

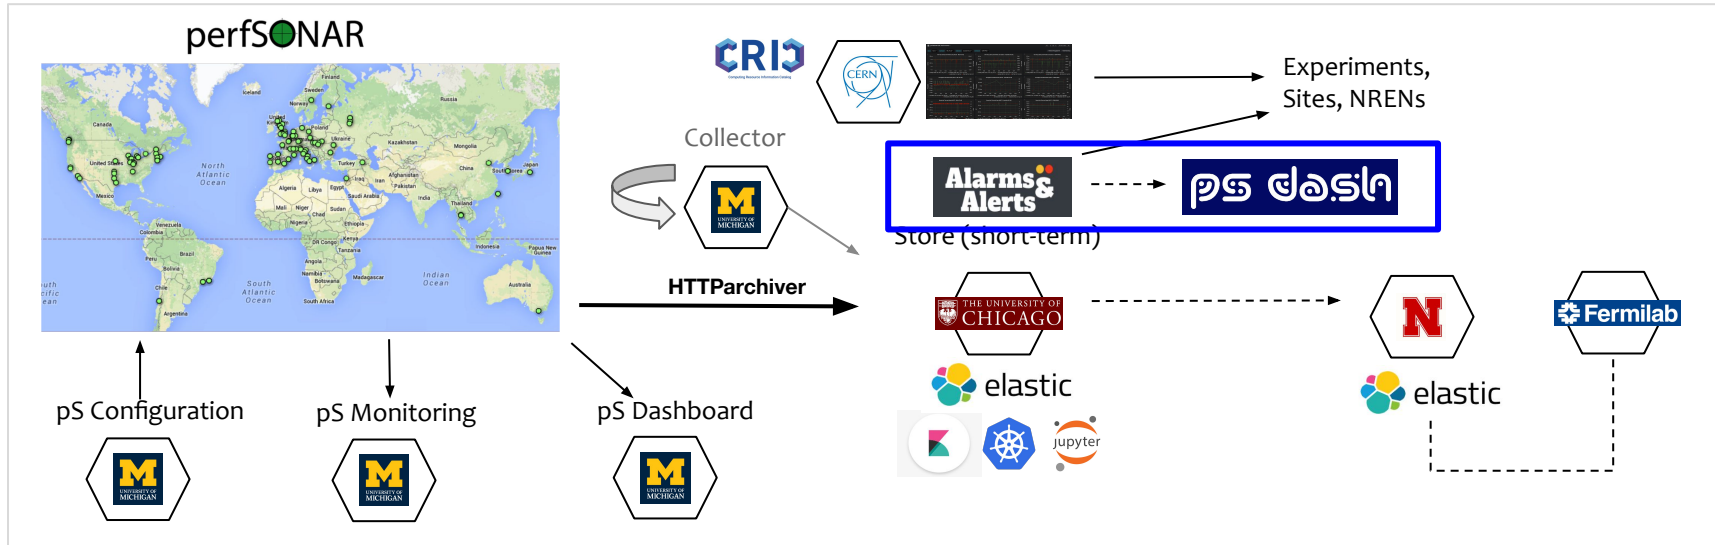
perfSONAR Network Analytics

#52 LHCOPN-LHCONE Meeting
IHEP Beijing CN
9 - 12 Oct 2024

Petya Vasileva, U Michigan
Marian Babik, CERN
Shawn McKee, U Michigan
Ilija Vukotic, U Chicago

