IPv6 Deployment in MyIX

Professor Dr. Sureswaran Ramadass

INTERNATIONAL UNIVERSITY OF MALAYA-WALES

What is an Internet Exchange

- Major providers connect their networks and exchange traffic.
- High-speed network or Ethernet switch.
- Simple concept – a place where providers come together to exchange traffic.

What happens at an Internet Exchange Point ?

- Multiple ISPs locate backbone IP router nodes in single building operated by co-location provider
- In-building connections
  - to shared interconnect fabric (using Ethernet LAN switching technology)
  - over point-to-point private interconnections
- Routing information, and hence traffic, is exchanged bi-laterally between ISPs
- Exchange operator may or may not be same organization as co-location provider
- Co-location provider will generally have other customers:
  - carriers, hosting, ASPs, content distributors

Inter-ISP Interconnect

- Peering:
  - two ISPs agree to provide access to each others' customers
  - commonly no money changes hands: "settlement free"
  - barter of perceived equal value
  - simple commercial agreements
- Public Interconnect:
  - Internet Peering Point ("IPP" or "IXP" or "NAP")
  - multiple parties connect to shared switched fabric
  - commonly Ethernet based
  - open, many-to-many connectivity
  - traffic exchange between consenting pairs of participants
- Other models exist

Importance of IXP Neutrality

- In most markets, IXPs are a natural monopoly
  - problem of trust between competitors
  - risks of abuse and conflicts of interest
- Successful IXPs are not usually:
  - owned, operated or housed by a single ISP or carrier
  - ISPs or wholesale IP transit providers
  - national or international backbones
- Co-location facility neutrality:
  - normally (mainly in Europe) these are buildings operated by independent commercial companies
  - though sometimes (mainly in US) co-los operate IXPs
  - IXPs tend not to be in carrier co-lo facilities

Some IXP Neutrality Principles

- Does not compete with its ISP members/customers
- Does not discriminate between its ISP members/customers
- Does not move traffic between cities or countries
- Does not make exclusive arrangements with:
  - ISPs
  - Carriers
  - Co-lo Providers
- Does not provide IP transit routing
- Does not take share of ISPs' transit revenues
- Only interconnects between metro area co-lo sites
- May be present at multiple co-lo sites and providers

Governance/Commercial Models

- Operated by public sector national academic network
  - BNIX, GIGAPIX
- Not-for-profit membership associations of participating ISPs (majority !)
  - LINX, AMS-IX
- Service within commercial co-location operator
  - Equinix, PAIX, IX Europe
- Companies whose shareholders are participating ISPs
  - MYIX, JPIX
- Independent neutral commercial companies
  - XchangePoint, JPNAP

Routing and Switching at IXPs

- ISPs perform Layer-3 IP routing over wide-area using routers connected by long-haul circuits
- IXPs perform layer-2 switching over local/metro area, usually using Ethernet
- ISPs interconnecting at IXPs exchange IP routing information using BGP (Border Gateway Protocol)

IXP Advantages

- Single large pipe to the IXP more efficient than many smaller pipes to many ISPs

