



NOTED: a framework to optimise network traffic via the analysis of data from FTS

Joanna Waczynska, Edoardo Martelli, Tony Cass (IT-CS); Edward Karavakis (IT-ST, FTS)

vCHEP21 - 18th May 2021

NOTED: Network Optimized Transfer of Experimental Data

NOTED is a software framework that aims to **detect the start of large data transfers and to predict their duration**, and then to trigger appropriate changes to optimise the network that delivers the transfers

The ultimate intention is **to reduce the transfer duration**, improve the effective transfer bandwidth for users, improve the efficient utilisation of available networks

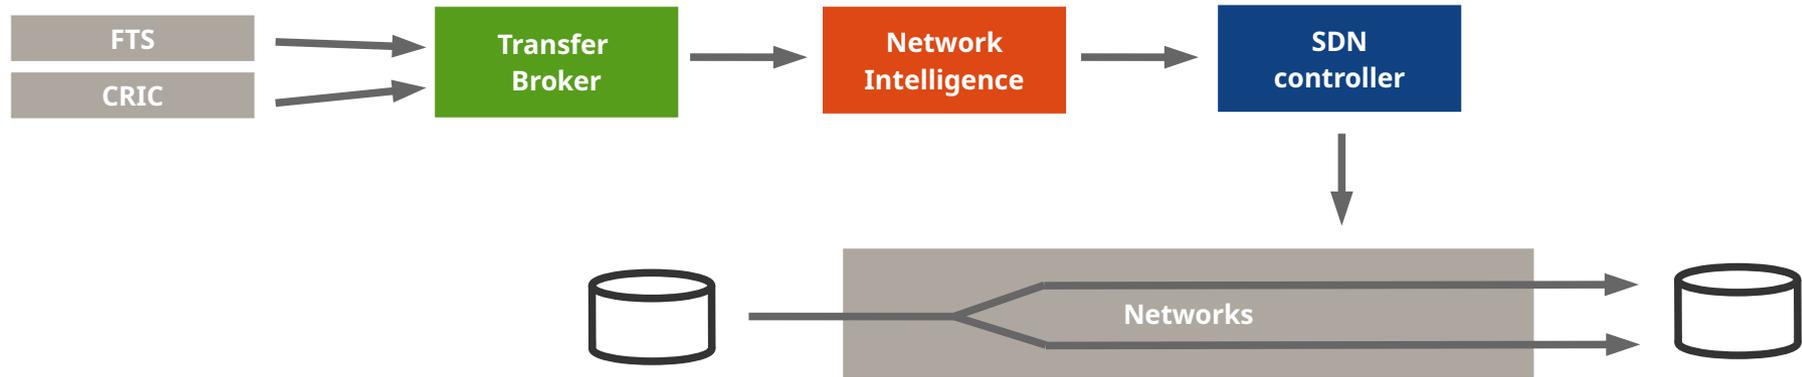
Architecture

There are three main components:

The **Transfer Broker (TB)** extracts transfer information from FTS and enriches it with network addresses from the CRIC database.

The **Network Intelligence (NI)** component interprets information provided by the TB to identify large transfers, and, using its knowledge of the network topology and how transfers will be routed, triggers corrective actions.

The **Software Defined Network Controller (SDNC)** is responsible for re-configuring the networks it controls to deliver better performance and to reduce the duration of file transfers



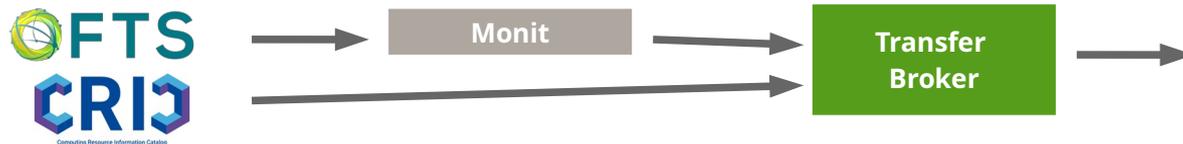
The Transfer Broker

The Transfer Broker (TB) is responsible for collecting and analysing data from FTS.

