Route Server Filtering

Barry O'Donovan

INEX Members' Meeting

March 31st, 2022, Dublin

An IXP Without Route Servers

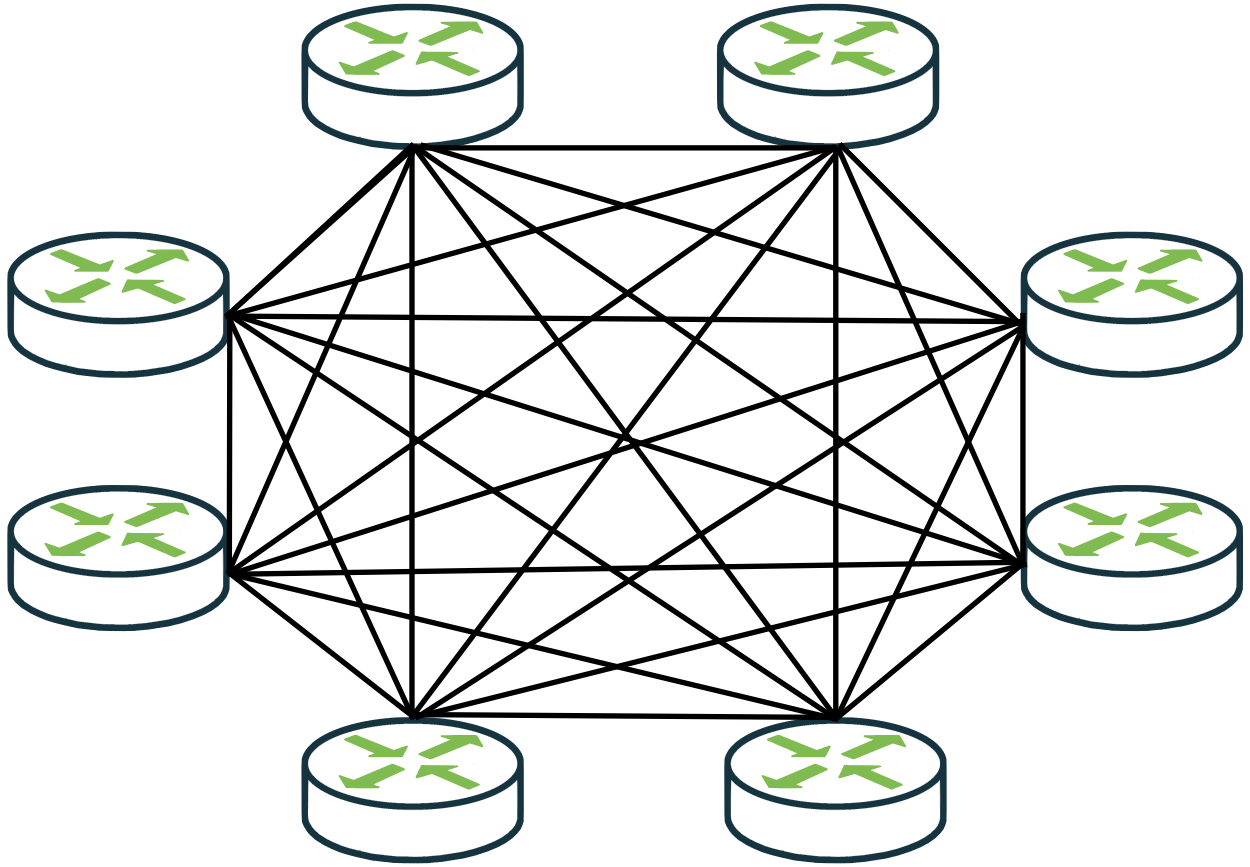

- $\frac{n(n-1)}{2}$ bilateral sessions
- 8 members - 28 sessions
- 100 members - 4,950 sessions

An IXP Without Route Servers

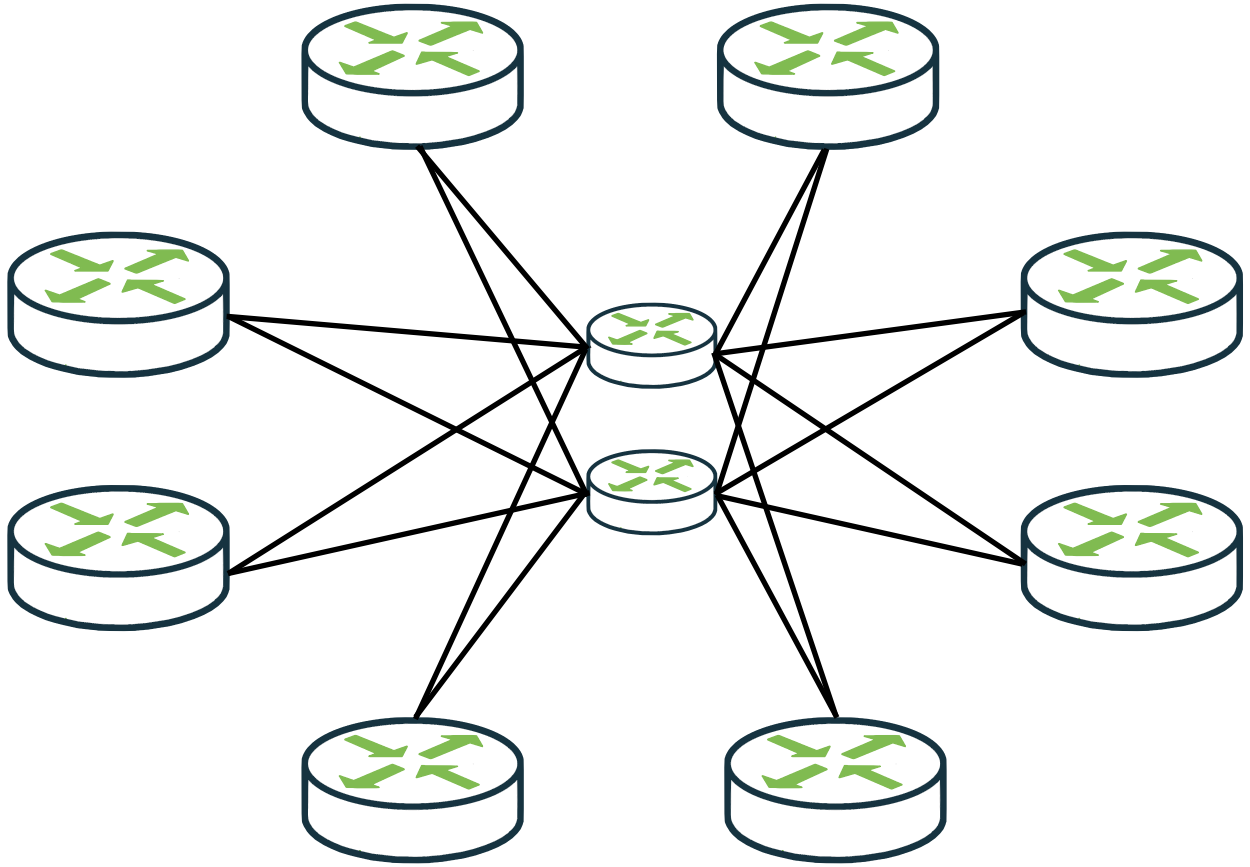

- RFC7947 - *multilateral interconnection using a third-party brokering system.*
- Unlike route reflectors, route servers operate with EBGP.
- Attribute and AS path transparency - ‘acts as if it’s doesn’t exist’.
- Critical infrastructure at an IXP.

Advantages of Route Servers

Immediate benefit of / for new members

Security and assurance for routes exchanged

Avoids email tennis, typos, mis-understandings etc.

No production changes, simpler configurations

At INEX: LAN1 - 88%. LAN2 - 91%. Cork - 100%.


INEX
The Need for Filtering

The Need for Route Server Filtering

- You are essentially “outsourcing” your routing policy
  - Usually fine - most networks peer openly at an IXP
- “I want to peer with everyone except X”
  - Route servers need some knobs to allow this
- Mostly standardised BGP community schema provides this
  - https://github.com/euro-ix/rs-workshop-july-2017/wiki/Route-Server-BGP-Community-usage
- Certainly standardised at nearly 200 exchanges using IXP Manager

(Large) Community Based Filtering

Action	Community
Prevent announcement of a prefix to a peer	`43760:0:peer-as`
Announce a route to a certain peer	`43760:1:peer-as`
Prevent announcement of a prefix to all peers	`43760:0:0`
Announce a prefix to all peers *(default)*	`43760:1:0`

NAMESPACE : ACTION : TARGET

Examples

- “Everyone except AS65501 and 65505”

```
43760:1:0 43760:0:65501 43760:0:65505
```

- “No one except AS65501 and 65505”

```
43760:0:0 43760:1:65501 43760:1:65505
```

NAMESPACE : ACTION : TARGET

Extra Filtering Options Large Communities

Action	Community
Prepend to peer AS once	`43760:101:peer-as`
Prepend to peer AS twice	`43760:102:peer-as`
Prepend to peer AS thrice	`43760:103:peer-as`

NAMESPACE : ACTION : TARGET

Stub Example - JunOS

```
edit policy-options
set community cl-inex-noadvertise-65501
    members large:43760:0:65501
set community cl-inex-noadvertise-65505
    members large:43760:0:65505

set policy-statement ps-export-v4-ix
   term what-i-want-to-send
        from …
        then community add cl-inex-noadvertise-65501
        then community add cl-inex-noadvertise-65505
```

Community Based Filtering in Practice

- Difficult at both ends of the network-size scale:
  - Small networks rarely touch their border routers
  - Large networks need cumbersome change control procedures

- Very complicated in a pinch
  - Community filtering is only half the story!
    - Still need to filter the routes you learn from the route servers

- DDoS events of Q2 2021
  - INEX Operations implemented route server filtering on an emergency basis for a number of members.

