
CERN

NOTED: An intelligent network controller to improve the throughput of large data transfers in File Transfer Services by handling dynamic circuits

CERN
IT Department CS Group

Data Challenge 2024 Workshop

Carmen Misa Moreira
Edoardo Martelli

Outline

Recapitulation

Motivation

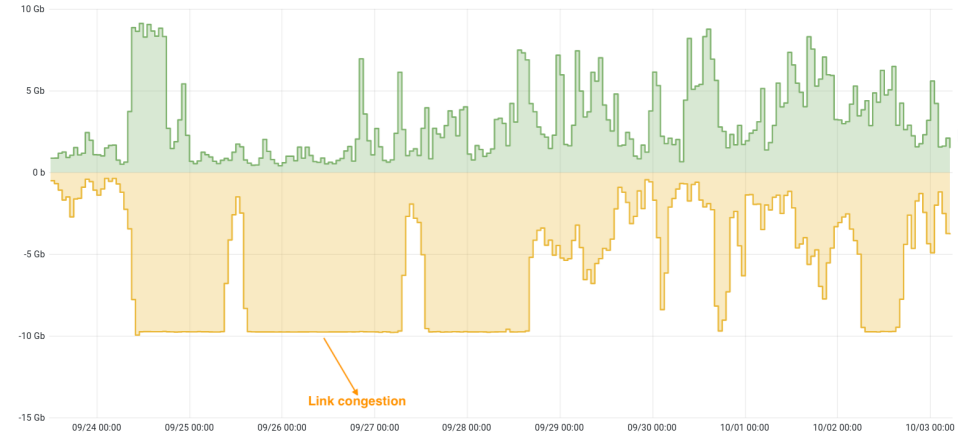

- Large data transfers can saturate network links while alternative paths may be left idle

Architecture

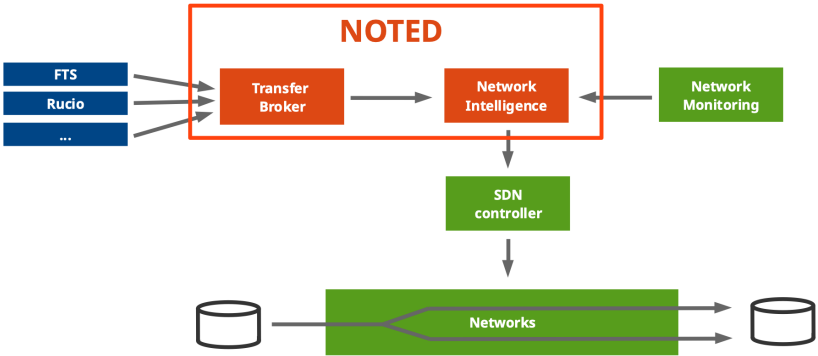

NOTED (Network Optimized Transfer of Experimental Data)

An intelligent network controller to improve the throughput of large data transfers in FTS (File Transfer Services) by handling dynamic circuits.

Elements

FTS (File Transfer Service):

- Analyse data transfers to estimate if any action can be applied to optimise the network utilization → get on-going and queued transfers.

CRIC (Computing Resource Information Catalog):

- Use the CRIC database to get an overview of the network topology → get IPv4/IPv6 addresses, endpoints, rcsite and federation.

FTS
File Transfer Service

elasticsearch

Package distribution and installation

Available in PyPI https://pypi.org/project/noted-dev/

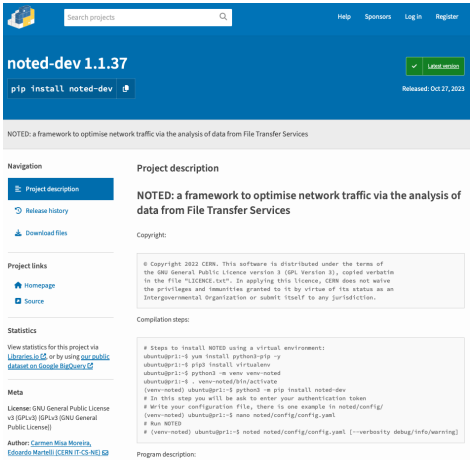

Common steps:

```
# Create a virtual environment:
$ pip3 install virtualenv
$ python3 -m venv venv-noted
$ .  venv-noted/bin/activate
```

Ubuntu installation:

```
# Install noted-dev
(venv-noted) $ python3 -m pip install noted-dev
# Write your configuration file
(venv-noted) $ nano noted/config/config.yaml
# Run NOTED
(venv-noted) $ noted noted/config/config.yaml
```