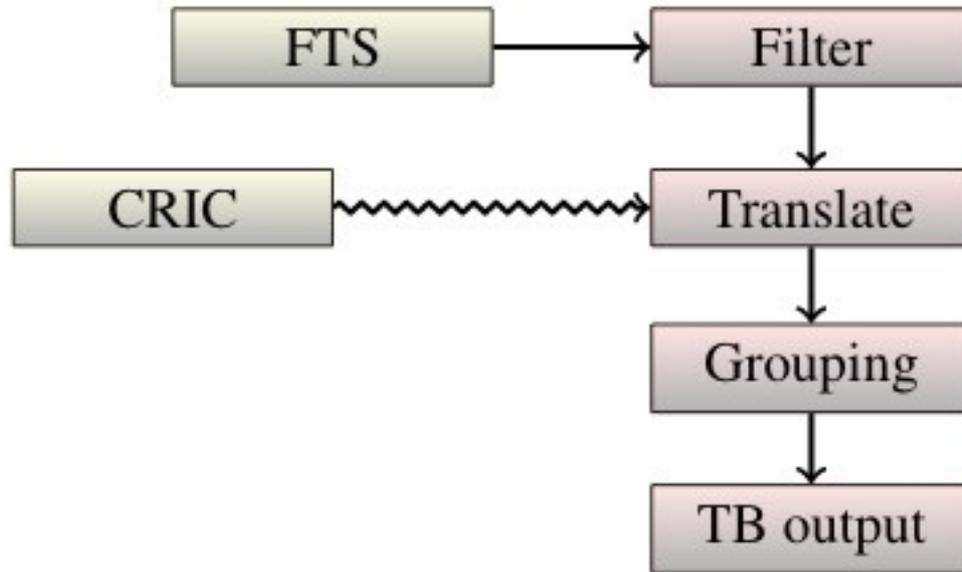
The Transfer Broker extracts FTS transfer information from MONIT and enriches it with network information (IP addresses of storage elements) from the CRIC database.

The TB checks the number of files in the submitted state (in the queue), the throughput, previous transfers already completed (within a certain interval), active and inactive transfers

FTS manages transfers independently one from the other. The TB aggregates the FTS information to identify bulk data transfers between pairs of sites



Transfer Broker stages



Only relevant transfers are considered (e.g. only LHCOPN transfers), to reduce the load

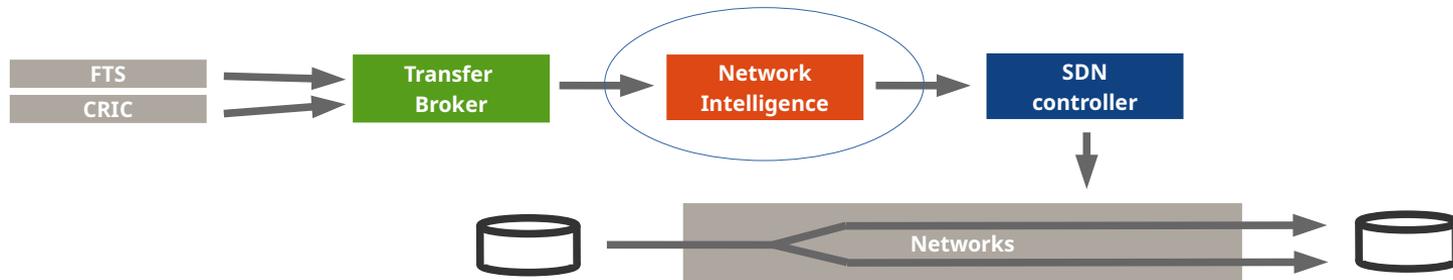
FTS data is enriched with the IP addresses of the storage elements

All transfers with the same src-dst are grouped together, because they will affect the same network path