Status and Plans

pS platform



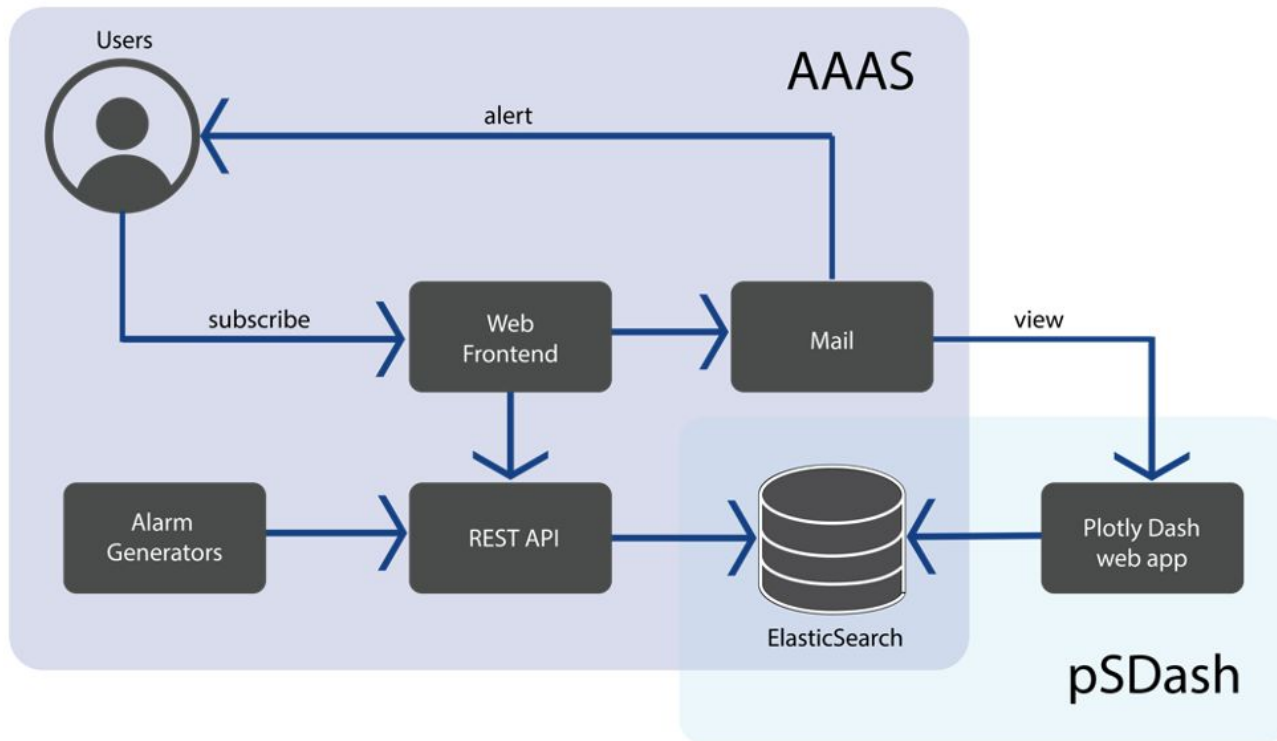
Analytics Tools

Covered in this presentation

● Alarms and Alerts

● pSDash

● ML implementation (in progress)



Alarming & Alerting Service: <https://psa.osg-htc.org/>

The perfSONAR Dashboard application: <https://ps-dash.uc.ssl-hep.org/>

How to subscribe for email alerts

Alarms

- Analytics
- Networking
 - Perfsonar
 - bad owd measurements
 - large clock correction
 - firewall issue
 - complete p
 - unresolvable
 - unresolvable
- Infrastructure
 - indexing
- Sites
 - destination cannot be reached from multiple
 - destination cannot be reached from any
 - bandwidth increased from/to multiple sites
 - bandwidth decreased from/to multiple sites
 - high packet loss
 - source cannot reach any
 - bandwidth increased
 - bandwidth decreased
- RENs
 - path changed
- Virtual Placement

Hover over an alarm to see details

Code running every 24h at UC k8s cluster, calculates average packet loss for all the src-dest pairs. Alarm is generated if the calculated value is equal to 100% and the site has complete packet loss to more than 10 other hosts. The code can be found here: <https://github.com/sand-ci/AlarmsAndAlerts/blob/main/ps-packetloss.py>

