Internet Exchanges in Asia Pacific

Katsuyasu Toyama

JPNAP / APIX association

Who is Katsuyasu Toyama?

- Chief Operating Officer (2015-)
- Director, General Manager of Technology Dept. (2007- 2015)

– Asia Pacific Internet Exchange (APIX) association
  - Chairperson

– IX Federation
  - Board member

– Global Peering Forum (GPF)
  - Board member

SVP of Internet and Mobile Services, NTT Communications (2012-2015)

Outline

- APIX association
- IXP Comparison in Asia Pacific, and Global
- IXP and Peering in Asia Pacific
- “Too much” quality of IXP operation in Japan?

APIX association

What is APIX association?

- “Asia-Pacific Internet Exchange”
  - An association of Internet Exchange Providers in Asia-Pacific region
  - Established in 2010, supported by APNIC
  - Two meetings per year, next meeting is 14th.
  - Objectives:
    To share information about
    technical, operational, and business issues and solutions regarding Internet Exchange

5 IXP associations in the world

Reference: Euro-IX Update at Euro-IX forum

IX-F

- “The Internet Exchange Federation (IX-F) acts as a platform for affiliated Internet Exchange Point Association (IXPAs).
  We are working together to build a global IXP community and help the development of IXPs throughout the world.”

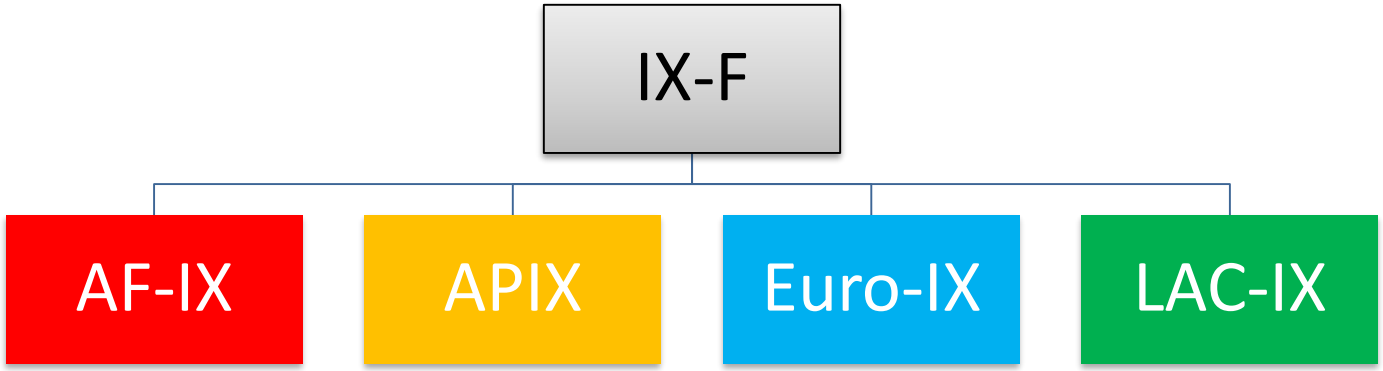

APIX members

- Membership Criteria
  - *Any entity, enterprise, or organization in the Asia-Pacific region, possessing legal personality under the laws of a country or economy where the entity resides, providing services regarding* ***Layer 2*** *Internet Exchange Points, sustaining* ***a neutral position****, and does* ***not enforce the purchase of transit services****. The Asia-Pacific region is defined as the economies where APNIC serves.*
    (from APIX bylaws)

APIX Members

- **23 IXPs, from 16 countries and economies**
  - Australia, Bangladesh, Cambodia, China, Hong Kong, India, Indonesia, Japan, Korea, Malaysia, Nepal, New Zealand, Singapore, Taiwan, Thailand, Vietnam

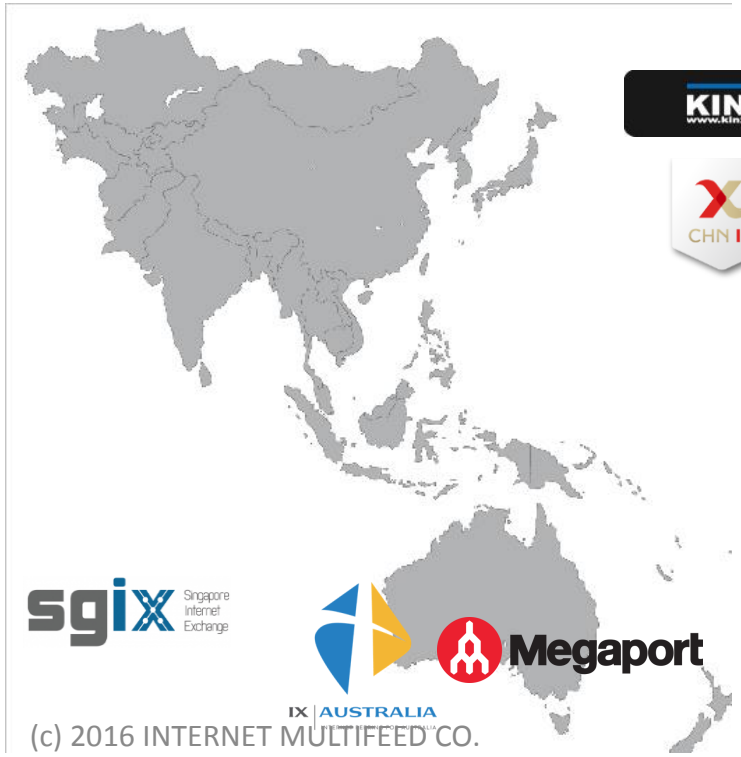

TPIX

CNX SOX VNIX IIX NZIX

APIX members

23 IXPs (3 members have IXPs in Australia)

- AMS-IX Hong Kong (HK)
- BBIX (JP)
- BDIX (BD)
- BKNIX (TH)
- CHN-IX (CN)
- CNX (KH)
- DIX-IE (JP)
- **Equinix** (HK, JP, SG, **AU**)
- HKIX (HK)
- IIX (ID)
- **IX-Australia** (**AU**, NZ)
- JPIX (JP)
- JPNAP (JP)
- KINX (KR)
- **Megaport** (AU, HK, SG, NZ)
- MyIX (MY)
- NIXI (IN)
- NP-IX (NP)
- NZIX (NZ)
- SGIX (SG)
- SOX (SG)
- TPIX (TW)
- VNIX (VN)

Topics discussed at previous APIX meeting

- **Technical session**
  - PeeringDB update (PeeringDB2.0)
  - Peering and interconnection market in Philippine
  - Netflix strategy for global expansion
  - Yet another route server software (GoBGP)

- **Administrative session**
  - Membership fee (just decided 100 USD per year per member IXP)
  - Election of the Steering Committee

Steering Committee

- Steering Committee members
  - Chairperson: Katsuyasu Toyama –JPNAP (JP)
  - Secretary: Ajai Kumar --NIXI (IN)
  - Treasurer: Gavin Tweedie -- Megaport (AU/HK/SG/NZ)
  - Meeting Arrangement: Che-Hoo Cheng --HKIX (HK)
  - Website Management: Brian Kim –Twitch (KR)

Katsuyasu Che-Hoo Gavin

Ajai

Brian

apix
The Asia Pacific Internet Exchange Association

IXP Comparison in Asia Pacific, and Global

Key Performance Indicators for IXPs in APAC region

1. **Number of ASN**
   – How many networks you can potentially peer there
2. **Peak traffic**
   – How much traffic you can potentially distribute there
3. **Average traffic per ASN**
   *= (**peak traffic**)/(**number of ASN**)*
   – How big eyeball and content providers gather there, or
   – How much broadband internet is spread in the market

Data from:

1. Number of ASN
   - **PeeringDB**
     - except for MSK-IX (Russia) and IX.br (Brazil), which seem too few in PeeringDB
2. Peak traffic
   - **Their web page on traffic stats**
     - Unfortunately some IXPs do not disclose data. They are excluded.
3. Average traffic per ASN

   = (peak traffic)/(number of ASN)

