



ESCAPE

European Science Cluster of Astronomy &
Particle physics ESFRI research Infrastructures

ESCAPE DIOS: WP2 Workshop

Dec 2020

Rosie Bolton

9 December 2020




ESFRIs



Our Mission

Prototype an infrastructure adapted to **Exabyte-scale** needs of large science projects

Common data infrastructure for
Astro-particle, Radio-astronomy, Gravitational
Waves, Cosmology and Particle Physics
Address **FAIR** data management principles



Partners



rijksuniversiteit
groningen



Welcome to the ESCAPE WP2 Workshop!

Today (now until 11:40):

Dress Rehearsal review - Experiments
view

11:55 - 13:00:

Dress Rehearsal review - Services
Perspective

Tomorrow: 09:00-10:00

Dress Rehearsal review - Feedback from sites

10:10 - 11:10:

FDR wrap-up: lessons learnt, identified issues, FDR
2.0 scopes and D2.2 deliverable preparations

Dress Rehearsal review - Services Perspective

11:25 - 13:00

2021 Program of Work: Timelines, future plans and
evolution of the pilot Data Lake towards the prototype
phase

Remote Workshop settings

All sessions are subject to the [SKA Code of Ethics and Conduct](#).

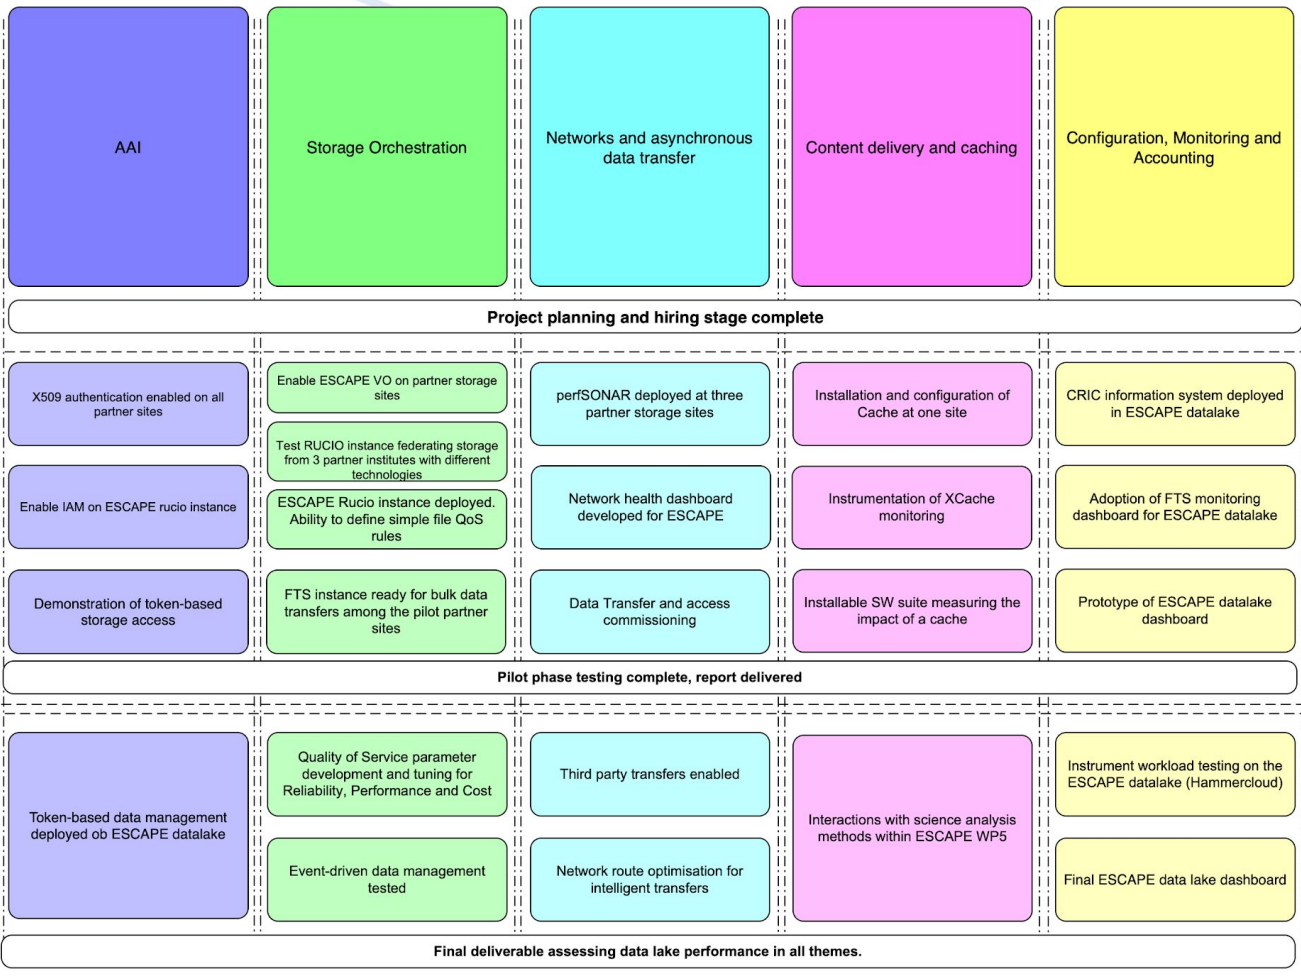
- Enter the **Zoom** room with your **full name**
- **Raise** your virtual **hand** to ask question - please do not jump in
- One conversation at a time
- Keep to the allocated time
- Stay **muted** to minimise background noises
- **When you are giving a talk, please turn on your camera if you can**
- **No applause please- it is painful on zoom.**
- We are recording this meeting (check - am I actually recording this?!)
- Closed-captioning has been enabled
- We will try to capture issues for further discussion




