



# Congestion driven NOTED

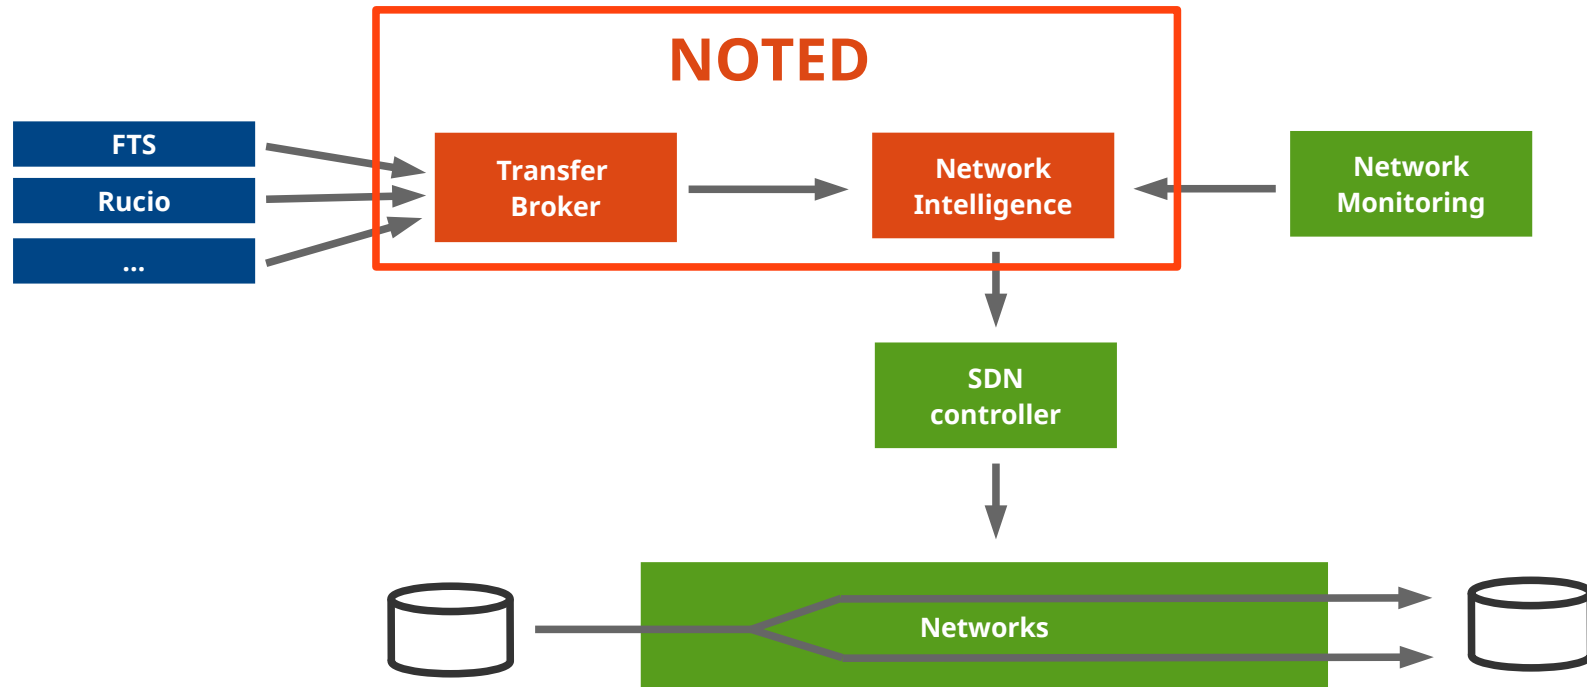
LHCONE meeting #51

19<sup>th</sup> October 2023 – Victoria CA

[edoardo.martelli@cern.ch](mailto:edoardo.martelli@cern.ch)