Number of ASN

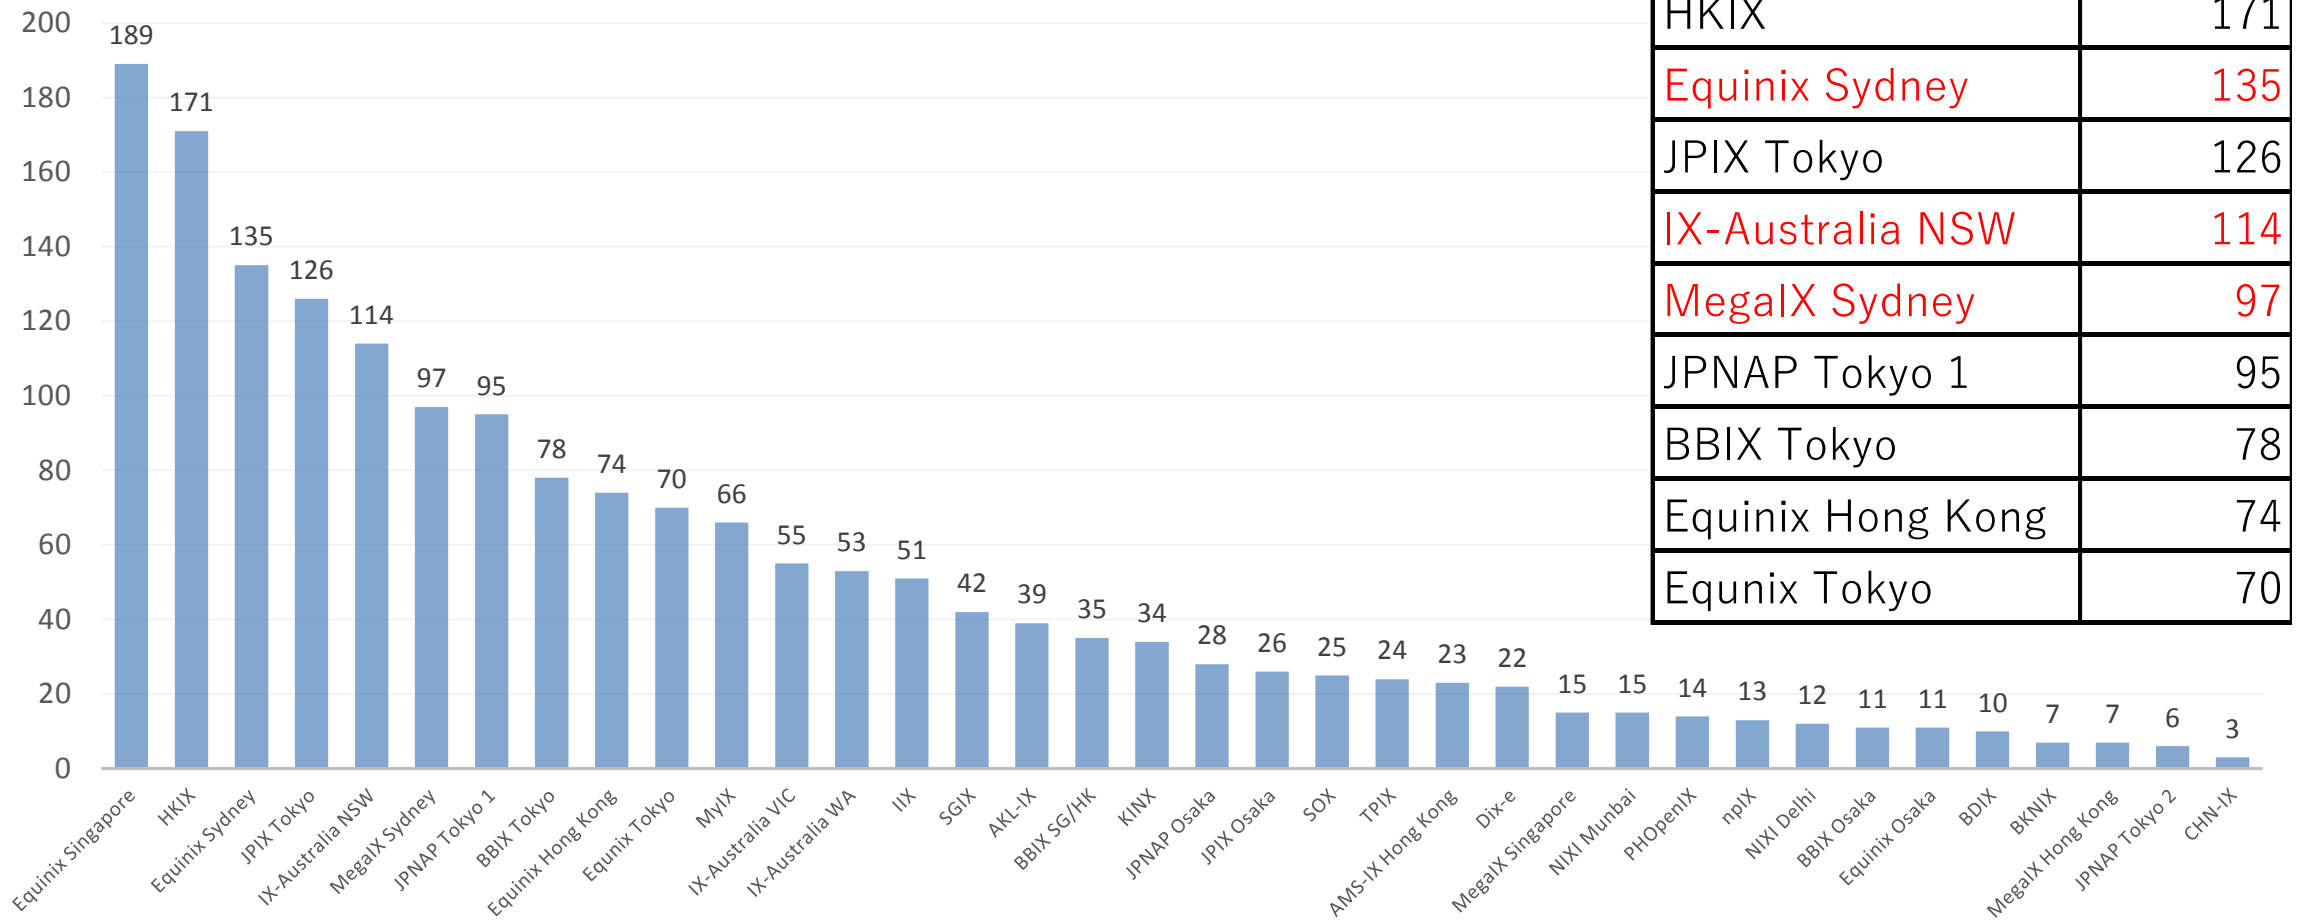

IXP	ASN
Equinix Singapore	189
HKIX	171
Equinix Sydney	135
JPIX Tokyo	126
IX-Australia NSW	114
MegaIX Sydney	97
JPNAP Tokyo 1	95
BBIX Tokyo	78
Equinix Hong Kong	74
Equnix Tokyo	70

Peak traffic

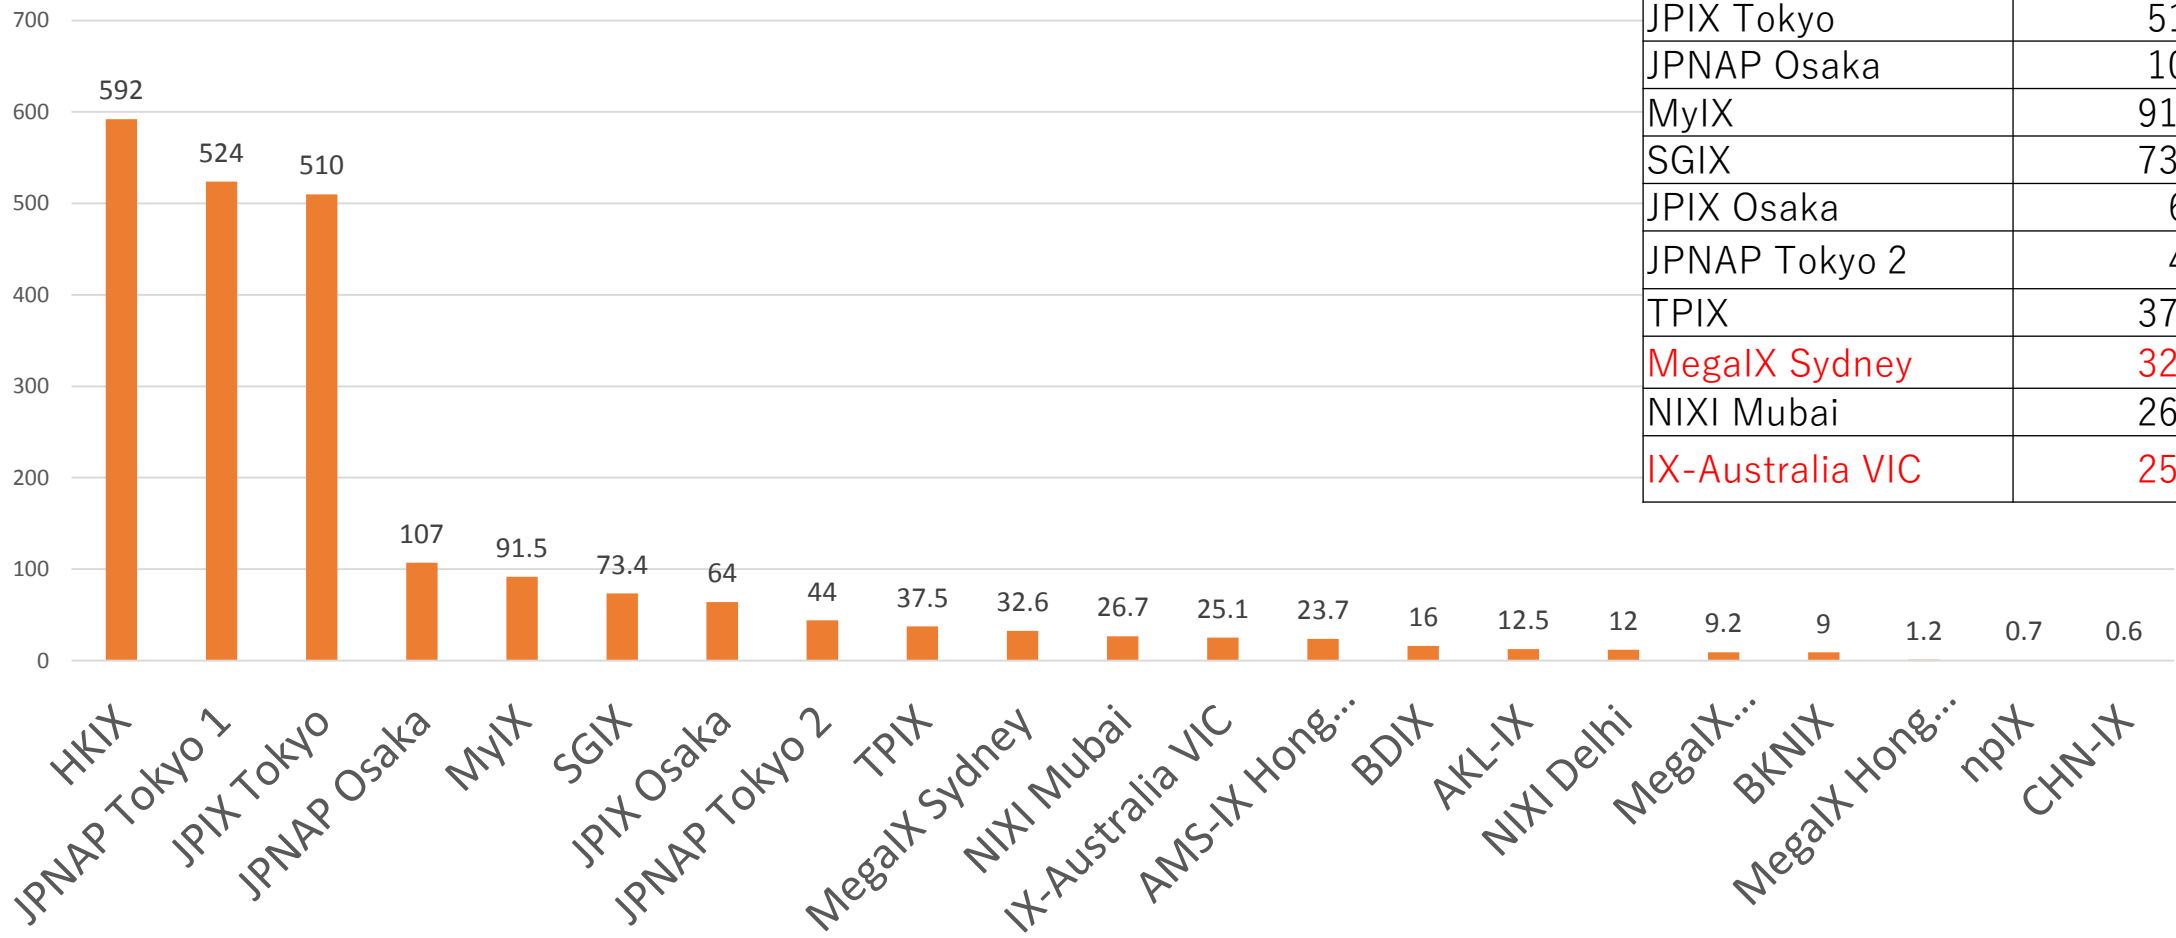

IXP	Peak
HKIX	592
JPNAP Tokyo 1	524
JPIX Tokyo	510
JPNAP Osaka	107
MyIX	91.5
SGIX	73.4
JPIX Osaka	64
JPNAP Tokyo 2	44
TPIX	37.5
MegaIX Sydney	32.6
NIXI Mubai	26.7
IX-Australia VIC	25.1

