



Notes on noted

Coralie Busse-Grawitz

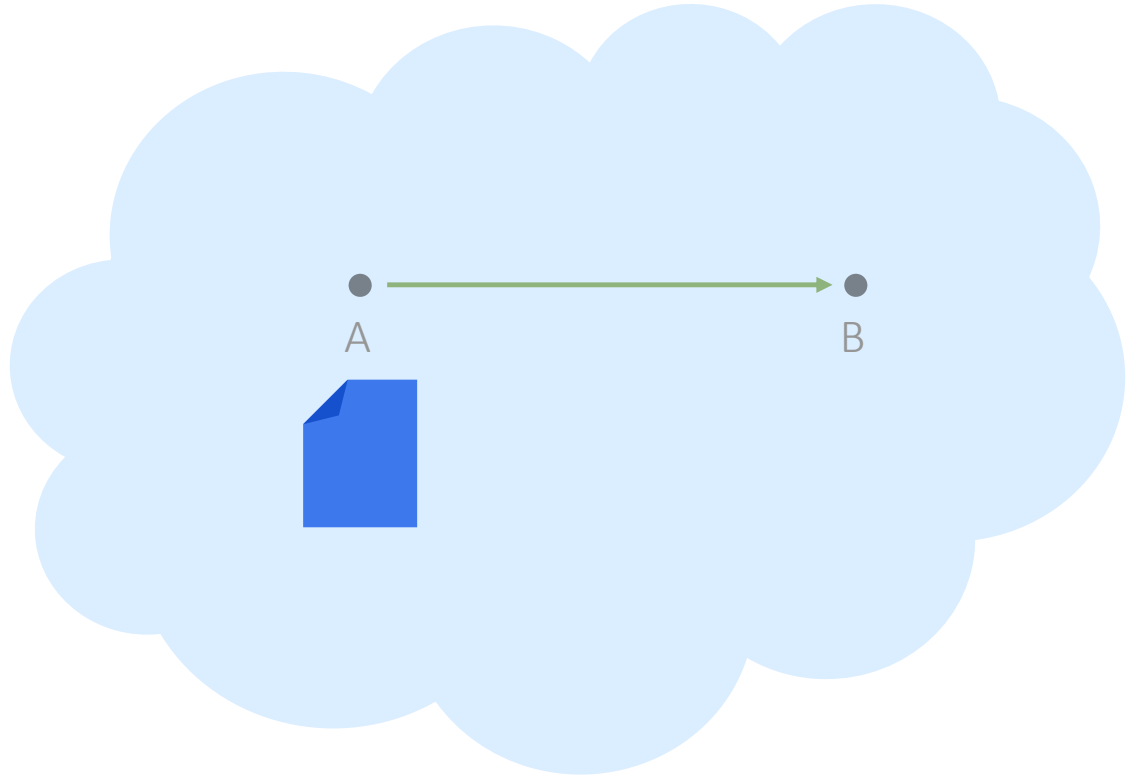
Supervisor: Edoardo Martelli



Fabiola

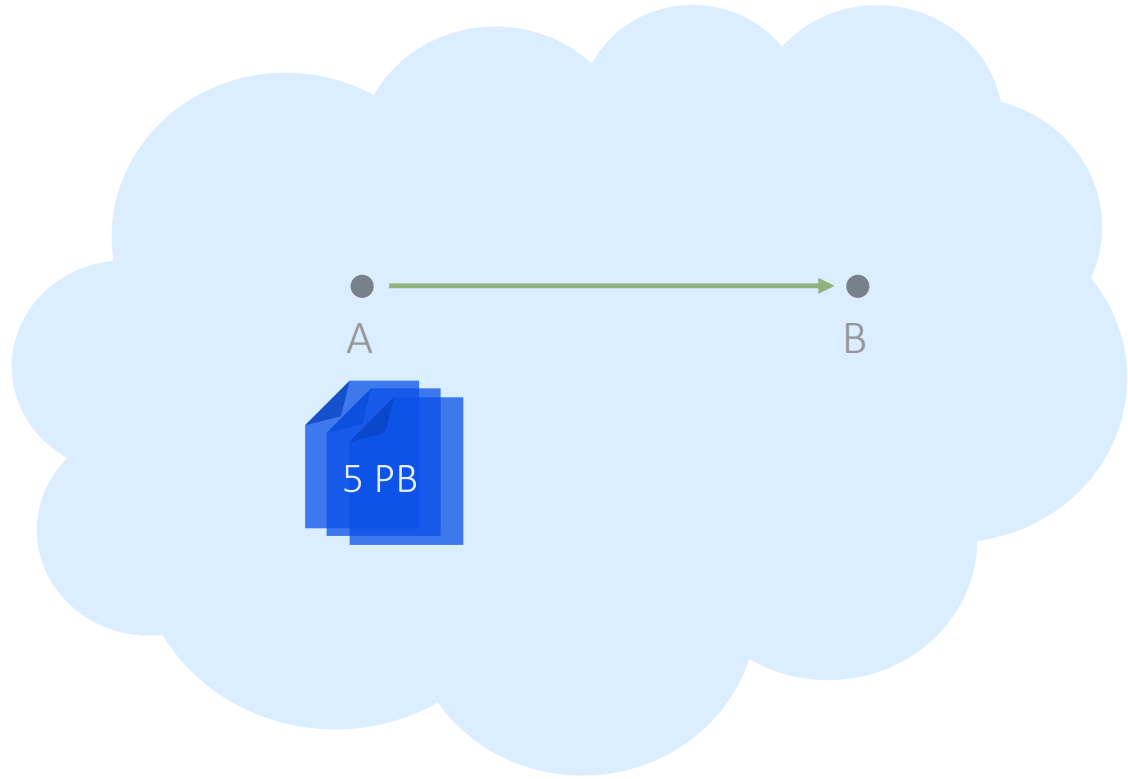


Fabiola






Fabiola
The Physicist






Progress






Progress

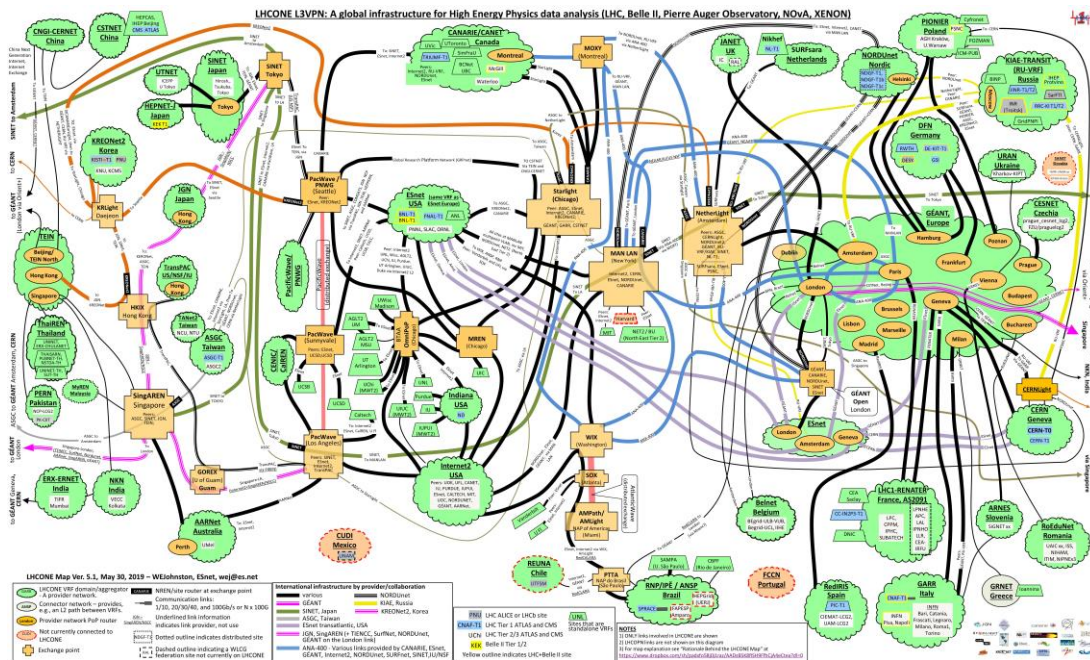




LHCONE..?





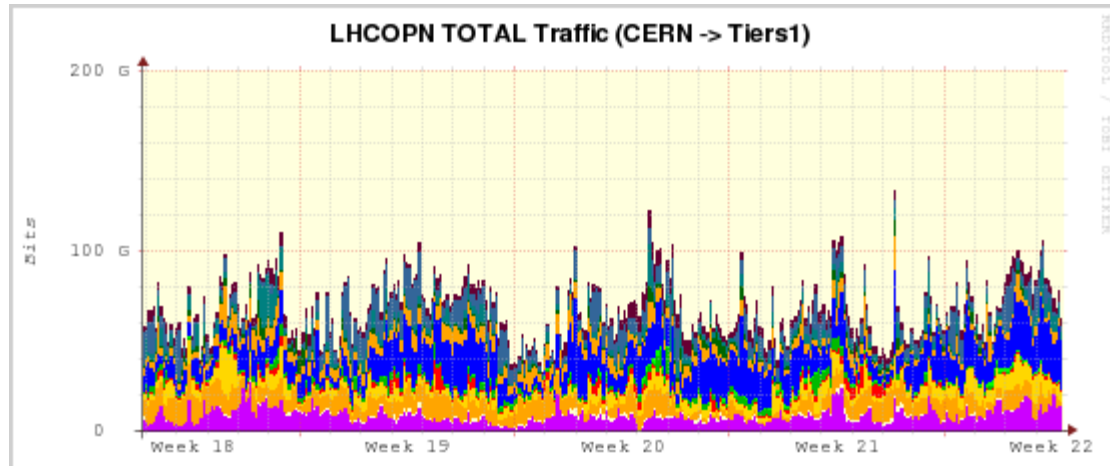


We could load-balance...

NOTED.




The luminosity upgrade causes
higher bandwidth usage (peaks)




https://netstat.cern.ch/monitoring/network-statistics/ext/cache/LHCOPN_00-Total-Traffic-1-LHCOPN_TOTAL-OUT_Yearly.png

Smooth and large data transfers
need network optimization