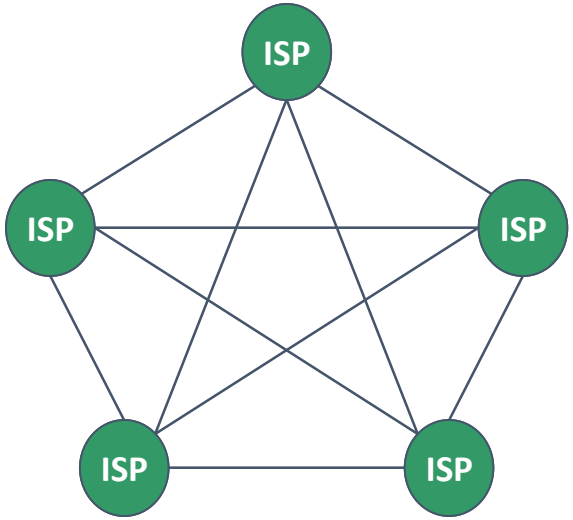

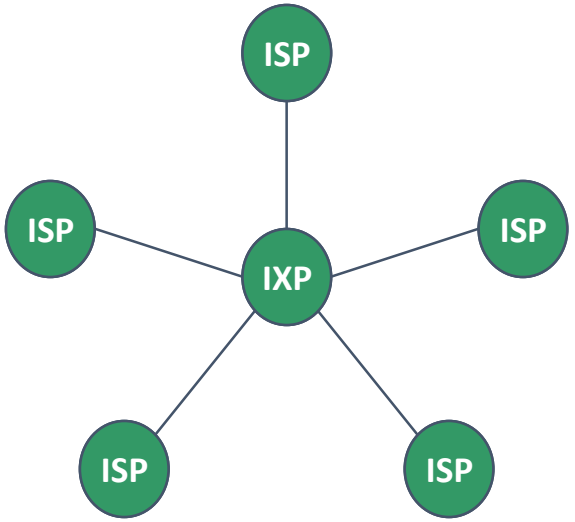

ISP = Internet Service Provider

IXP = Internet eXchange Point

IXP Advantages

- Keeps domestic traffic within a country/region without having to take indirect international route
- Typically 20-35% of traffic can be domestic
- Reduced bandwidth costs
- Improved throughput and latency performance
- Economies of scale
- Commercial basis of traffic exchange between ISPs across IXP usually via cost-saving peering
- Critical mass of ISPs in a single location creates competitive market in provision of capacity, transit and services

Malaysian Internet Exchange

Malaysia Internet Exchange(MyIX)

- ASN: 55822
- Traffic Profile: Balanced
- Traffic Volume: 80 Gbps
- Peering policy : Open, # Peers 70
- Peering Locations: Kuala Lumpur
- PeeringDB Entry: as55822.peeringdb.com
- Contact: raja.mohan@myix.my

INTERNATIONAL UNIVERSITY OF MALAYA-WALES

INTERNATIONAL UNIVERSITY OF MALAYA-WALES

MyIX: History

- Established in Dec 2006.
- Funded by the government
- Supported by regulator and the Malaysian Communication and Multimedia Commission (MCMC)
- Started with 16 members, 3 sites (KL)
- Managed and operated by industry
- Non – profit
- Layer 2, Open peering
- Supports Unicast IPv4 and IPv6

INTERNATIONAL UNIVERSITY OF MALAYA-WALES

Peering Nodes (Central)

TM

CSF1
Jalan Teknokrat 6
CyberJaya , Selangor

40 Gbps
Ring

AIMS

1st Floor, Menara Aik Hua,
Cangkat Raja Chulan,
Kuala Lumpur

JARING
come together

Ground Floor Telco Room,
Bangunan INTAC
Technology Park Malaysia,
Bukit Jalil, Kuala Lumpur

INTERNATIONAL UNIVERSITY OF MALAYA-WALES

Peering Nodes (Regional)