CentOS installation:

```
# Download noted-dev.tar.gz
(venv-noted) $ wget url_pypi_repo_tar_gz
# Install noted-dev
(venv-noted) $ tar -xf noted-dev-1.1.62.tar.gz
(venv-noted) $ pip install noted-dev-1.1.62/
# Run NOTED
(venv-noted) $ noted noted/config/config.yaml
```

Package distribution and installation

Available in Docker https://hub.docker.com/r/carmenmisa/noted-docker

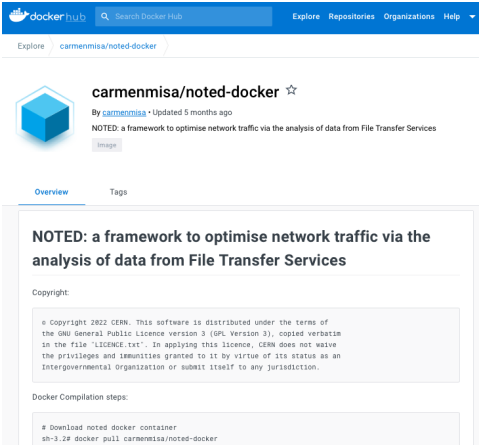

Installation:

```
# Download noted docker container:
$ docker pull carmenmisa/noted-docker

# Run docker container:
$ docker run --detach --entrypoint /sbin/init
--network="host" --privileged --name noted_controller
carmenmisa/noted-docker

# Copy your configuration file into the container:
$ docker cp src/noted/config/config.yaml
noted_controller:/app/noted/config

# Run commands in the container from outside:
$ docker exec noted_controller noted -h
$ docker exec noted_controller
/app/src/noted/scripts/setup.sh mail

# Run NOTED
$ docker exec noted_controller noted config/config.yaml &
```

LHCOPN, LHCONE version

New features

CUSTOM

NOTED is working based on the parameters written in a config.yaml file by the network administrator to monitor FTS data transfers

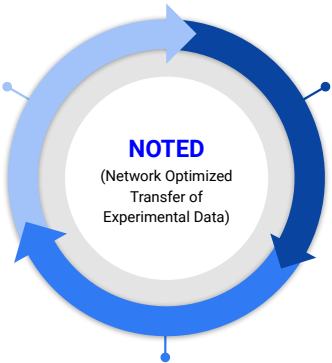

LHCOPN

When CERN NMS raises an alarm on an interface in one of the LHCOPN border routers, NOTED identifies the Tier 1 and starts to monitor FTS data transfers → automatically!

LHCONE

When CERN NMS raises an alarm on an interface in one of the LHCONE border routers, NOTED identifies the Tier 2, Tier 3 and starts to monitor FTS data transfers → automatically!

- Much more complex for LHCONE since a single path is shared by multiple sites $\sim$ 100.

Network monitoring (LHCOPN, LHCONE version)