UI Based Filtering in IXP Manager

UI Based Filtering in IXP Manager

- Purpose: move the complexity from member router to route server (RS)

- Mechanism is unchanged - just where it happens moves:
  - RS tags your routes in ingress rather than you doing it on egress
  - RS filters routes to be advertised to you on egress rather than you on ingress

1. Intended for relatively simple routing policies
2. Consider how you order your rules
3. You are responsible for your own routing policy - if in doubt, just contact INEX Operations

Route Server Filtering for AS112

You have no filters in production.

No route server filters have been defined. Create Route Filter

Route Server Filtering

IXP Manager supports the industry standards for community based route server filtering. You can find the official documentation here. Using the BGP-community mechanism can be difficult to implement where a network engineer is not familiar with BGP communities or where a network may have arduous change control processes for altering a router's configuration.

This purpose of this tool is to allow IXP participants to implement the exact same mechanism but rather than tagging your routes on egress from your router / manipulating routes on ingress to your router, the IXP's route servers perform the equivalent tagging / route manipulation as they accept your routes or send you routes from other networks.

Please note the following important points:

1. This tool is intended to help you make relatively simple routing policies.
2. When processing routes, please consider the ordering of your rules and ensure to put more specific rules first.
3. You are responsible for your own routing policy and ensuring any rules you set here have the desired effect. If in doubt, feel free to contact our operations team.

Create Route Server Filter for AS112

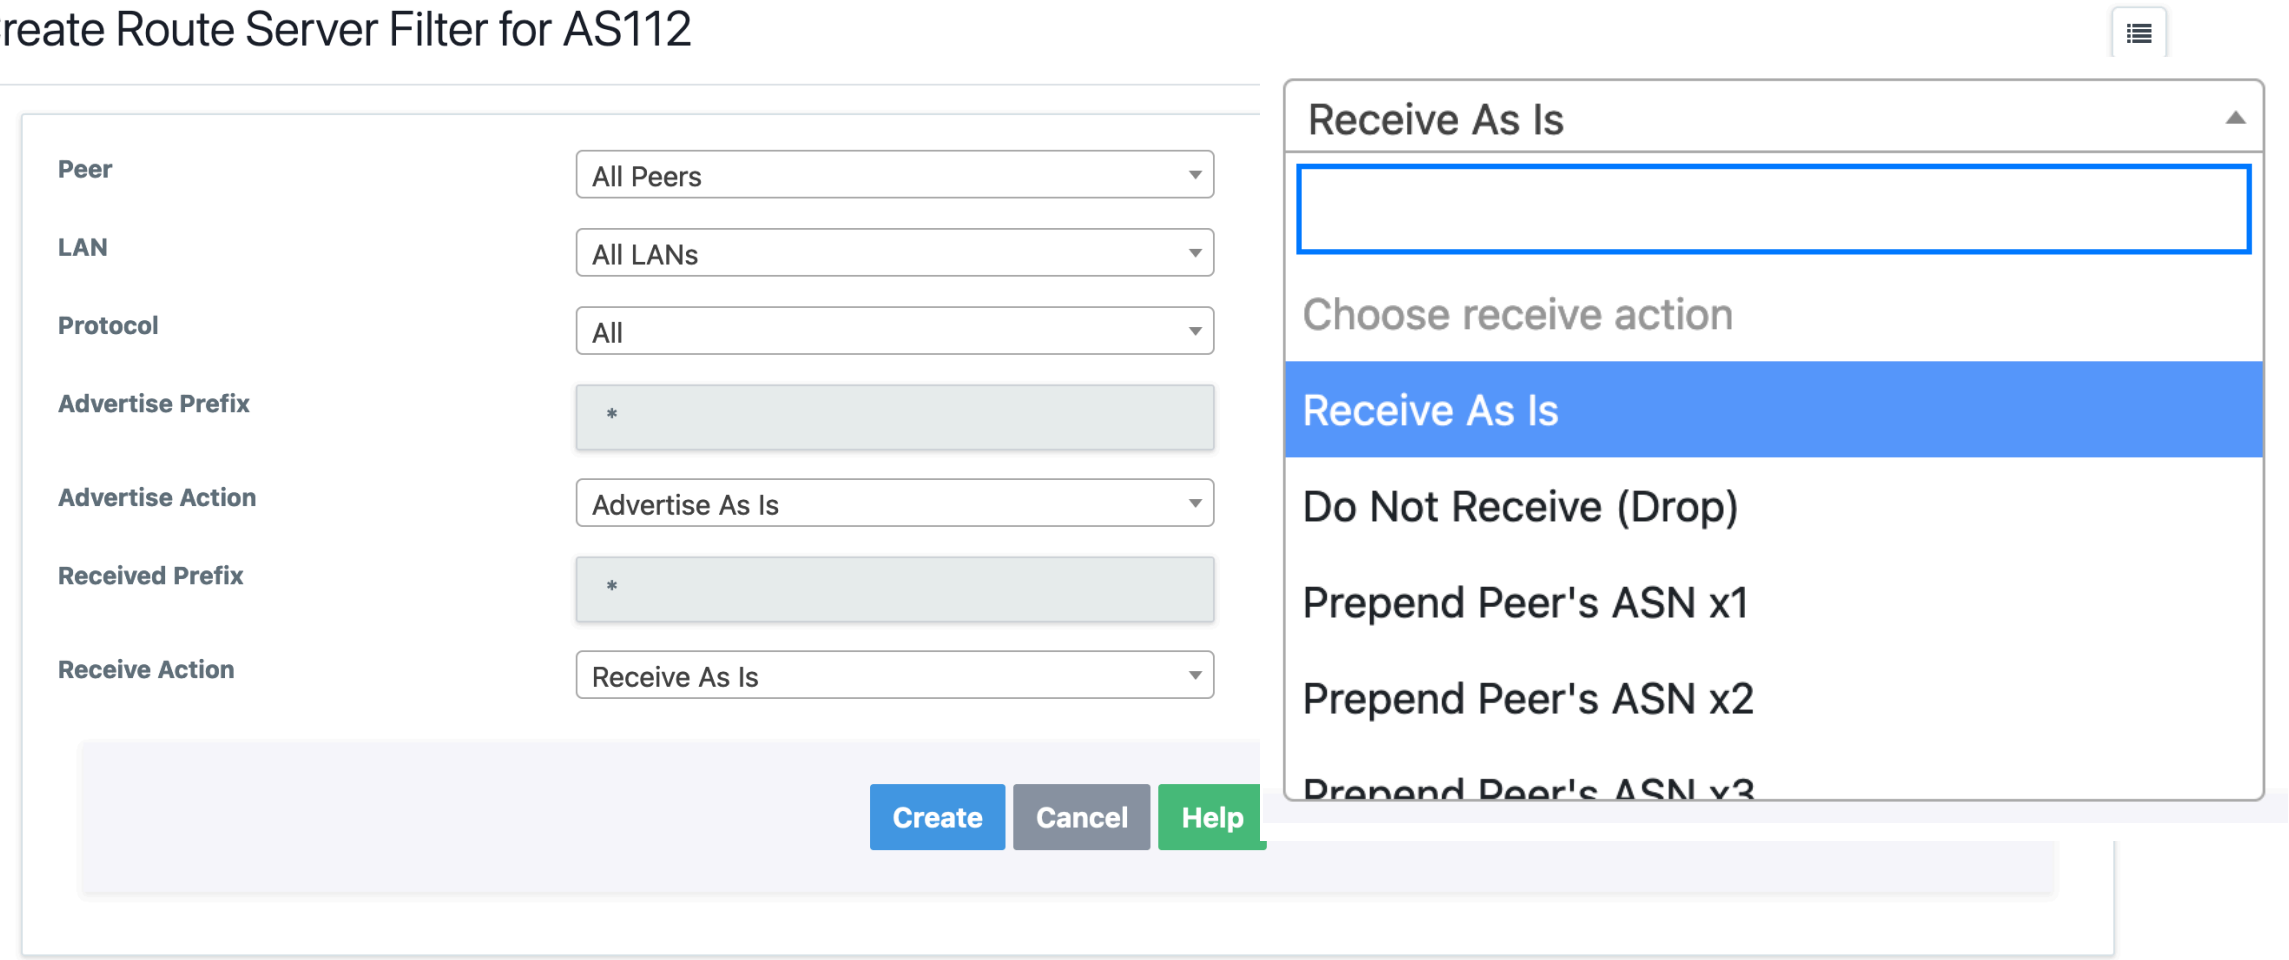

Route Server Filtering for AS112

Your filters are not in sync with our production configuration. You can continue editing or: Revert Commit

Staged Rules (Deploy via Commit above)

Peer	LAN	Protocol	Advertised Prefix	Advertise Action	Received Prefix	Receive Action	Enabled	Order	Actions
HEAnet	Peering LAN 1	Both	*	Do Not Advertise	*	Do Not Receive (Drop)	Yes	1	

Rules in Production

There are no filters in production.