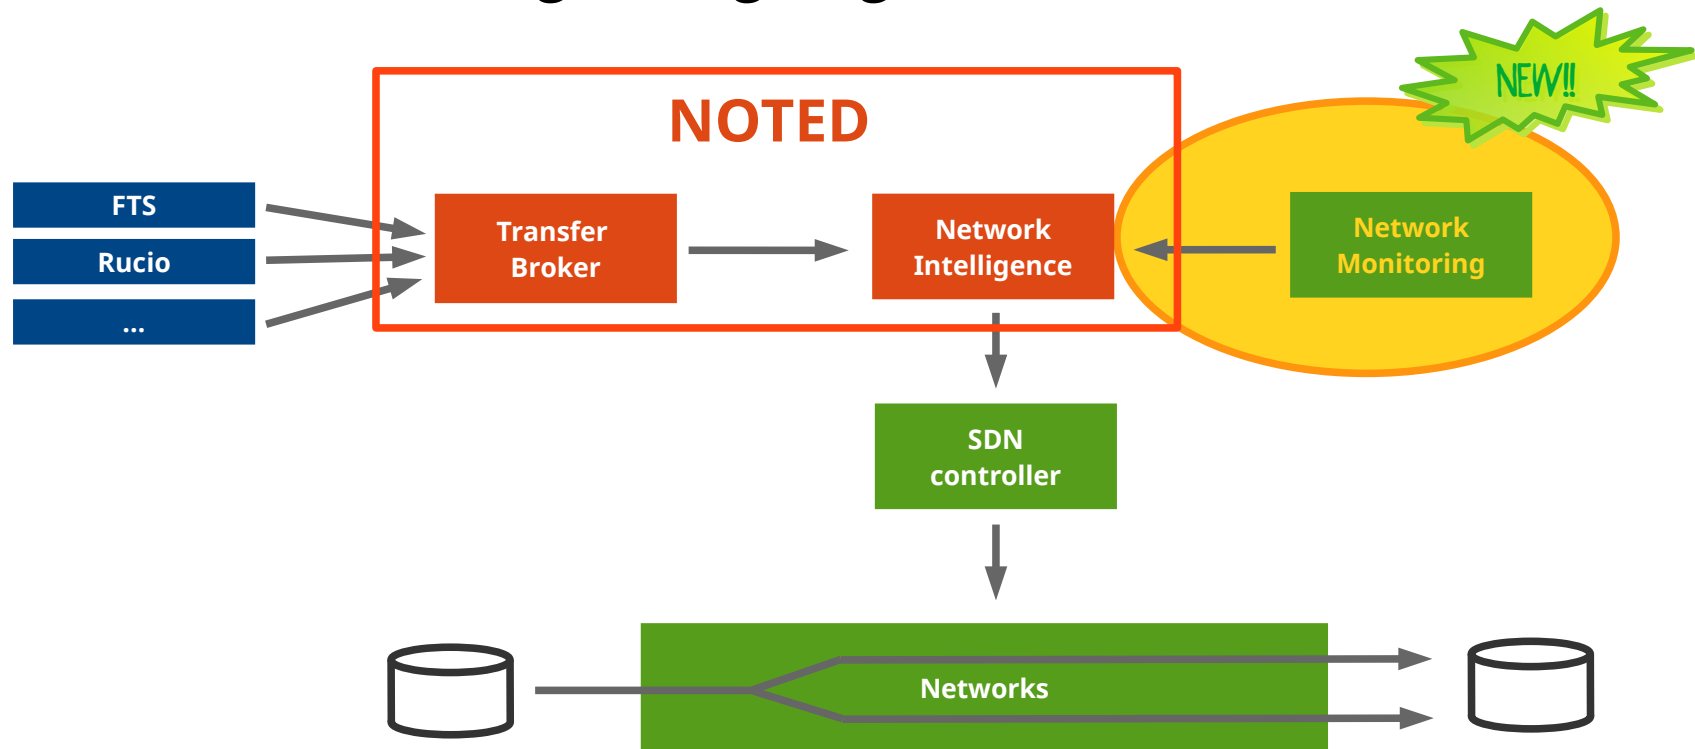
# Recap on NOTED

**NOTED:** framework that dynamically improves network performances for large, on-going data transfers



# Recap on NOTED

**NOTED:** framework that dynamically improves network performances for large, on-going data transfers



# Wait for link saturation alarms

- Poll the alarms **LOAD THRESHOLD EXCEDEED** generated by the CERN NMS

Alarms / Spectrum Alarms ☆ 🔊

Instance Production ▾ Entity name Enter variable value Cause ID Enter variable value Exclude secstring --- ▾

IT/CS Alarm History

Severity	Occ	Entity name	Type	Class	Alarm name	Ack	Start at	Cleared at	Duration
MINOR	1	<a href="#">1513-e-rjup1-1_irb.2126</a>	Gen_IF_Port	Port	OUT LOAD THRESH...	No	2023-10-06 08:59:23		
MINOR	1	<a href="#">1513-e-rjup1-1_irb.3530</a>	Gen_IF_Port	Port	OUT LOAD THRESH...	No	2023-10-06 08:43:05		
MINOR	1	<a href="#">1513-v-rjuxl-12_xe-1_0_11</a>	Gen_IF_Port	Port	OUT LOAD THRESH...	No	2023-10-06 08:21:10		
MINOR	1	<a href="#">1513-e-rjup1-1_irb.3530</a>	Gen_IF_Port	Port	OUT LOAD THRESH...	No	2023-10-06 08:08:03	2023-10-06 08:22:50	00:14:47
MINOR	1	<a href="#">1513-e-rjup1-1_irb.3530</a>	Gen_IF_Port	Port	IN LOAD THRESHO...	No	2023-10-06 07:28:02	2023-10-06 07:32:47	00:04:45
MINOR	1	<a href="#">1513-e-rjup1-1_irb.3530</a>	Gen_IF_Port	Port	IN LOAD THRESHO...	No	2023-10-06 06:58:02	2023-10-06 07:07:47	00:09:45
MINOR	1	<a href="#">1513-v-rjuxl-12_xe-1_0_12</a>	Gen_IF_Port	Port	OUT LOAD THRESH...	No	2023-10-06 06:46:00	2023-10-06 07:36:14	00:50:14
MINOR	1	<a href="#">1513-e-rjup1-1_irb.2126</a>	Gen_IF_Port	Port	OUT LOAD THRESH...	No	2023-10-06 06:34:23	2023-10-06 08:44:08	02:09:45
