IPv6 adoption at IIJ

Matsuzaki ‘maz’ Yoshinobu

<maz@iij.ad.jp>

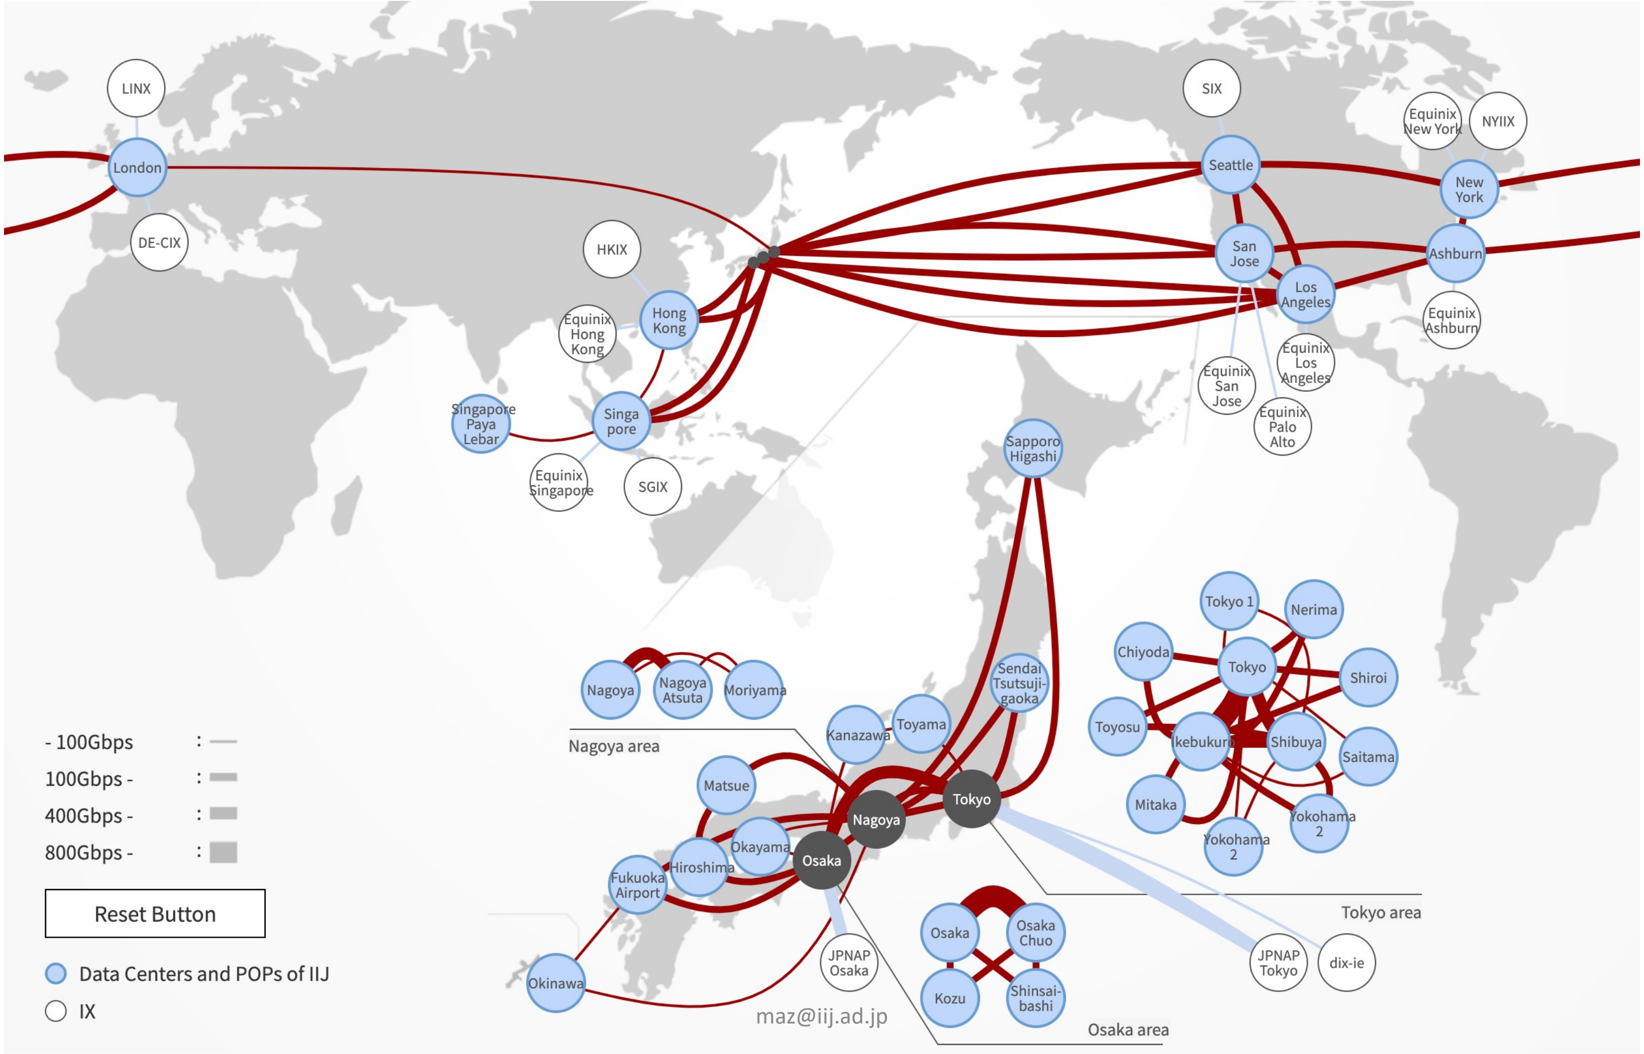
LINX
London
DE-CIX
HKIX
Equinix Hong Kong
Hong Kong
Singapore Paya Lebar
Singa pore
Equinix Singapore
SGIX
SIX
Seattle
Equinix New York
NYIIX
New York
San Jose
Ashburn
Los Angeles
Equinix Ashburn
Equinix Los Angeles
Equinix San Jose
Equinix Palo Alto
Sapporo Higashi
Nagoya
Nagoya Atsuta
Moriyama
Nagoya area
Sendai Tsutsuji-gaoka
Kanazawa
Toyama
Matsue
Tokyo
Nagoya
Okayama
Osaka
Hiroshima
Fukuoka Airport
Okinawa
JPNAP Osaka
Tokyo 1
Nerima
Chiyoda
Tokyo
Shiroi
Toyosu
Ikebukuro
Shibuya
Saitama
Mitaka
Yokohama 2
Yokohama 2
Tokyo area
JPNAP Tokyo
dix-ie
Osaka
Osaka Chuo
Kozu
Shinsai-bashi
Osaka area
- 100Gbps :
100Gbps - :
400Gbps - :
800Gbps - :
Reset Button
Data Centers and POPs of IIJ
IX
maz@iij.ad.jp