Current Subscriptions

Category	Subcategory	Event	Tags
Networking	Sites	high packet loss on multiple links	*
Networking	Perfsonar	bad owd measurements	*
Networking	Perfsonar	firewall issue	*
Networking	Perfsonar	complete packet loss	*
Networking	Perfsonar	unresolvable host	*
Networking	Perfsonar	unresolvable host	*
Networking	Sites	bandwidth decreased from/to multiple sites	MWT2
Networking	Perfsonar	large clock correction	*

Showing 1 to 8 of 8 entries

Update Subscription

Track a specific site or a host

Sat, 05 Oct 2024 03:02:07 Networking/Perfsonar/firewall issue Firewall issue

tags: GSI-LCG2-LHCONE

Site: GSI-LCG2-LHCONE, perfsonar node: dclxdlperfsonar2.gsi.de seems to have a firewall issue. Packet loss is 100% from at least 10 hosts. Affected sites: ['CERN-PROD-LHCOPNE', 'NLT1-SARA-LHCOPNE', 'PRAGUELCG2-LHCONE', 'PIC-LHCOPNE', 'IN2P3-CC-LHCOPNE']. More information could be found in pSDash:

<https://ps-dash.uc.ssl-hep.org/loss-delay/bdf4ce3c56b604ace5944490946673b1e218b376fb15b4230afaf9aa>

What you'll receive

In the past 12 hours, path between 12 pairs diverged and went through ASN 174 owned by COGENT-174, US. The change affected the following sites ['SPRACE-REDNESP', 'CA-WATERLOO-T2-LHCONE', 'UKI-LT2-QMUL', 'RRC-KI-T1-LHCOPNE', 'CERN-PROD-LHCOPNE', 'DESY-ZN-LHCONE', 'INFN-T1-LHCOPNE', 'MWT2_IU', 'UKI-SOUTHGRID-RALPP-LHCONE', 'JP-KEK-CRC-02', 'CBPF-LHCONE', 'PIC-LHCOPNE', 'IN2P3-CC-LHCOPNE']. The code can be found on pS-Dash: <https://ps-dash.uc.ssl-hep.org/paths/c648363dca4131c09b9daee028215540ac6bc3bc8329c175961b74aa>

Tue, 02 Jul 2024 04:08:11 Networking/Sites/bandwidth decreased bandwidth decreased

tags: JINR-LCG2-LHCONE, UKI-SOUTHGRID-RALPP-LHCONE

Bandwidth decreased for the ipv6 links between sites JINR-LCG2-LHCONE and UKI-SOUTHGRID-RALPP-LHCONE. Current throughput is 17 MB, dropped by -98% with respect to the 21-day-average. More information could be found in pS-Dash:

<https://ps-dash.uc.ssl-hep.org/throughput/7cedc9ccf2c135d32acae685b34d714837fc6e3579b9143a452b3b4f>

Understanding the paths (via pSDash)

BASELINE PATH

Taken in 98% of time

Always reaches destination: NO



ALTERNATIVE PATHS

Taken in 1.0% of time

Always reaches destination: YES



Taken in 0.0% of time

Always reaches destination: NO



Taken in 0.0% of time

Always reaches destination: YES



Taken in 0.0% of time

Always reaches destination: NO




TW-ASGC → AGLT2_UM



The plot shows the AS numbers for every hop and the frequency of their occurrences at each position (source and destination not included). The dark blue values of 1 mean the ASN was always used at that position; Close to 0, means the ASN rarely appeared; OFF indicates the device did not respond at the time of the traceroute test; 0 is when there was a response, but the ASN was unknown.