MINOR	1	<a href="#">1513-v-rjuxl-12_xe-1_0_13</a>	Gen_IF_Port	Port	OUT LOAD THRESH...	No	2023-10-06 06:15:58	2023-10-06 07:55:57	01:39:59
MINOR	1	<a href="#">1513-e-rjup1-1_irb.3530</a>	Gen_IF_Port	Port	OUT LOAD THRESH...	No	2023-10-06 05:53:02	2023-10-06 06:27:47	00:34:45



# Identify the prefixes routed via the interface

Identify the prefixes routed via the alarmed interface

## - Find the IP of the next hop:

```
BORDER-ROUTER> show interfaces irb.3530 terse
Interface          Admin Link Proto  Local                               Remote
irb.3530           up    up    inet   172.24.18.9/30
                   inet6  2001:1458:302:38::1/64
```

## - Find the routed prefixes:

```
BORDER-ROUTER> show route next-hop 2001:1458:302:38::2

2a00:139c::/45      *[BGP/170] 2d 23:16:51, MED 10, localpref 100
                   AS path: 58069 I, validation-state: unverified
                   > to 2001:1458:302:38::2 via irb.3530
```

# Find the corresponding site in CRIC

- **Lookup routed prefixes in CRIC to identify the destination site**

This action also confirms that the alarm is related to a LHCOPN/ONE link

The screenshot displays the CRIC interface for NetworkRoute: FZK-LCG2-LHCOPNE. It lists various network attributes such as Network Site (DE-KIT), ASN (58069), and a list of IP prefixes. The prefix 2a00:139c::/45 is highlighted with a red box. Below this, the LHCONE section shows a bandwidth limit of 200 and collaborations including WLCG, BelleII, PierreAugerObservatory, and XENON. The LHCOPN section also shows a bandwidth limit of 200. Red arrows from the text on the left point to the highlighted prefix and the LHCOPN section.