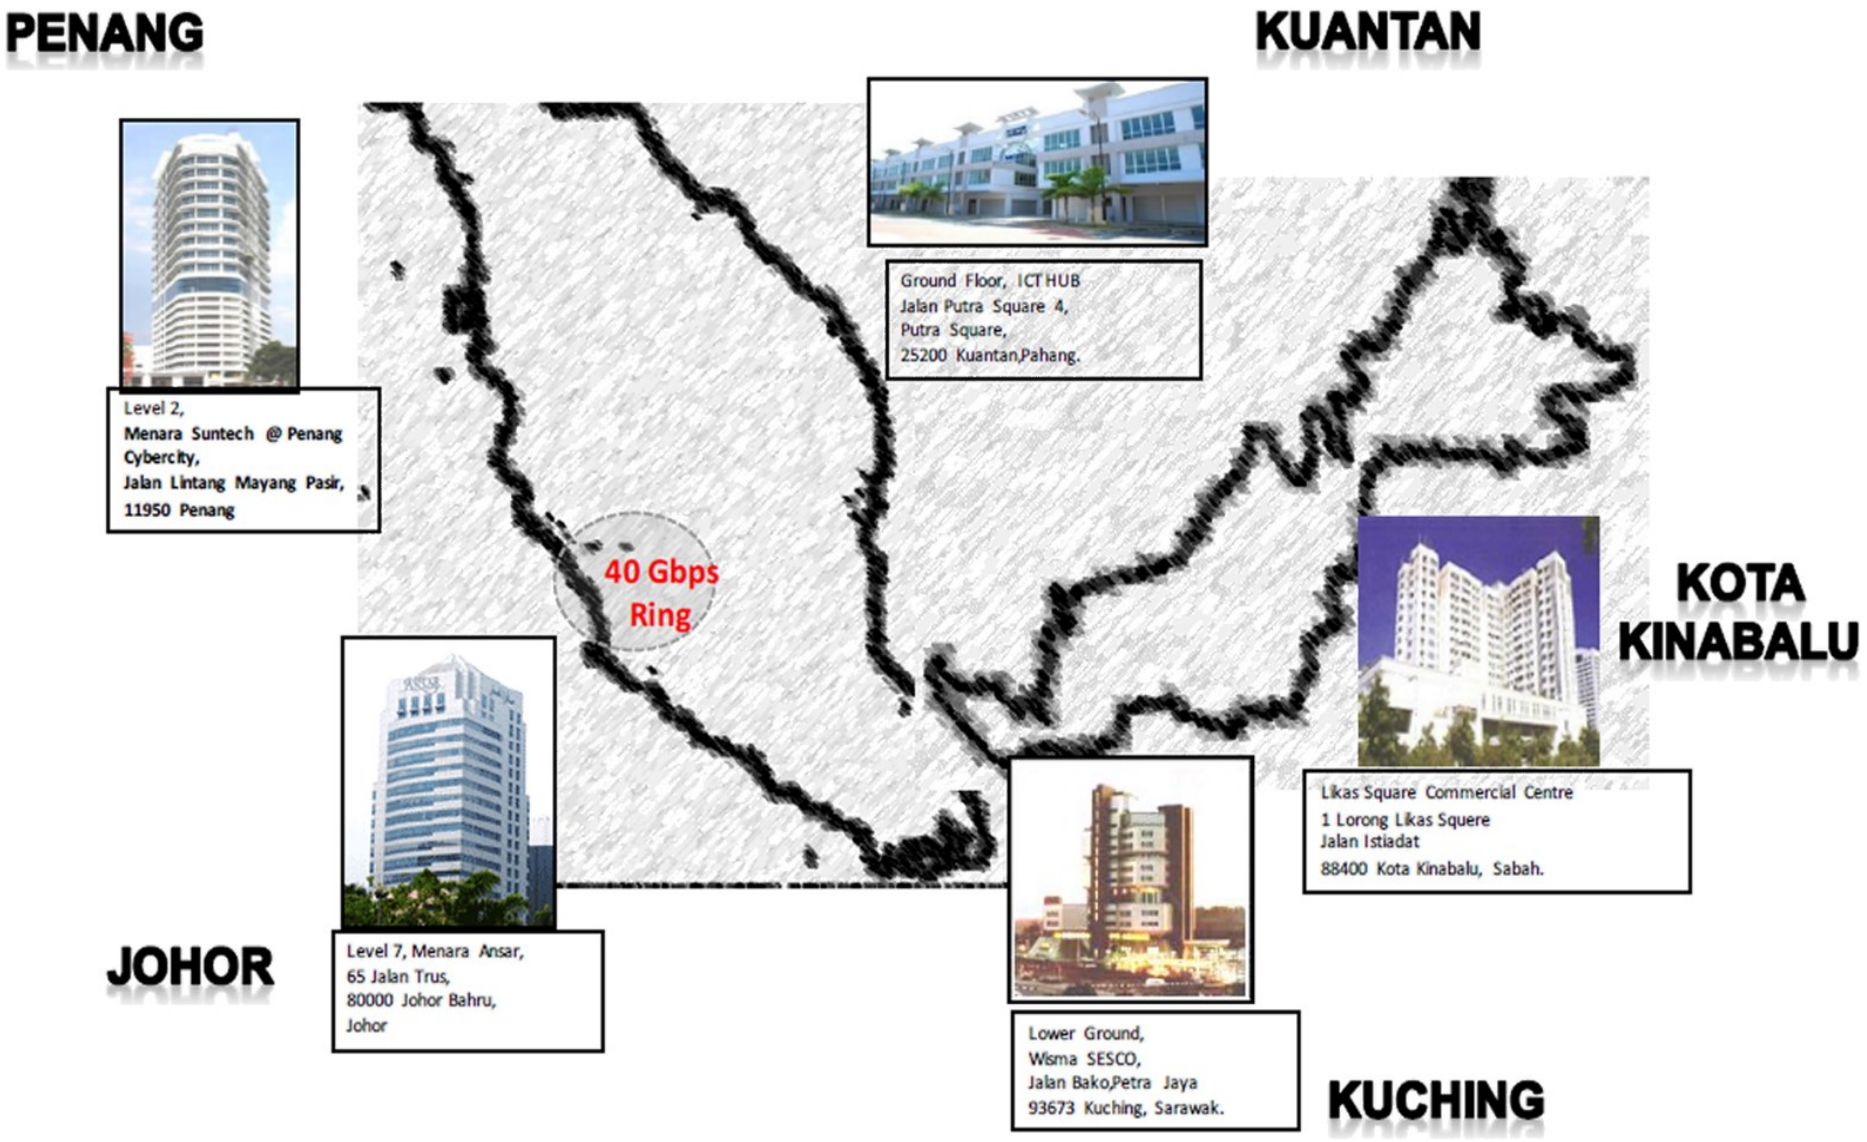

INTERNATIONAL UNIVERSITY OF MALAYA-WALES

MyIX Membership

- Ordinary Member
  - Local NSP
  - ASN
  - Voting Right
- Associate Member
  - Local company
  - ASN
  - Non-Voting Member
- Peering Partner
  - ASN