Site Status Dashboard

Summary



Highest number of alarms from site
OSG-WLCCG (DE): 26

Highest number of alarms from country
United Kingdom: 83

SITE	STATUS	NETWORK	INFRASTRUCTURE	OTHER	URL
filter data...					
INFN-CATANIA	●	0	0	0	See latest alarms
RRC-KI	●	0	0	0	See latest alarms
RO-NIPNE-LHCONE	●	0	0	0	See latest alarms
RO-16-UAIC	●	0	0	0	See latest alarms
RO-14-ITIM	●	0	0	0	See latest alarms
RO-11-NIPNE	●	0	0	0	See latest alarms
RO-07-NIPNE	●	0	0	0	See latest alarms
RO-03-UPB	●	0	0	0	See latest alarms
RO-02-NIPNE	●	0	0	0	See latest alarms
RAL-LCG2-ECHO	●	0	0	0	See latest alarms
RAL-LCG2	●	0	0	0	See latest alarms
PURDUE-LHCONE	●	0	0	0	See latest alarms
PURDUE	●	0	0	0	See latest alarms
PRAGUELCG2	●	0	0	0	See latest alarms
PIC	●	0	0	0	See latest alarms
OU_OCHEP_SWT2	●	0	0	0	See latest alarms

« < 1 / 9 > »

How was the status determined?



Types of alarms

NETWORK

- destination cannot be reached
- source cannot reach any destination
- bad owd measurements
- large clock correction
- complete packet loss
- firewall issue

INFRASTRUCTURE

- bandwidth decreased on multiple sites
- path changed between sites
- unresolvable hosts
- no tests for a host

Other


- bandwidth increased
- bandwidth decreased (on a single pair)
- high packet loss

The goal is to
proactively discover network issues



How to combine pS tests?



Tests' rate of execution varies by type



■ traceroute test every 10 min


● throughput test every 6-24 hours

Correlate network tests



Trends on routers

2001:630:0:9011::189




Each **point** represents the throughput values collected when the node was on the path

Challenges


- 50% of the paths are incomplete
- some addresses are private

Path from JINR-T1-LHCOPNE to BEIJING-LCG2



To build **reliable topology** models for identifying weak points on the network, we need to **reconstruct the paths**





What is the most probable C, given it's between A and B?

Possible intermediates:

r237 $P(A \rightarrow C) = 0.008$ and $P(C \rightarrow B) = 0.038$

r265 $P(A \rightarrow C) = 0.056$ and $P(C \rightarrow B) = 0.009$

r536 $P(A \rightarrow C) = 0.176$ and $P(C \rightarrow B) = 0.551$

r792 $P(A \rightarrow C) = 0.072$ and $P(C \rightarrow B) = 0.008$

r838 $P(A \rightarrow C) = 0.008$ and $P(C \rightarrow B) = 0.01$

The most probable intermediate router between r792 and r237 is r536 with a probability of 0.097

There are multiple possibilities for C. What is the correct node that lies between A and B depends more on the surrounding nodes rather than on highest probability value

