Best Practices for HKIX Peering

ISP Symposium 2017

Kenneth CHAN

Team Lead, HKIX

www.hkix.net

18 Dec 2017

What is HKIX?

- Established in Apr 1995, Hong Kong Internet eXchange (HKIX) is the main layer-2 Internet eXchange Point (IXP) in Hong Kong where various autonomous systems interconnect with one another and exchange traffic
- HKIX is now owned and operated by the Hong Kong Internet eXchange Limited (a wholly-owned subsidiary of The Chinese University of Hong Kong Foundation Limited) in collaboration with Information Technology Services Centre of The Chinese University of Hong Kong
- HKIX serves both commercial networks and R&E networks
- The original goal is to keep intra-Hong Kong traffic within Hong Kong

Help Keep Intra-Asia Traffic within Asia

- We have almost all the Hong Kong networks
- So, we can attract participants from Mainland China, Taiwan, Korea, Japan, Singapore, Malaysia, Thailand, Indonesia, Philippines, Vietnam, India and other Asian countries
- We now have more non-HK routes than HK routes
- We do help keep intra-Asia traffic within Asia
- In terms of network latency, Hong Kong is a good central location in Asia
- HKIX does help HK maintain as one of the Internet hubs in Asia
- HKIX supports both domestic and international traffic