Effect on Route Server Configuration

```
# UI Based Filtering (import rules)

# Filter id:38 created:2022-03-30 14:28:28 updated:2022-03-30 14:28:28
# NO_ADVERTISE to HEAnet
bgp_large_community.add( 43760:0:1213 );



# UI Based Filtering (export rules)

# Filter id:38 created:2022-03-30 14:28:28 updated:2022-03-30 14:28:28
if ( bgp_path.first = 1213 ) then {
    # NO_ADVERTISE - do not advertise this to the route server client
    reject;
}
```

INEX Operations Update March 2022

Barry O'Donovan

INEX Members' Meeting

March 31st, 2022, Dublin

INTERCONNECTING NETWORKS AND PEOPLE FOR OVER 25 YEARS

Connected Edge Capacity

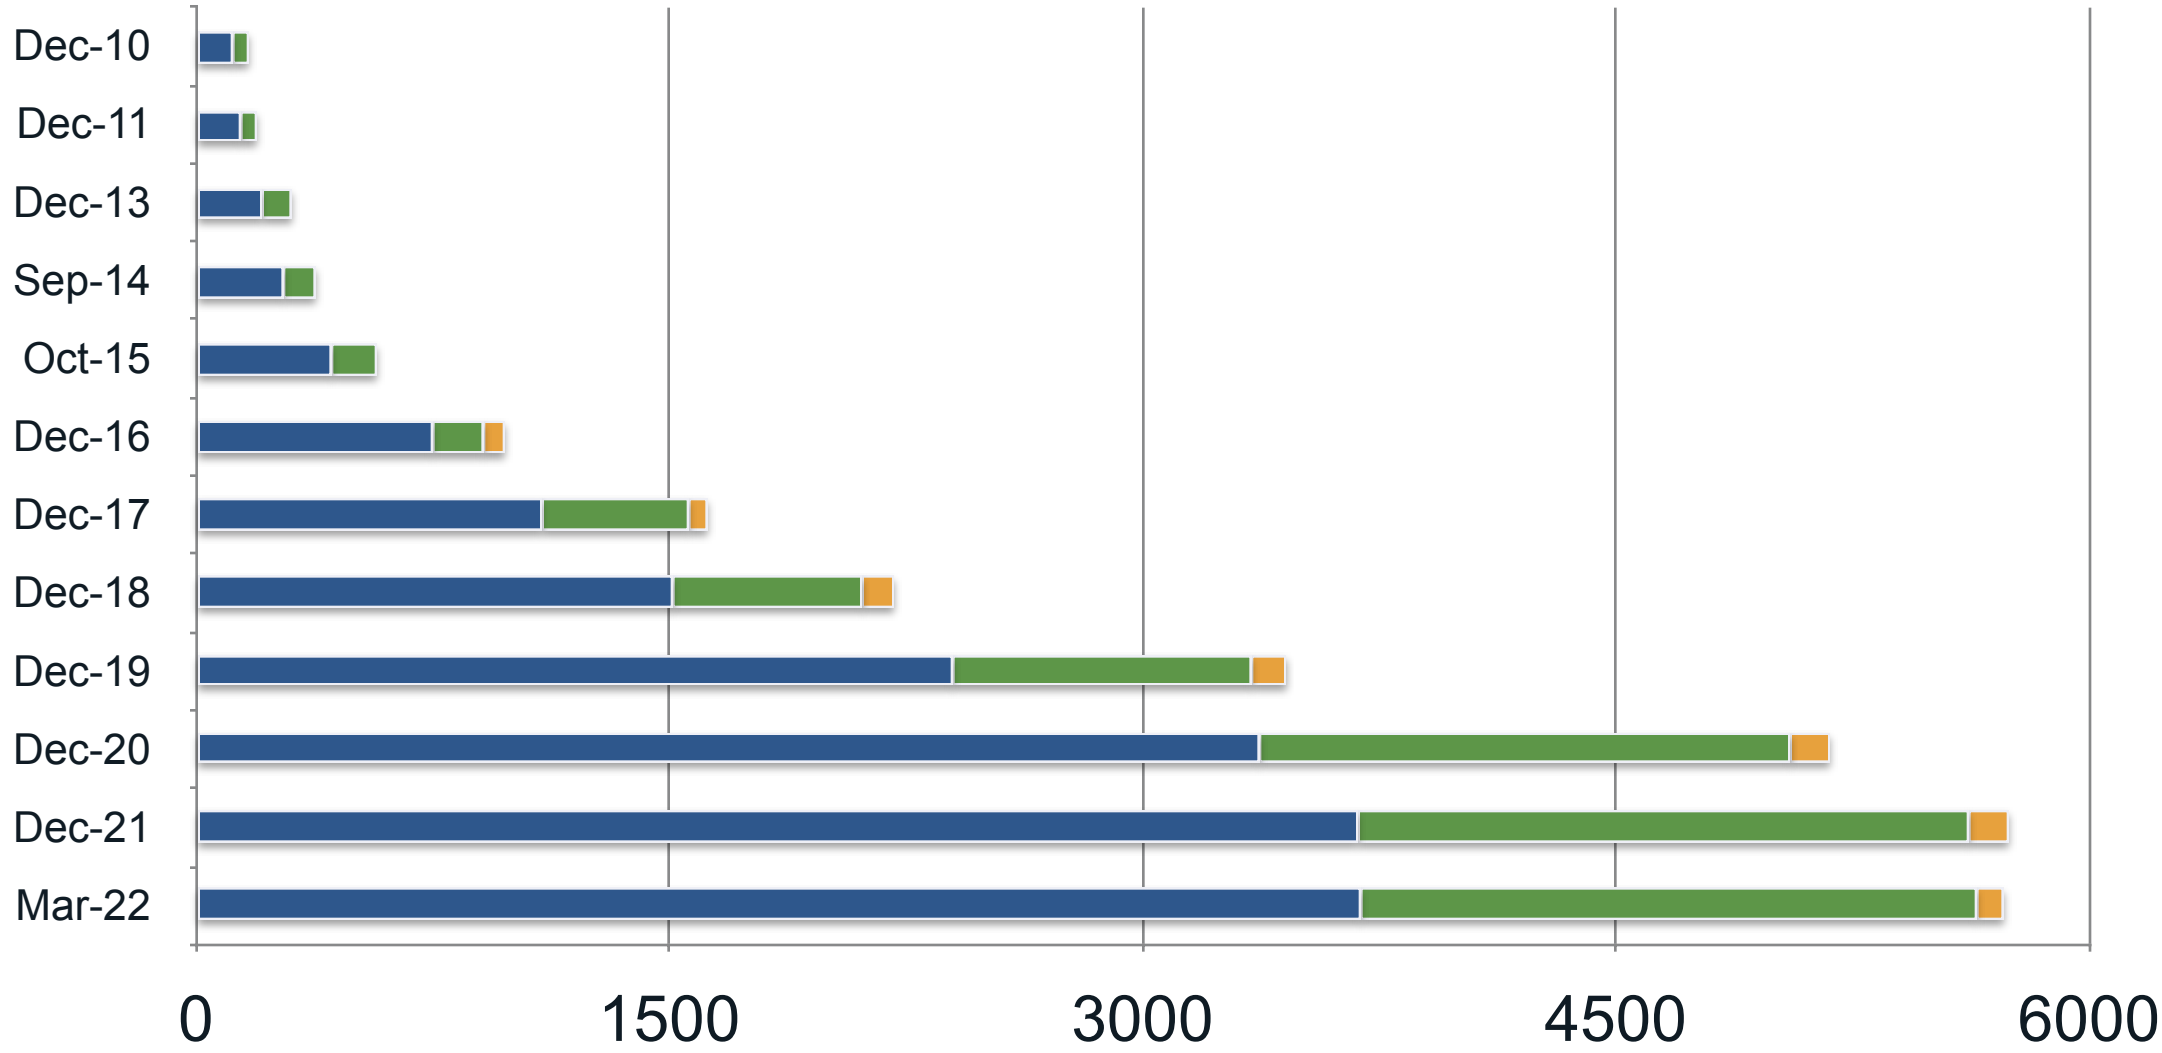

Traffic Peaks :: 1996 - 2021

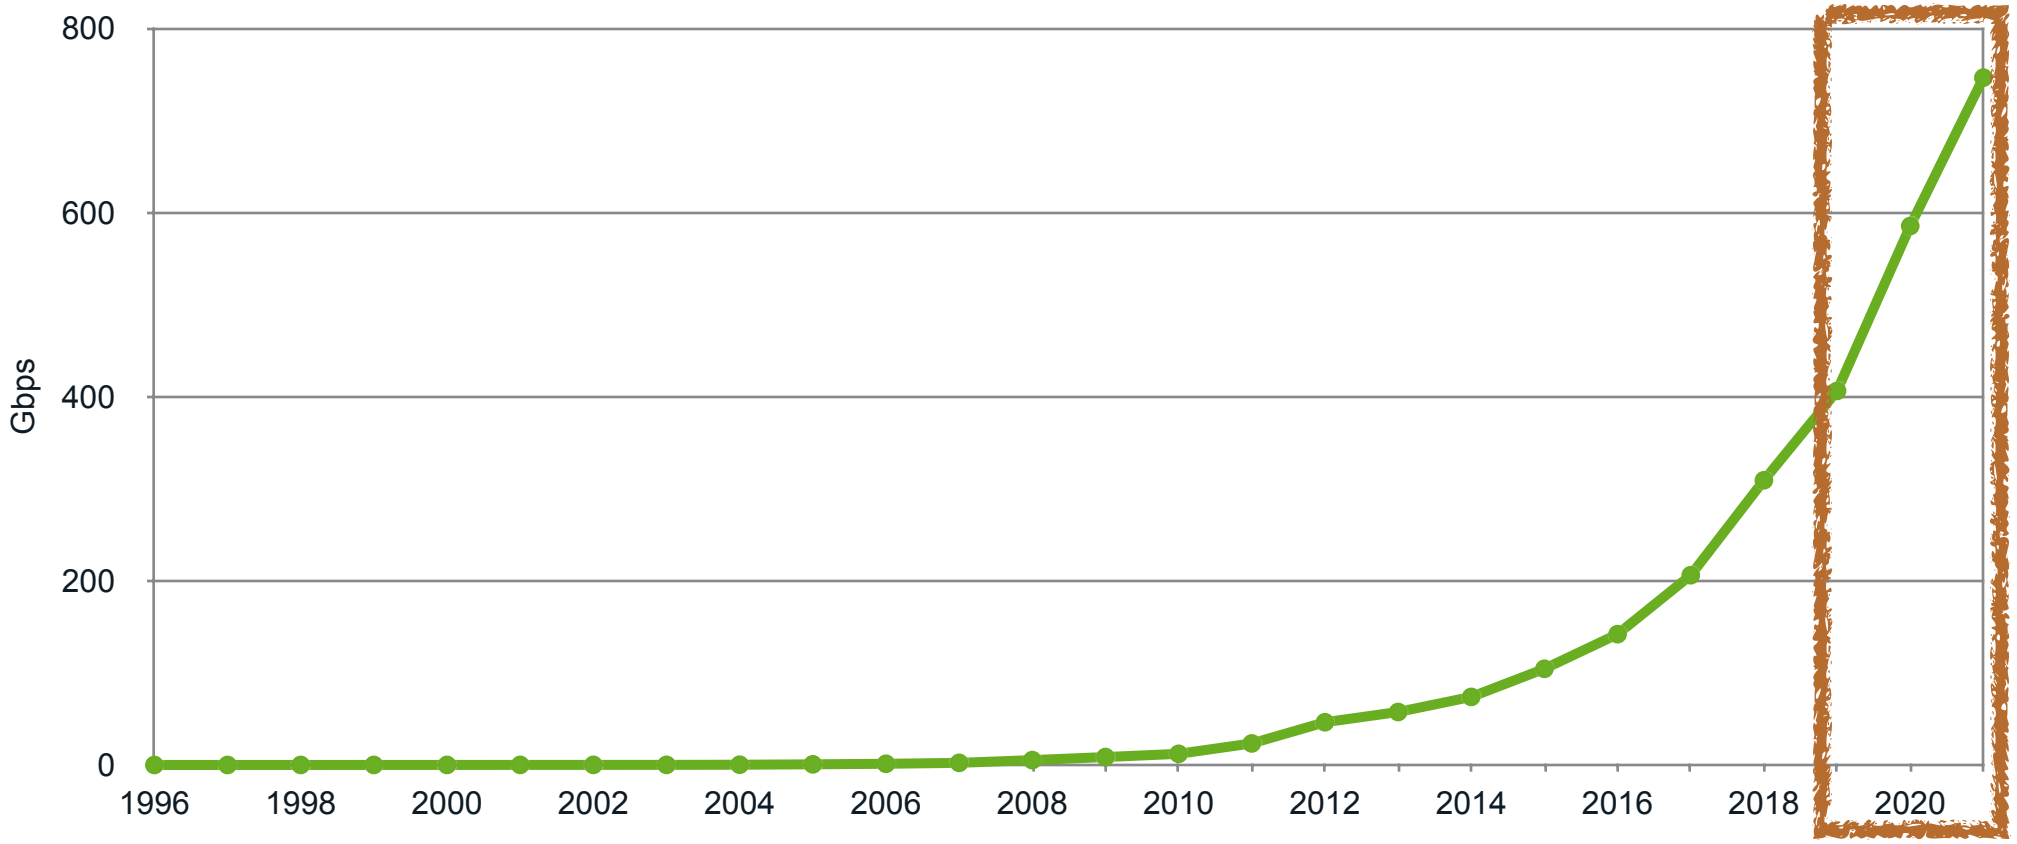

Network Overview

- Traffic peaks continue to grow (2020: 585; 2021: 756 Gbps)
- Supply chain issues very much part of our planning now
- All INEX LAN core / edge capacity in good shape
- Currently deploying 400G-capable core
- In talks with vendors about 400G switches
- Starting to EOL 1Gb ports
- Route server upgrades imminent
- New management servers in production
