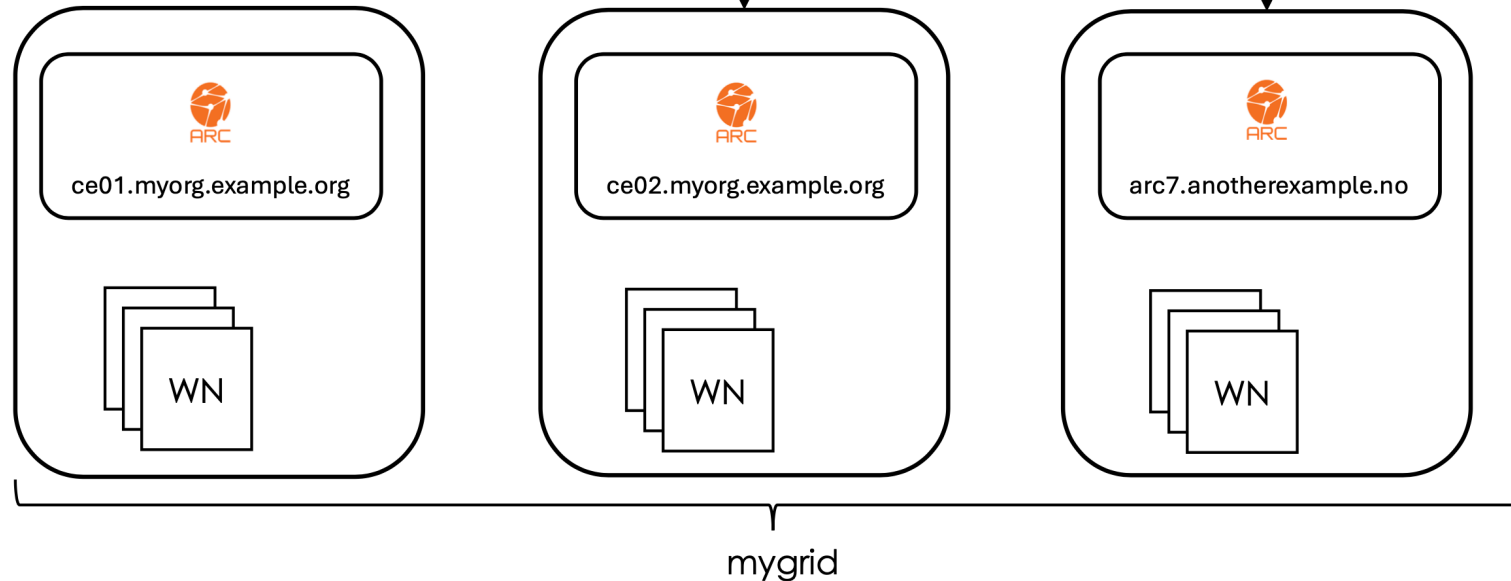
Remote ARC client

```
arcsub --registry mygrid.example.org -T arcrest job.xml
```

INFO

Arcsub includes selecting (matchmaking) resources based on the job description.

ARCHERY



Get in touch and links


- Email: nordugrid-discuss@nordugrid.org ([Subscribe](#))
- Skype [channel dedicated to ARC support](#)
- ARC 7 documentation – WIP! Things still not fully updated
<http://www.nordugrid.org/arc/arc7>
- This tutorial:
<http://www.nordugrid.org/arc/arc7/admins/tutorial/tutorial.html>
- Ansible playbook for ARC: <https://github.com/nordugrid/arc-ansible>