Average traffic per ASN

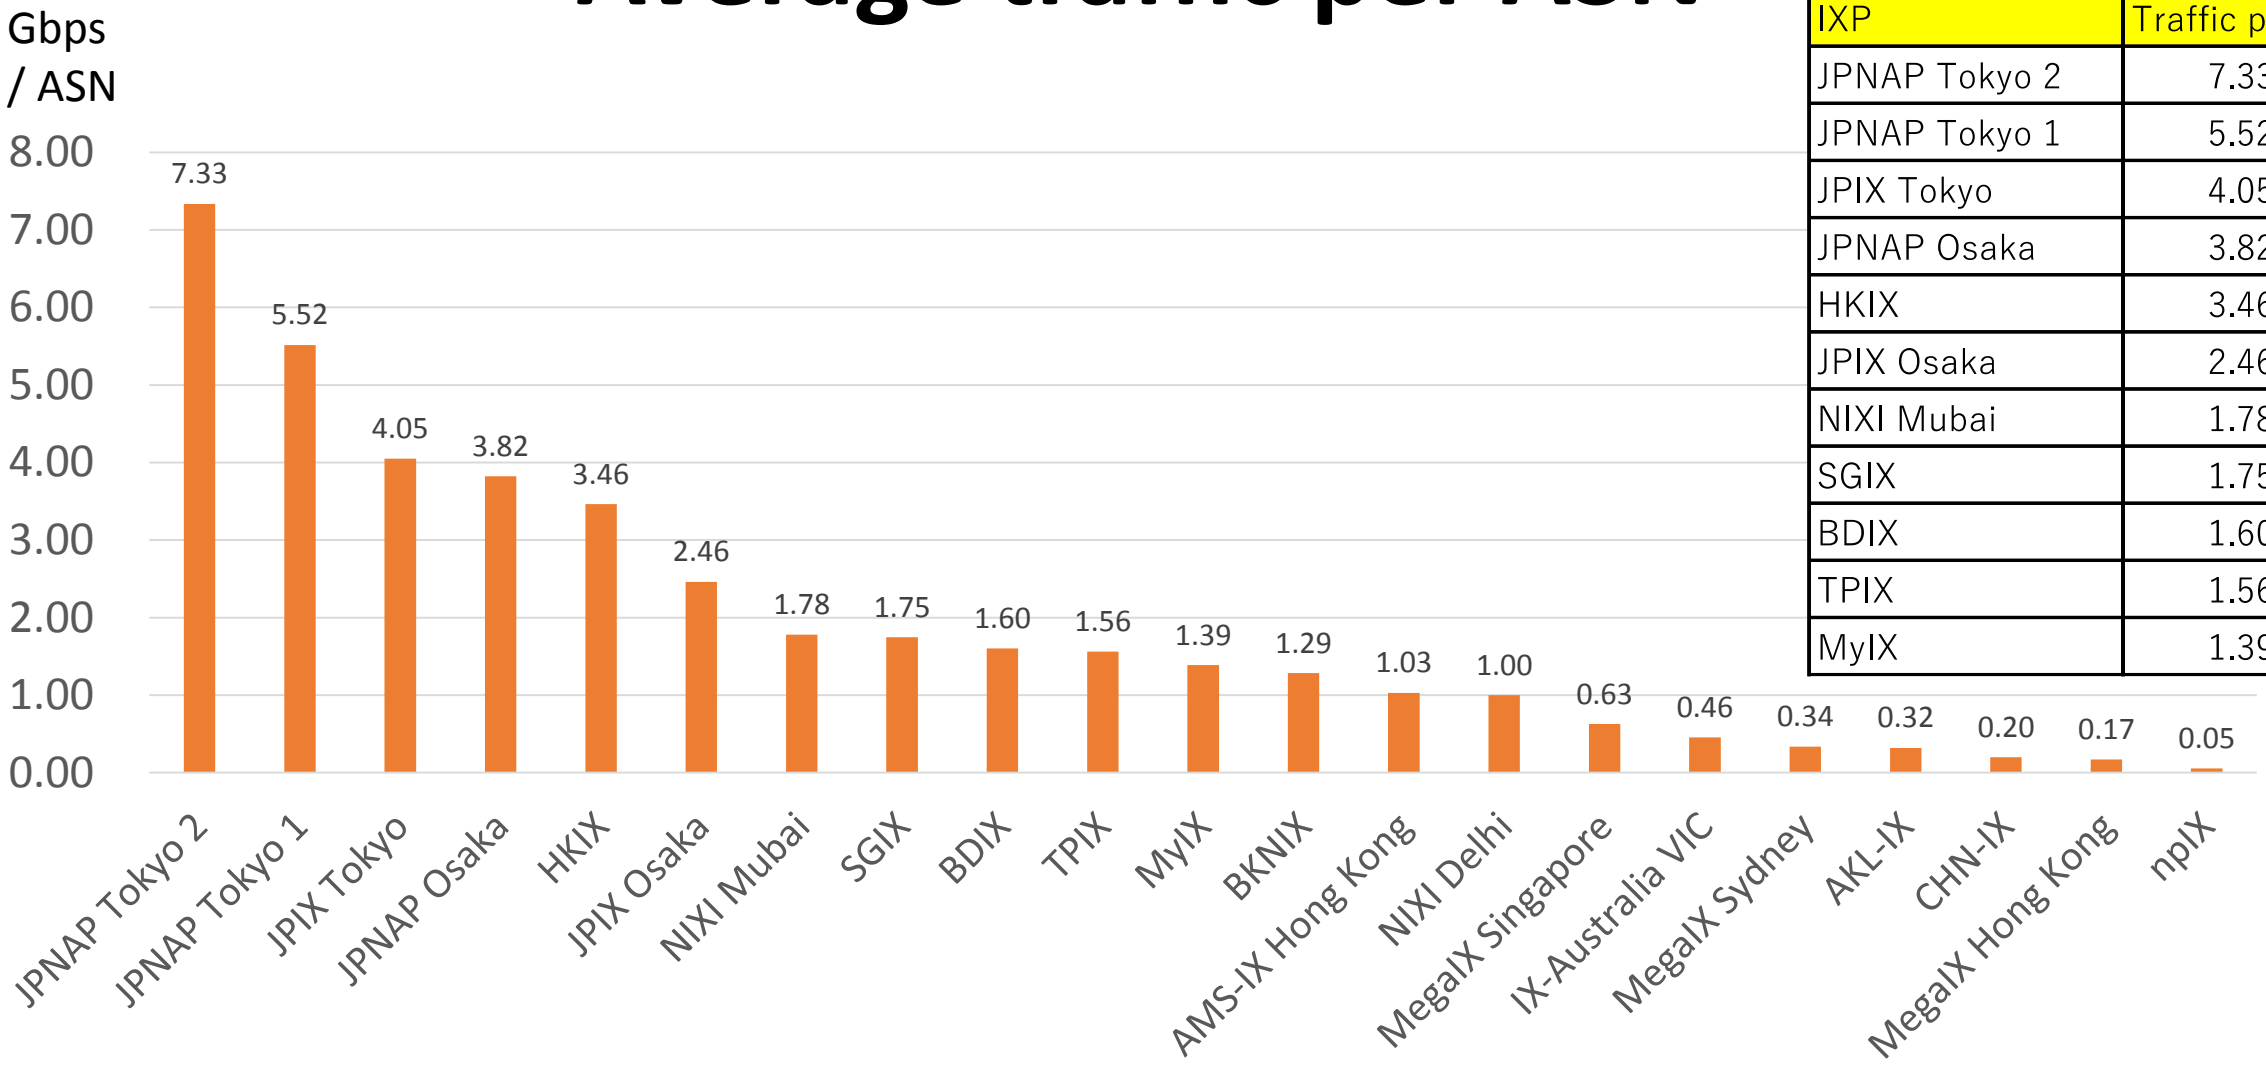

IXP	Traffic per
JPNAP Tokyo 2	7.33
JPNAP Tokyo 1	5.52
JPIX Tokyo	4.05
JPNAP Osaka	3.82
HKIX	3.46
JPIX Osaka	2.46
NIXI Mubai	1.78
SGIX	1.75
BDIX	1.60
TPIX	1.56
MyIX	1.39

Comparing Asia and the rest of the world

- Europe:
  - AMS-IX Amsterdam, Netherland
  - DE-CIX Frankfurt, Germany
  - LINX London, United Kingdom
  - MSK-IX Moscow, Russia
- Latin America:
  - IX.br Sao Paulo, Brazil