  - Replaced servers will be recycled to INEX Cork

DWDM Upgrades

Behind the Scenes of a Maint. Window

[inex-tech] INEX scheduled maintenance :: 2022-03-03 :: metro core upgrades

INEX is scheduling a maintenance window on 2022-03-03 between 00:00 and 04:00 UTC to upgrade DWDM transmission capacity on the metro link between Park West and Citywest.

The purpose of this upgrade is to augment capacity on both LAN1 and LAN2 between these two points of presence. **The existing muxes and 200G transponders will be swapped out for wider-band muxes and 600G transponder cards.**

This operation will affect both INEX LAN1 and LAN2, but is not expected to disrupt access service on either network.

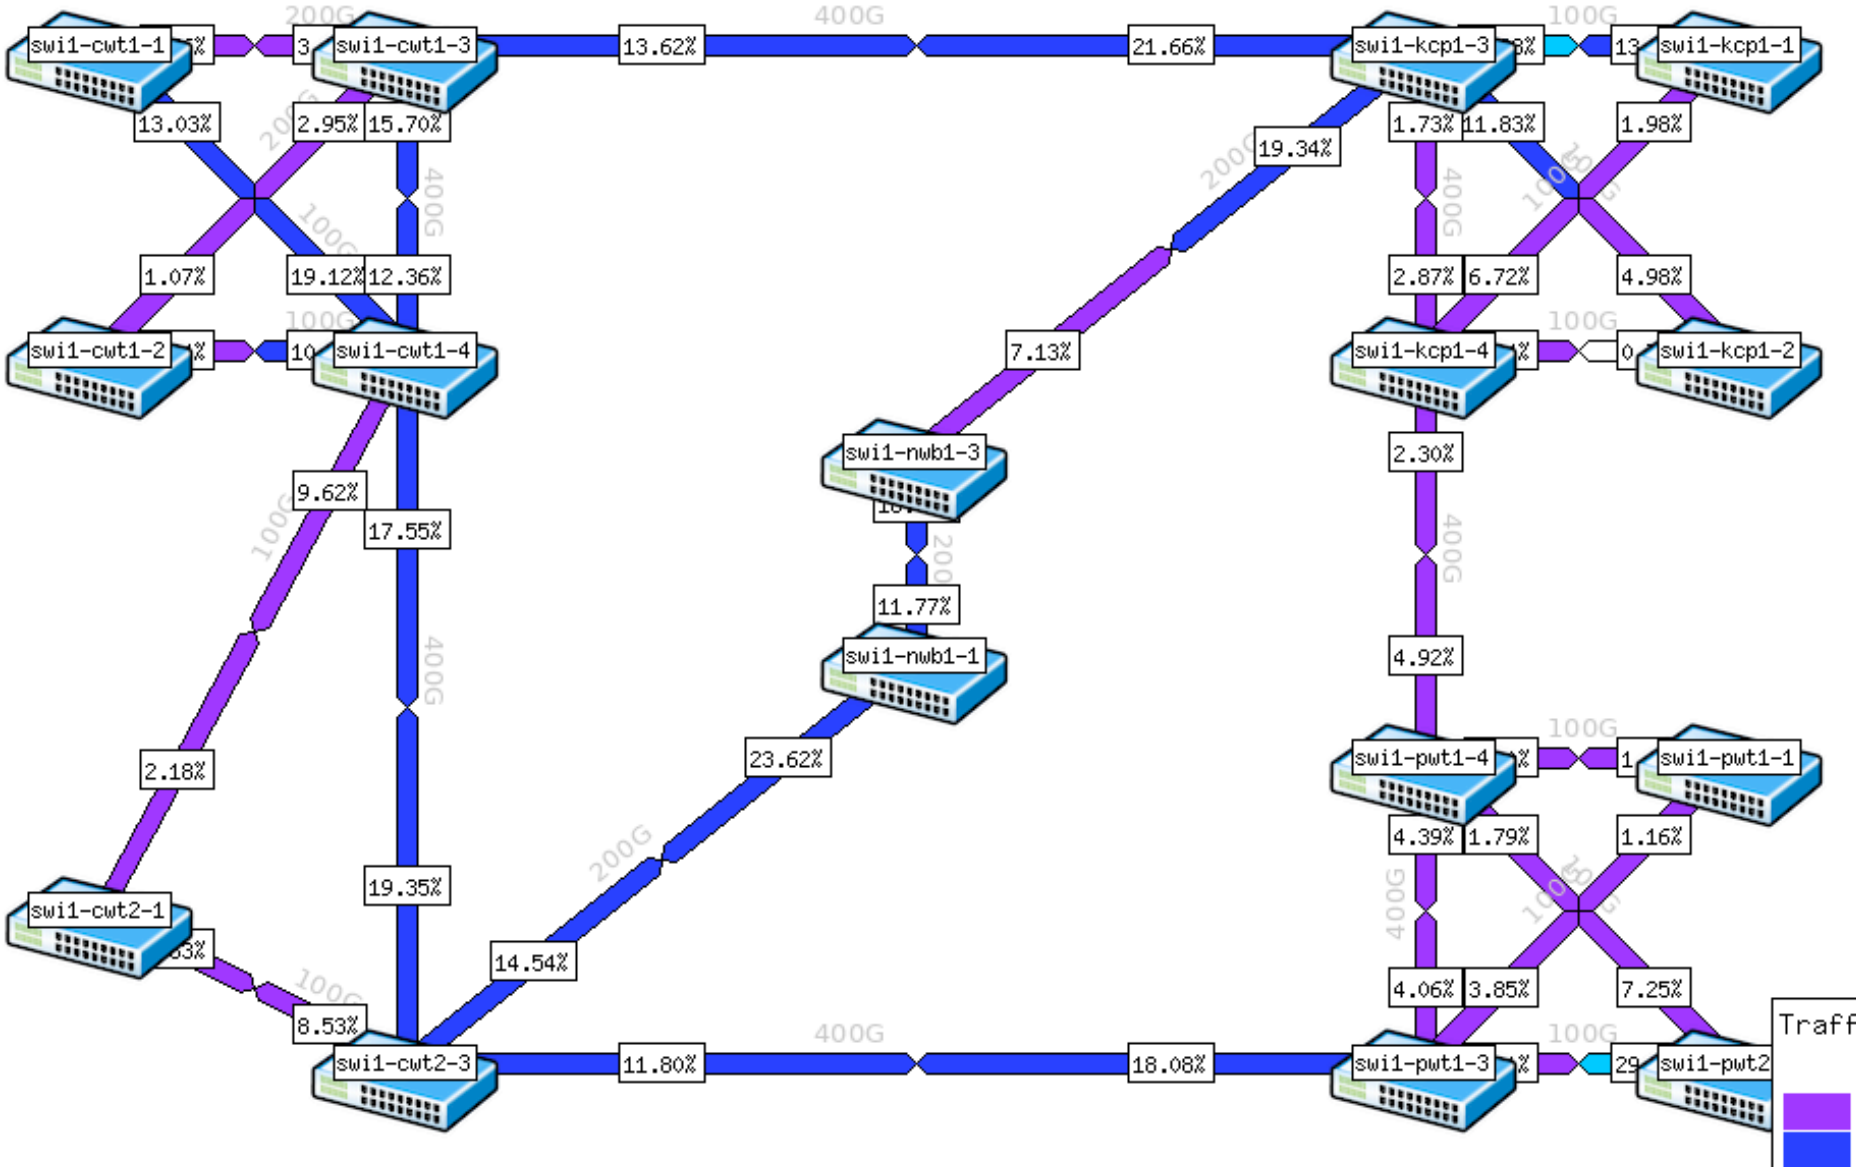
swi1-cwt1-1
swi1-cwt1-3
swi1-kcp1-3
swi1-kcp1-1
swi1-cwt1-2
swi1-cwt1-4
swi1-kcp1-4
swi1-kcp1-2
swi1-nwb1-3
swi1-nwb1-1
swi1-pwt1-4
swi1-pwt1-1
swi1-cwt2-1
swi1-cwt2-3
swi1-pwt1-3
swi1-pwt2
13.62%
21.66%
13.03%
2.95%
15.70%
1.73%
11.83%
1.98%
19.34%
1.07%
19.12%
12.36%
2.87%
6.72%
4.98%
7.13%
2.30%
9.62%
17.55%
11.77%
4.92%
2.18%
23.62%
4.39%
1.79%
1.16%
19.35%
14.54%
4.06%
3.85%
7.25%
8.53%
11.80%
18.08%
Traff