History of IIJ’s IPv6 services

- 1999
  - IPv6 over IPv4 tunnel service (experimental)
- 2000
  - IPv6 native service (leased line, experimental)
- 2001
  - IPv4/IPv6 dual stack service (leased line)
  - IPv6 data center service
- 2011
  - IPv6 PPPoE service over NTT FLETS (FTTH)
- 2012
  - IIJmio IPv4/IPv6 dual stack service (mobile)

IPv6 mobile service (MVNO)

- Offering SIM / eSIM
  - MVNO of NTTdocomo and KDDI au
- IP addressing
  - A public IPv6 (/64) + private IPv4 (/32)
  - Depending on users' APN PDP type configuration

Business Model

- IIJ sells bandwidth, or the ability to transfer data
  – Basically, we don't care what is carried, I mean IPv4 or IPv6
- Our customers can select protocol(s), or
  – IPv4 only
  – IPv4/IPv6 dual stack
  – IPv6 only

History of IIJ’s IPv6 backbone

- Initially we started with dedicated backbone
- 1998
  - PC based router(kame stack)
  - Tunnel and Ethernet
- 2000
  - Cisco c72xx series
  - Tunnel, Ethernet and T1 line
- 2005
  - Started to migrate to IPv4/IPv6 dual stack backbone