Smooth and large data transfers need network optimization

Network
Optimized
Transfer of
Experimental
Data



To load-balance dynamically,
we tackle these **key challenges**

when to do

Obtain transfer information

what, and

Load-balance only during the transfer

how, and

Find load-balancing mechanism

how well

Measure impact

To load-balance dynamically,
we tackle these **key challenges**

Obtain transfer information

Load-balance only during the transfer

Find load-balancing mechanism

Measure impact

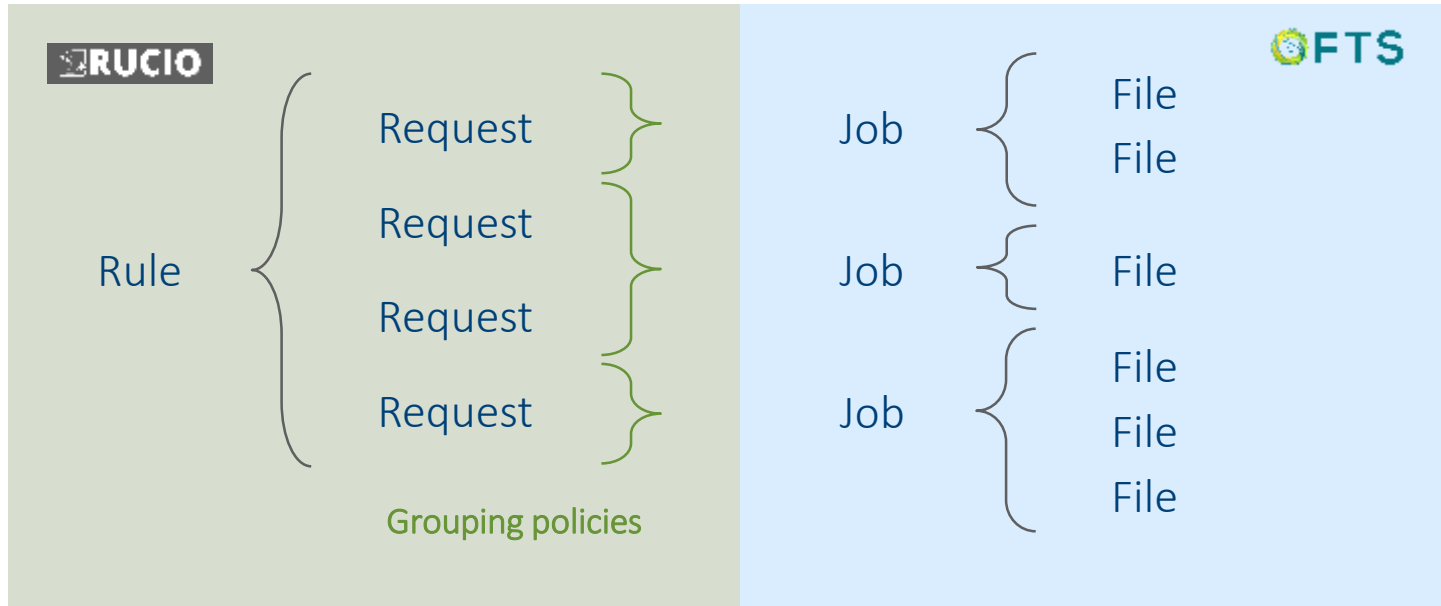
To obtain transfer information,
we look inside **Fabiola's transfer tool**