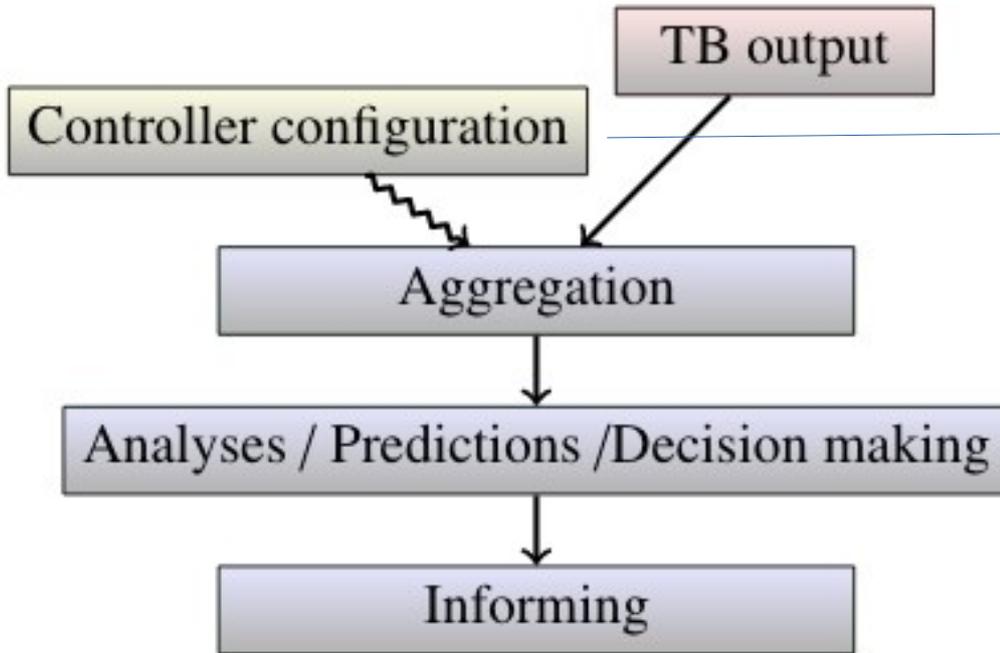
The Network Intelligence component

The Network Intelligence (NI) component is a topology-aware system that **interprets information from the Transfer Broker and can signal the start and stop of large aggregated transfers** affecting a given link/path.

An NI is configured to analyse all the transfers affecting a given network link/path that can be improve. To optimize different network links, multiple NI instances would be needed



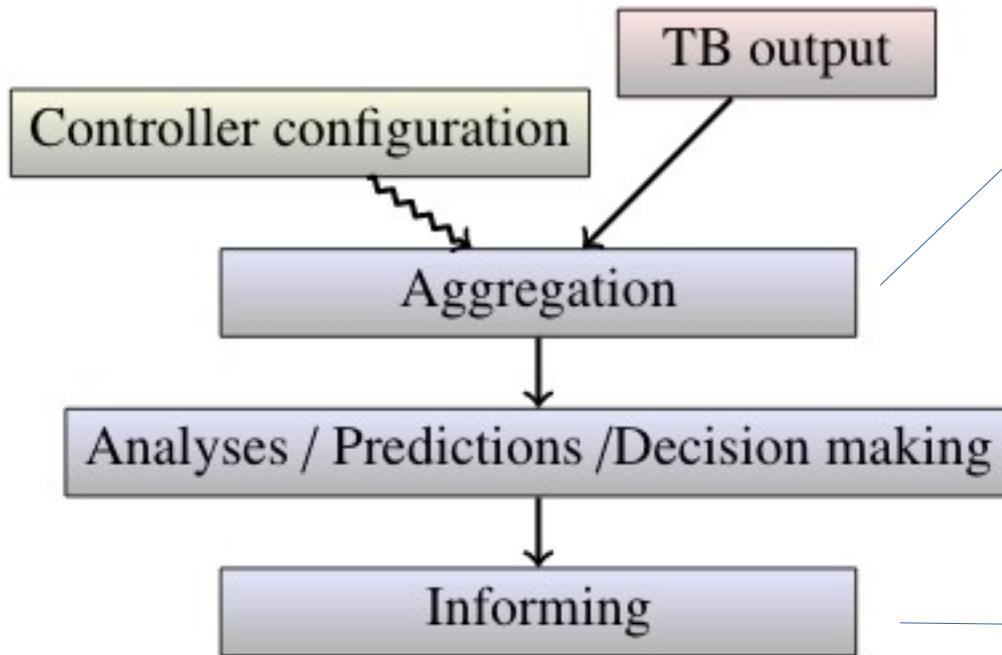
Components of the NI - input



the Transfer Broker output with all information about transfers and sites

the controller configuration which sets up the control parameters used for the analysis, decision making and predictions for the different network segments of interest

Components of the NI - Processes



Aggregation of all transfers affecting a common link or path

process that analyses all the transfers, detects when the transfer volume is above/below a certain threshold, estimates the Time-to-complete, decides when a transfer has started and stopped