Dual stack backbone

- All our backbone nodes are IPv6 enabled

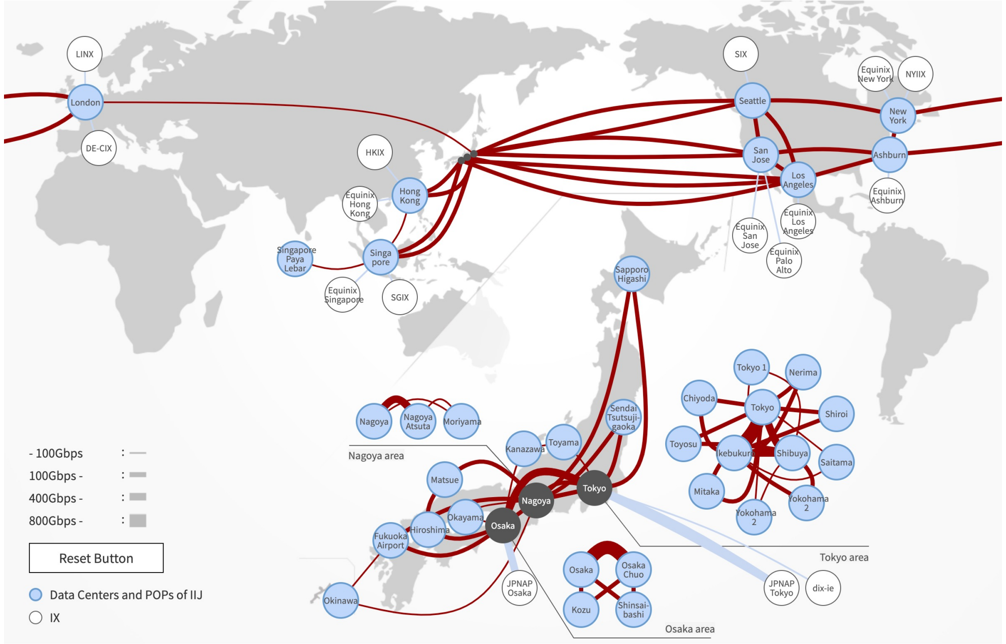

Addressing

- /128 for loopback interfaces
- /64 for links
  - /127 is used on several inter-router links
- Enterprises
  - /48
- Consumers
  - /48, /56
- See also ripe-690

  https://www.ripe.net/publications/docs/ripe-690

IPv6 Interface ID

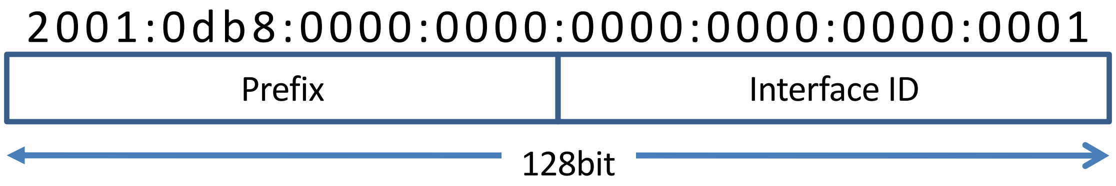

- We assign unique 64bit ID to every router
  - A router uses the ID as its interface id on every interface
  - No correlate with IPv4 address
- Exception
  - Customer facing interface
  - /127

Global Unicast Address of Routers

- Inter-router link does not require public address inside AS
  - OSPFv3 uses link-local address to exchange LSAs
  - Only loopback interface needs to have a public address
- But we configure public address on every interface
  - To be able to ping to check the interface availability

link-local Unicast Address of Routers

- fe80::/64
- AS IS
  - We don't touch
  - Most routers use Modified EUI-64 format address
- A virtual address for vrrp/hsrp is another story
  - For user convenience, we use a simple address like fe80::1