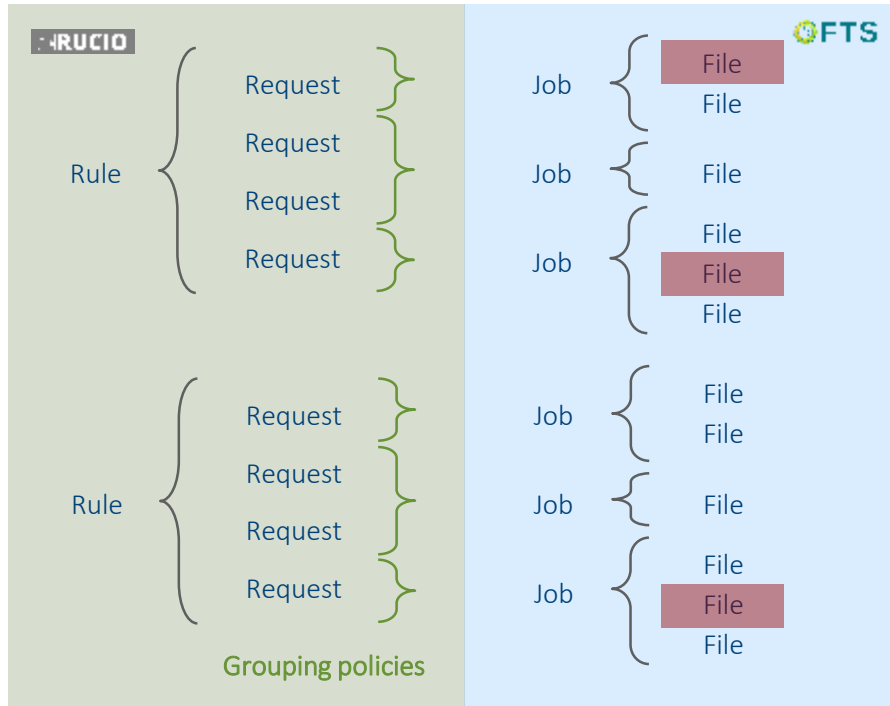
To obtain transfer information,
we look inside **Rucio** and **FTS**



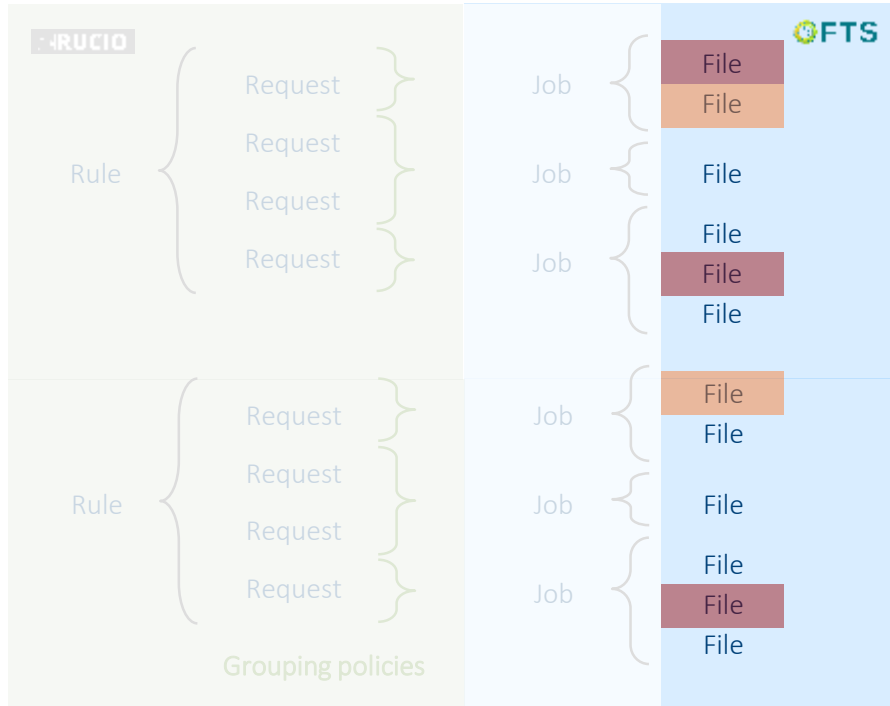
To obtain transfer information,
we look inside **Rucio** and **FTS**