HKIX Model — MLPA over Layer 2 + BLPA

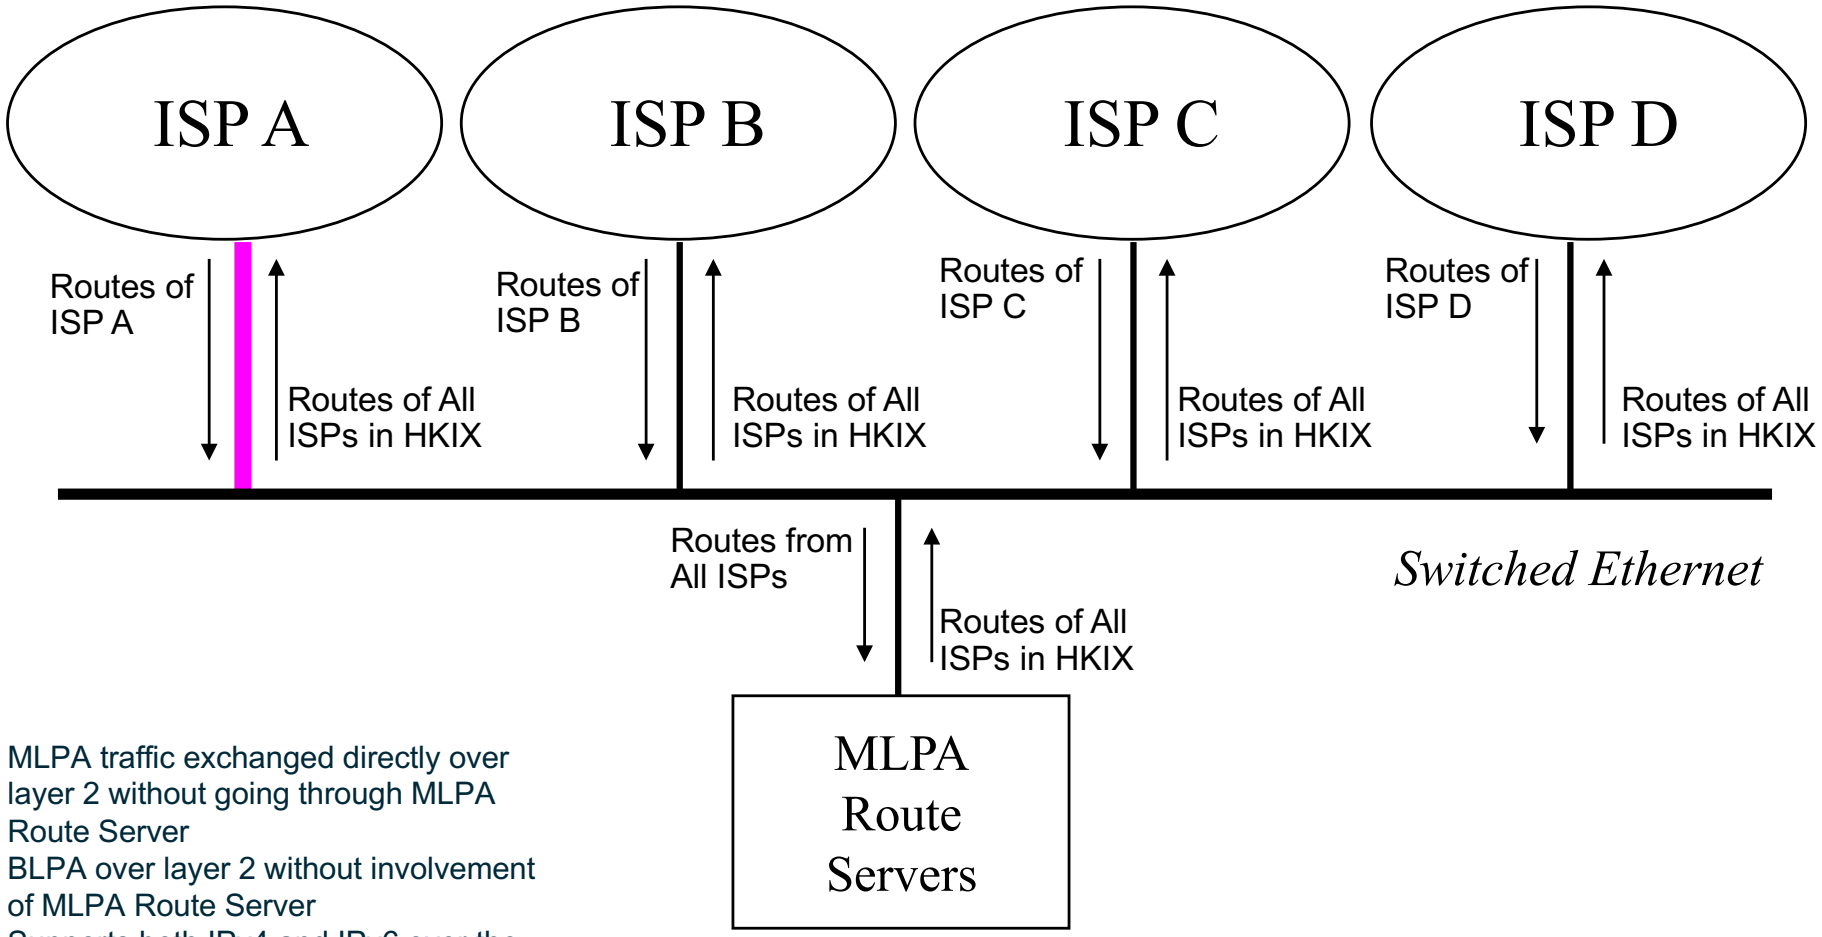

- MLPA traffic exchanged directly over layer 2 without going through MLPA Route Server
- BLPA over layer 2 without involvement of MLPA Route Server
- Supports both IPv4 and IPv6 over the same layer 2 infrastructure