INTERNATIONAL UNIVERSITY OF MALAYA-WALES

Member & Traffic Growth

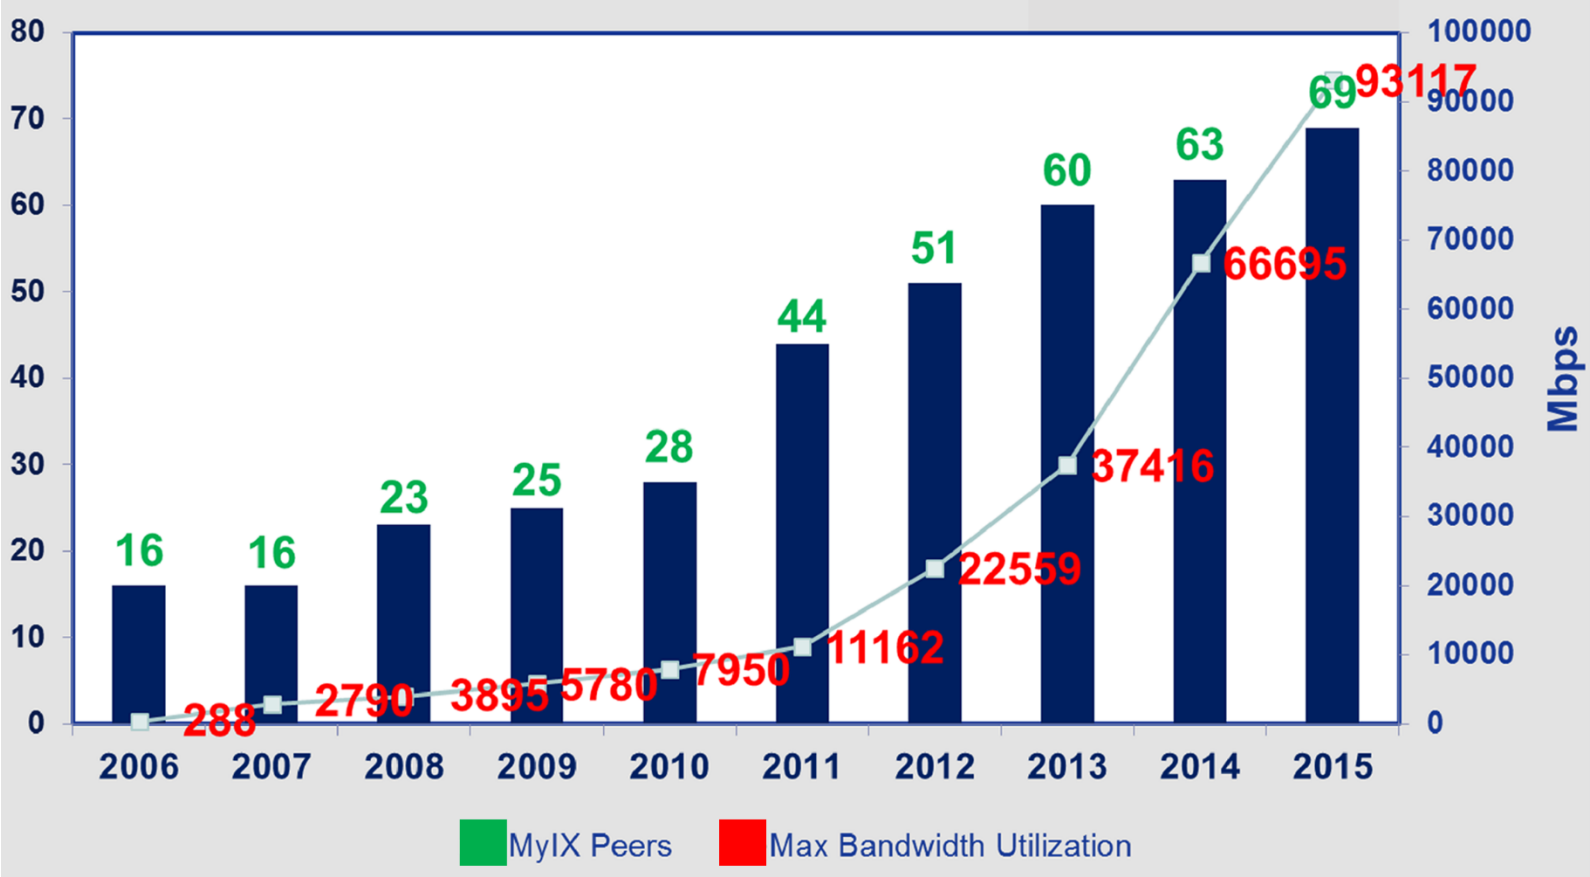

MyIX Peers : IPv4 *218.100.44.0/24

1. Acme Commerce Sdn Bhd
2. ACODA
3. AIMS Data Centre Sdn Bhd
4. Akamai Technologies
5. Aktif Setegap Sdn Bhd
6. Altel Communications Sdn Bhd
7. Amazon.com
8. Arcnet
9. ASIAN BROADCASTING NETWORK (M) SDN BHD
10. BIGBand
11. Biznet Networks
12. Celcom (M) Berhad
13. CloudFlare
14. Colocation Hosting Sdn Bhd
15. DiGi Telecommunications Malaysia
16. Everworks IDC Sdn Bhd
17. ExaBytes Network
18. Extreme Broadband
19. Facebook
20. Fibrecomm Network (M) Sdn Bhd
21. Formis Development
22. Global Transit
23. Communications
24. Google Inc.
25. HGC
26. Hitachi Sunway Data Centre Services Sdn Bhd
27. IP ServerOne Solutions Sdn Bhd
28. ISC
29. JARING Communications Sdn Bhd
30. KKIPC - SSDC
31. LightsUp Network Solution Sdn. Bhd.
32. Macro Lynx Sdn Bhd
33. Malaysian Research & Education Network (MYREN)
34. Maxis Communications Bhd
35. Microsoft
36. MNC Wireless Berhad
37. Monash University Sunway Campus
38. MYCorous360
39. MyKRIS ASIA KL
40. Net Onboard Sdn Bhd
41. Net2One Sdn Bhd
42. Netnod
43. OCESB
44. Pacific Link Telecom (M) Sdn Bhd
45. Packet Clearing House
46. Packet-One Networks Sdn Bhd