Different rules can act on the same source-destination pair



Hence, we get transfer information from the **FTS** queue of submissions




Files in
SUBMITTED
state



To obtain transfer information,
we interface with Rucio and FTS



Files in
SUBMITTED
state



To load-balance dynamically,
we tackle these **key challenges**

Obtain transfer information

Load-balance only during the transfer

Find load-balancing mechanism

Measure impact

To load-balance dynamically,
we tackle these **key challenges**

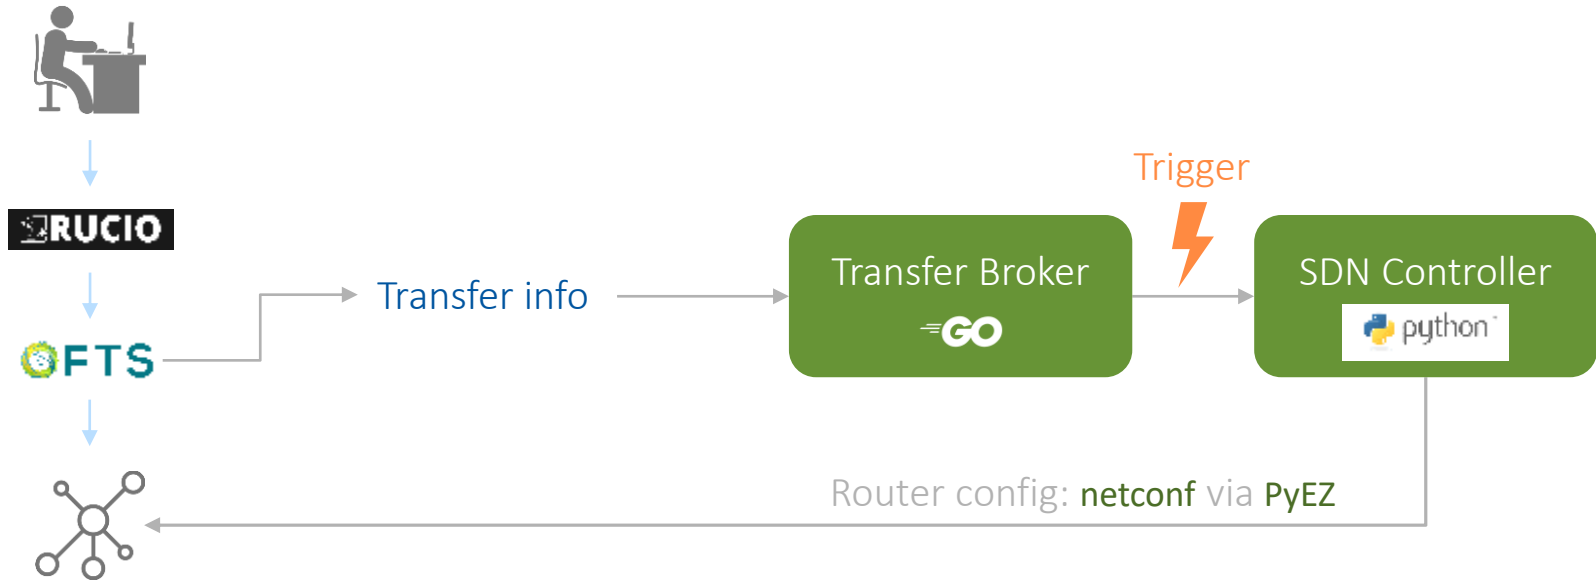
Obtain transfer information

Load-balance *only during* the transfer

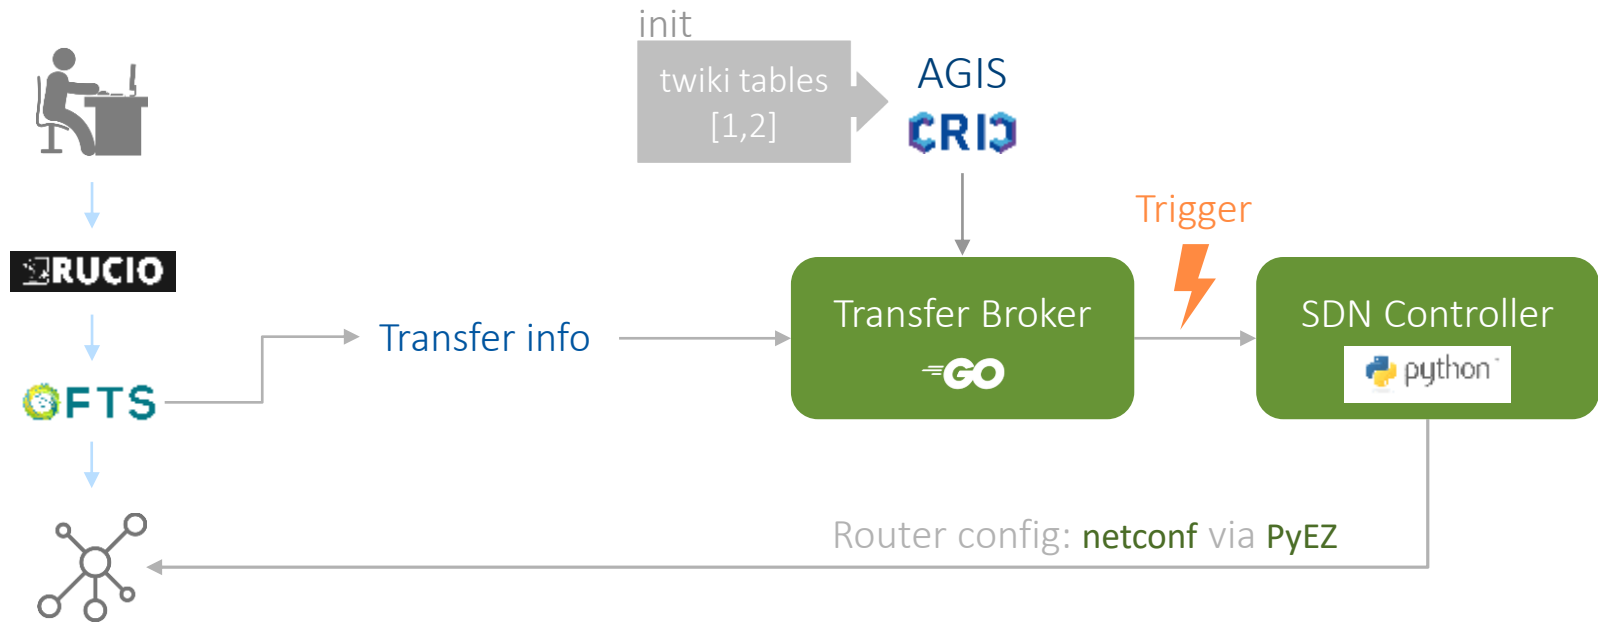
Find load-balancing mechanism

Measure impact

When a large transfer starts,
we trigger a SDN controller



We construct the SDN action
with IP prefixes from AGIS/CRIC