Raise Hand

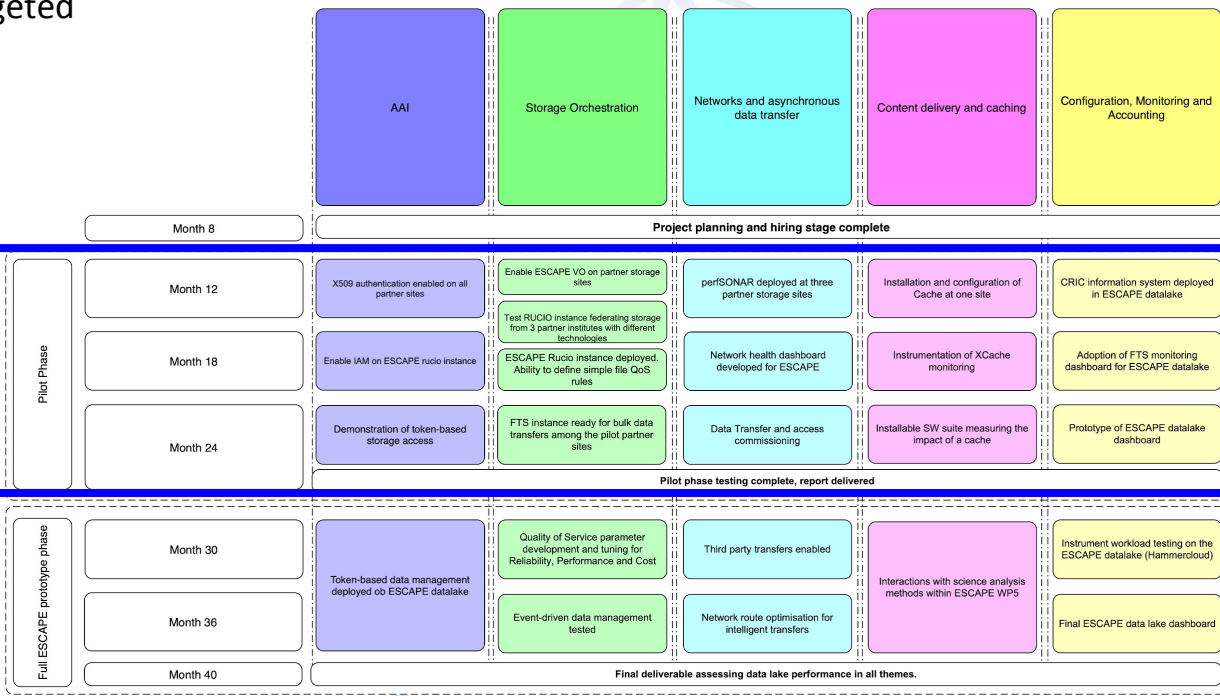
Participants Panel or
Shortcut: **alt/cmd + Y**