Route aggregation inside AS

- Not so aggressive at this moment
  - We allocate /56 to every POP for links inside POP
  - But currently, # of IPv6 prefix << # of IPv4 prefix

```
IPv6 routing table name is default(0) global scope - 186092 entries
Number of prefixes:
   /8: 1, /16: 1, /19: 1, /20: 14, /21: 3, /22: 7, /23: 7, /24: 31
   /25: 8, /26: 16, /27: 20, /28: 203, /29: 4293, /30: 603, /31: 315, /32: 23047
   /33: 3160, /34: 2709, /35: 1043, /36: 6259, /37: 975, /38: 1755, /39: 1551, /40: 14559
   /41: 940, /42: 3456, /43: 1078, /44: 15929, /45: 1994, /46: 3416, /47: 4046, /48: 86450
   /49: 15, /51: 4, /52: 32, /53: 1, /54: 1, /56: 93, /57: 2, /60: 243
   /61: 1, /64: 2548, /112: 9, /120: 6, /126: 161, /127: 144, /128: 4942
```

routing protocols

IPv4

- OSPFv2
  - Mostly area 0
  - Max-metric on startup
- BGP4
  - Peer over IPv4
  - Route-Reflector

IPv6

- OSPFv3
  - Area 0 only
  - Max-metric on startup
- BGP4+
  - Peer over ipv6 global
  - Route-Reflector (the same as IPv4)

Router ID

- Routing protocols usually require 32bit ID
  – Even a routing protocol for IPv6
  – We use IPv4 address of loopback interface as its router ID
- Every routers has IPv4 address in our network

OSPFv3 link cost

- We set the same link cost value as IPv4's.
  - The network topology is almost same.
  - Working fine ☺
- When we had to use RIPng as IGP (no choice at that time), these were so much trouble.

BGP4+ nexthop attribute

- AS IS
  - We keep the nexthop attribute heard from eBGP peers
  - No nexthop-self
  - The same policy as our IPv4's
- We are importing prefixes of IXs and PNIs into OSPFv3
  - The same policy as our IPv4's

BGP4+ multipath

- Well-known load sharing methods
  - Relax BGP best path selection a bit to allow multiple paths to be used at the same time
  - We have been using IPv4 for a long time and are using IPv6 as well
- Still confusing multipath behavior
  - Actually non-standard
  - Vendor-specific

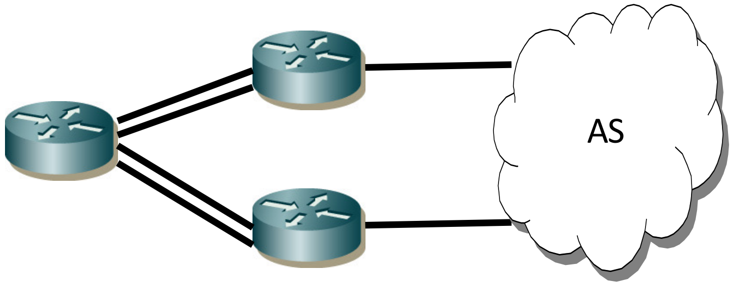

Management

- Remote access to routers' vty
  - IPv4/IPv6 dual stack
  - Permit access only from a stepping host
- Other services
  - Mostly IPv4 only
  - AAA, snmp, syslog, ntp, flow export

Availability Check

- PING!
  - Both IPv4 and IPv6
  - Dual stack routers receive ping twice as much

iij.ad.jp DNS

```
iij.ad.jp.        IN   NS     dns0.iij.ad.jp.
iij.ad.jp.        IN   NS     dns1.iij.ad.jp.

dns0.iij.ad.jp.   IN   A        210.130.0.5
dns0.iij.ad.jp.   IN   AAAA   2001:240::105
dns1.iij.ad.jp.   IN   A        210.130.1.5
dns1.iij.ad.jp    IN   AAAA   2001:240::115
```

iij.ad.jp SMTP

```
iij.ad.jp.          IN   MX  10  omgi.iij.ad.jp.

omgi.iij.ad.jp.     IN   A       202.214.79.36
omgi.iij.ad.jp.     IN   A       202.32.225.116
omgi.iij.ad.jp.     IN   AAAA  2001:240:bb46:8143::64
omgi.iij.ad.jp.     IN   AAAA  2001:240:bb5f:86c::1:64
```

iij.ad.jp WEB

```
www.iij.ad.jp.    IN    A       202.232.2.180
www.iij.ad.jp.    IN    AAAA    2001:240:bb81::10:180
```