To load-balance dynamically,
we tackle these **key challenges**

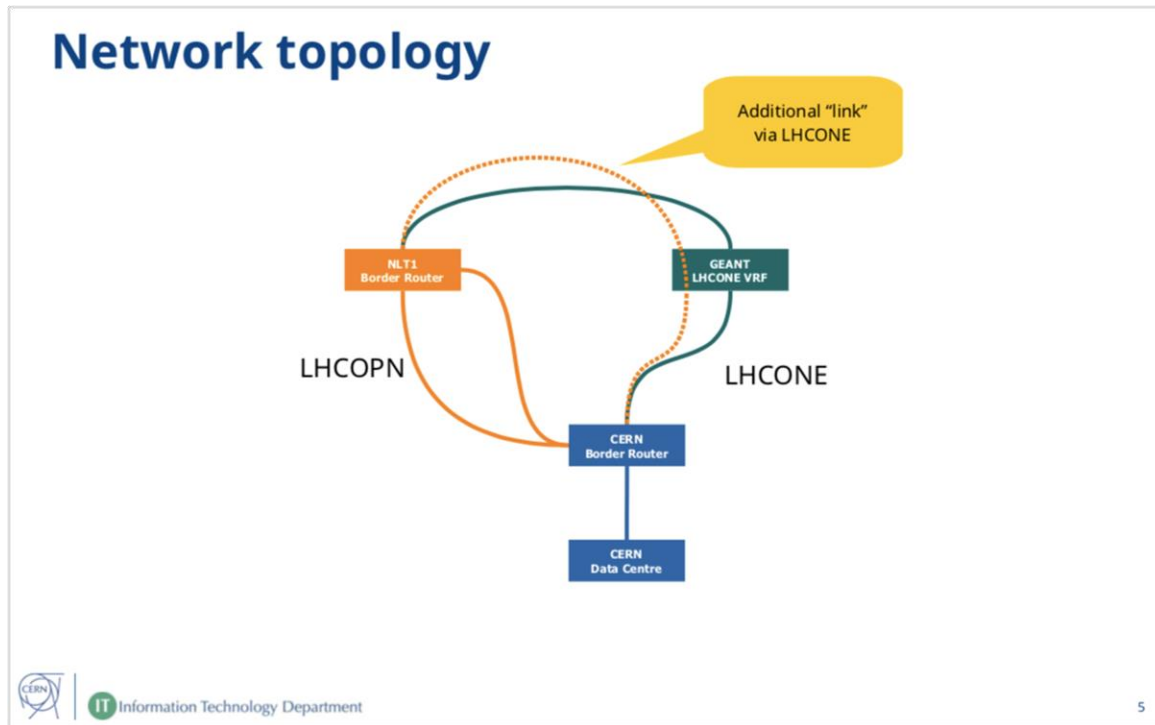
Obtain transfer information

Load-balance *only during* the transfer

Find load-balancing mechanism

Measure impact

As a first load-balancing mechanism,
we automated BGP multipath to NL-T1 and DE-KIT



To load-balance dynamically,
we tackle these **key challenges**

Obtain transfer information

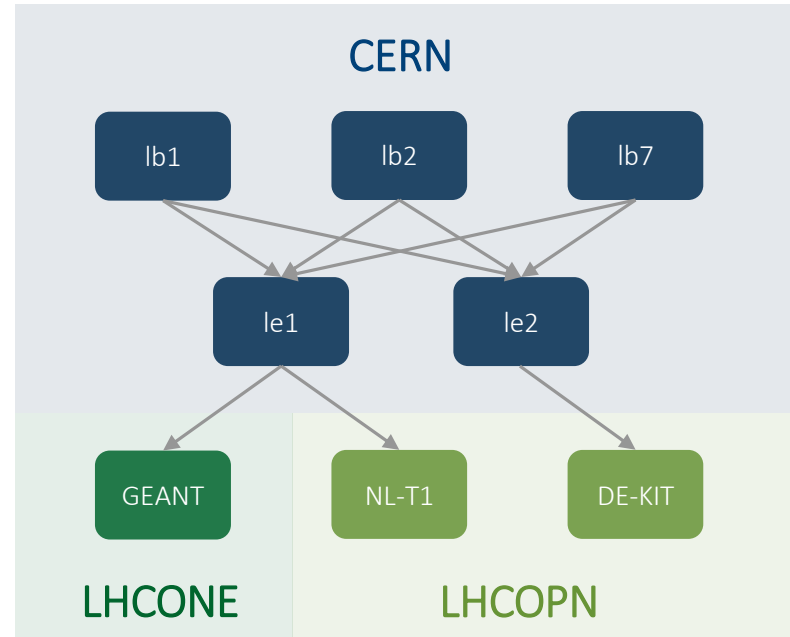
Load-balance *only during* the transfer

Find load-balancing mechanism

Measure impact

We measured the impact in a proof-of-concept data transfer

- A) 20 TB from CERN to NL-T1
- B) 20 TB from CERN to DE-KIT

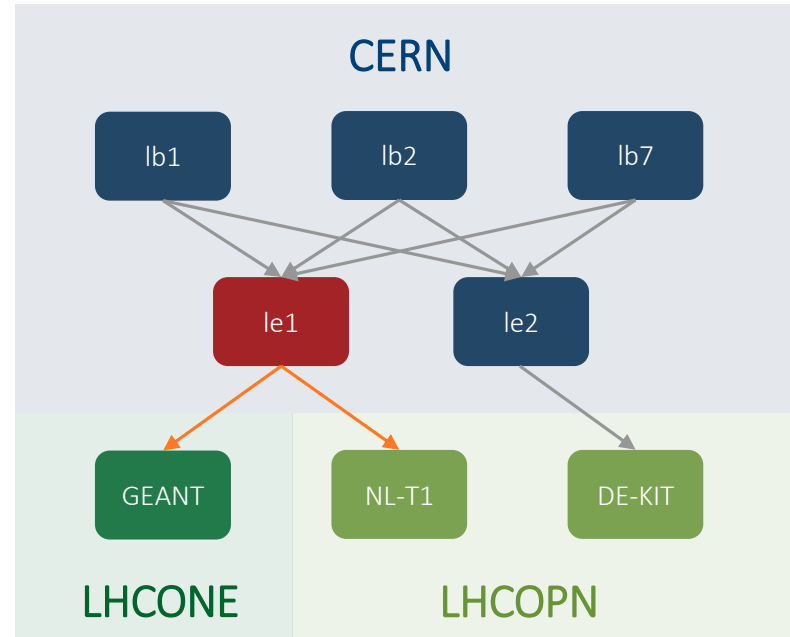


We measured the impact in a proof-of-concept data transfer

- A) 20 TB from CERN to NL-T1
- B) 20 TB from CERN to DE-KIT

Routers to configure...