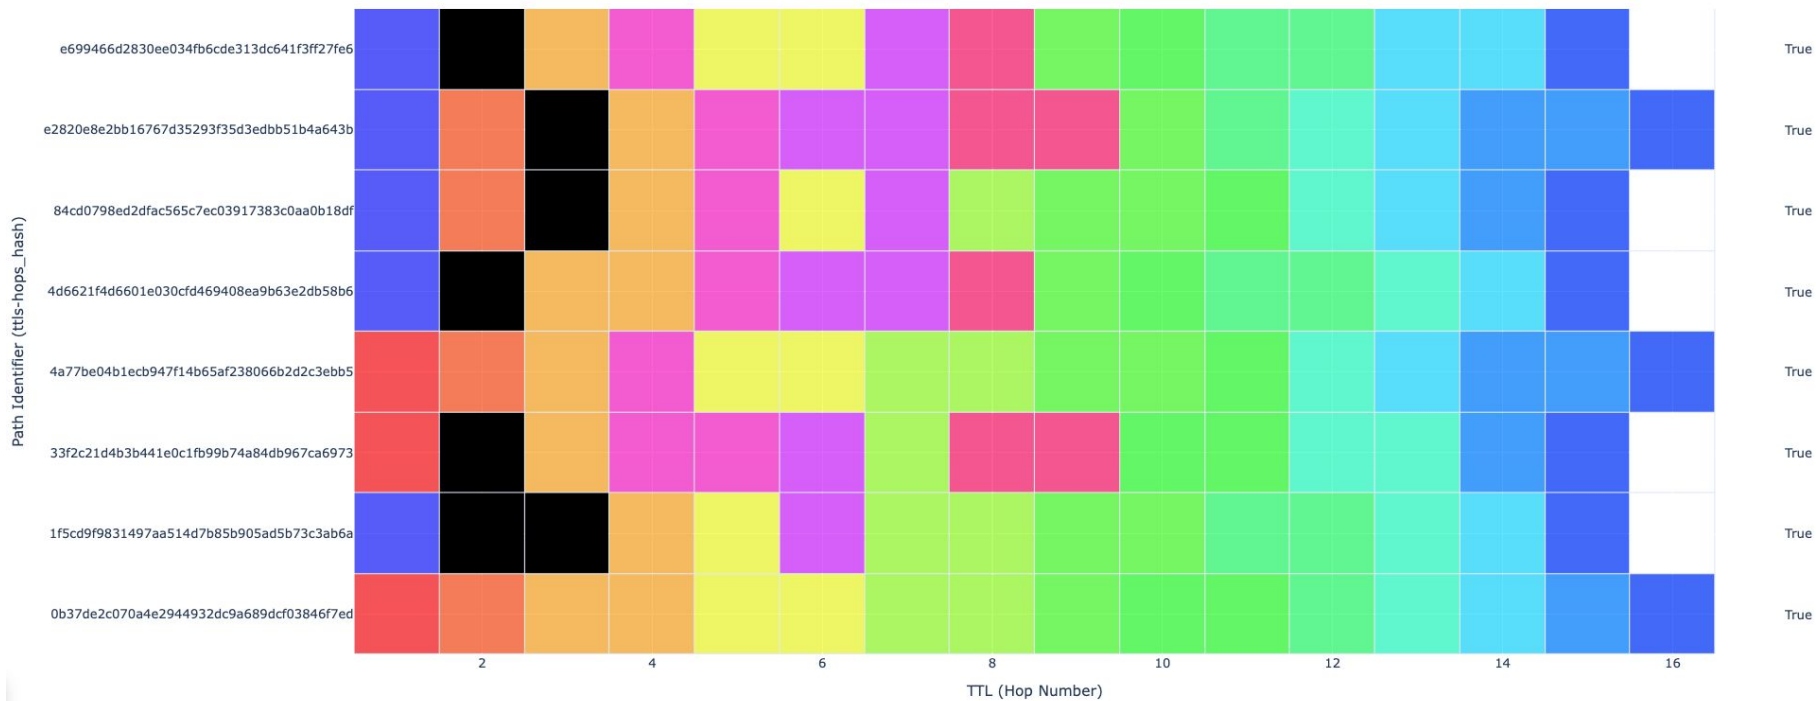
Site to site path signature

Unknown IP

Each color is a different IP

Path signature between UKI-NORTHGRID-MAN-HEP and JINR-LCG2


Destination Reached



Path reconstruction via ML


- Transformer model
- Pays attention to the nodes preceding the one in question
- Currently excludes unknown nodes as targets

The model performs well and converges quickly



Path signature between UKI-LT2-IC-HEP and T2_US_WISCONSIN

Destination Reached




Unknown IP

Each color is a different IP

Predicted next node: r1015

Corner cases



Unknown IP

Each color is a different IP

Contextual node or retrying until it gets to the final node?

Next steps

Include the **corner cases** by:

- Creating “confidence score” to be used into the loss function
- Collapsing the multiple consecutive unknown nodes into a single one

Implement another sequence model to compare the results



Once the **topology** is **fixed**, we can proceed by building more **complex models** that incorporate other metrics such as **loss, bandwidth** or **file transfer** statistics

Acknowledgements

We would like to thank the **WLCG**, **HEPiX**, **perfSONAR** and **OSG** organizations for their work on the topics presented.