Scatter plot: ASN x Peak traffic

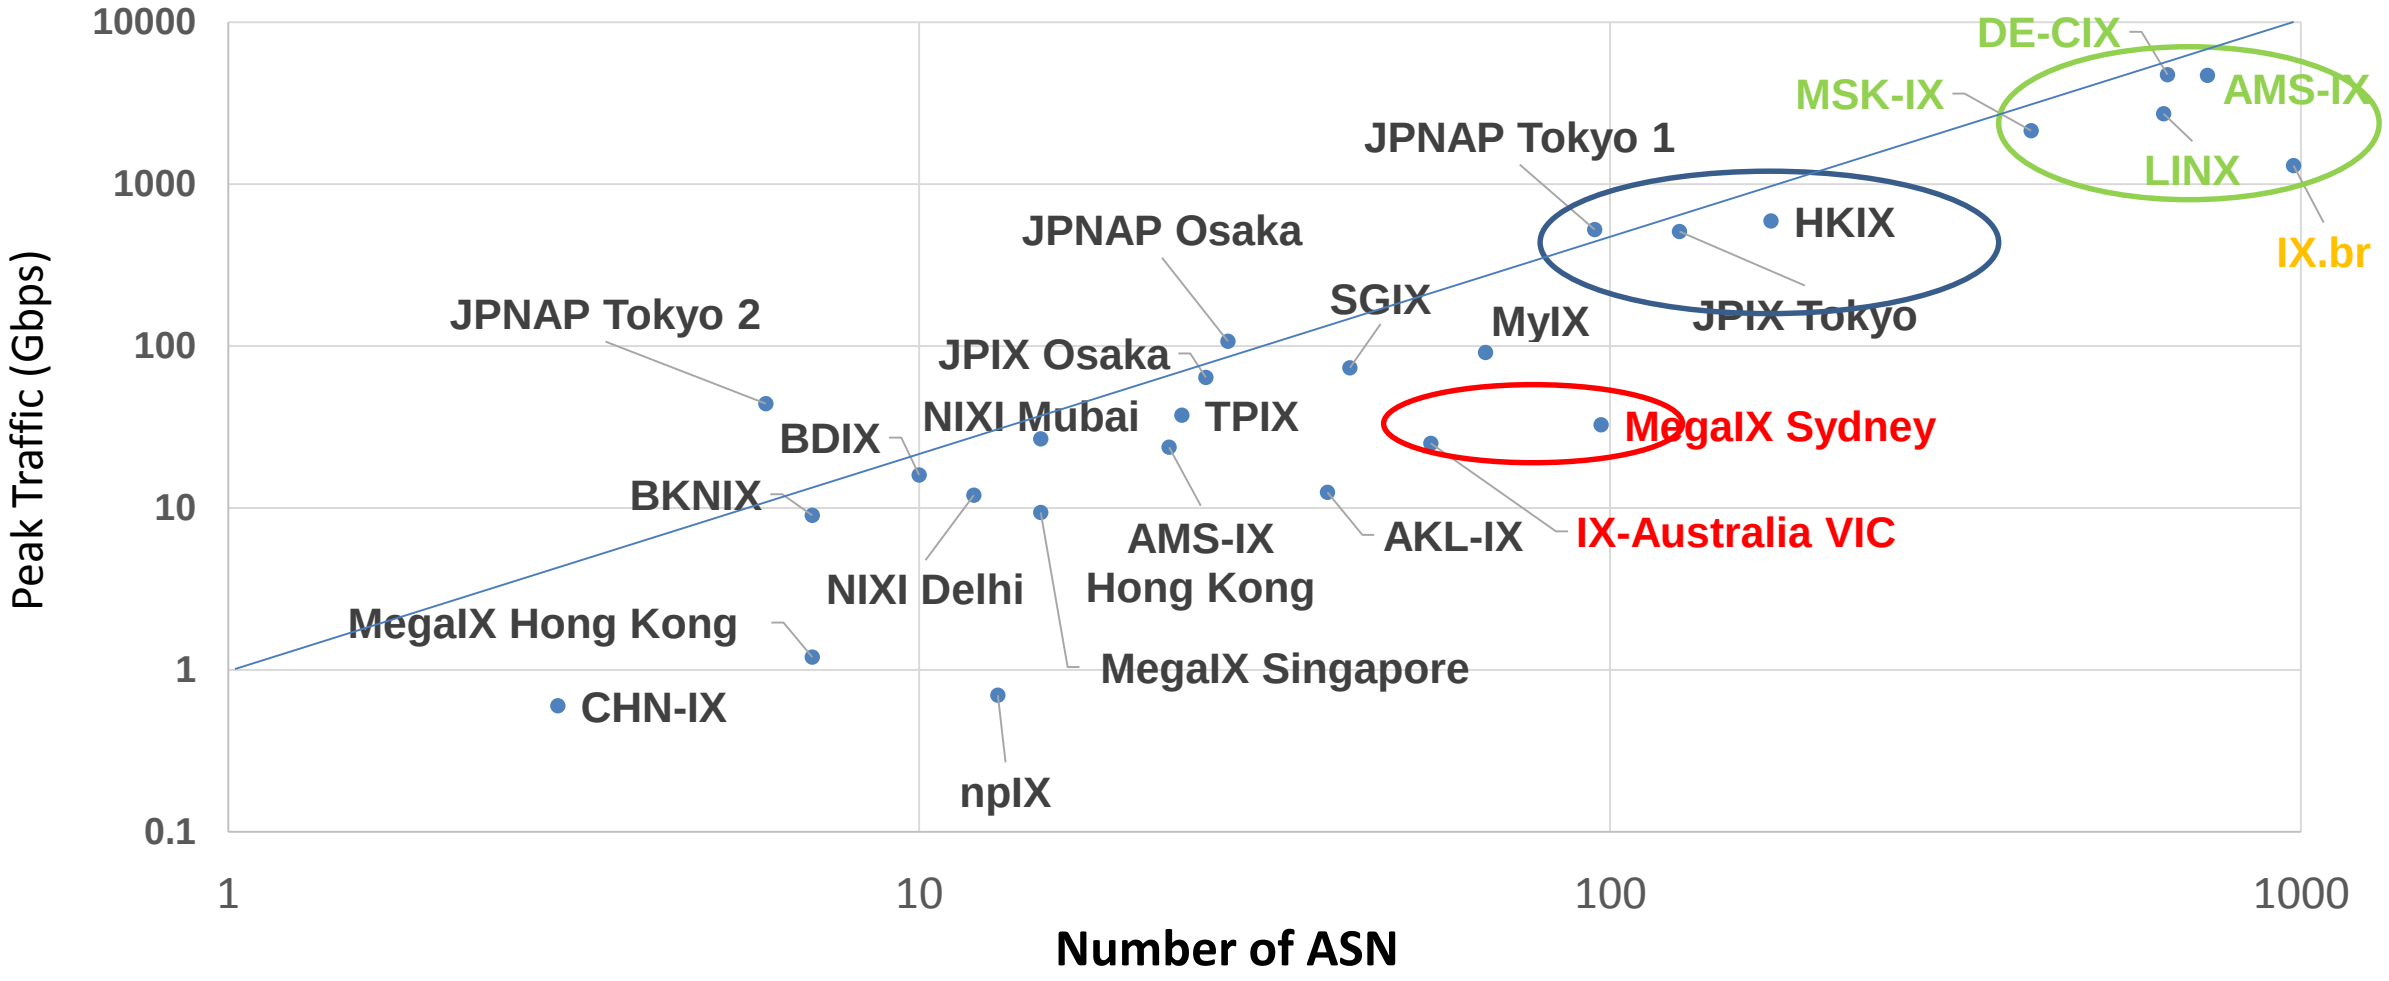

Scatter plot: ASN x Traffic/ASN

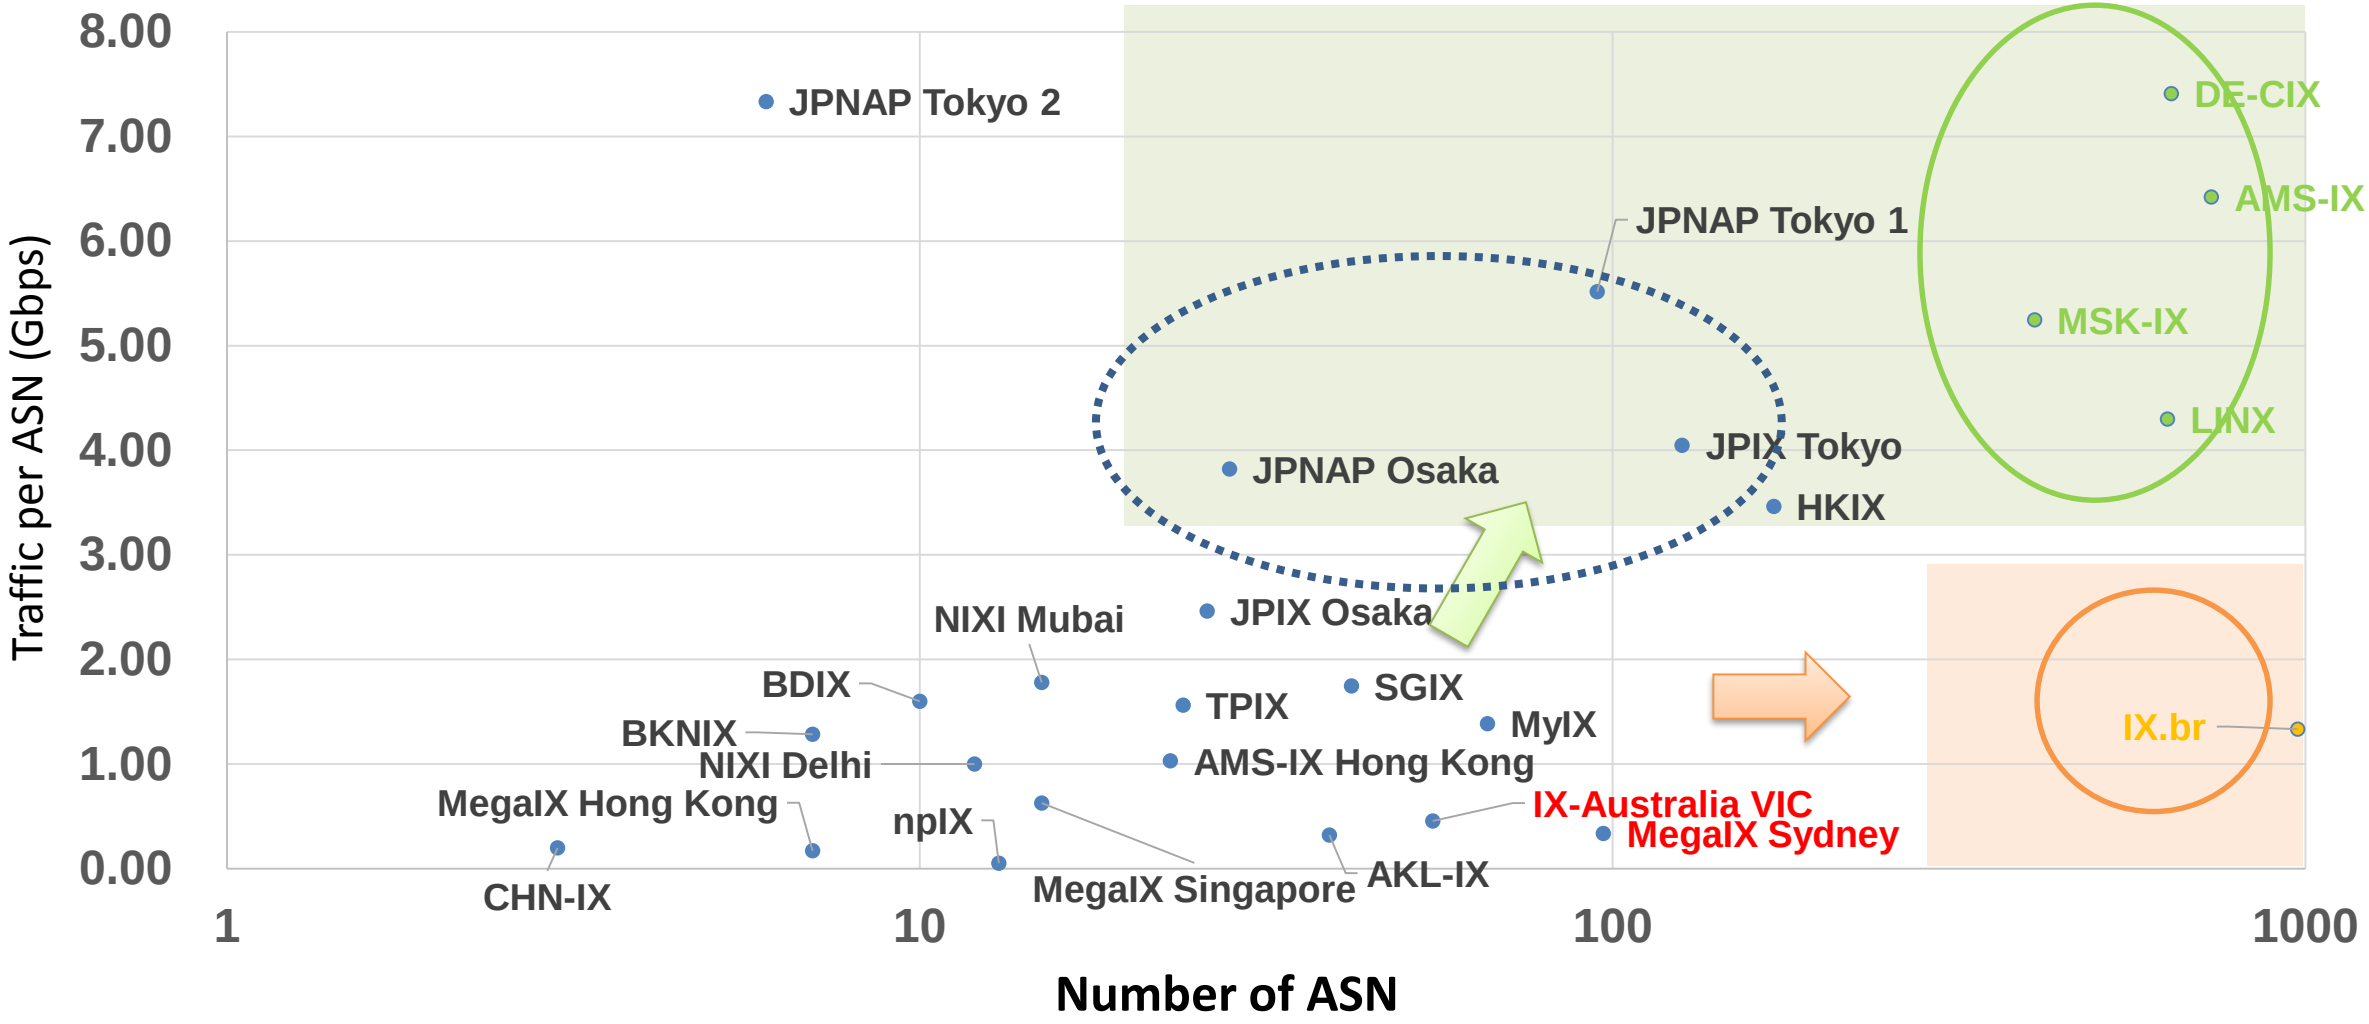

apix
The Asia Pacific Internet Exchange Association

IXP and Peering in Asia Pacific

IXP scene in Asia

- Asia
  - Hong Kong, Singapore are the hubs of south east & south Asia
  - Tokyo has big traffic due to broadband penetration;
    Tokyo is gradually attracting the attentions from Asian ISPs due to Japanese content
  - Sydney, many ASN at IXP but relatively small traffic. Why?