... to obtain equal cost

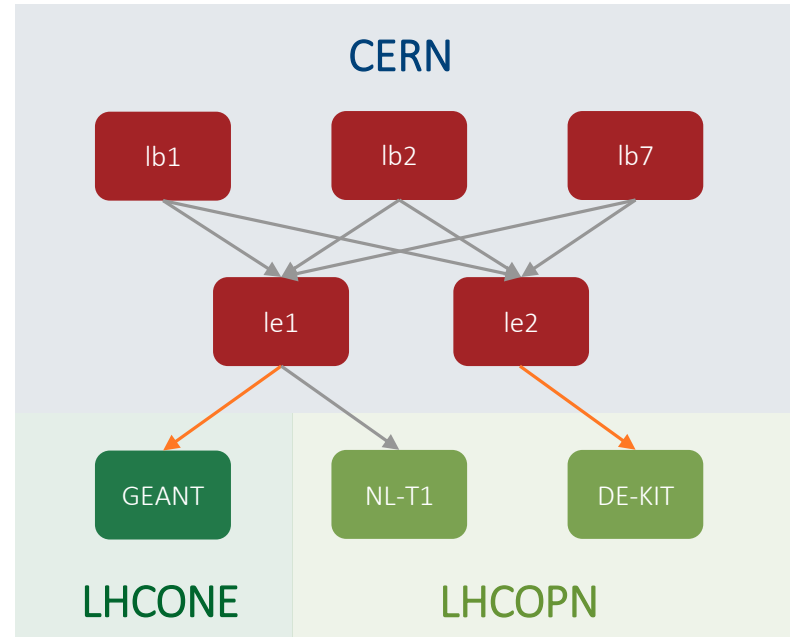


We measured the impact in a proof-of-concept data transfer

- A) 20 TB from CERN to NL-T1
- B) 20 TB from CERN to DE-KIT

Routers to configure...

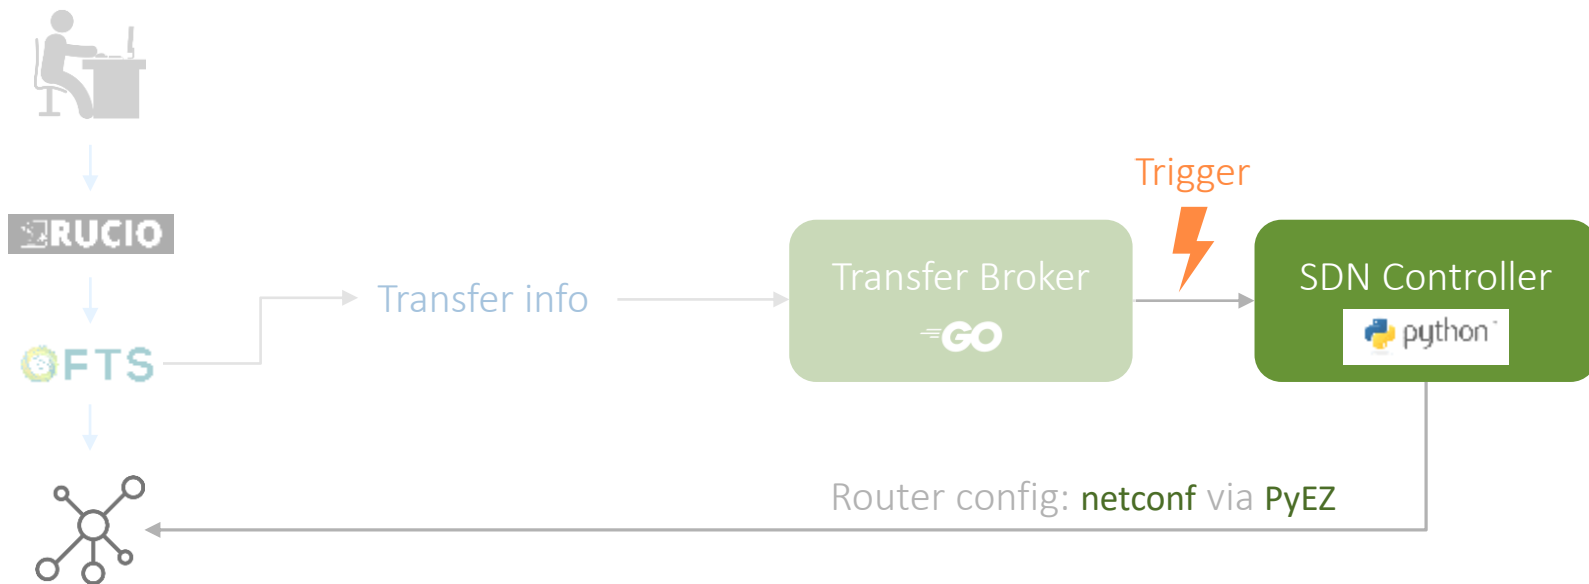
... to obtain equal cost



We see the transfer
in the **FTS** queue size



We **trigger** the load-balancing...



We also see the actions in the SDN controller behind the scenes

2019-12-11 16:15:05,349 - CERN-PROD---DE-KIT: start loadbalancing

2019-12-11 16:15:05,349 - host le1: prepend policy: PS-NOTED-LHCONE-V4-IN-DE-KIT to: PS-LHCONE-V4-IN in group: LHCONE in neighbor: xxx

2019-12-11 16:15:10,620 - host le1: prepend policy: PS-NOTED-LHCONE-V6-IN-DE-KIT to: PS-LHCONE-V6-IN in group: LHCONE in neighbor: xxx

2019-12-11 16:15:15,907 - host le2: prepend policy: PS-NOTED-LHCOPN-V4-IN-DE-KIT to: PS-LHCOPN-V4-IN in group: LHCOPN in neighbor: xxx

2019-12-11 16:15:22,700 - host le2: prepend policy: PS-NOTED-LHCOPN-V6-IN-DE-KIT to: PS-LHCOPN-V6-IN in group: LHCOPN in neighbor: xxx

2019-12-11 16:15:29,501 - host lb1: set multiple-as

2019-12-11 16:15:33,387 - host lb2: set multiple-as

2019-12-11 16:15:37,304 - host lb7: set multiple-as

We also see the actions in the SDN controller behind the scenes

2019-12-11 17:06:05,349 - CERN-PROD---DE-KIT: stop loadbalancing

2019-12-11 17:06:05,349 - host le1: load policies: PS-LHCONE-V4-IN in group: LHCONE in neighbor: xxx