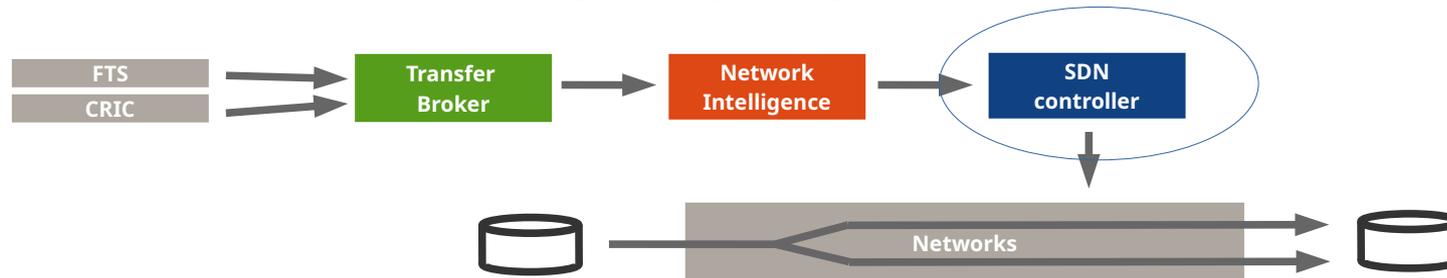
A start-of-transfer and end-of-transfer signal is sent out

Software Defined Network Controller

The Software Defined Network Controller (SDNC) is the component that **actually modifies the configuration of network devices, implementing and removing special routing policies** at the times suggested by the NI to improve the performance of data transfers.

For example, a data transfer could be load-balanced over the primary link and a second link which may have a different purpose, but it is under utilised at that moment. Alternatively, the NI could request a temporary network circuit of a higher capacity to be provisioned between the source and the destination of the transfer.

There can be multiple SDNC, in case of links spanning multiple network domains



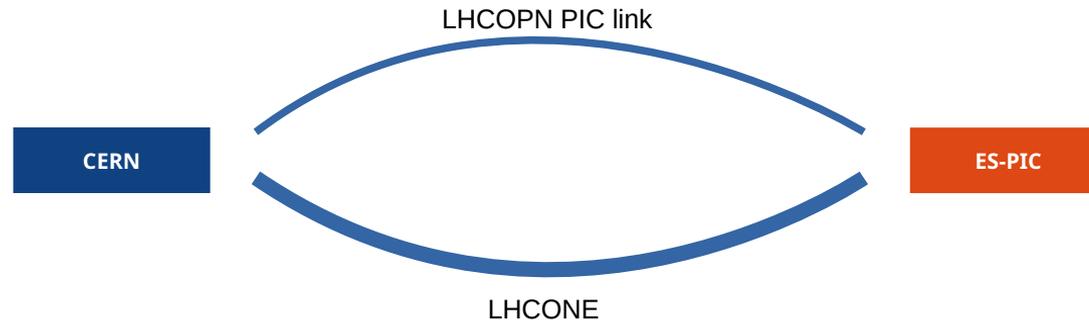
NOTED applications

NOTED has been tested to identify real transfers and apply real network optimizations.