- Poll the alarms IN/OUT LOAD THRESHOLD EXCEEDED generated by the CERN NMS

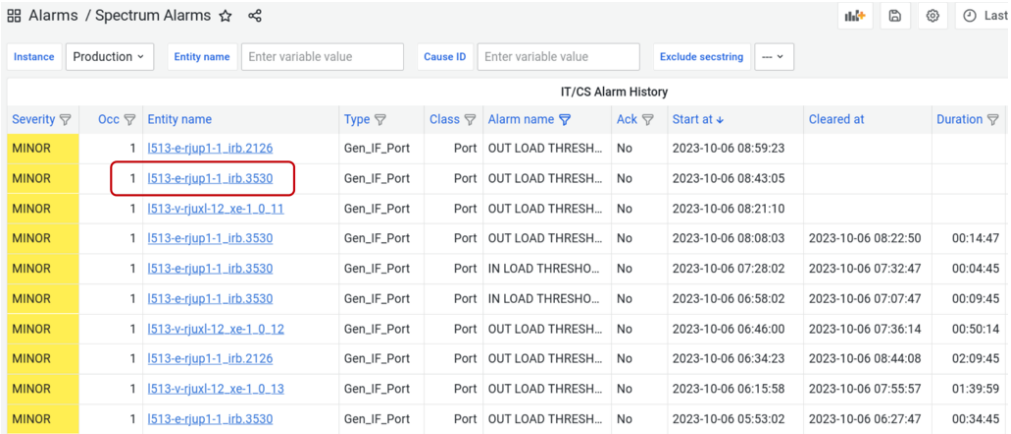

Severity	Occ	Entity name	Type	Class	Alarm name	Ack	Start at	Cleared at	Duration
MINOR	1	513-e-rjup1-1_irb.2126	Gen_IF_Port	Port	OUT LOAD THRESH...	No	2023-10-06 08:59:23		
MINOR	1	513-e-rjup1-1_irb.3530	Gen_IF_Port	Port	OUT LOAD THRESH...	No	2023-10-06 08:43:05		
MINOR	1	513-v-rjuxl-12_xe-1_0_11	Gen_IF_Port	Port	OUT LOAD THRESH...	No	2023-10-06 08:21:10		
MINOR	1	513-e-rjup1-1_irb.3530	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	513-e-rjup1-1_irb.3530	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	513-e-rjup1-1_irb.3530	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	513-v-rjuxl-12_xe-1_0_12	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	513-e-rjup1-1_irb.2126	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	513-v-rjuxl-12_xe-1_0_13	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	513-e-rjup1-1_irb.3530	Gen_IF_Port	Port	OUT LOAD THRESH...	No	2023-10-06 05:53:02	2023-10-06 06:27:47	00:34:45

Border router forwarding table (LHCOPN, LHCONE version)