Network Site	DE-KIT
ASN	58069
More specific	False
Monitoring URL	Not set
Networks	<ul style="list-style-type: none"><li>157.180.228.0/22</li><li>157.180.232.0/22</li><li>192.108.45.0/24</li><li>192.108.46.0/23</li><li>192.108.68.0/24</li><li>2a00:139c::/45</li></ul>
Last modified	2022-01-19 12:27:47.145947

LHCONE

Bandwidth limit	200
Collaborations	WLCG, BelleII, PierreAugerObservatory, XENON

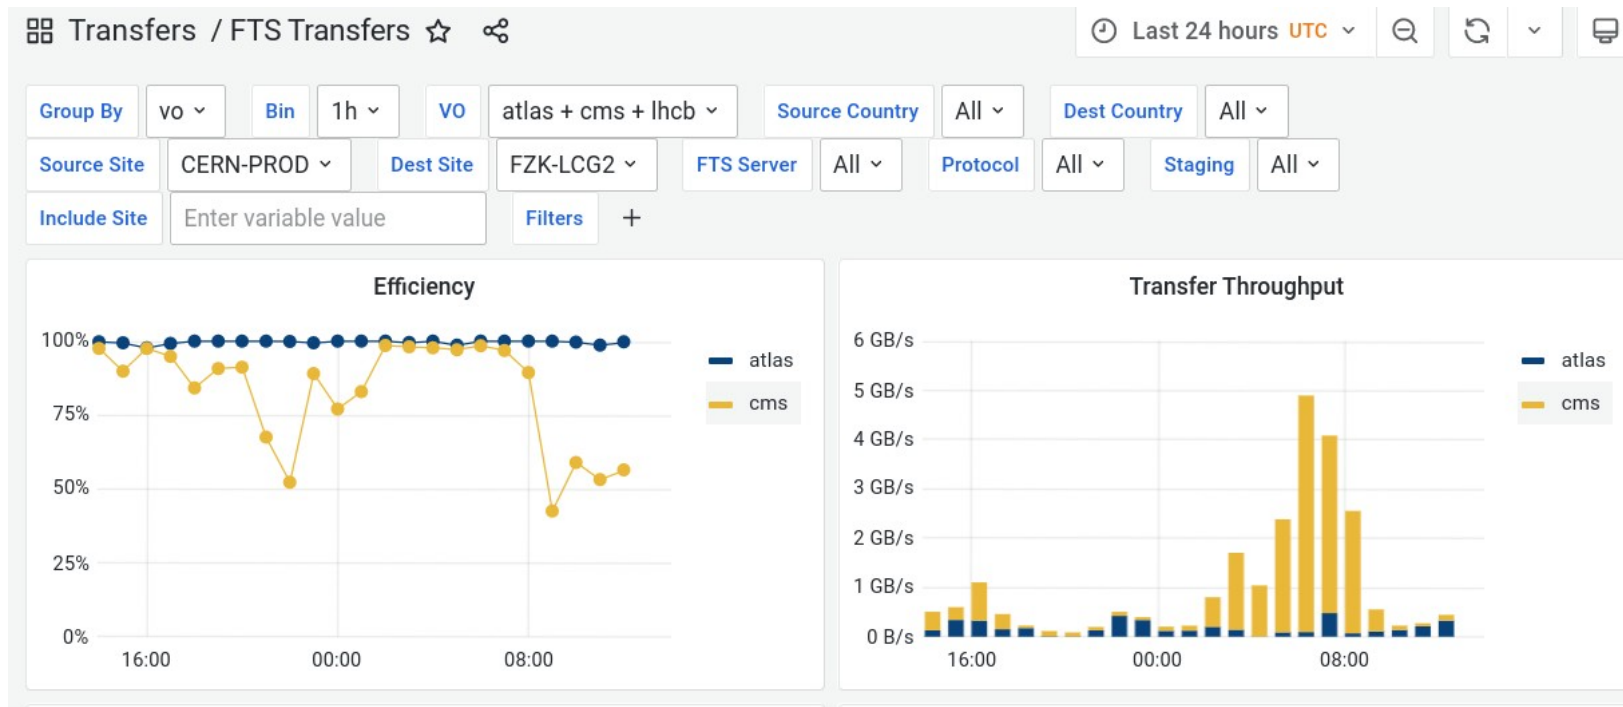
LHCOPN

Bandwidth limit	200
-----------------	-----

CRIC WLCG Web Portal server: wlcg-cric-2.cern.ch

# Look for FTS transfers in FTS Monit

- Get confirmation that it is a FTS transfer that is causing the congestion



# Compile NOTED configuration and run it

- Compile a NOTED configuration to monitor and estimate the duration of the overloading transfers
- Launch NOTED to manage network corrective actions for the overloaded link
- Publish the NOTED actions in Monit Grafana

NOTED Alarms ⓘ

ID ↓	Alarm name	Version	NOTED status	NOTED action	SDN status	Max FTS Throughput [Gb/s]	Interface	NOTED alarm start
5	CH-CERN to DE-KIT	LHCOPN	Action	On-going SDN. FTS throughput [Gb/s]: 28.72	Provided	24.9	<a href="#">I513-e-rjup1-1_irb.3530</a>	2023-10-11 08:27:56
4	CH-CERN to DE-KIT	LHCOPN	Running	Spectrum generated an alarm: NOTED is inspecting FTS.			<a href="#">I513-e-rjup1-1_irb.3512</a>	2023-10-10 18:03:00
3	CH-CERN to DE-KIT	LHCOPN	Running	Spectrum generated an alarm: NOTED is inspecting FTS.			<a href="#">I513-e-rjup1-1_irb.3530</a>	2023-10-10 18:02:43
2	CH-CERN to CA-TRIUMF	LHCOPN	Running	Spectrum generated an alarm: NOTED is inspecting FTS.			<a href="#">I513-e-rjup1-1_irb.2126</a>	2023-10-10 18:02:28
1	CH-CERN to FR-IN2P3-LAPP	LHCONE	Running	Spectrum generated an alarm: NOTED is inspecting FTS.			<a href="#">I513-e-rjup1-1_irb.111</a>	2023-10-10 18:02:13



# NOTED actions

**Decision-making:** NOTED has started and it is looking for FTS transfers to potentially execute an action or not

**Running:** NOTED is running but there are no transfers in FTS so NOTED is waiting and running until the link-saturation alarm is cleared.

**Monitoring:** NOTED is running and there are on-going FTS transfers, but they are below the defined bandwidth threshold that we establish