IXP scene in Asia

- Carrier neutral IXP in Asian countries are increasing
  - Thailand (BKNIX) and Mainland China (CHN-IX)

- Remote peering is slowly spread into Asia

BKNIX

Bangkok Neutral Internet Exchange

- Carriers or big ISPs own their “IX”, but they are layer-3 (routers)
- BKNIX is a “neutral” internet exchange in Thailand

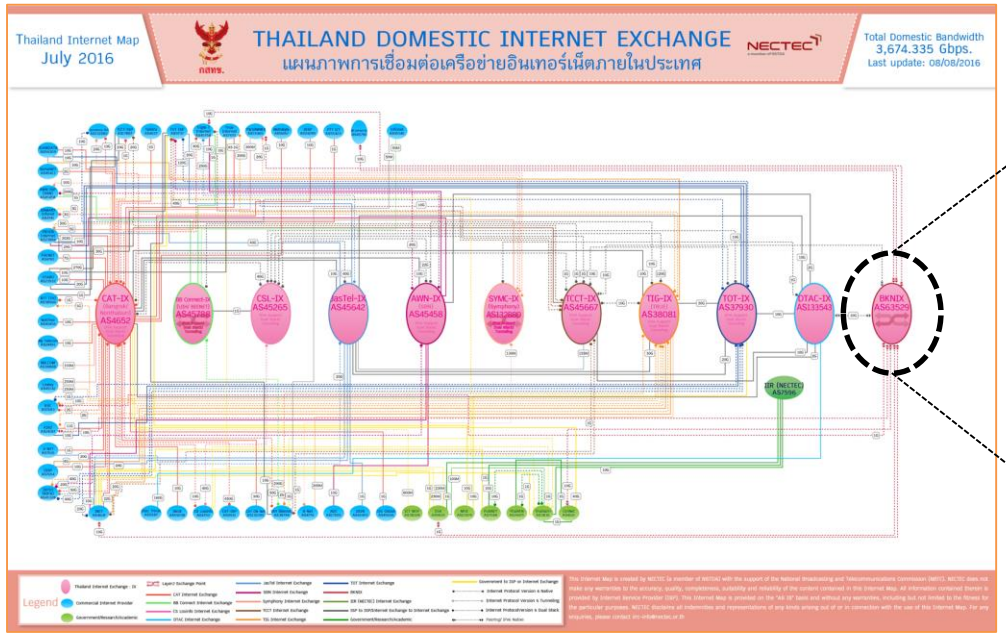

http://internet.nectec.or.th/webstats/internetmap.current.iir?Sec=internetmap_current

https://bknix.co.th/en/index.php?module=service&content=1

CHN-IX

- CHN-IX has been established by ChinaCache's initiative
  - Beijing, Guangzhou, Shanghai
- Initial target is ICP and "Tier-2" ISP

Remote peering

- ISPs extend their network to another country or region, to peer with ISPs there.
  - For instance, eyeball tries to peer with content such as video or game
- Without colos or routers in remote side

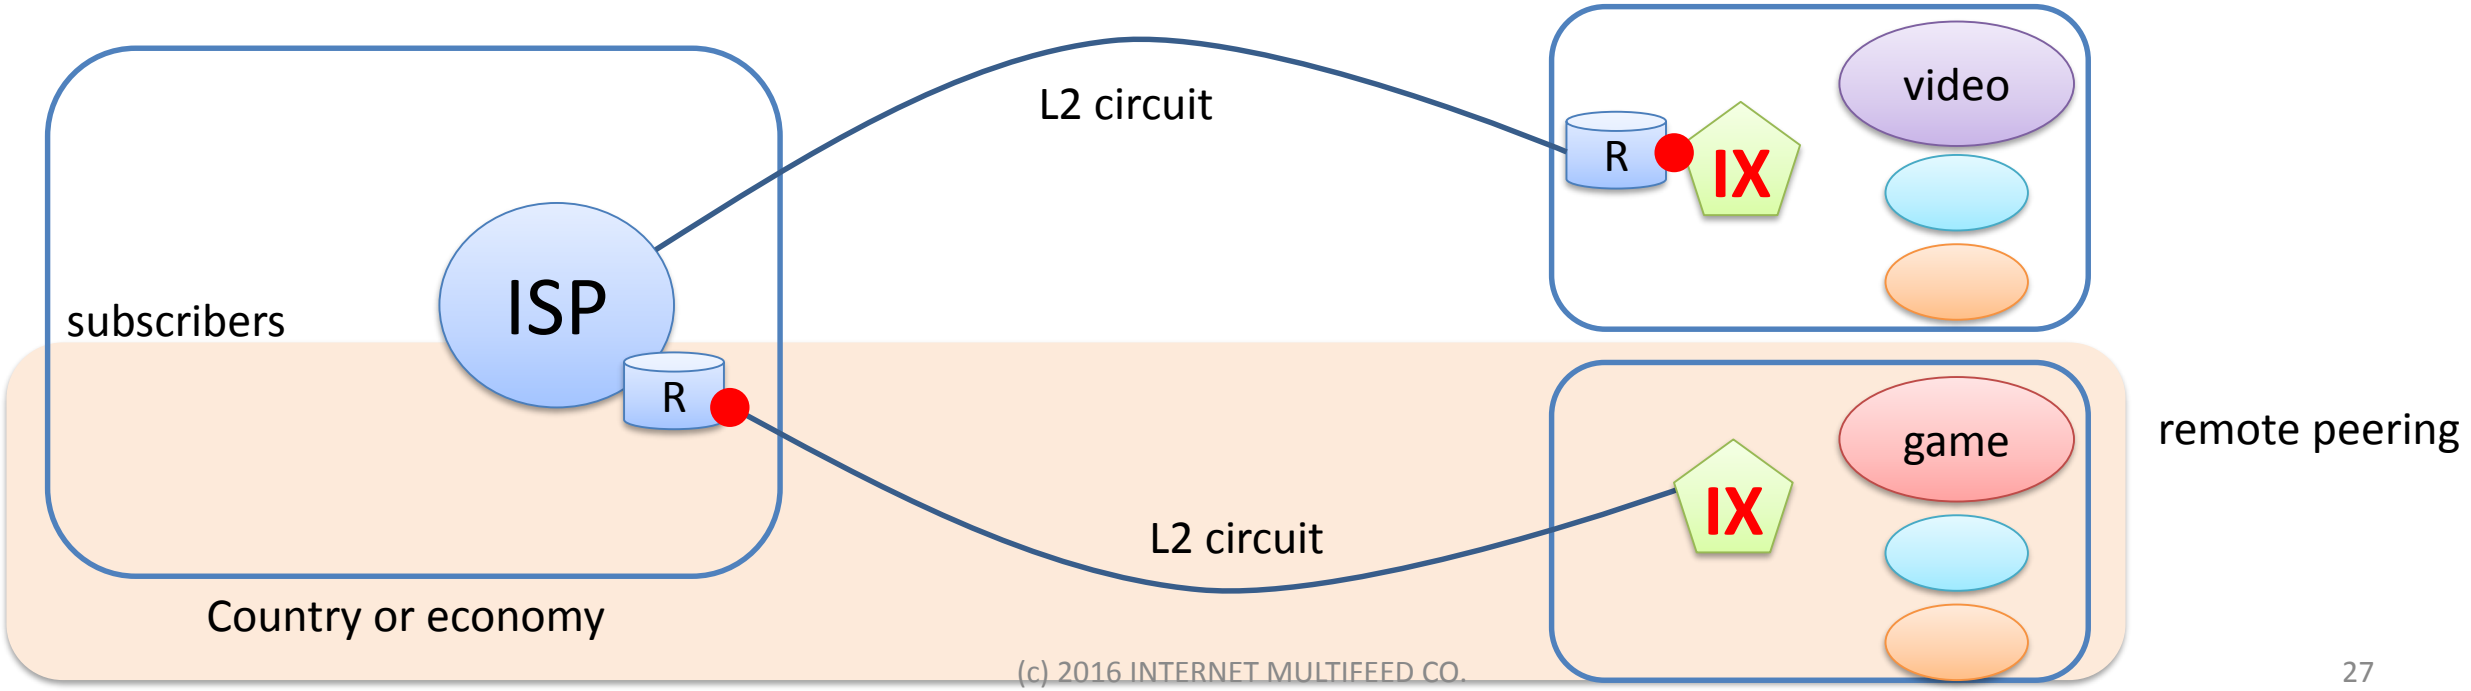

Remote peering

- Motivation or benefits
  - to fulfil subscribers' requests such as latency to servers, not by transit providers but by themselves
- This works fine if:
  - "**transit** fee" > "**circuit** fee + **IX port** fee"

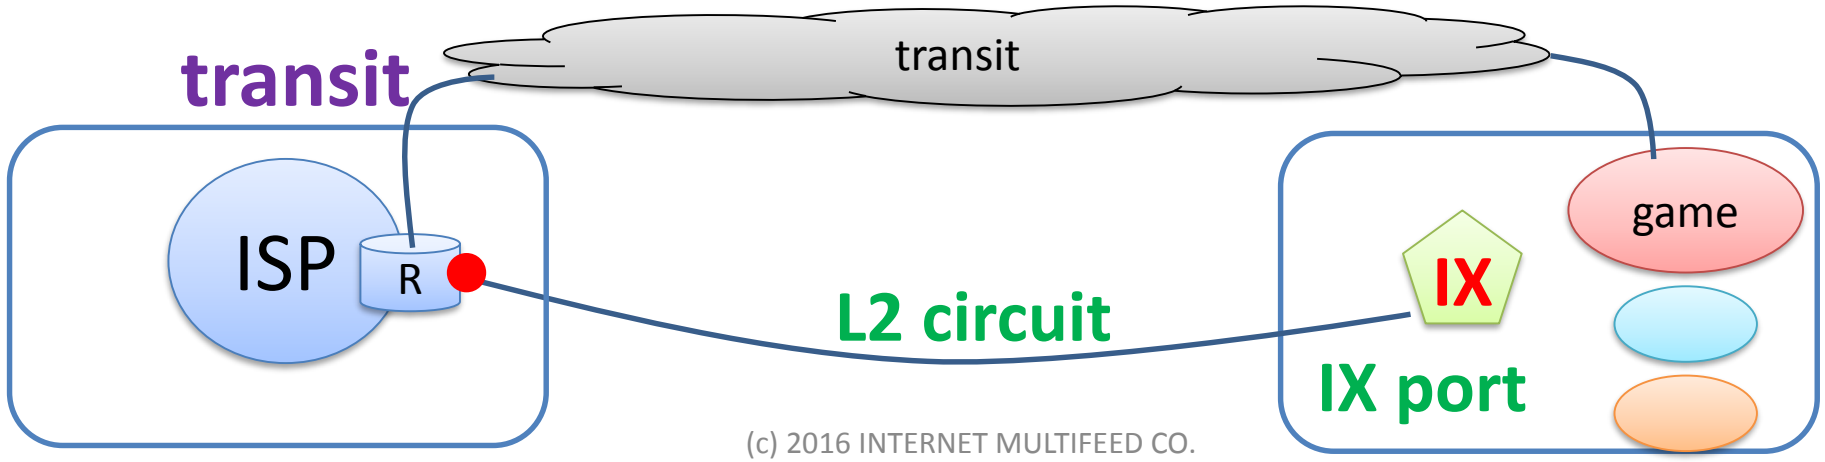

Remote peering

- In Europe,
  "transit fee" > "**circuit fee** + IX port fee" sometimes holds
  - Fibers and L2 circuits are obtained at reasonable price, because land lines can be built much cheaper than sea cables, and also L2 providers compete with each other