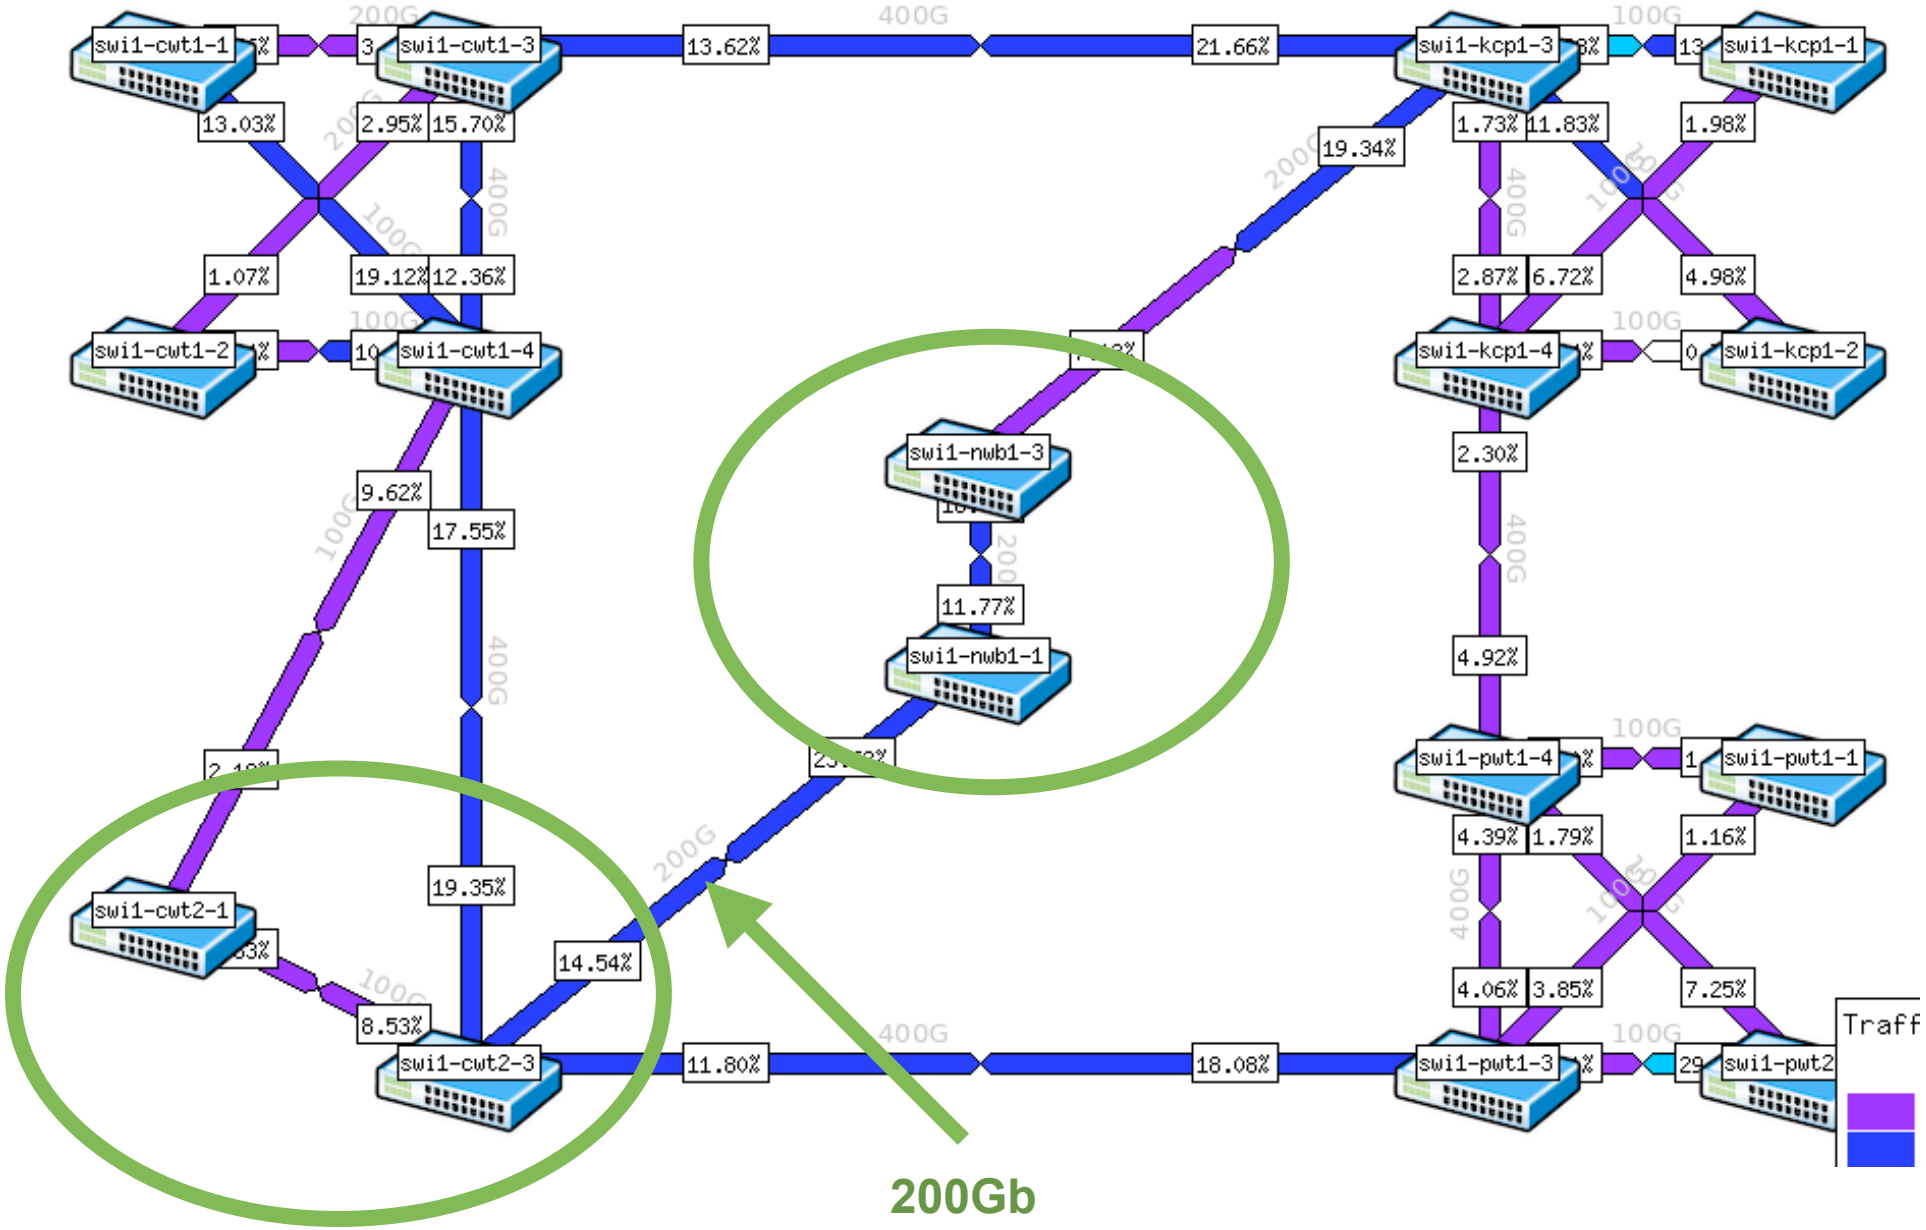
swi1-cwt1-1
swi1-cwt1-3
swi1-cwt1-2
swi1-cwt1-4
swi1-cwt2-1
swi1-cwt2-3
swi1-nwb1-3
swi1-nwb1-1
swi1-kcp1-3
swi1-kcp1-1
swi1-kcp1-4
swi1-kcp1-2
swi1-pwt1-4
swi1-pwt1-1
swi1-pwt1-3
swi1-pwt2
13.62%
21.66%
13.03%
2.95%
15.70%
1.07%
19.12%
12.36%
19.34%
1.73%
11.83%
1.98%
2.87%
6.72%
4.98%
2.30%
9.62%
17.55%
11.77%
4.92%
19.35%
14.54%
8.53%
11.80%
18.08%
4.39%
1.79%
1.16%
4.06%
3.85%
7.25%
400G
200G
100G
200Gb