2019-12-11 17:06:09,963 - host le1: load policies: PS-LHCONE-V6-IN in group: LHCONE in neighbor: xxx

2019-12-11 17:06:14,527 - host le2: load policies: PS-LHCOPN-V4-IN in group: LHCOPN in neighbor: xxx

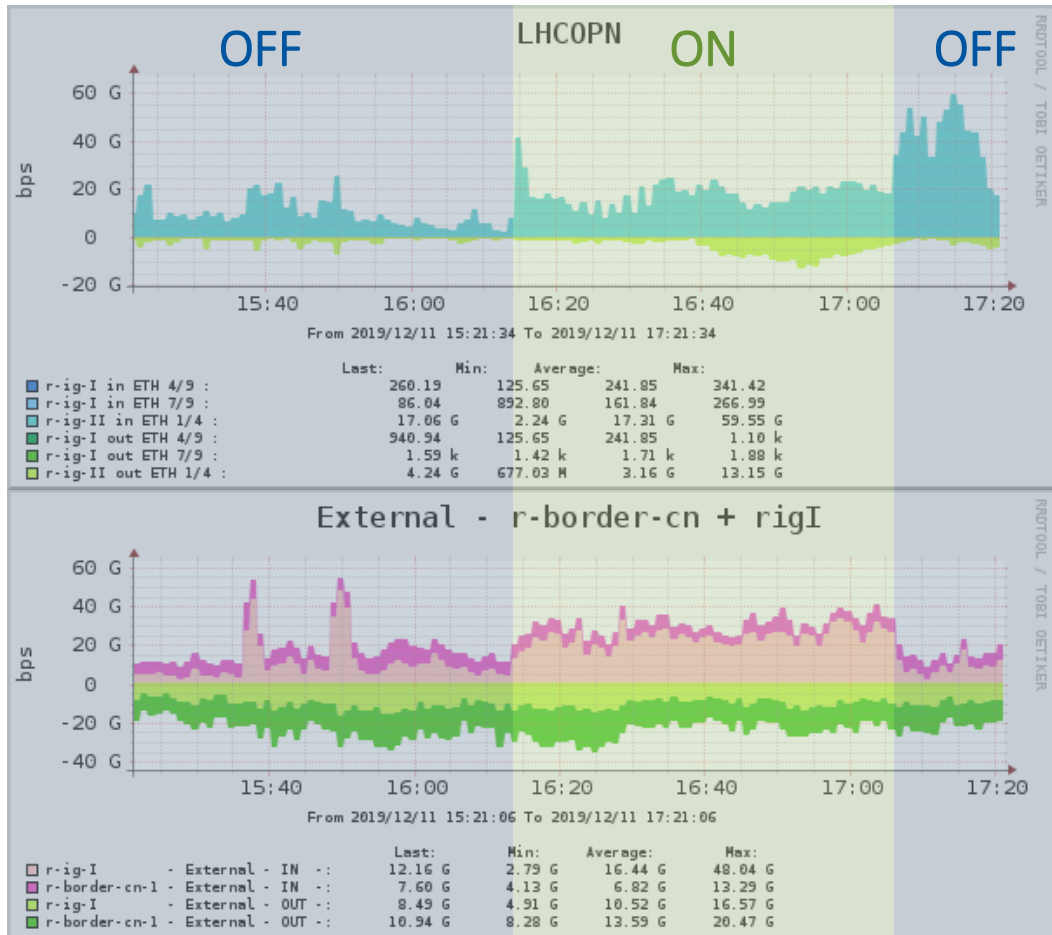
2019-12-11 17:06:18,795 - host le2: load policies: PS-LHCOPN-V6-IN in group: LHCOPN in neighbor: xxx