- In Asia
  - Gradually L2 connection providers are penetrating, according to price down of sea cables

IXP competition in Japan

- Major IXPs in Japan
  - JPNAP
  - JPIX
  - BBIX
  - EIE (Equinix)

  - They are all commercial, and neutral internet exchanges
  - Neutral, but major carriers are behind them, and only use specific IXPs
  - Content providers connects all the IXPs,
    but small eyeballs connect one or two of them

IXP competition in Japan

	JPIX	JPNAP	BBIX	EIE (Equinix)
	Commercial, Neutral	Commercial, Neutral	Commercial, Neutral	Commercial, Neutral
	KDDI group	NTT/IIJ group	Softbank group	US company
Tokyo				
ASN	126	95	78	70
Traffic	510G	524G	Not Disclosed	Not Disclosed
Osaka				
ASN	26	28	11	11
Traffic	64G	107G	Not Disclosed	Not Disclosed
Features				
	Number of Networks is biggest in Tokyo	Largest Traffic in Tokyo Biggest in Osaka	Also has IXs in Singapore and Hong Kong	US content providers

Peering Trend in Japan
Big domestic eyeballs extend to abroad

- **Domestic eyeballs (not carriers) extend their reach to abroad**
  - Biglobe (AS2518) Singapore, Hong Kong, Los Angeles, San Jose
  - K-opticom(AS17511) Amsterdam, Los Angeles

  - *For bigger eyeballs, peering cost (IX port + Leased Line) becomes less than transit by negotiation.*
- **Remote peering will be more popular in near future**

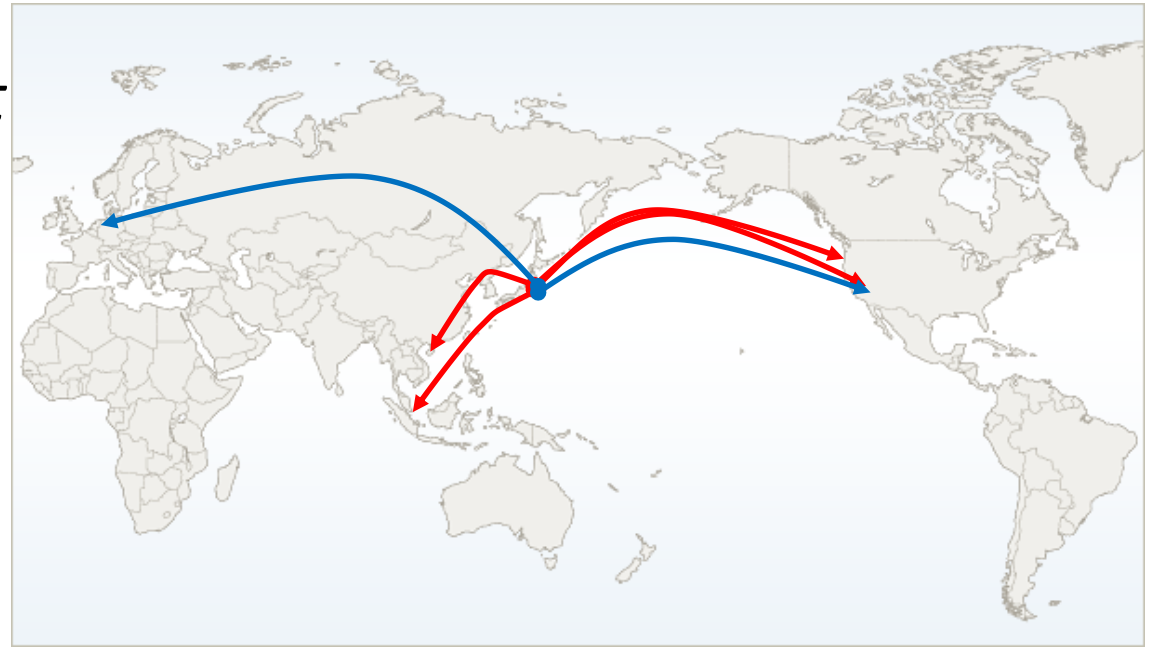

“Too much” quality of IXP operation in Japan?

Internet Operation in Japan

- High reliability and high quality
  - Sometimes good, sometimes bad

- Internet is “best effort” service
  - usually means “do our best as far as we (service provider) can”
  - in Japan sometimes “do your best until we (customer) are satisfied”

Japanese customers: strong request for service quality

- Our customer ISPs, and their customer enterprises are very severe for service quality.

- If a trouble happens and network is down, ISP orders us to report its root causes and their countermeasures immediately.
  - It is strongly requested by their customers (enterprises)

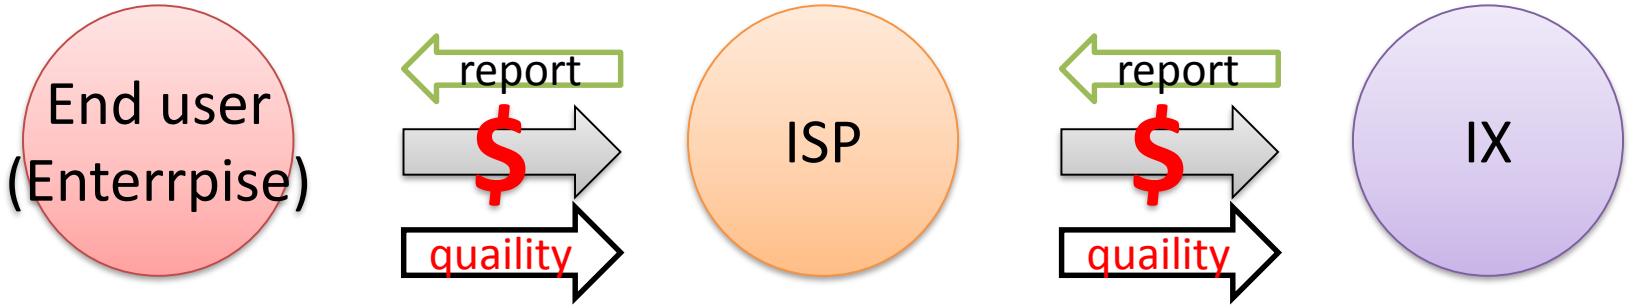

Japanese customers: strong request for service quality

- Operation center of our customer is always watching logs and traps.
- If they notice anomalies in such monitoring data, they will call us to confirm if something bad is happening or not.
  - For instance, if their bgp session is down, they will call us whether the down is caused by IXP network, such as switch down, interface down, or backbone circuit is down.

- Accordingly, JPNAP built a reliable platform using optical switches and redundant ethernet switch network.
  - The result of competition, Japanese IXPs have adopted such redundant network architecture.