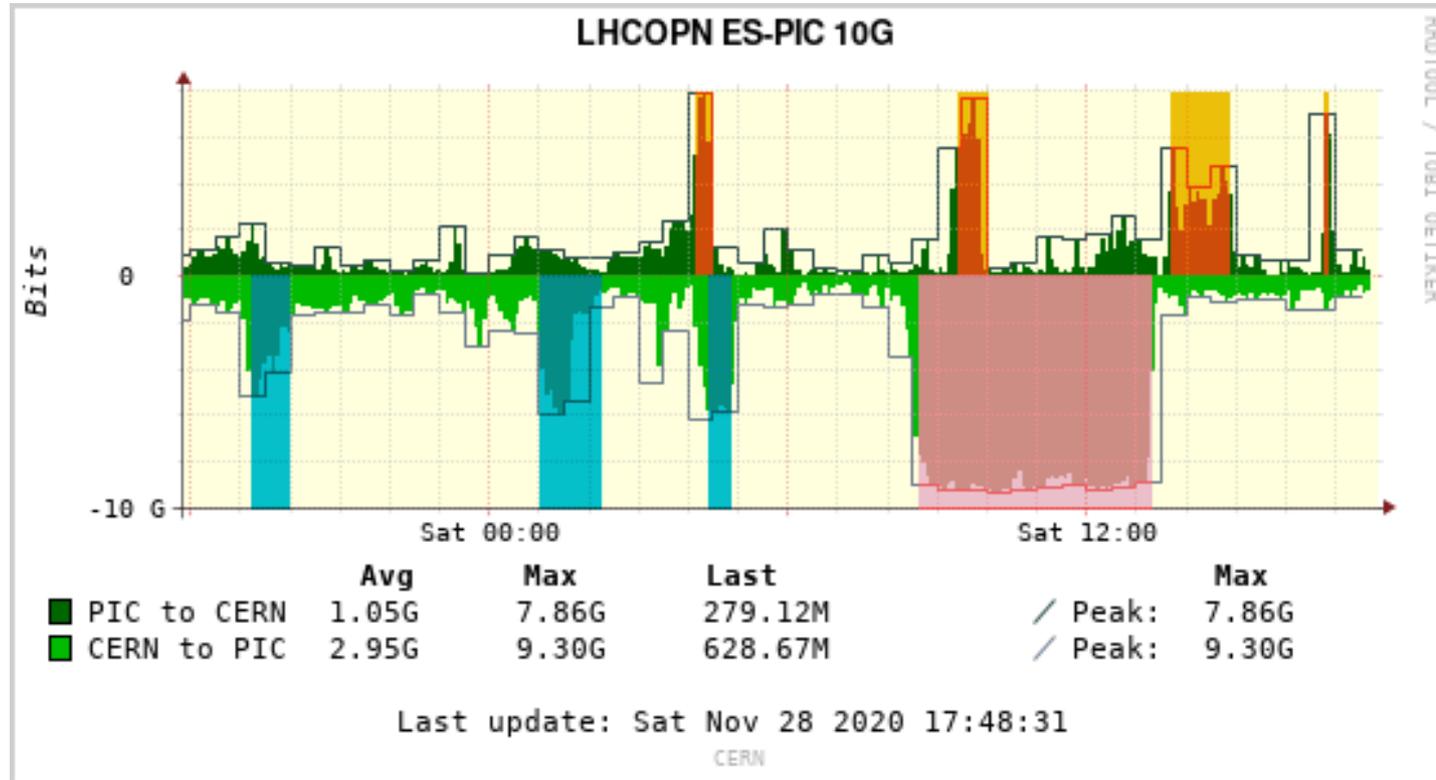
To make sure the tests happened during time agreed with WLCG operations, test transfers of real data were setup by ATLAS using Rucio.

Application: CERN-PIC LHCOPN+LHCONE

CERN-PIC: ES-PIC has a 10G LHCOPN link to CERN. During this test the SDNC could activate BGP load balancing over the 100G LHCONE connection (link shared by all the Tier1s)



