

# INEX Operations Update December 2022

Barry O'Donovan  
INEX Members' Meeting  
December 15th, 2022, Dublin




INTERCONNECTING NETWORKS AND PEOPLE FOR OVER 25 YEARS




# Traffic & Trends



# Connected Edge Capacity




# Traffic Peaks :: 1996 - Now






# Traffic Peaks :: 1996 - Now - YoY Growth



# Annual Trend / Growth - INEX

## Day Graph




	Max	Average	Current
In	1.035 Tbits	678.955 Gbits	429.351 Gbits
Out	1.020 Tbits	677.840 Gbits	429.337 Gbits



# Operations Update

## Core Network

# Optical Line Equipment - Infinera GX G30



## CHM2T

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


## CHM1G

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





## INEX LAN2 - Weathermap



Created: 2022-12-14 20:35:01



# Core Network Upgrades in 2023

- Increase core capacity of LAN1 to 600Gb
  - Mostly optics and RTU licences
- Augment PWT1/PWT2 LAN1 capacity to 2 x [2 x 100Gb LAGs]
- Augment PWT1/PWT2 LAN2 capacity to [4 x 10Gb LAG]
- Estimate LAN2 core to 400Gb not required until 2024





# PoP Rebuilds

CWT1 - Equinix DB1


KCP1 - Equinix DB2



# PoP Rebuilds

- Existing cabinets are not fit for purpose
  - Old style 1990s telco design
  - 600/800mm; solid doors; bottom-to-top ventilation.
  - No possibility of vertical cable management
  - Zero-U PDUs impede racked kit
  - No top of rack access for structured cabling
- Structured cabling shows measurable degradation due to age
  - Physical wear and tear; deplasticisation; ferrule degradation








# PoP Rebuilds

- In place cabinet replacement?
  - Insufficient physical space in DB2 to accommodate
  - Not supported by Equinix
  - Carries greater risk and potential down time
- Equinix DB1
  - New cabinet alongside existing two
  - No requirement for temporary structured cabling
  - Ultimately reducing two cabinets to one






# PoP Rebuilds


- Equinix DB2
  - Two new cabinets in a new location, side by side
  - Requirement for temporary structured cabling
  - Ultimately reducing three cabinets to two
- Plan agreed and supported by Equinix teams
  - PoPs migrated consecutively (single point of focus)
  - Supported by Equinix Migrations Team for cross connects
  - Service-affecting components during weekend daytime hours




Equinix MMR











Migrate


- member x-connects
- switches
- management equipment
- DWSM
- etc....





# Route Servers & Filtering


# 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.

# 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 over 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 ( <i>default</i> )	43760:1:0

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

peer-as == 0 ==> ALL PEERS

NAMESPACE : ACTION : TARGET

# 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. **If in doubt, just contact INEX Operations**

## Your INEX - IXP Manager Dashboard

- Peering Manager
- Route Server Filtering
- Looking Glass
- Peering Matrix



A document store, to facilitate transparency of documents, with the members, by *Information* menu item above or by [clicking here](#). This document store will be populated with basis.

# 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	<div></div>

## Rules in Production

 There are no filters in production.



# Route Server Resilience

# Background

- Route servers considered a critical production service
  - Deployed in pairs on each LAN
- Run on Dell hypervisors treated as production network appliances
  - No other function
  - Dual PSU, hardware RAID, iDRAC
  - Deployed in different PoPs
  - Full array of monitoring / checks













# Previous Configuration Update Method

- Via script ~4 times per day, offset against its resilient route server
- Sanity check - don't install a rs2 config on rs1 ('rs1-lan1-ipv4')
- Check retval for curl api call for config
- Check downloaded file exists and is non-zero size
- Check there are member BGP config stanzas in file (grep)
- Use Bird to parse the config file and check retval
- Backup old config file and replace with new
- Is Bird running? Either start or reconfigure as appropriate
- If reconfiguring and it fails, revert to old configuration file
- API call to IXP Manager to signal update complete



# New Configuration Update Method

- On demand and no longer offset. Script augmented:
- IXP Manager database updated so that all routers now have named pairs
- Script now obtains a lock before starting the update process
  - A locked router prevents its pair from updating
- Additionally, old config and new config are diff'd
  - Bird only reloaded if configuration has actually changed
- API signal to mark update as complete also releases lock
- I.e. A failed update run will prevent the paired router from updating
- Stress tested/ing on the route collectors

Handle	Name	Vlan	Pair	Peering IP	ASN	Last Updated	Actions
rs1-cork-ipv4	RS1 - Cork - IPv4	INEX Cork	rs2-cork-ipv4	185.1.69.8	43760	5 minutes ago	 
rs1-cork-ipv6	RS1 - Cork - IPv6	INEX Cork	rs2-cork-ipv6	2001:7f8:18:210::8	43760	3 minutes ago	 
rs1-lan1-ipv4	RS1 - LAN1 - IPv4	INEX LAN1	rs2-lan1-ipv4	185.6.36.8	43760	36 minutes ago	 
rs1-lan1-ipv6	RS1 - LAN1 - IPv6	INEX LAN1	rs2-lan1-ipv6	2001:7f8:18::8	43760	36 minutes ago	 
rs1-lan2-ipv4	RS1 - LAN2 - IPv4	INEX LAN2	rs2-lan2-ipv4	194.88.240.8	43760	36 minutes ago	 
rs1-lan2-ipv6	RS1 - LAN2 - IPv6	INEX LAN2	rs2-lan2-ipv6	2001:7f8:18:12::8	43760	36 minutes ago	 

# IXP Manager





# IXP Manager

- [www.ixpmanager.org](http://www.ixpmanager.org) -> Community -> Statistics
  - 126 of 205 IXPs using their IX-F Member Export API.
  - A further 72 IXPs using PeeringDB.
  - The remaining 7 are not currently polled.
- Provided with a NCC Group security audit of IXP Manager
  - Fixes in v6.3.
- Release v6.4 will include the route server filtering feature.





# Thank you



**INEX**

INTERCONNECTING NETWORKS  
AND PEOPLE FOR OVER 25 YEARS