Intelligent Optical Switch System

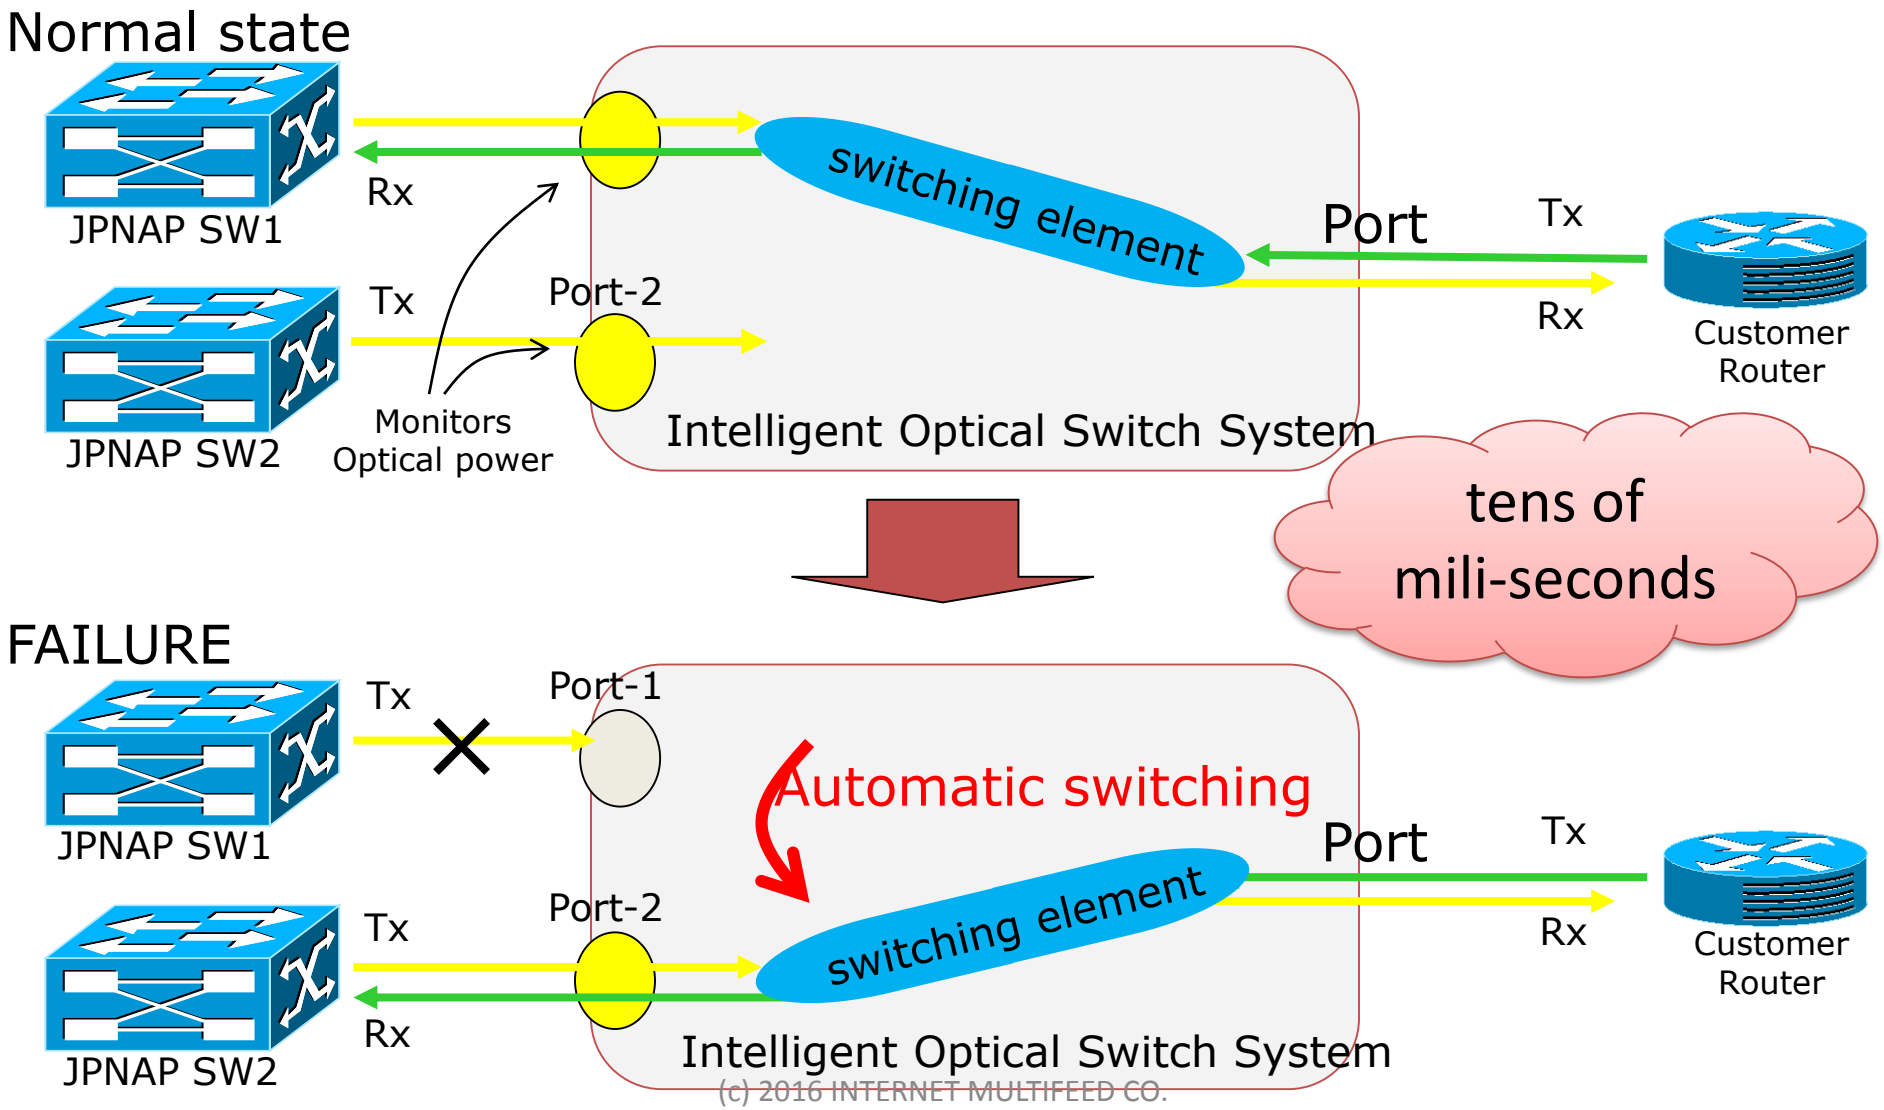

JPNAP: Redundant network architecture

- Backbones and Ethernet switches are redundant
- Optical switch connects Customer ISP/ICP router to two Ethernet switches

- Down time could be minimum

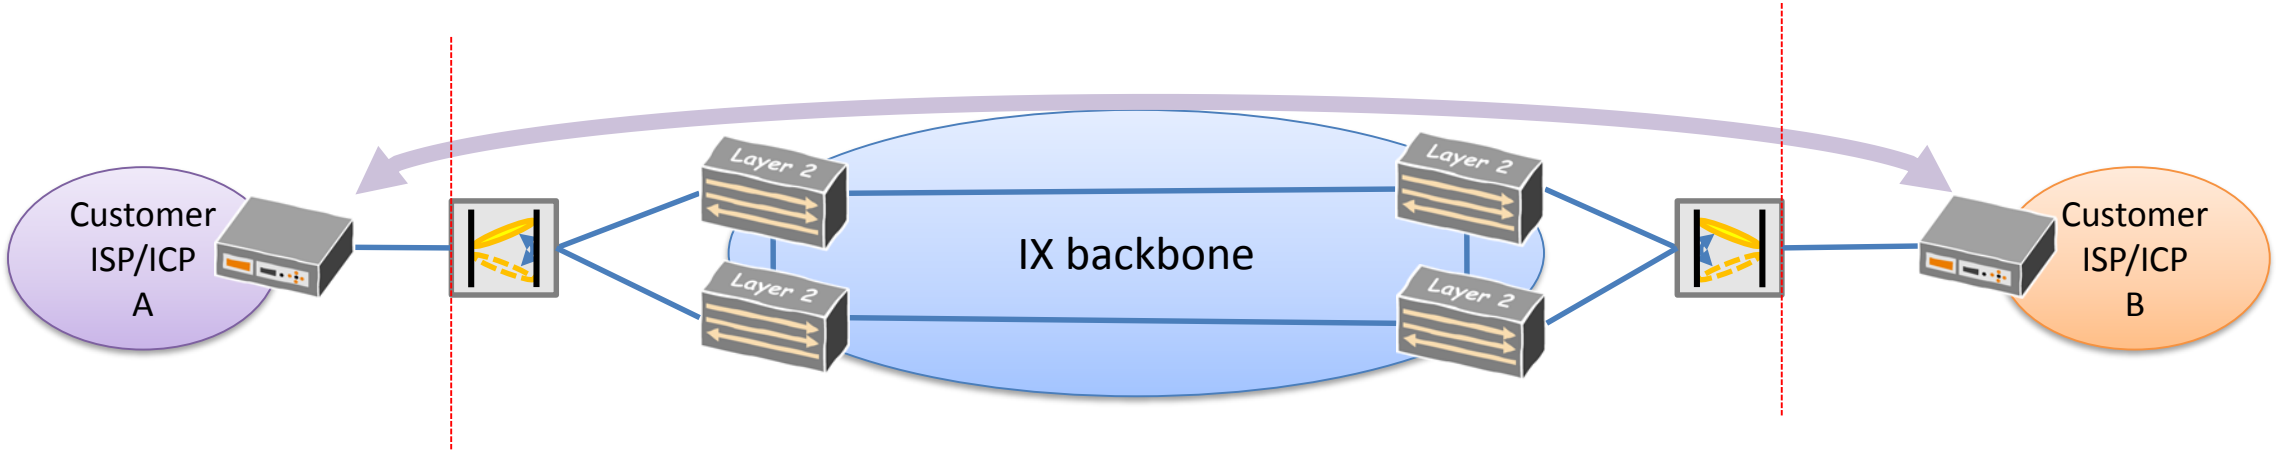

JPNAP: Redundant network architecture

- In case of failure, automatically failover to the backup switch.
- BGP session is kept

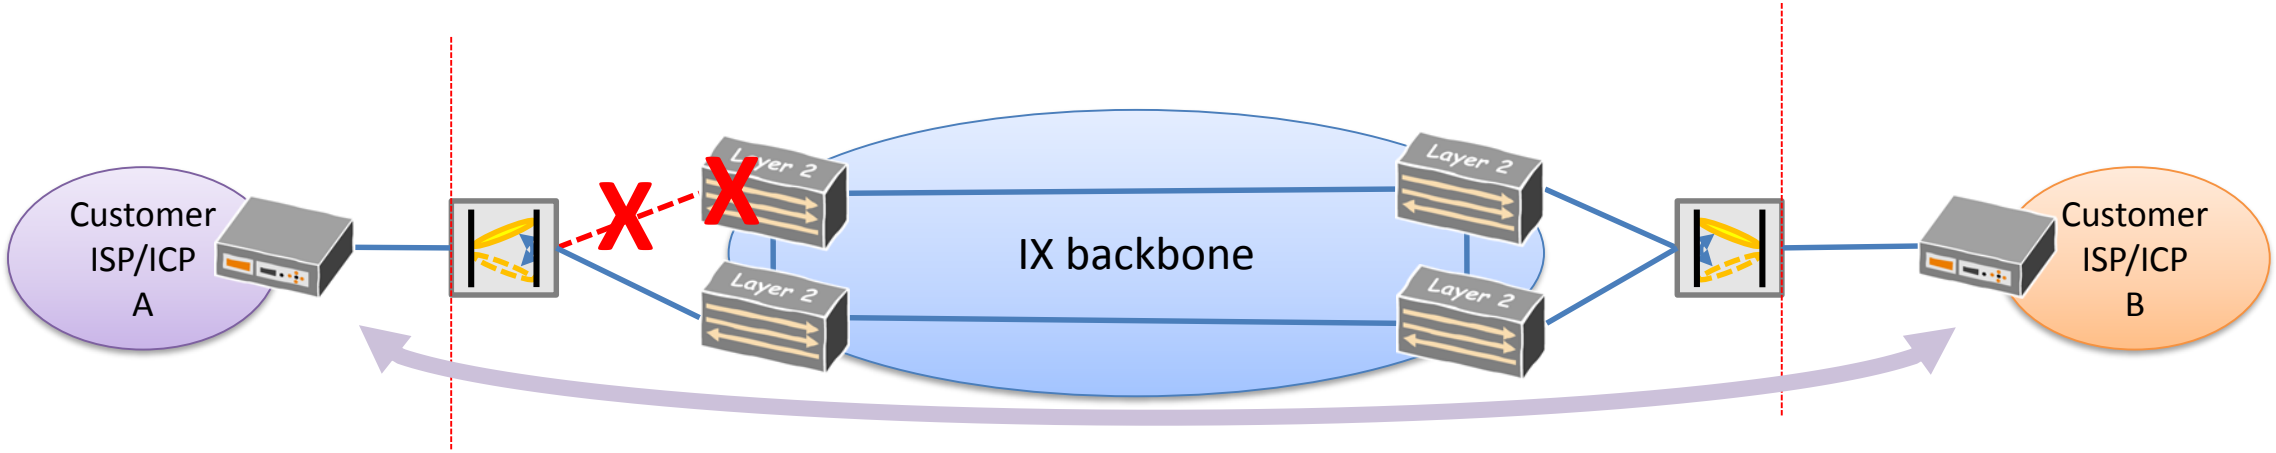

JPNAP: Improving failure detection time

- To detect network down as soon as possible, we monitor;
  - both customers' and ours' interface by Ping
  - Link down / up traps from our ethernet switches
  - bgp sessions between customers' bgp speakers and ours