47. Piradius Sdn. Bhd
48. Redtone-CNX Broadband Sdn Bhd
49. Symphonet SB
50. Tg Agas Technology Sdn Bhd
51. TheGigabit
52. TIME DotCom Berhad
53. TM
54. Tune Music
55. Twitter, Inc.
56. U Mobile Sdn. Bhd.
57. VADS Berhad
58. VTELECOMS BERHAD
59. Webvisions Network Services Sdn Bhd
60. YTL Communications Sdn Bhd

*Source: peeringdb.com

INTERNATIONAL UNIVERSITY OF MALAYA-WALES

MyIX Peers : IPv6 *2001:de8:10::/112

1. ACODA
2. AIMS Data Centre Sdn Bhd
3. Akamai Technologies
4. Aktif Setegap Sdn Bhd
5. Altel Communications Sdn Bhd
6. Arcnet
7. ASIAN BROADCASTING NETWORK (M) SDN BHD
8. Biznet Networks
9. CloudFlare
10. DiGi Telecommunications Malaysia
11. Facebook
21. Global Transit Communications
22. Google Inc.
23. HGC
24. Hitachi Sunway Data Centre Services Sdn Bhd
25. IP ServerOne Solutions Sdn Bhd
26. JARING Communications Sdn Bhd
27. Macro Lynx Sdn Bhd
21. Malaysian Research & Education Network (MYREN)
22. Maxis Communications Bhd
23. Microsoft
24. Monash University Sunway Campus
25. MyKRIS ASIA KL
26. Net2One Sdn Bhd
27. Netnod
44. Packet-One Networks Sdn Bhd
45. Redtone-CNX Broadband Sdn Bhd
46. Symphonet SB
47. Tg Agas Technology Sdn Bhd
48. TIME DotCom Berhad
49. TM
50. VADS Berhad
51. VTELECOMS BERHAD
52. YTL Communications Sdn Bhd