**Action:** NOTED is running and has triggered a SDN action to provide more bandwidth

**Stopped:** NOTED has stopped because there are no transfers in FTS and the link-saturation alarm has cleared

# NOTED actions

Development / NOTED Alarms ★



Last 24 hours ▾



1m ▾



## NOTED Alarms

ID <span>▾</span>	Alarm name	NOTED status	NOTED action	SDN status	Max FTS Throughput [Gb/s]	Interface	NOTED alarm start	NOTED alarm end	NOTED duration
26	CA-TRIUMF to CH-CERN	Decision-making	An action on the link may be required: number of events: 1. Throughput [Gb/s]: 8.31			<a href="#">I513-e-rjup1-1_irb.2126</a>	2023-09-01 15:46:30		
25	CH-CERN to DE-KIT	Monitoring	No transfers found in FTS. NOTED is still running until Spectrum clears the alarm.			<a href="#">I513-e-rjup1-1_irb.3530</a>	2023-09-01 15:42:25		
24	CH-CERN to DE-KIT	Monitoring	No transfers found in FTS. Spectrum alarm is cleared	Released	0.740	<a href="#">I513-e-rjup1-1_irb.3530</a>	2023-09-01 15:25:34	2023-09-01 15:30:43	00:05:09
23	CA-TRIUMF to CH-CERN	Stopped	The large data transfer is finished.	Released	17.0	<a href="#">I513-e-rjup1-1_irb.2126</a>	2023-09-01 15:15:49	2023-09-01 15:26:04	00:10:15
22	CH-CERN to DE-KIT	Monitoring	No transfers found in FTS. Spectrum alarm is cleared	Released	1.61	<a href="#">I513-e-rjup1-1_irb.3530</a>	2023-09-01 14:57:50	2023-09-01 15:16:17	00:18:27