Application: CERN-PIC



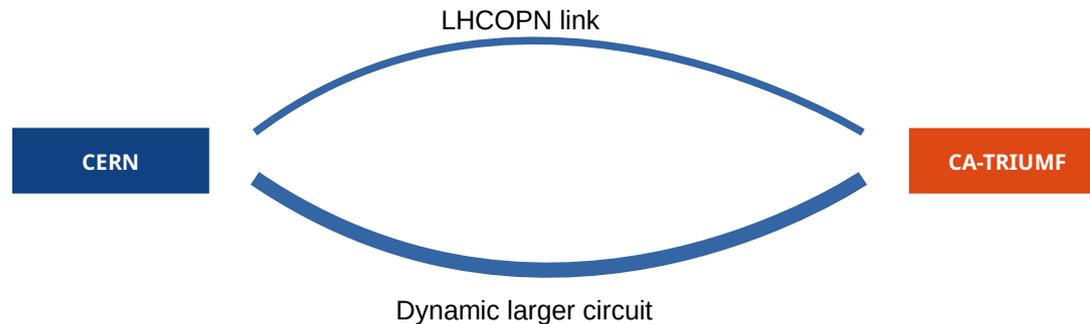
Transfers detected by NOTED

Application: CERN-TRIUMF dynamic circuit

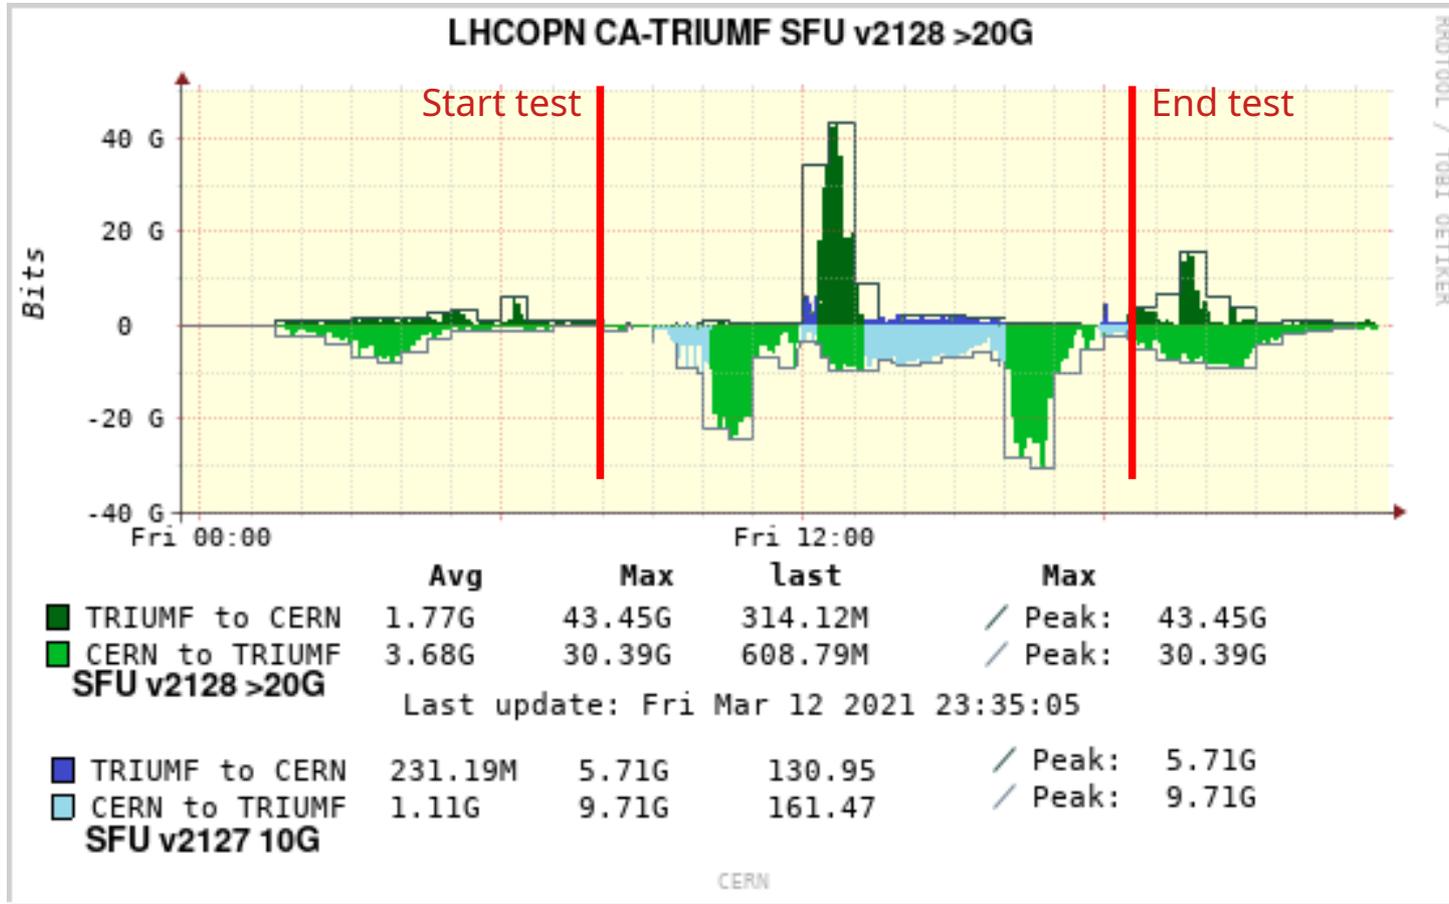
CERN-TRIUMF: CA-TRIUMF has a primary >20G LHCOPN link to CERN and a backup 10G link.

During the test the 10G link was set as primary link and NOTED could request to move the traffic to the >20G link when a large transfer was detected.

This test wanted to simulate the possibility to use a dynamic circuit provided by a on-demand service



Application: CERN-TRIUMF



Conclusions

The NOTED project has demonstrated that analysing data from transfer services such as FTS can **allow the understanding and the prediction of network traffic**

NOTED has also demonstrated that it can use these predictions to **reconfigure the network and deliver faster data transfers** and more efficient network utilization

*More information on Joanna's plenary talk tomorrow at 15:30:
Convolutional LSTM models to estimate network traffic*

Thank you for your attention

Questions?