*Source: peeringdb.com

INTERNATIONAL UNIVERSITY OF MALAYA-WALES

IPv6 Peering in IXP

- Individual IPv4 and IPv6 BGP peering sessions is established between all BGP neighbors
- Since IPv4 and IPv6 are not interoperable, IPv6 topology of the IXP should be in parallel and independent of IPv4.
- Prevents black holing of IPv6 traffic in the event of a protocol outage because the IPv6 session goes down when IPv6 reachability is lost

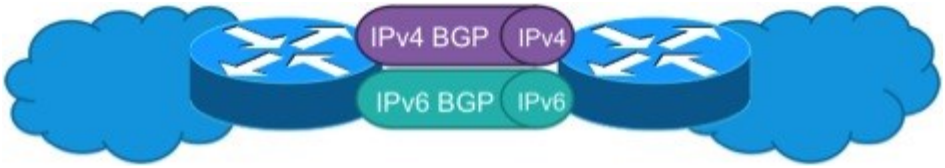

INTERNATIONAL UNIVERSITY OF MALAYA-WALES

IPv6 Deployment in IXP

- Since most Internet Exchange Points (IXPs) work at the Layer 2 level, making the deployment of IPv6 an easy task.

- IXPs normally implement additional services such as statistics gathering, route servers, looking glasses, and broadcast controls that may be impacted by the implementation of IPv6.

Switch Fabric Configuration

- An Ethernet-based IXP switch fabric implements IPv6 over Ethernet.
  - Switching of IPv6 traffic happens in the same way as in IPv4.
  - Functions such as SNMP, or flow analysis may require IPv6 as an underlying layer.
- There are two common configurations of IXP switch ports to support IPv6:
  - Dual-stack LAN
    - When both IPv4 and IPv6 traffic share a common LAN. No extra configuration is required in the switch.
  - Independent VLAN
    - When an IXP logically separates IPv4 and IPv6 traffic in different VLANs.

IXP IPv6 Addressing Plan

- RIRs have specific address policies to assign Provider Independent (PI) IPv6 addresses to IXPs.
  - Usually /48 or shorter prefixes but may vary depending on the IXP.
  - MyIX IPv6 Address Block is 2001:de8:10::/112.
- Depending on the country and region of operation, address assignments may be made by NIRs.
- IPv6 ULAs are normally not used in an IXP LAN.
- IXPs normally use manual address configuration
  - allows IXP participants to replace network interfaces with no need to reconfigure BGP sessions' information, and it also facilitates management tasks.
- When selecting the use of static Interface Identifiers (IIDs), there are different options on how to fill its 64 bits.

RIR IXP IPv6 Allocation Policies

RIR	Category	Policy	Comment
AFRINIC	Size	/48 minimum.	Part of the "Provider Independent (PI) Assignment for End-Sites" policy
	Eligibility	– Minimum number of three peers connected – Open policy for anyone to connect/peer.	
APNIC	Size	/48 minimum.	
	Eligibility	APNIC members with IPv4 resources assigned under the IPv4 IXP policy, but with no IPv6 resources, automatically qualify for an IPv6 /48.Members that do not hold an IPv4 critical infrastructure assignment from APNIC, that have existing IPv6 resources, or that wish to request more than /48 should meet the following requirement: The IXP must have a clear and open policy for others to join and must have at least three members.	
LACNIC	Size	/48 minimum, /32 maximum.	
	Eligibility	The IXP must have a clear and open policy for others to join and must have at least three members. It must also provide documentation showing that it is an IXP, list of participants, structure diagram, numbering plan and a utilization plan for the following three and six months.	
ARIN	Size	/48 minimum.	
	Eligibility	Exchange point operators must provide justification for the allocation, including: connection policy, location, other participants (minimum of two total), ASN, and contact information.	
RIPE NCC	Size	/64 or /48.	
	Eligibility	The IXP must have a clear and open policy for others to join and must have at least three members.	