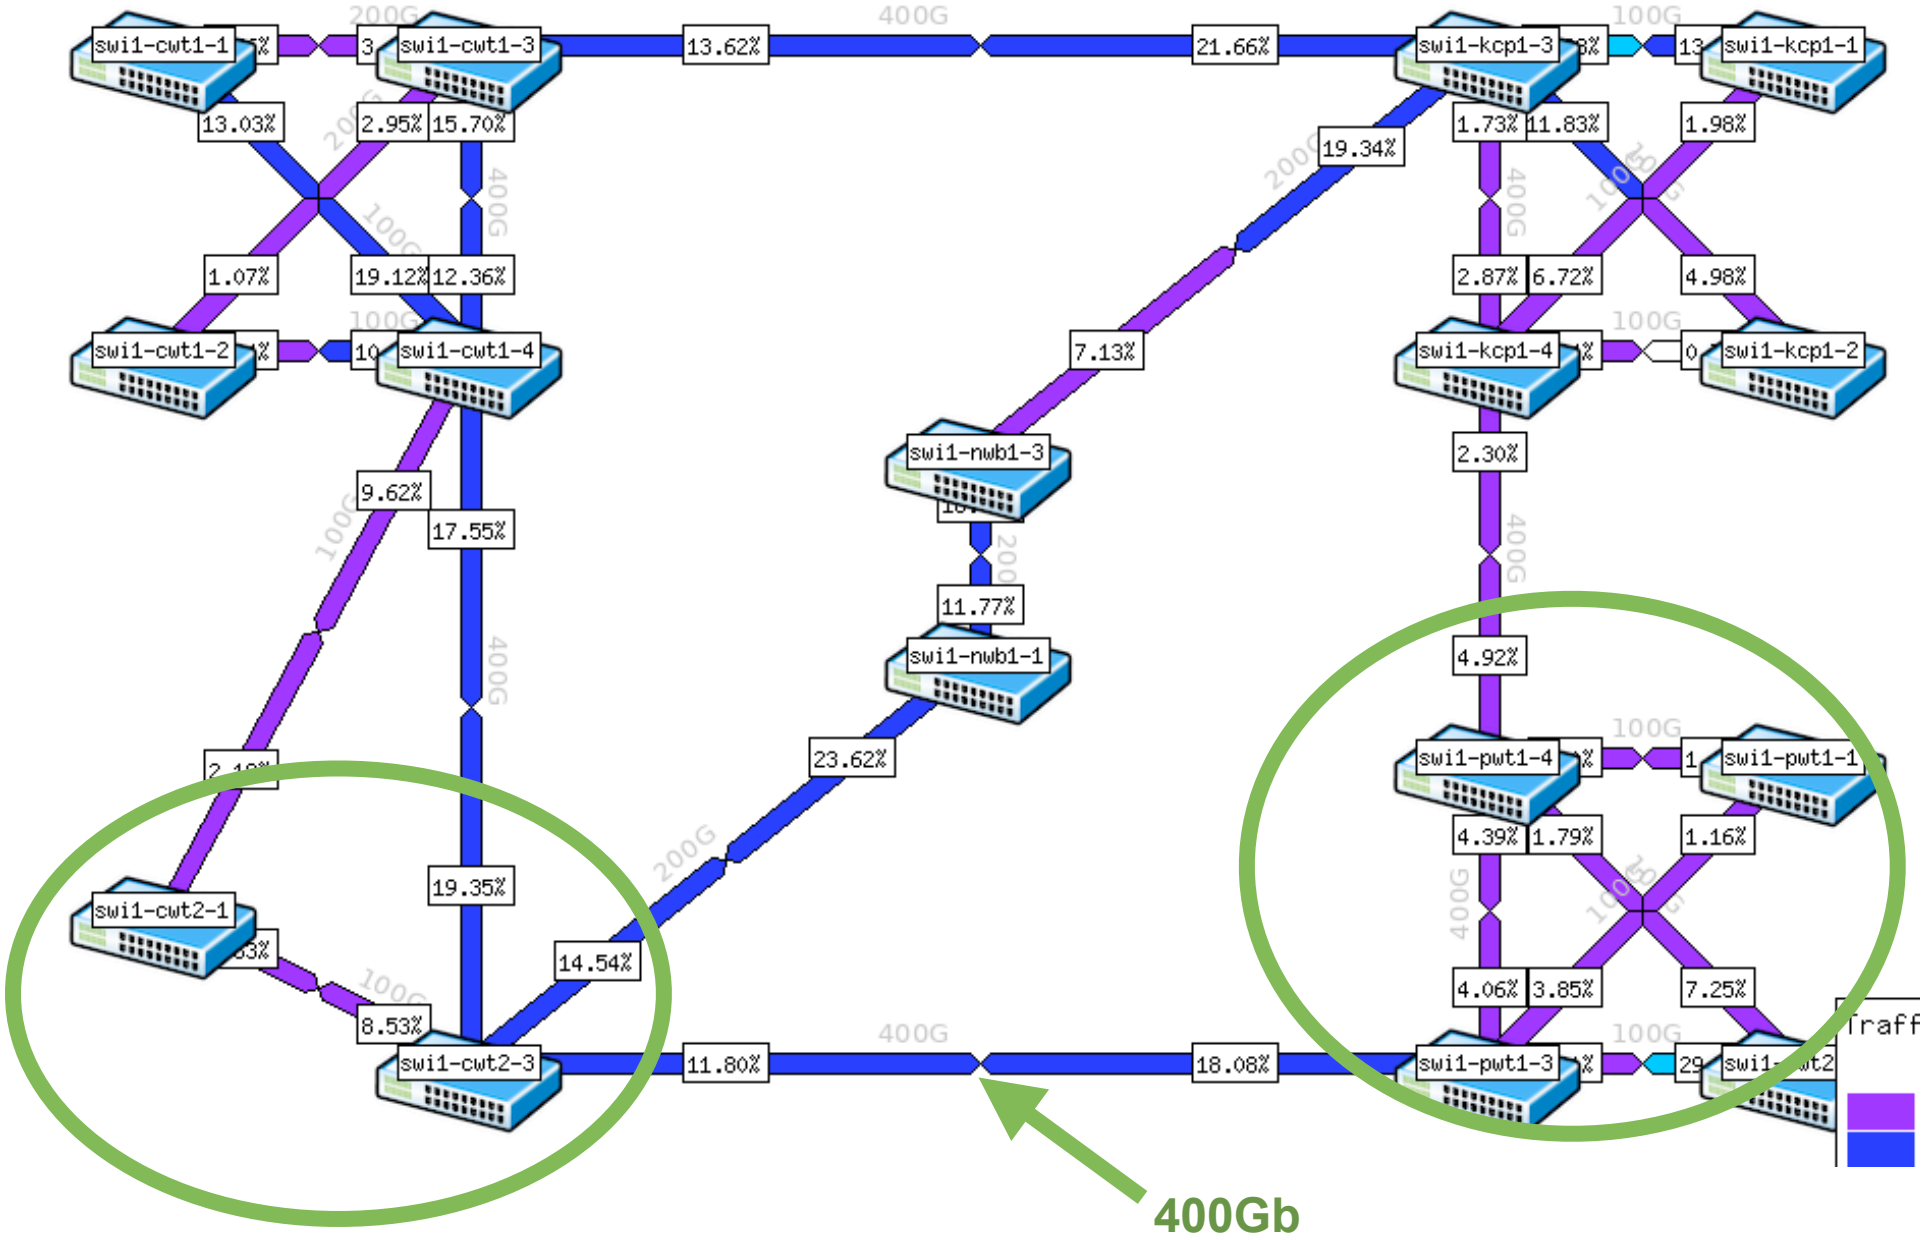
swi1-cwt1-1
swi1-cwt1-3
swi1-cwt1-2
swi1-cwt1-4
swi1-cwt2-1
swi1-cwt2-3
swi1-kcp1-3
swi1-kcp1-1
swi1-kcp1-4
swi1-kcp1-2
swi1-nwb1-3
swi1-nwb1-1
swi1-pwt1-4
swi1-pwt1-1
swi1-pwt1-3
13.62%
21.66%
13.03%
2.95%
15.70%
1.07%
19.12%
12.36%
19.34%
1.73%
11.83%
1.98%
2.87%
6.72%
4.98%
7.13%
2.30%
9.62%
17.55%
11.77%
4.92%
23.62%
19.35%
14.54%
8.53%
11.80%
18.08%
4.39%
1.79%
1.16%
4.06%
3.85%
7.25%
400G
200G
100G
400Gb