2019-12-11 17:06:23,081 - host lb1: delete multiple-as

2019-12-11 17:06:27,011 - host lb2: delete multiple-as

2019-12-11 17:06:30,938 - host lb7: delete multiple-as

For the DE-KIT transfer,
we see the traffic shifting



Plots by Bruno Hoefft



To load-balance dynamically,
we tackle these **key challenges**



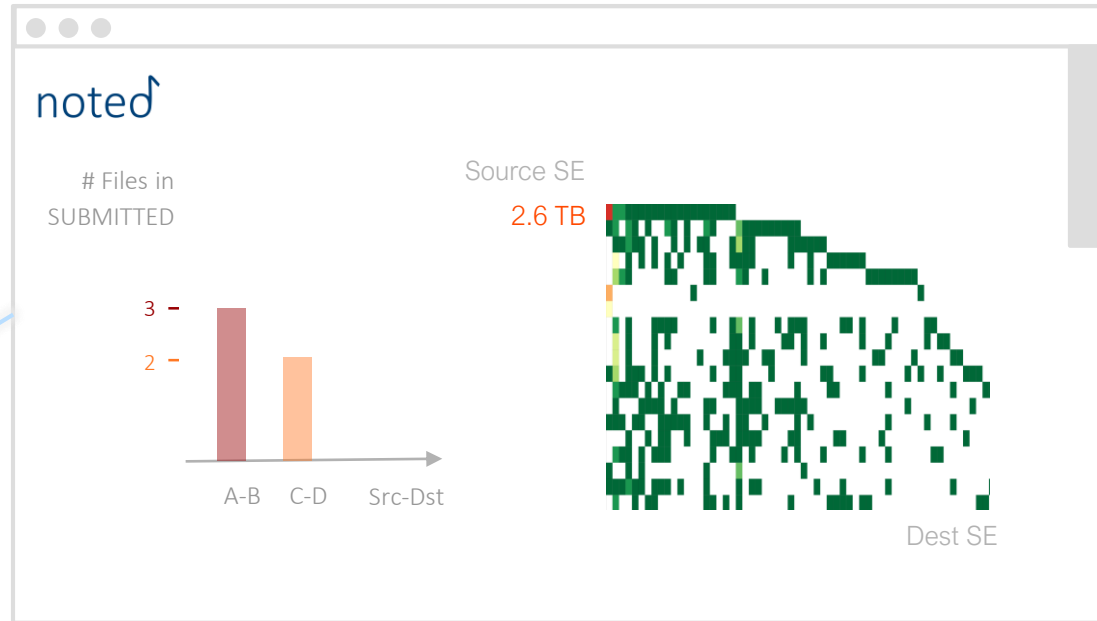
Obtain transfer information

Load-balance *only during* the transfer

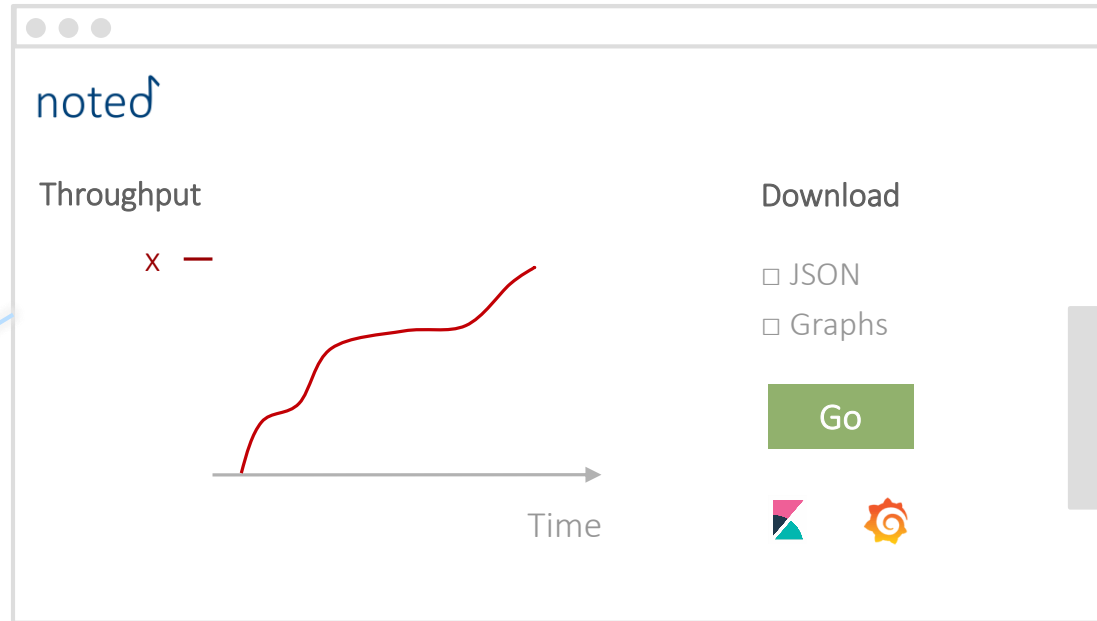
Find load-balancing mechanism

Measure **& show** impact

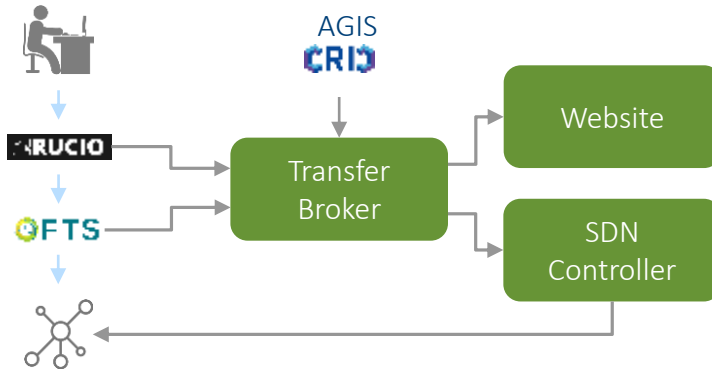
We can make the Transfer Broker
information publicly available *for you*



We can make the Transfer Broker
information publicly available *for you*



To load-balance dynamically,
we tackled these key challenges



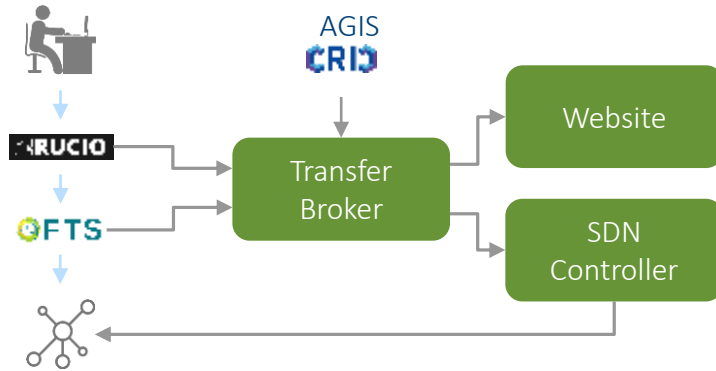
Obtain transfer information

Load-balance *only during* the transfer

Find load-balancing mechanism

Measure & show impact

To load-balance dynamically,
we tackled these key challenges



Obtain transfer information


Load-balance *only during* the transfer

Find load-balancing mechanism

Measure & show impact

We would love to hear your feedback!

Future work...



Load-balancing:
segment routing in MPLS


Future work...



Load-balancing:
segment routing in MPLS

Use StackStorm as
SDN controller

Future work...



Load-balancing:
segment routing in MPLS

Use StackStorm as
SDN controller

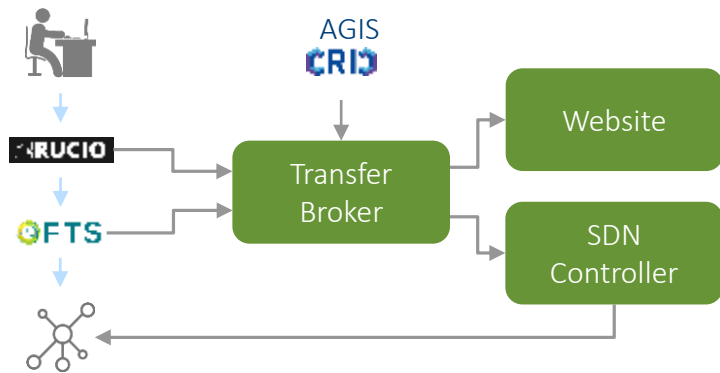
Also cover remoteIO/WMagent
(some jobs produce ~1PB load,
heavily used by e.g. CMS)

Questions?

edoardo.martelli@cern.ch

coralie.busse-grawitz@cern.ch

To load-balance dynamically,
we tackled these key challenges



Obtain transfer information

Load-balance *only during* the transfer

Find load-balancing mechanism

Measure & show impact

We would love to hear your feedback!

Rucio does **not** always
group requests by location