In addition we want to explicitly acknowledge the support of the **National Science Foundation** which supported this work via:

- OSG: NSF MPS-1148698
- IRIS-HEP: NSF OAC-1836650

**Thank
you!**

**Any
questions?**

Contact us @ net-discuss@umich.edu

Backup slides

Get your site name

from pSDash: <https://ps-dash.uc.ssl-hep.org/search-alarms>

perfSONAR Toolkit Information

Kibana: Packet Loss in OSG/WLCG

Kibana: Packet Loss Tracking

MEPHi Tracer: Traceroute explorer


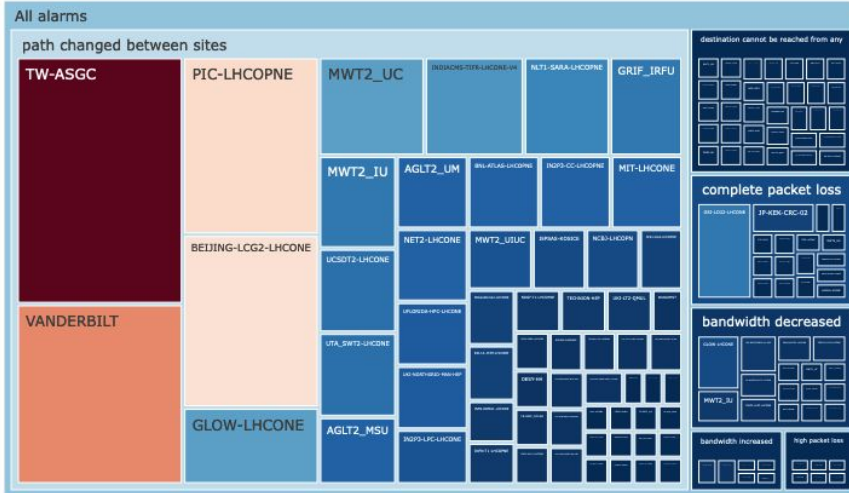


SITES OVERVIEW

SEARCH ALARMS

EXPLORE PATHS

MAJOR ALARMS



Search & Explore the Networking Alarms

10/08/2024 → 10/09/2024

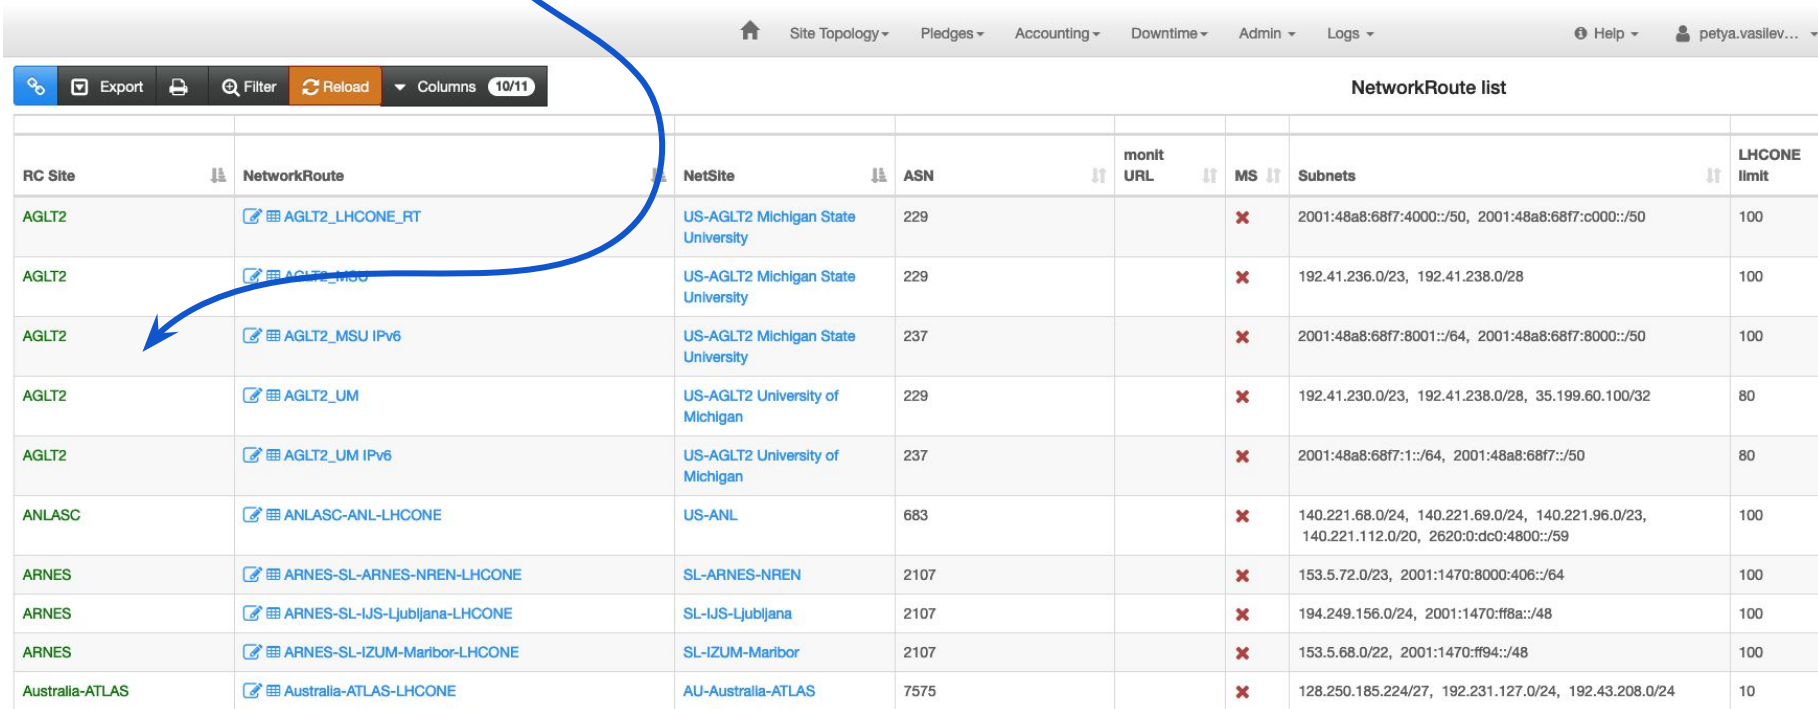
Rounded to the day

- Search for a site
- AGLT2_MSU
 - AGLT2_UM
 - BEGRID-ULB-VUB
 - BEIJING-LCG2-LHCONE
 - BNL-ATLAS-LHCOPNE
 - BUDAPEST

List of alarms

Get your site name

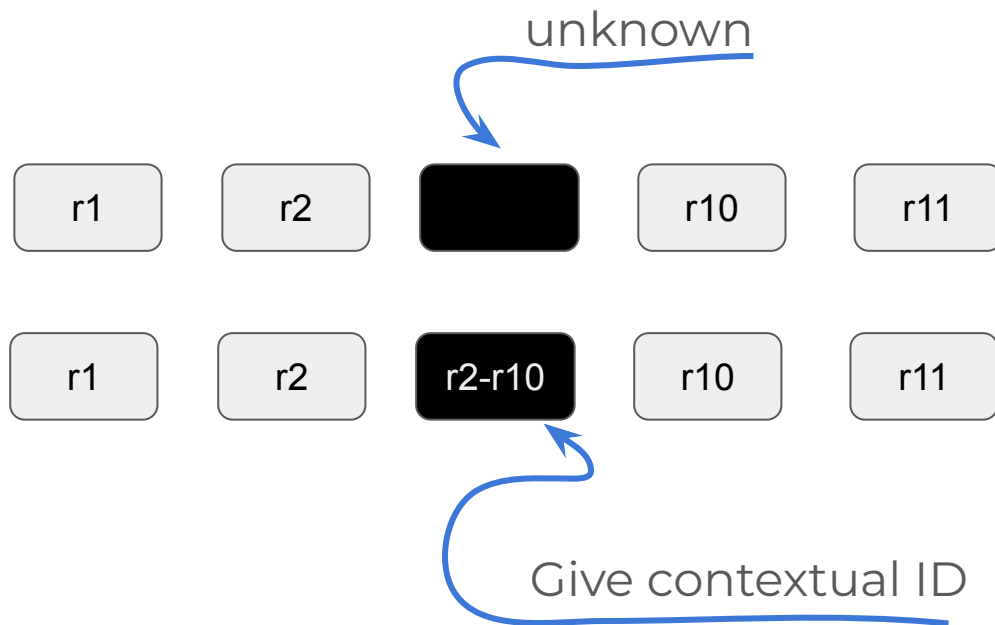
from CRIC: <https://wlcg-cric.cern.ch/core/networkroute/list/>