Project Plan - recap



Project Plan - recap

In the proposal, our pilot phase targeted having 3 sites for assessment.



**Pilot Phase
(up to now)**

**Full Prototype
(starts 2021)**



"D2.2" - 2nd WP2 Deliverable

"Assessment of the **Pilot** Datalake"

SKAO is responsible for this deliverable - **it is due very soon** -15th Jan

Emphasis on:

- the on-boarding experiences of experiment reps

- useful features of the datalake in context of the use case for each experiment

- functionality of each datalake service

- future directions, improvements and innovations (else no point assessing it)



The Full Dress Rehearsal, FDR1, 17 November 2020

- Our Datalake is not a production system, but we "dream big"
- Weekly Deployment and Operations meetings have helped develop a team of experts and space to share knowledge
- Datalake with 12 (now 13) Rucio Storage Elements working together
 - Identify 17th Nov 2020 for a "Golden Window" of good behaviour!
- Series of "Data Injection" meetings over the past few months for experiment teams to gain confidence using rucio
- By FDR day, 12 additional "users" representing experiments interacted with rucio
 - On 17th Nov 2020, we asked all users to perform tests to exercise aspects of the datalake relevant for their experiments or local setups.



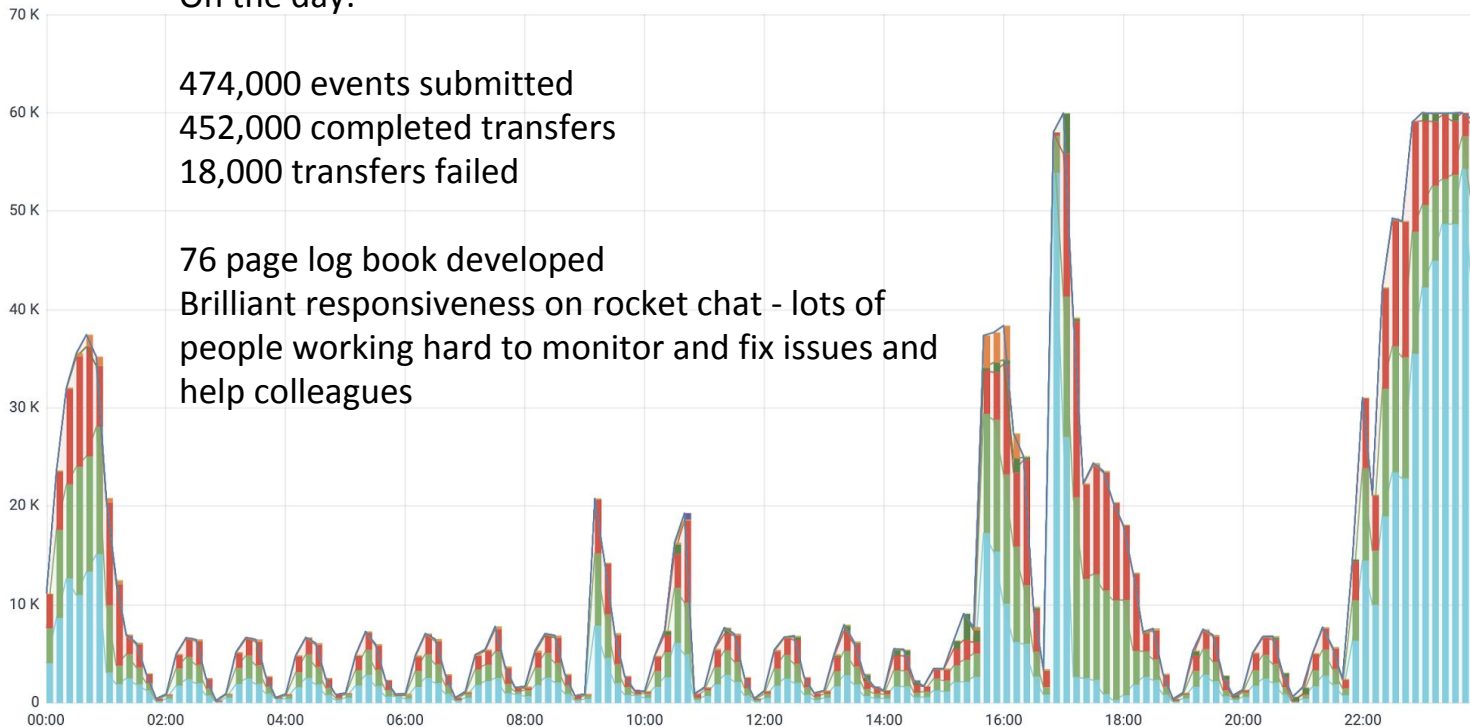
Bin 10m ▾ scopes All ▾ file name * src All ▾ dst All ▾