Development / NOTED Alarms ★



Last 2 days ▾



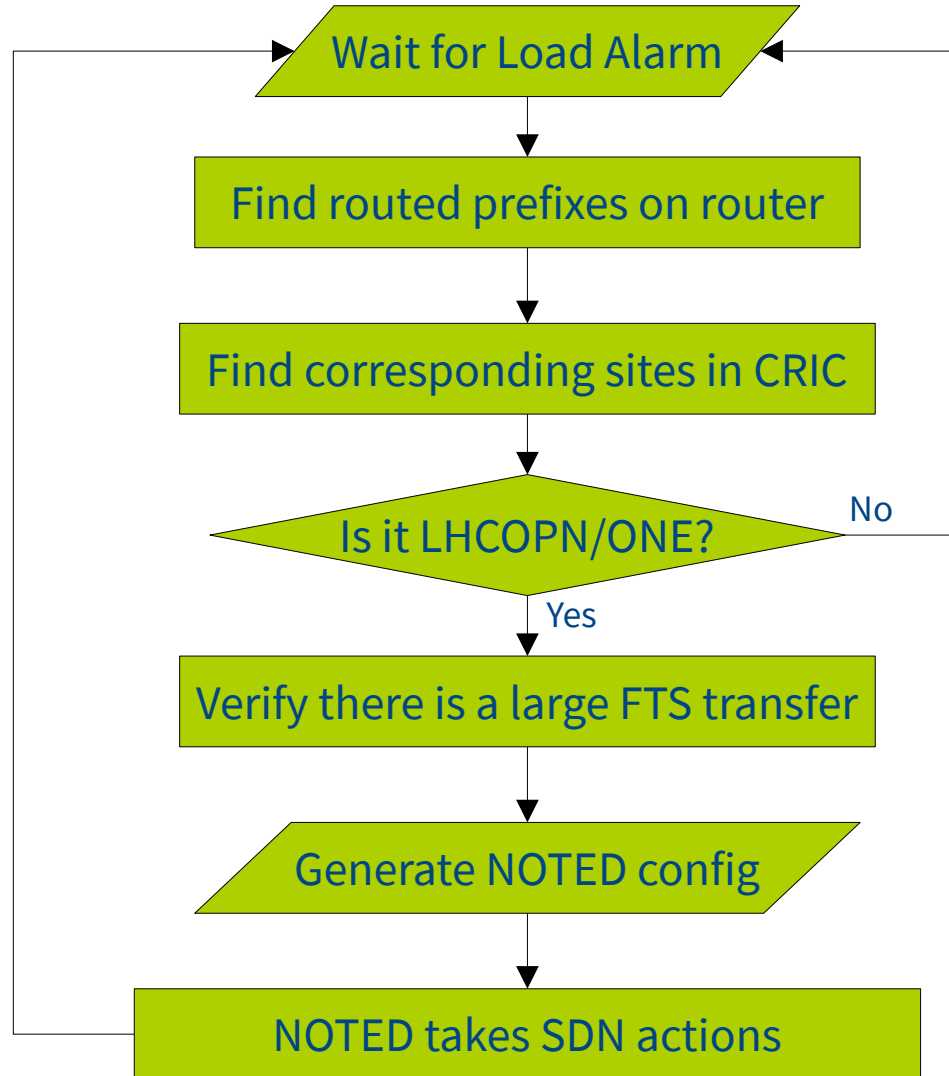
1m ▾



## NOTED Alarms

ID <span>▾</span>	Alarm name	NOTED status	NOTED action	SDN status	Max FTS Throughput [Gb/s]	Interface	NOTED alarm start	NOTED alarm end	NOTED duration
26	CA-TRIUMF to CH-CERN	Action	On-going SDN. FTS throughput [Gb/s]: 6.22	Provided	8.31	<a href="#">I513-e-rjup1-1_irb.2126</a>	2023-09-01 15:46:30		
25	CH-CERN to DE-KIT	Monitoring	No transfers found in FTS. NOTED is still running until Spectrum clears the alarm.			<a href="#">I513-e-rjup1-1_irb.3530</a>	2023-09-01 15:42:25		
24	CH-CERN to DE-KIT	Monitoring	No transfers found in FTS. Spectrum alarm is cleared	Released	0.740	<a href="#">I513-e-rjup1-1_irb.3530</a>	2023-09-01 15:25:34	2023-09-01 15:30:43	00:05:09
23	CA-TRIUMF to CH-CERN	Stopped	The large data transfer is finished.	Released	17.0	<a href="#">I513-e-rjup1-1_irb.2126</a>	2023-09-01 15:15:49	2023-09-01 15:26:04	00:10:15
22	CH-CERN to DE-KIT	Monitoring	No transfers found in FTS. Spectrum alarm is cleared	Released	1.61	<a href="#">I513-e-rjup1-1_irb.3530</a>	2023-09-01 14:57:50	2023-09-01 15:16:17	00:18:27
21	CA-TRIUMF to CH-CERN	Monitoring	No transfers found in FTS. Spectrum alarm is cleared	Released	31.8	<a href="#">I513-e-rjup1-1_irb.2126</a>	2023-09-01 14:04:05	2023-09-01 14:52:00	00:47:55