IPv6 related Publication

- IIJ had published our deployment schedule on its web site
  - http://www.iij.ad.jp/service/IPv6schedule/ (not exists anymore)
  - Status and schedule of IPv6 adoption for each service
- This helped our customers to plan their IPv6 adoption

Japanese Broadband Subscribers

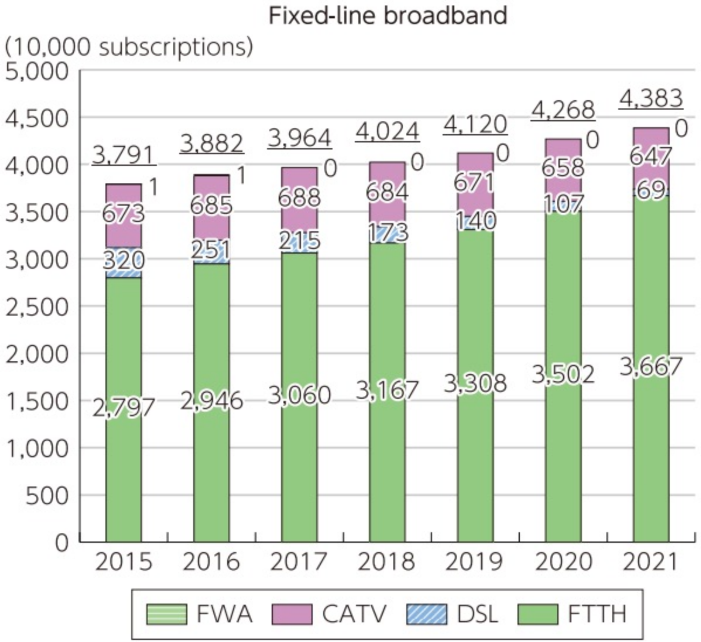

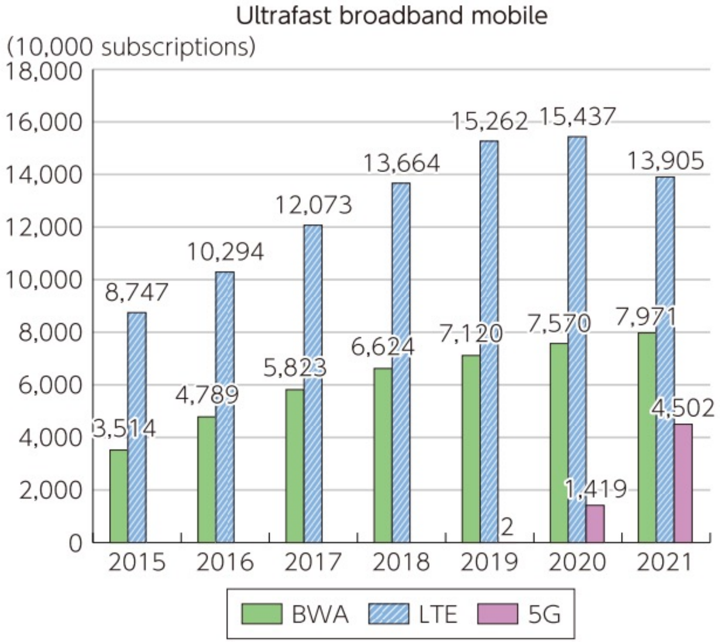

https://www.soumu.go.jp/johotsusintokei/whitepaper/eng/WP2022/2022-index.html

- Mostly FTTH
- Rapid increase of 5G subscribers

Measurements of IPv6 adoption

- Cooperate with providers to publish data
  - TF was formed by telecom operators, government, and academia to address IPv4 address exhaustion
- Encouraged each operator to adopt IPv6