Events by type over time (only scope filtered)

On the day:

474,000 events submitted
452,000 completed transfers
18,000 transfers failed

76 page log book developed
Brilliant responsiveness on rocket chat - lots of people working hard to monitor and fix issues and help colleagues



	total ▾
transfer-queued	796.0 K
transfer-submitted	474.2 K
transfer-done	452.3 K
deletion-done	25.5 K
transfer-failed	18.5 K
create_dts	10.0 K
detach	688
register_cnt	661
erase	17
transfer-submission_failed	9
create_cnt	4



The Full Dress Rehearsal, FDR1, 17 November 2020

On the day:

Mix of file sizes, dataset structures and rule types

Use of specific placement (at target RSEs) and automated placement based on QoS - which gave good resilience to downtime

Some transfers failed but rucio showed itself to be very resilient, able to recover

Rucio Stats ☆ 🔊

2020-11-17 00:00:00 to 2020-11-17 23:59:59

source: Monit ESCAPE (long term) | rses: All | scopes: All | experiments: All | bin: 1m

DIDs per Experiment (replica=1)

Experiment	Number of DIDs	Number of files	Number of datasets	Number of containers	Average Filesize ↓
LOFAR	25.3 K	25.2 K	5	0	1.666 GB
FAIR	194	192	2	0	1.036 GB
CMS	401	398	3	0	1.026 GB
MAGIC	13.5 K	824	12.6 K	18	573 MB
ATLAS	7.604 K	6.952 K	652	0	235 MB
LSST	350 K	350 K	13	0	18.5 MB
CTA	564 K	563 K	1.458 K	0	9.273 MB
SKA	2.736 Mil	2.703 Mil	33.0 K	25	3.259 MB
VIRGO	15.6 K	15.6 K	10	0	86.4 kB