Static IID Selection Method

1. Some IXPs like to include the decimal encoding of each participant's ASN (Autonomous System Number) inside its correspondent IPv6 address:

   - IXP LAN prefix: 2001:db8::/64
   - ASN: 64496
   - IPv6 Address: 2001:db8:0000:0000:0000:0006:4496:0001/64 or its equivalent representation 2001:db8::6:4496:1/64

Static IID Selection Method

2. Although BCD encoding is more "human-readable", some IXPs prefer to use the hexadecimal encoding of the ASNs number as the upper part of the IID:

   - IXP LAN prefix: 2001:db8::/64
   - ASN: 64496 (DEC) or fbf0 (HEX)
   - IPv6 Address: 2001:db8:0000:0000:0000:0000:fbf0:0001/64 or its equivalent representation 2001:db8::fbf0:1/64

Static IID Selection Method

3. A third scheme for statically assigning IPv6 addresses on an IXP LAN could be to relate some portions of a participant's IPv6 address to its IPv4 address:

   - IXP LAN prefix: 2001:db8::/64
   - IPv4 Address: 192.0.2.123/23
   - IPv6 Address: 2001:db8:2::123/64

Static IID Selection Method

4. A fourth approach might be based on the IXPs ID for that participant.
5. These four methods are non-exhaustive and the IXP may decide to employ other mechanisms for selecting their IID.

IXP IPv6 Multicast

- IXPs typically control broadcast traffic across the switching fabric in order to avoid broadcast storms by only allowing limited ARP.
- IPv6 does not support broadcast, but IXPs can control multicast traffic in each LAN instead using ICMPv6 Neighbor Discovery.
- IPv6 Multicast traffic exchange, an IXP can use
  - Same LAN being used for unicast IPv6 traffic exchange
  - Same LAN being used for IPv4 Multicast traffic exchange
  - Dedicated LAN for IPv6 Multicast traffic exchange.

Other IPv6 Elements in IXP

- Reverse DNS
  - If reverse DNS is configured, DNS servers should be reachable over IPv6 transport for complete IPv6 support.
- Route-Server
  - To provide Multi-Lateral Peering Agreements (MLPA), looking-glass, or route-collection service, the equipment should be able to transport IPv6 traffic and to support MP-BGP extensions for IPv6 address family
- External and Internal Support
  - Each service, e.g. web, that is currently accessed through IPv4 or that handle IPv4 addresses should be evaluated for IPv6 support.

Other IPv6 Elements in IXP

- IXP Policies and IPv6
  - IXP policies and contracts should be revised as any mention of IP should be clarified if it refers to IPv4, IPv6, or both.
- Security Considerations
  - IPv6 implementation should be scrutinized to ensure there is no security vulnerabilities such as bogus addresses.
  - Limit IPv6 DoS attacks to the IXP switch fabric by not globally announce the IXP LANs prefix.

IXPs Role in Driving IPv6 Adoption

- There is some growth in IPv6 ASNs since World IPv6 Launch but the growth has not been significant.
- IXPs can play a big role in encouraging and supporting IPv6 connectivity among its peers.
- Peer with other IPv6 networks as much as possible.

	Globally	AFRINIC	APNIC	ARIN	Lacnic	RIPE NCC
World IPv6 Launch	13.7% (5650 out of 41473 ASNs)	11.6% (65 out of 559 ASNs)	18.0% (893 out of 5003 ASNs)	10.3% (1561 out of 15184 ASNs)	15.2% (333 out of 2195 ASNs)	15.0% (2725 out of 18126 ASNs)
Today	16.1% (7168 out of 44470 ASNs)	14.0% (91 out of 651 ASNs)	20.0% (1091 out of 5477 ASNs)	12.4% (1947 out of 15655 ASNs)	15.7% (428 out of 2722 ASNs)	18.1% (3525 out of 19484 ASNs)

INTERNATIONAL UNIVERSITY OF MALAYA-WALES

IPv6 Human Resource Development

- IPv6 is not as simple as it sounds, and is not an extension of taking the 32bit IPv4 address and stretching it to 128 bits.

- So the first step before rolling out or pushing IPv6 into a organisation would be the need to train the ISP/IXP engineers to be IPv6 qualified

- The Global IPv6 Forum has recognised and certified numerous such programs around the world.

THANK YOU

PROFESSOR DR. SURESWARAN RAMADASS
sures@nltvc.com