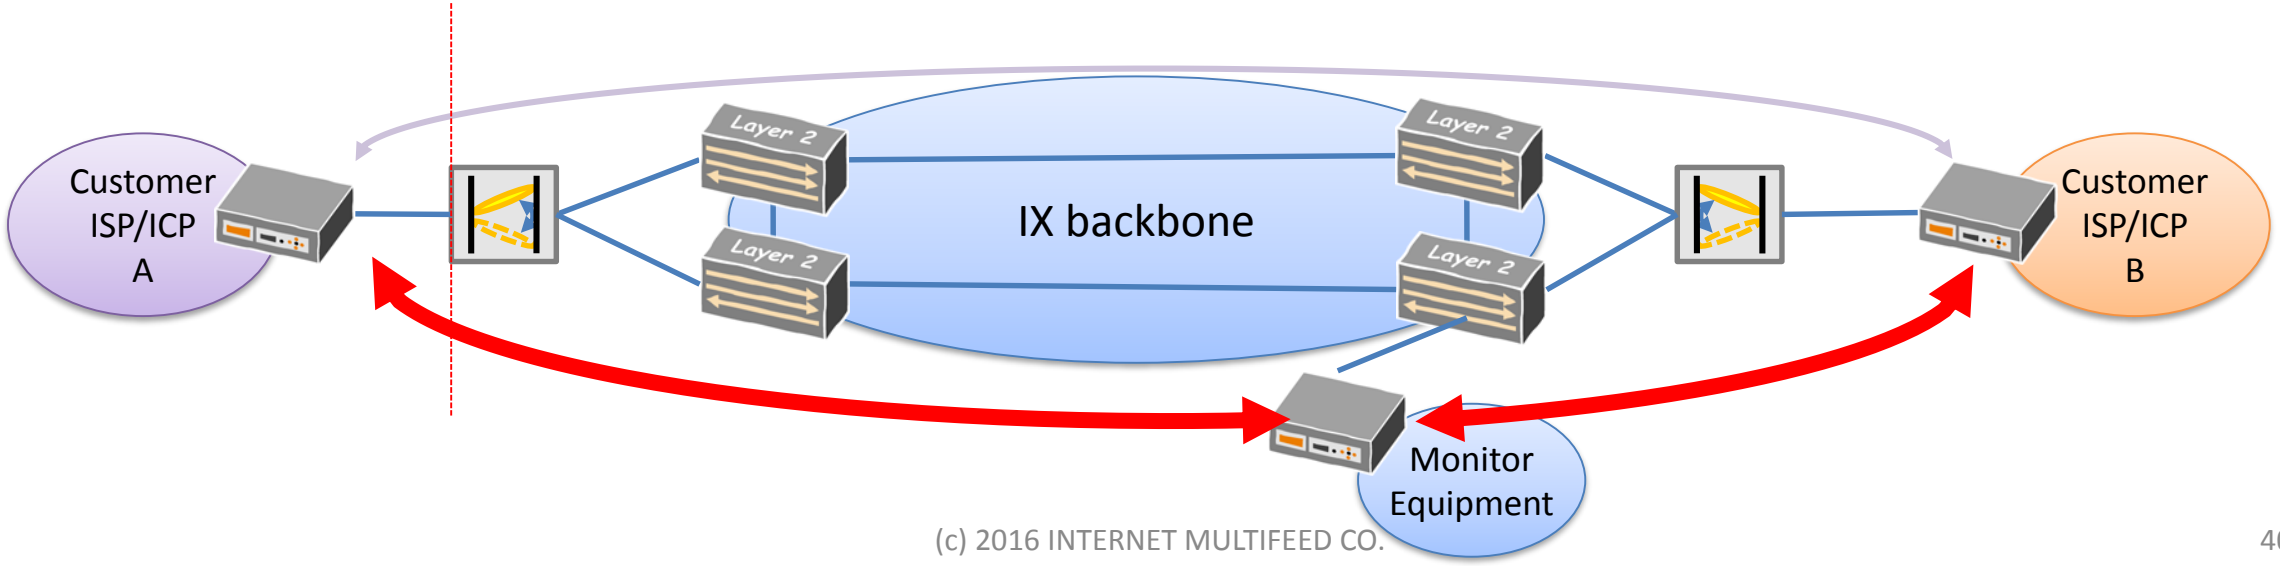

JPNAP: Improving failure detection time

Usually traffic fairly balanced on LAG

-> We use the characteristic to detect the failure such as LINK up but frames disappear

Past		Current
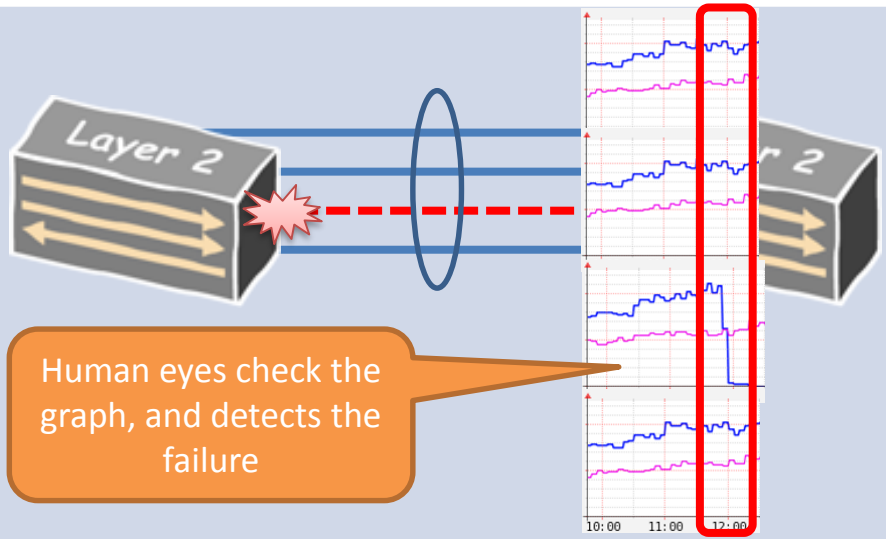		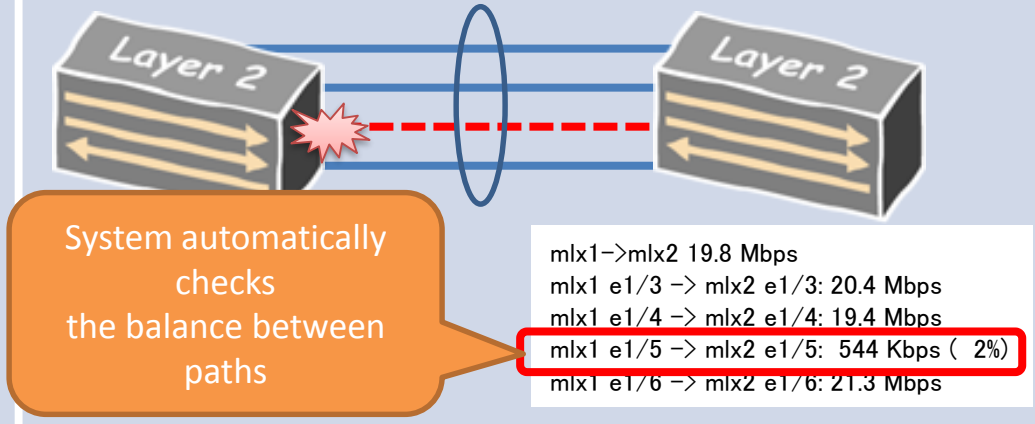
First alert from Ping failure or BGP down	→	Our monitoring system checks automatically
10 minutes at maximum to detect by human, from traffic graph updated every 5 minutes	→	**20 second at maximum by system**, using SNMP polling every 10 second
Human disables the port	→	System disables the port automatically
More than 10 minutes to recover	→	**Less than 20 second** to recover

Next step?

- My personal opinion, “too much” operation should be reduced.
- But this is a kind of Japanese culture; seems difficult to change.
- Also IXPs are afraid to lose customers, they cannot step out from the current situation. ☹

And one more topic from JPNAP…

Route server platform

- Route server platform is one of outstanding issues at IXP
  - Current popular platforms: BIRD, OpenBGPd, Quagga, Cisco
    - BIRD: recently becoming very popular platform
    - OpenBGPd: multithreaded, works on BSD
    - Quagga: old and monolithic architecture, maintenance is poor
    - Cisco: expensive
  - Quagga has become old and maintenance is poor.
  - New platform to replace Quagga is necessary.

GoBGP

- Open source BGP daemon
  - https://github.com/osrg/gobgp
  - developed by NTT Laboratories
- Go Language (developed by Google)
  - Using multi-core, faster processing in parallel
- API-first
  - GoBGP is designed to be controlled through API
  - Command Line interface is implemented into a software to kick the APIs

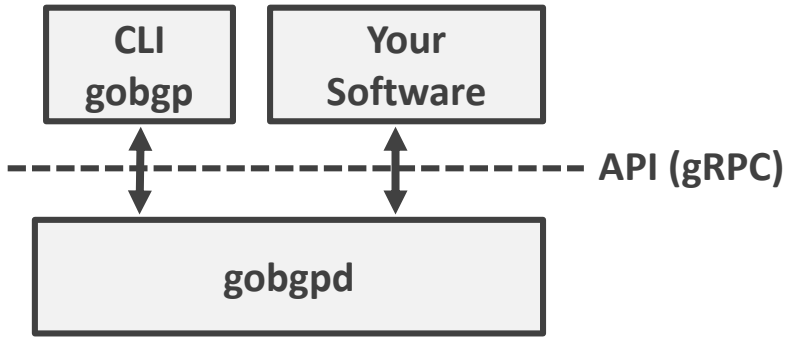

Performance comparison in laboratory with other BGP daemons

- Much better performance than Quagga, an older and monolithic BGP daemon
- GoBGP is just a little better than BIRD, widely deployed BGP daemon

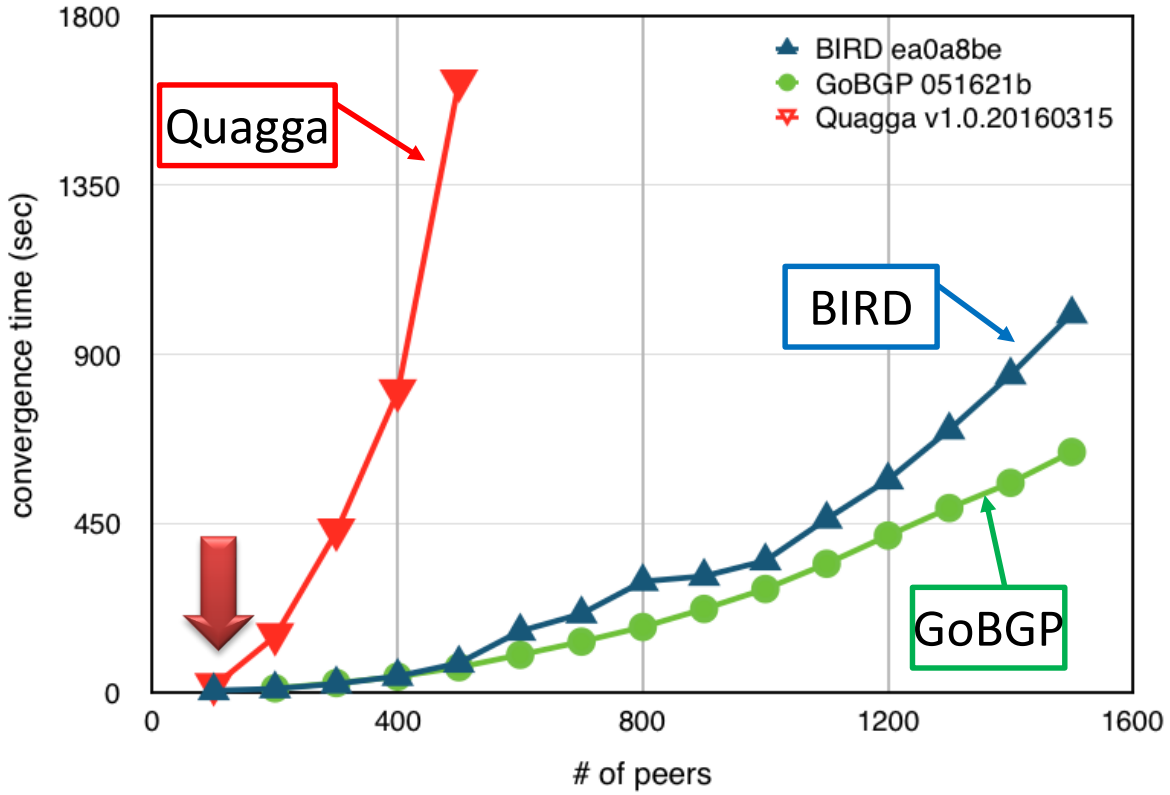

- Advertise 100 prefixes per peer
- Without route filters

GoBGP Route Server deployed at JPNAP

- Deployed at JPNAP Osaka on Aug. 3, 2016
  in place of Quagga route server
- So far no troubles ;-)
- Next, deploying at JPNAP Tokyo 1

Thank you!

Contact

APIX: sc _at_ apix.asia

JPNAP: info _at_mfeed.ad.jp