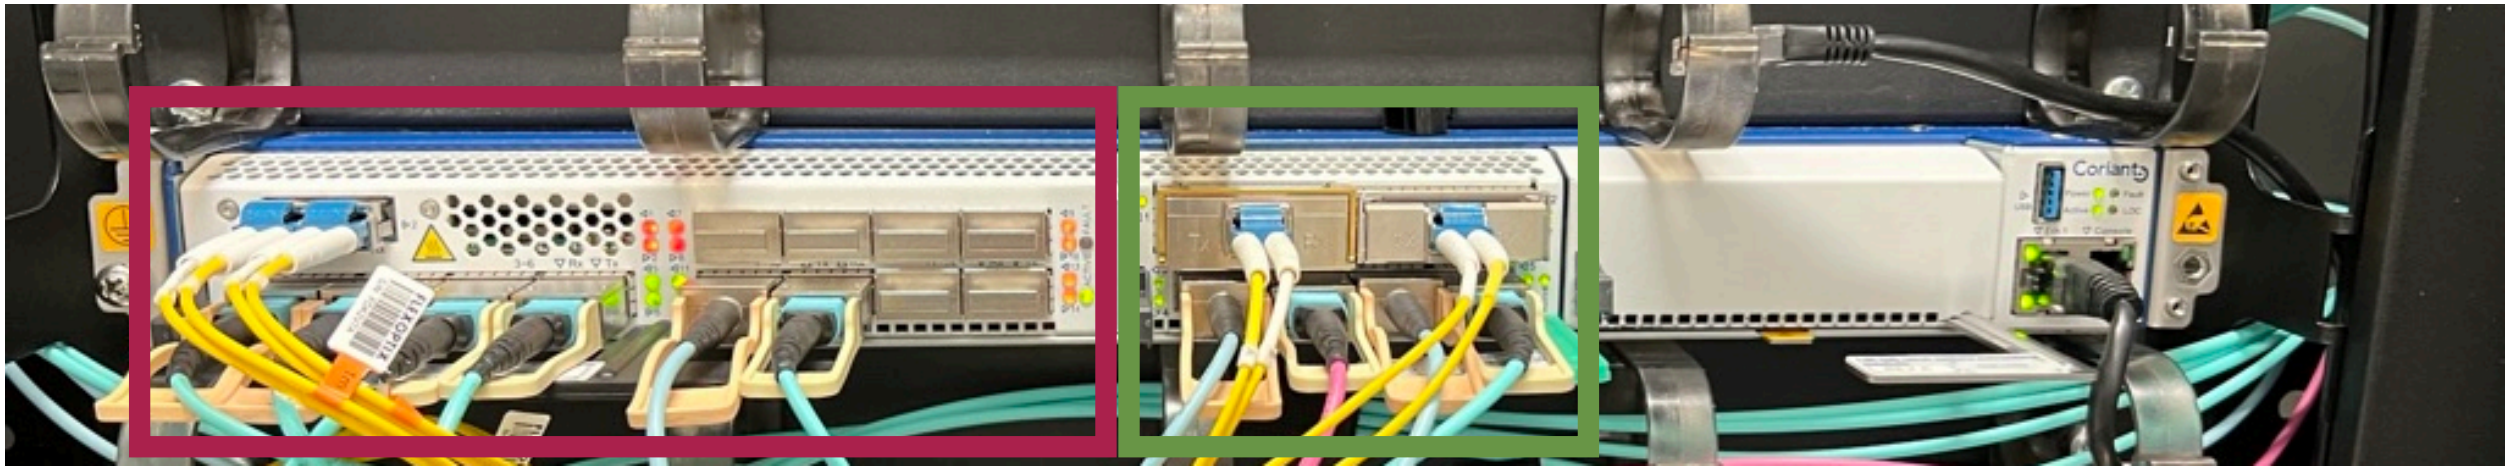

CHM2T

Line: 2 x 600Gb
Client: 12 x 100Gb
Or 3 x 400Gb

CHM1G

Line: 2 x 200Gb
Client: 4 x 100Gb

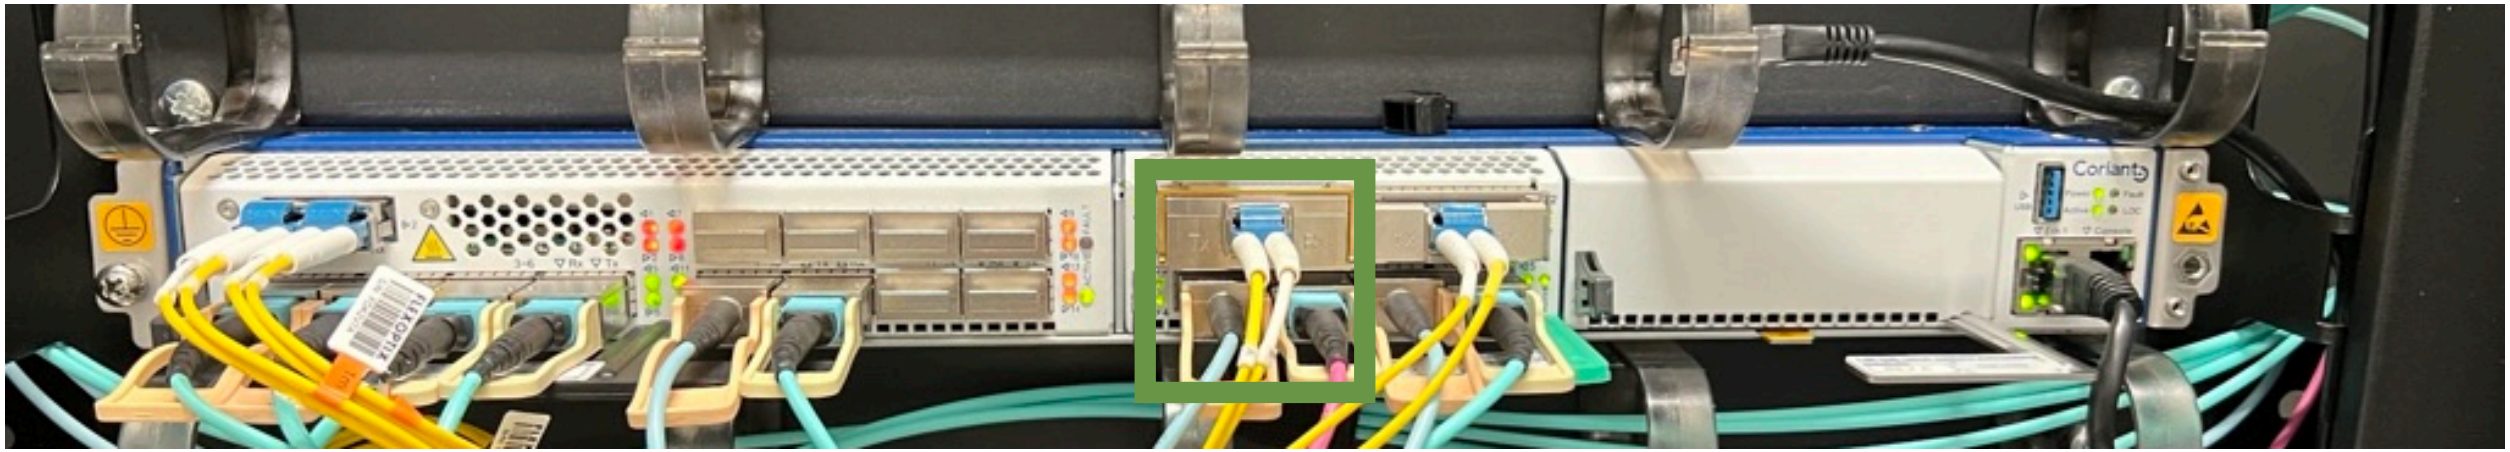

NWBP
INEX LAN1

2 x 100Gb

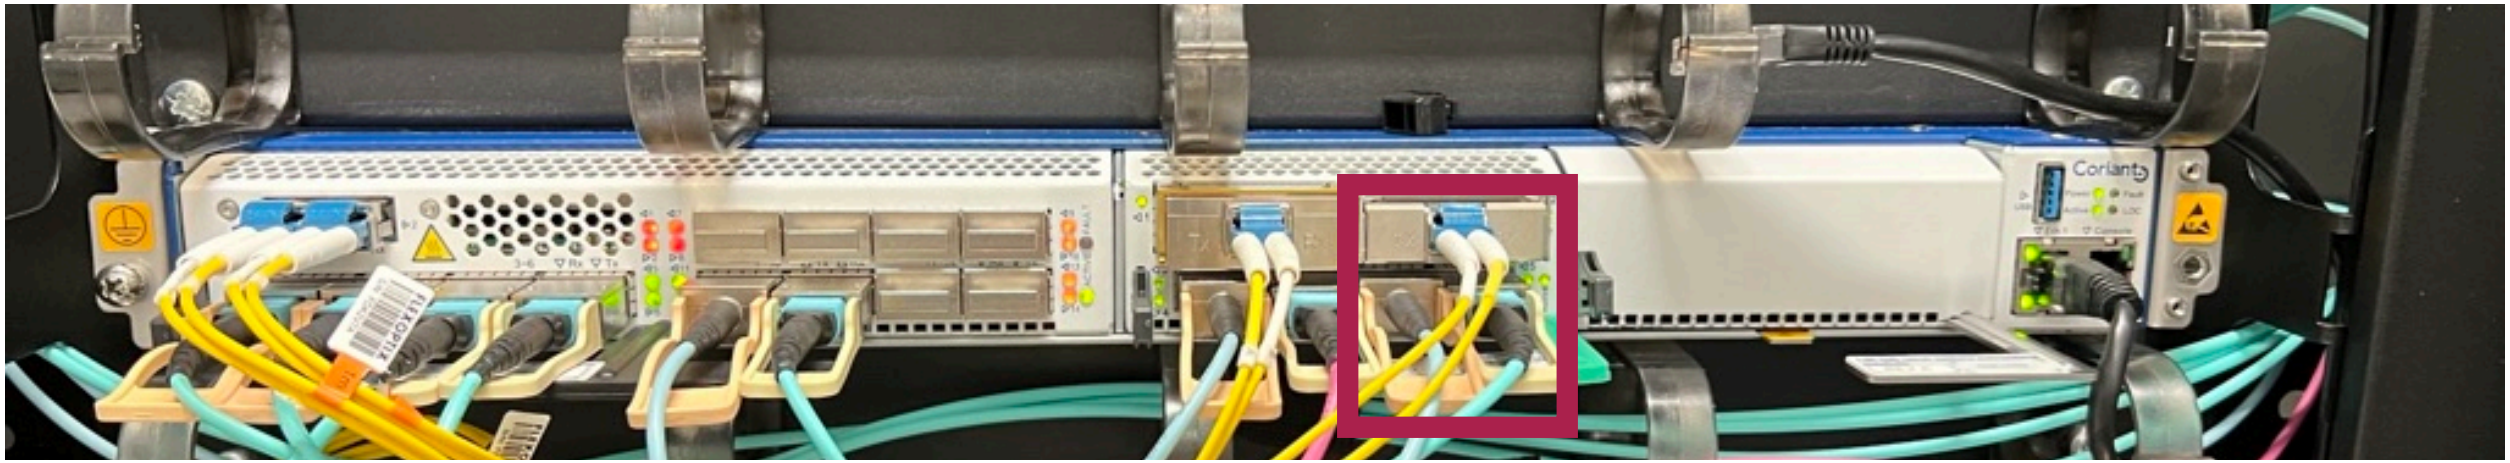

NWBP
INEX LAN2

2 x 100Gb

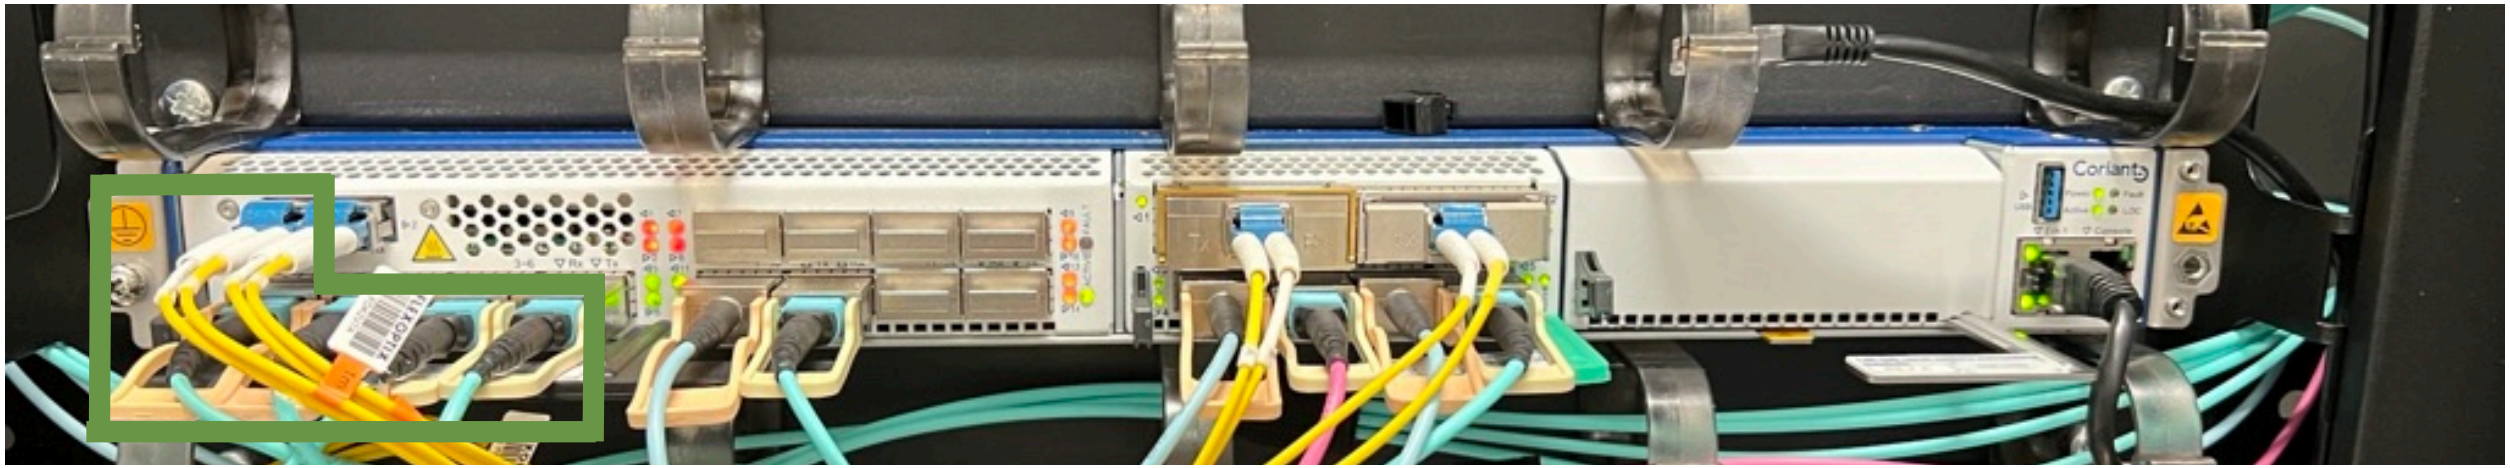

PWT1
INEX LAN1

4 x 100Gb

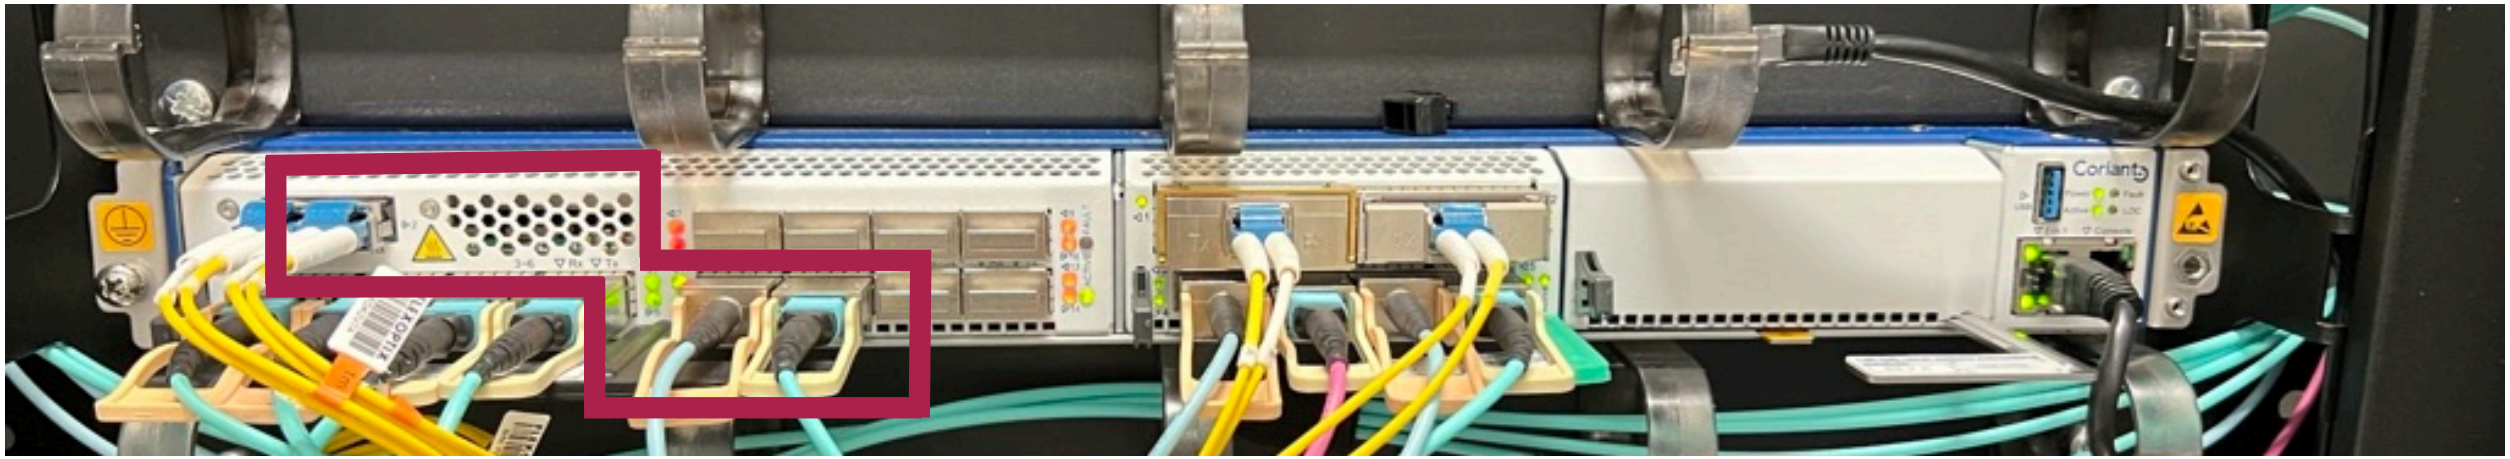

PWT1
INEX LAN2

2 x 100Gb

Other Items of Note

Route Server Refresh

- “Route Server” servers treated like network appliances
  - Critical service, critical infrastructure

- Two physical servers per LAN
  - Deployed in different PoPs

- Standard config includes:
  - Dual PSU
  - OOB management (iDRAC)
  - Hardware RAID (PERC)

Route Server Refresh

- Virtualisation employed as an operational tool
  - Build new stacks side by side
  - Swap in by power down / power up
  - Snapshot for same VM upgrades

- Typically Ubuntu LTS, Bird, Birdseye Looking Glass, mgmt
  - Not part of SaltStack - avoid ‘fat fingered’ mistakes

- New servers will be going live over the coming weeks
  - Dell PowerEdge R350 x 2

Issues

- March 19th - swi1-nwb1-1 unexpected switch reboot
  - No apparent cause
  - Will watch and RMA if necessary

- Route collector BGP session flaps
  - Observed some flaps recently
  - CPU / memory bound
  - Will resolve with a maintenance window

ISO 27001

- Aim still to achieve ISO 27001 compliance in 2022
- Provides 90 - 95% of the requirements of NISD / NISD2
- ISO 27001:2013 gap analysis complete - the INEX team:
  - has a very good understanding of information security in general
  - have implemented several effective information security measures within the organisation
  - there appears to be a clear commitment within the organisation to improve security on an ongoing basis
- Task: existing practices need to be formalised by defining them in policy / procedure.

IXP Manager

IXP Manager

- New release imminent - route server filtering.
- Yann Robin moved back to France late 2021
- Beginning the process of recruiting a new PHP developer
  - This is funded via sponsorship


Thank you
INEX
INTERCONNECTING NETWORKS
AND PEOPLE FOR OVER 25 YEARS