The screenshot shows a web interface for managing network routes. At the top, there is a navigation bar with links for Site Topology, Pledges, Accounting, Downtime, Admin, and Logs. A user profile for 'petya.vasilev...' is visible on the right. Below the navigation bar is a toolbar with buttons for Export, Filter, Reload, and Columns (10/11). The main content area is titled 'NetworkRoute list' and contains a table with the following columns: RC Site, NetworkRoute, NetSite, ASN, monit URL, MS, Subnets, and LHCONE limit. A blue arrow points from the URL in the text above to the 'AGLT2_MSU' row in the table.

RC Site	NetworkRoute	NetSite	ASN	monit URL	MS	Subnets	LHCONE limit
AGLT2	AGLT2_LHCONE_RT	US-AGLT2 Michigan State University	229		✗	2001:48a8:68f7:4000::/50, 2001:48a8:68f7:c000::/50	100
AGLT2	AGLT2_MSU	US-AGLT2 Michigan State University	229		✗	192.41.236.0/23, 192.41.238.0/28	100
AGLT2	AGLT2_MSU IPv6	US-AGLT2 Michigan State University	237		✗	2001:48a8:68f7:8001::/64, 2001:48a8:68f7:8000::/50	100
AGLT2	AGLT2_UM	US-AGLT2 University of Michigan	229		✗	192.41.230.0/23, 192.41.238.0/28, 35.199.60.100/32	80
AGLT2	AGLT2_UM IPv6	US-AGLT2 University of Michigan	237		✗	2001:48a8:68f7:1::/64, 2001:48a8:68f7::/50	80
ANLASC	ANLASC-ANL-LHCONE	US-ANL	683		✗	140.221.68.0/24, 140.221.69.0/24, 140.221.96.0/23, 140.221.112.0/20, 2620:0:dc0:4800::/59	100
ARNES	ARNES-SL-ARNES-NREN-LHCONE	SL-ARNES-NREN	2107		✗	153.5.72.0/23, 2001:1470:8000:406::/64	100
ARNES	ARNES-SL-IJS-Ljubljana-LHCONE	SL-IJS-Ljubljana	2107		✗	194.249.156.0/24, 2001:1470:ff8a::/48	100
ARNES	ARNES-SL-IZUM-Maribor-LHCONE	SL-IZUM-Maribor	2107		✗	153.5.68.0/22, 2001:1470:ff94::/48	100
Australia-ATLAS	Australia-ATLAS-LHCONE	AU-Australia-ATLAS	7575		✗	128.250.185.224/27, 192.231.127.0/24, 192.43.208.0/24	10

Transformer model: Preprocessing example



The dataset

path	r1	r2	r2-r10	r10	r11		
Input	r1					Target	r2
	r1	r2					r2-r10
	r1	r2	r2-r10				r10
	r1	r2	r2-r10	r10			r11

Skip unknown nodes as targets