20	CA-TRIUMF to CH-CERN	Stopped	The large data transfer is finished.	Released	10.2	<a href="#">I513-e-rjup1-1_irb.2126</a>	2023-09-01 13:01:12	2023-09-01 13:13:38	00:12:26
19	CH-CERN to DE-KIT	Stopped	The large data transfer is finished.	Released	2.50	<a href="#">I513-e-rjup1-1_irb.3530</a>	2023-09-01 10:19:05	2023-09-01 13:01:54	02:42:49
18	CH-CERN to DE-KIT	Running	Monitoring transfers. FTS throughput [Gb/s]: 0.01			<a href="#">I513-e-rjup1-1_irb.3530</a>	2023-09-01 09:34:47		
17	CH-CERN to CA-TRIUMF	Stopped	The large data transfer is finished.	Released	1.13	<a href="#">I513-e-rjup1-1_irb.2126</a>	2023-09-01 08:31:13	2023-09-01 08:52:40	00:21:27
16	CH-CERN to DE-KIT	Stopped	The large data transfer is finished.	Released	44.0	<a href="#">I513-e-rjup1-1_irb.3530</a>	2023-09-01 08:28:12	2023-09-01 09:36:33	01:08:21

# Workflow



# Plan

- Implemented for LHCOPN links
- Working on version for LHCONE links (more complex because of many possible destinations on the same link)
- Dry-run during DC24 (unless there is the real possibility to relief heavily congested links)

*Questions?*

*edoardo.martelli@cern.ch*  
*carmen.misa@cern.ch*