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
```

Identify WLCG destination site (LHCOPN, LHCONE version)

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

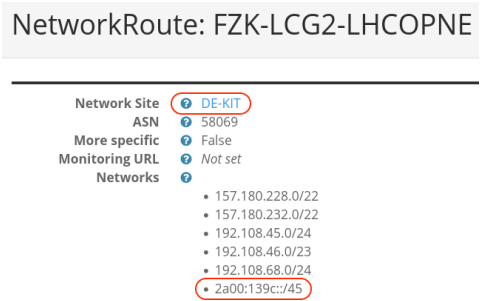

- Look for FTS transfers and make a network decision if it is causing congestion:

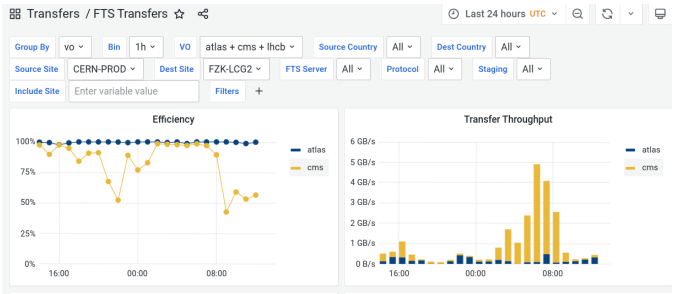

NOTED actions

- Decision-making: NOTED is making the network decision 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 an SDN action to provide more bandwidth.
- Stopped: NOTED has stopped because there are no transfers in FTS and the link-saturation alarm has cleared.

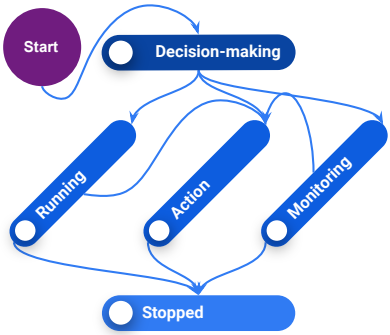

NOTED alarms in MONIT Grafana [Link to the dashboard]

NOTED Alarms

ID	Alarm name	Version	NOTED status	NOTED action	SDN status	Max FTS Throughput [Gb/s]	Interface
184	CH-CERN to CA-TRIUMF	CUSTOM	Running	Spectrum generated an alarm: NOTED is inspecting FTS.	Not provided	0	
187	DE-KIT to CA-TRIUMF	CUSTOM	Action	On-going SDN. FTS throughput [Gb/s]: 5.56	Provided	9.94	
211	CH-CERN to FR-CCIN2P3	CUSTOM	Monitoring	No transfers found in FTS. NOTED is still running until Spectrum clears the alarm.	Not provided		
219	DE-KIT to CA-TRIUMF	CUSTOM	Stopped	The large data transfer is finished.	Released	22.3	
73	ES-ATLAS-T2 to CH-CERN	LHCONE	Decision-making	An action on the link may be required: number of events: 1. Throughput [Gb/s]: 4.12	Not provided		I513-e-rjup1-1_irb.111
83	FR-CCIN2P3 to CH-CERN	LHCONE	Action	On-going SDN. FTS throughput [Gb/s]: 4.94	Provided	7.52	I513-e-rjup1-1_irb.111
84	RO-LCG to CH-CERN	LHCONE	Stopped	The large data transfer is finished.	Released	10.3	I513-e-rjup1-1_irb.111
85	ES-PIC to CH-CERN	LHCONE	Action	On-going SDN. FTS throughput [Gb/s]: 5.94	Provided	12.6	I513-e-rjup1-1_irb.111
107	FR-GRIF to CH-CERN	LHCONE	Monitoring	No transfers found in FTS. NOTED is still running until Spectrum clears the alarm.	Not provided		I513-e-rjup1-1_irb.111
108	IT-INFN-T2 to CH-CERN	LHCONE	Stopped	The large data transfer is finished.	Released	27.9	I513-e-rjup1-1_irb.111
116	UK-SouthGrid to CH-CERN	LHCONE	Running	Spectrum generated an alarm: NOTED is inspecting FTS.	Not provided		I513-e-rjup1-1_irb.111
29	AU-ATLAS to CH-CERN	LHCOPN	Stopped	The large data transfer is finished.	Released	8.79	I513-e-rjup1-1_irb.3530
30	CH-CERN to CA-TRIUMF	LHCOPN	Action	On-going SDN. FTS throughput [Gb/s]: 7.45	Provided	31.5	I513-e-rjup1-1_irb.2126
31	CH-CERN to DE-KIT	LHCOPN	Stopped	The large data transfer is finished.	Released	17.7	I513-e-rjup1-1_irb.3530
32	CH-CERN to DE-KIT	LHCOPN	Monitoring	No transfers found in FTS. NOTED is still running until Spectrum clears the alarm.	Not provided	0	I513-e-rjup1-1_irb.3530
36	NL-T1 to CH-CERN	LHCOPN	Decision-making	An action on the link may be required: number of events: 1. Throughput [Gb/s]: 6.48	Not provided		I513-e-rjup1-1_irb.3530
37	DE-KIT to CH-CERN	LHCOPN	Running	Spectrum generated an alarm: NOTED is inspecting FTS.	Not provided		I513-e-rjup1-1_irb.3530

NOTED demo at DC24

Pre-testing at SC23 (LHCONE, LHCOPN and custom versions)

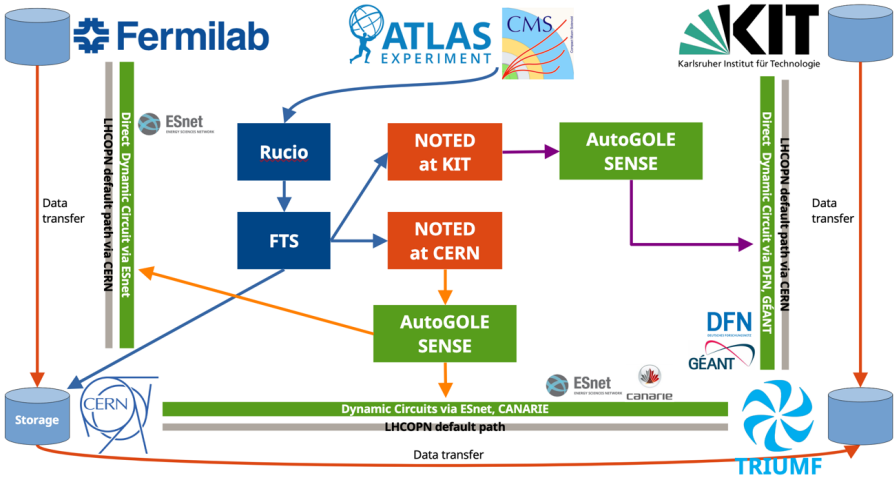

Pre-testing at SC23

Components:

- 3x NOTED controllers and FTS at CERN.
  - 2x custom version for TRIUMF and Fermilab.
  - 1x LHCOPN/LHCONE version.
- 1x NOTED custom controller at KIT.
- Data storage at CERN, TRIUMF, KIT and Fermilab.
- AutoGOLE/SENSE circuits between CERN-TRIUMF, CERN-Fermilab and KIT-TRIUMF.
  - SENSE circuits are provided by ESnet, CANARIE, DFN and GÉANT.

Participants:

Plans for DC24

- Monitoring of LHCONE and LHCOPN links at CERN.
- Dry-run mode: no real actions are taken.
- In case there is the possibility to relief any heavily congested link, NOTED can be used with real SDN actions.

Thanks for your attention!

NOTED: An intelligent network controller to improve the throughput of large data transfers in File Transfer Services by handling dynamic circuits

CERN
IT Department CS Group

Data Challenge 2024 Workshop

Carmen Misa Moreira
Edoardo Martelli


CERN