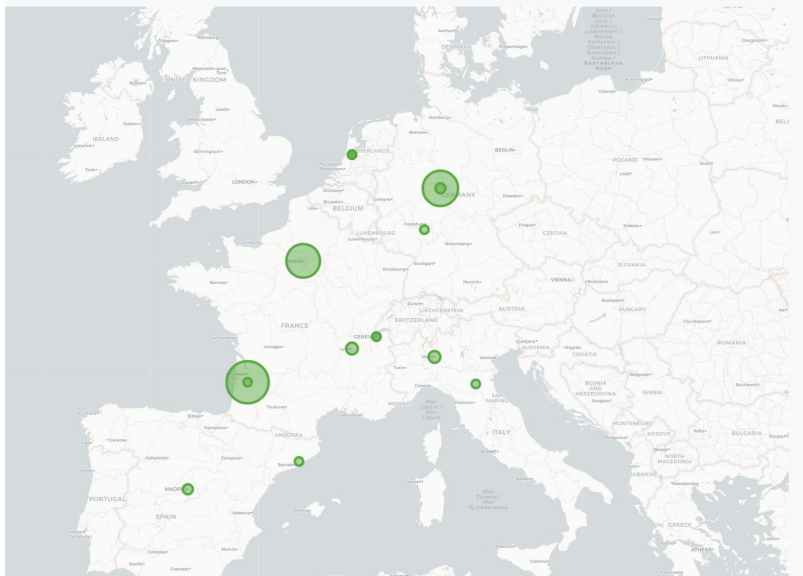


We are running an international datalake!

FDR1 Transfers map by source

FS Transfers (map-experimental-rizart) ☆ 🔊

Number of transfers by source



Production / Rucio Events ☆ 🔊

2020-11-17

Transfer Matrix - Replica Creation

src	dst	DESY-DCACI	SARA-DCACI	PIC-DCACI	EULAKE-1	LAPP-DCACI	IN2P3-CC-DCACHE	CNAF-STORM	ALPAMED-DCACHE	GSIR-ROOT	INFN-NA-DPM	LAPP-WEBDV	INFN-NA-DPM-FED	INFN-ROMA1
DESY-DCACI	DESY-DCACI	NO DATA	100%	51%	100%	104%	100%	100%	93%	35%	98%	100%	100%	NO DATA
SARA-DCACI	DESY-DCACI	100%	NO DATA	98%	100%	100%	100%	98%	88%	25%	98%	98%	95%	NO DATA
PIC-DCACI	DESY-DCACI	100%	100%	NO DATA	99%	100%	100%	100%	100%	23%	100%	100%	95%	NO DATA
EULAKE-1	DESY-DCACI	100%	75%	47%	NO DATA	100%	100%	100%	100%	42%	100%	100%	100%	NO DATA
LAPP-DCACHE	DESY-DCACI	100%	100%	99%	100%	NO DATA	98%	100%	98%	16%	98%	94%	95%	NO DATA
IN2P3-CC-DCACHE	DESY-DCACI	100%	100%	89%	100%	100%	NO DATA	100%	91%	35%	98%	100%	100%	NO DATA
CNAF-STORM	DESY-DCACI	100%	100%	98%	100%	100%	97%	NO DATA	100%	18%	100%	100%	100%	NO DATA
ALPAMED-DPM	DESY-DCACI	28%	94%	100%	100%	100%	100%	100%	NO DATA	49%	92%	100%	100%	NO DATA
GSIR-ROOT	DESY-DCACI	100%	99%	94%	100%	99%	100%	100%	89%	NO DATA	100%	97%	95%	NO DATA
INFN-NA-DPM	DESY-DCACI	100%	100%	100%	100%	99%	100%	99%	90%	45%	NO DATA	98%	NO DATA	NO DATA
LAPP-WEBDV	DESY-DCACI	100%	100%	100%	100%	100%	100%	98%	100%	100%	100%	NO DATA	100%	NO DATA
INFN-NA-DPM-FED	DESY-DCACI	100%	100%	96%	100%	93%	100%	96%	81%	40%	NO DATA	96%	NO DATA	NO DATA
INFN-ROMA1	DESY-DCACI	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA

Failed Transfers

Failed Transfers

timestamp ↑ scope name size (bytes) src



Next

FDR1 was a great success, now we should...

- Ensure that we collate appropriate output for D2.2
- Take this opportunity to **capture lessons learned**
 - How can we improve our monitoring, testing, reactions (operations)
 - How can we improve the services themselves?
 - What more do experiment teams need in order to assess a datalake that more closely matches their use cases?
- Pause and ask where priorities are going forward
- See who needs the second "golden window" identified for next Tuesday, and what support they need