Vega EuroHPC ARC use-case

Vega/EuroHPC - IZUM, Maribor, Slovenia



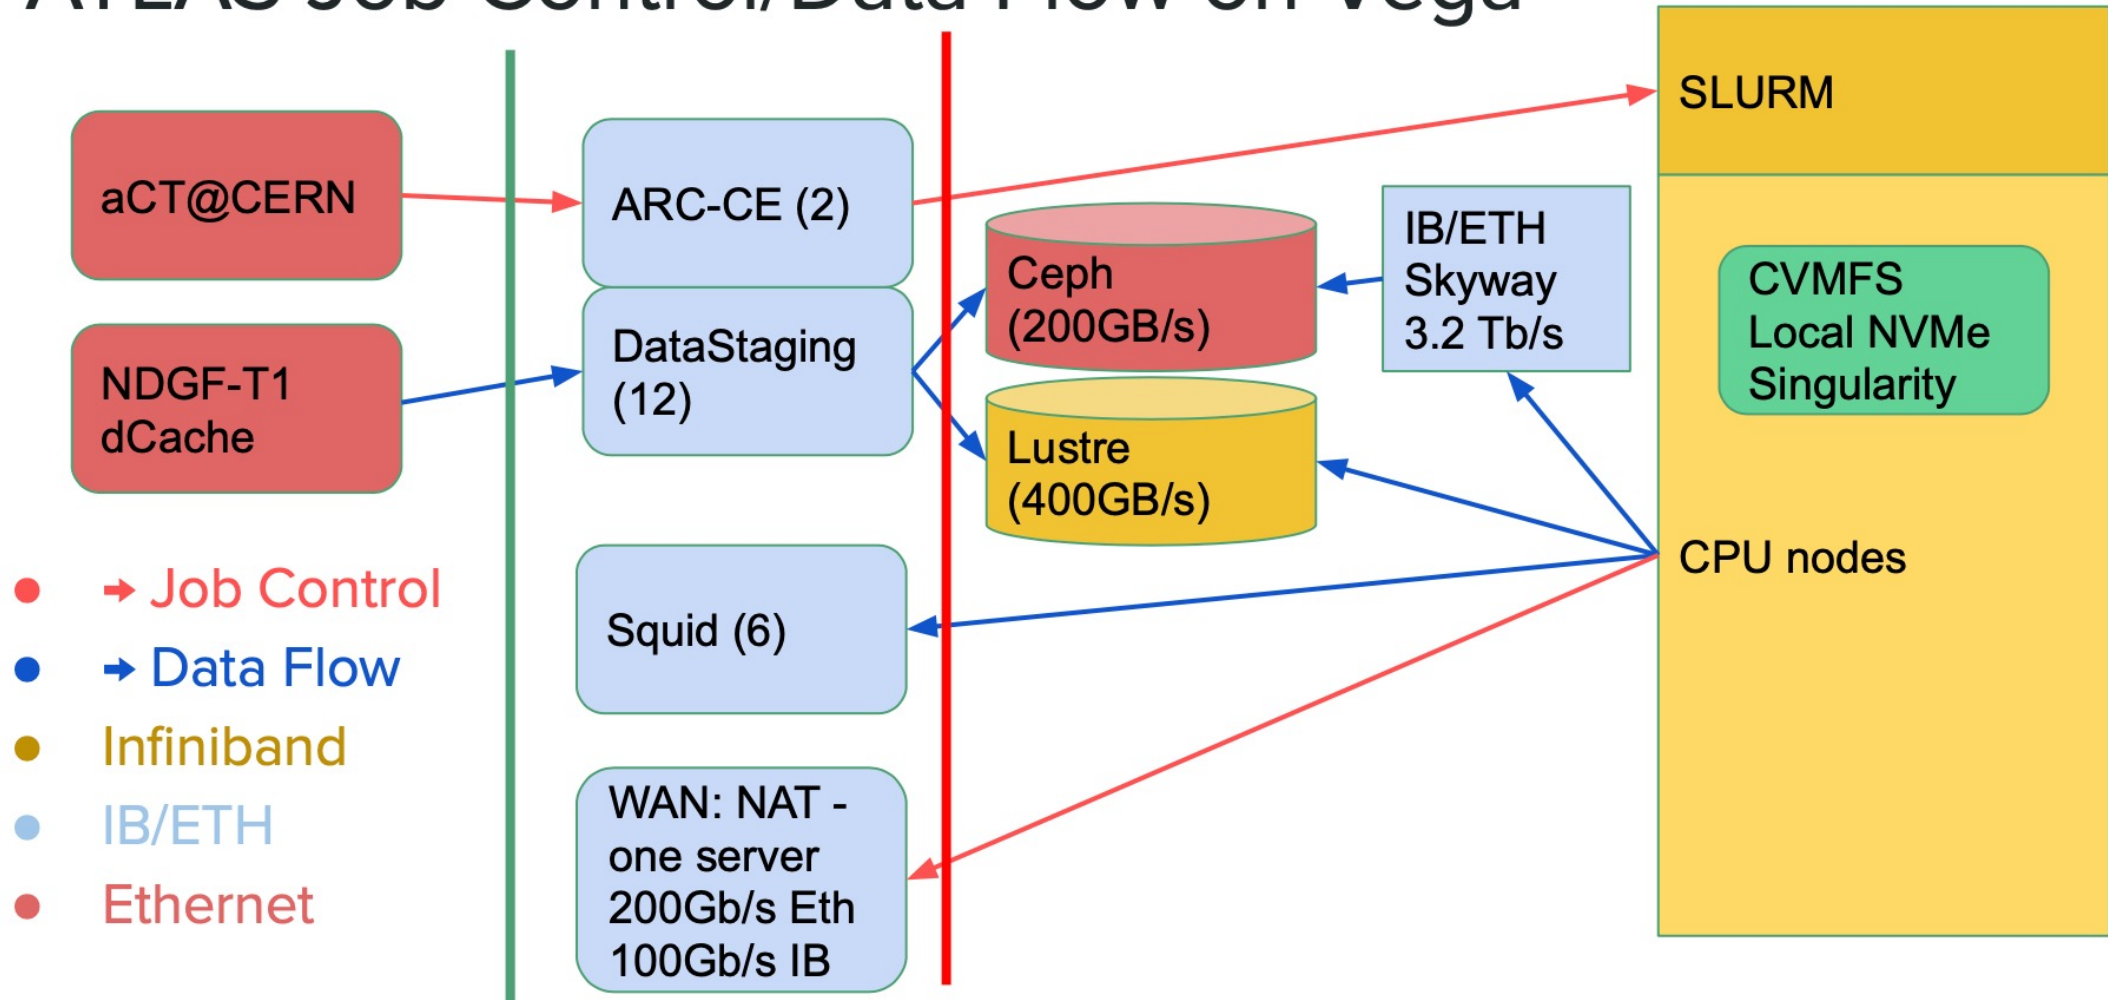
Number of Nodes	1020
Number of login nodes	4 CPU login nodes and 4 GPU login nodes
Compute partitions	CPU partition: 960 (768 standard CPU nodes and 192 high memory CPU nodes * GPU partition: 60 nodes
Storage capacity	1 PB high-performance NVMe storage, 23 PB raw large-capacity storage (18 PB usable)
Sustained Performance	6,9 petaflops
Peak Performance	10.1 petaflops







Why ARC-CE on Vega?

- ATLAS, Belle-2, Pierre Auger, Vera Rubin:
 - Slovenia pledges compute resources to these experiments
 - Central distribution of production jobs - remote job submission
 - Transparent transfer of input/output files to experiment storage
- Other data-demanding communities on Vega
 - Cryomicroscopy
 - Genomics, DNA sequencing
 - AI image processing (generative networks development)
- Many individual users also use ARC-CE
 - Managing jobs on local PC
 - submission to other SI clusters
- Some communities have private service VMs (eg FAIR data, custom executors)

ATLAS Job Control/Data Flow on Vega





ARC for any user-community

Developed for WLCG and HPC in particular

- But can run in connection with any compute resource that runs a batch system (not yet Kubernetes, to come in later releases) (or fork 😊)
- Suitable for any research community with access to multiple compute resources and where data can be accessed over the internet (or is already in place locally at the compute resource).
- New initiatives for user-friendliness for non-LHC user-communities
<https://indico.naic.no/event/259/contributions/1060/subcontributions/4>
 - New in-development integration with the Galaxy web-portal job-submission interface
 - Easy-deployment of ARC+Slurm on EGI Infrastructure Manager