New HKIX Dual-Core Two-Tier Spine-and-Leaf Architecture For 2014 and Beyond

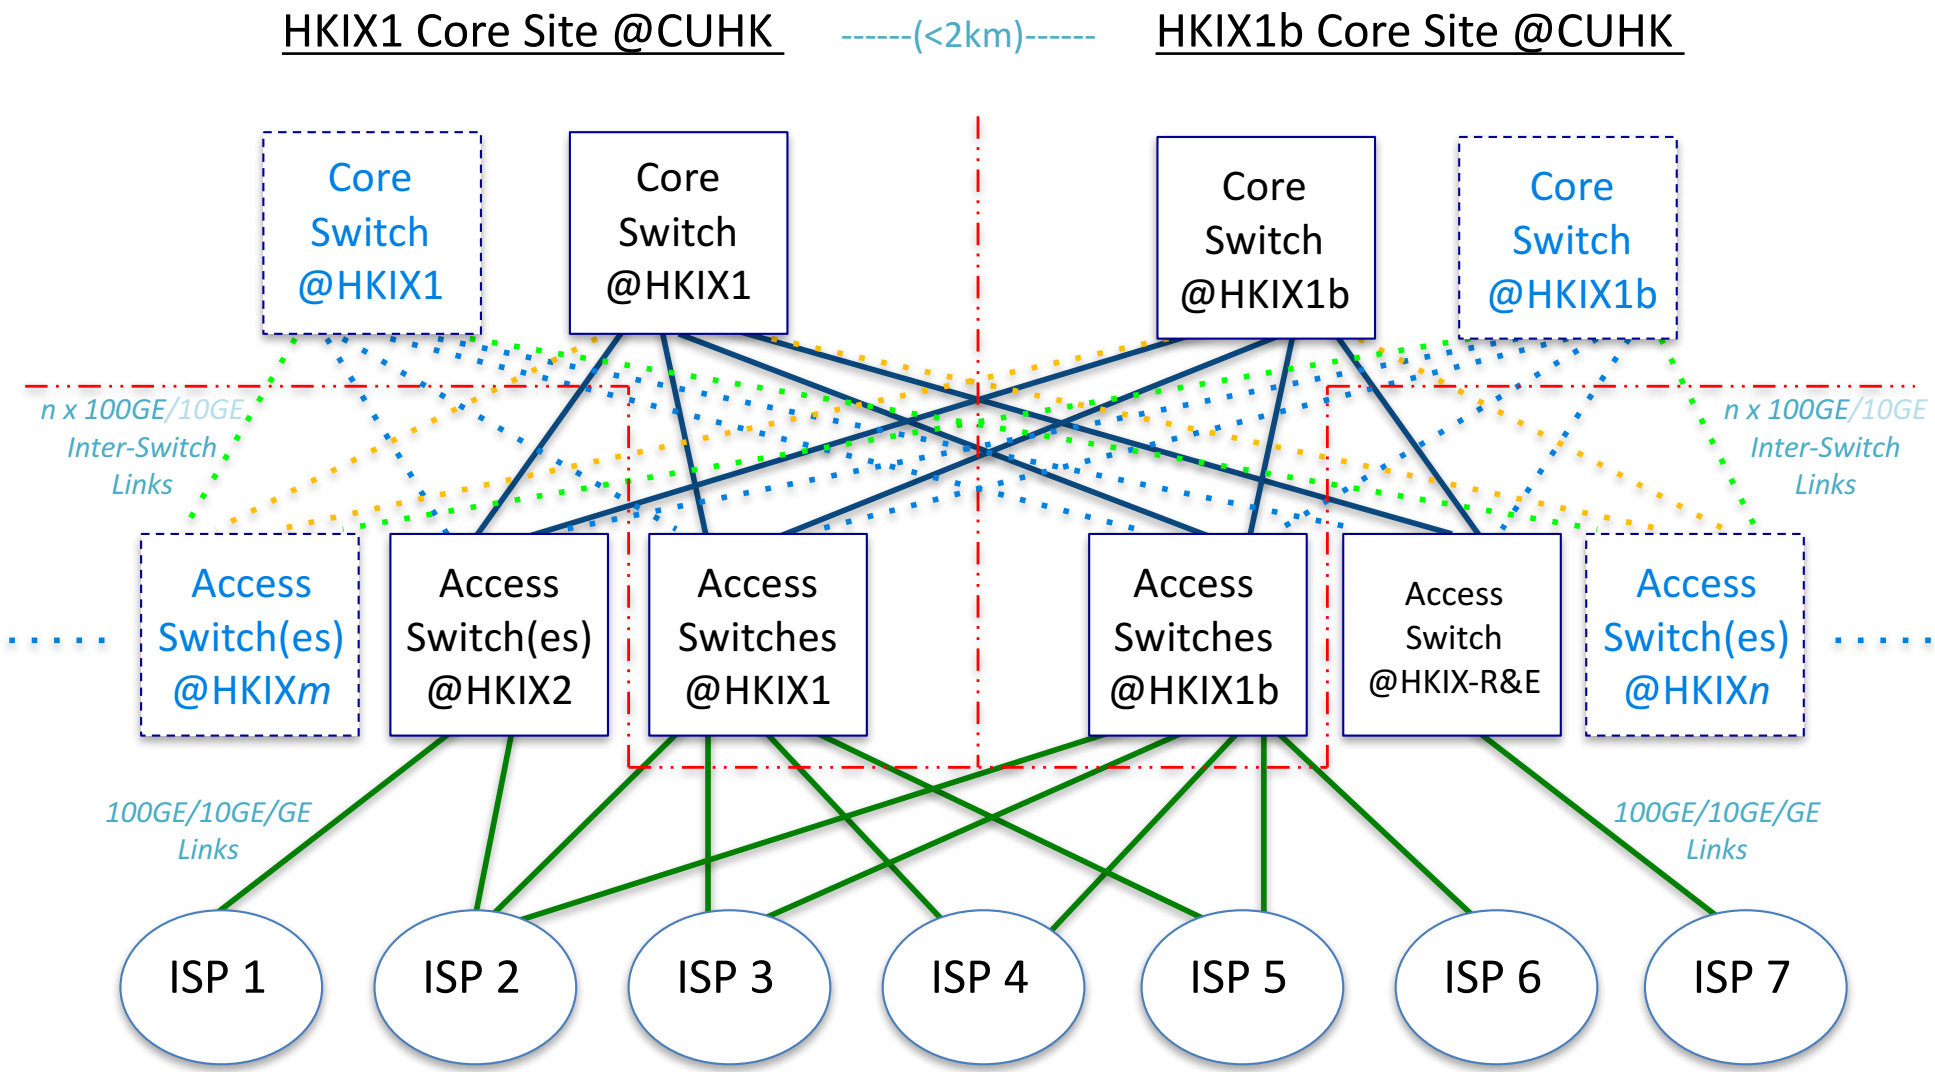

HKIX Traffic in 2007

Daily Graph (5 Minute Average)

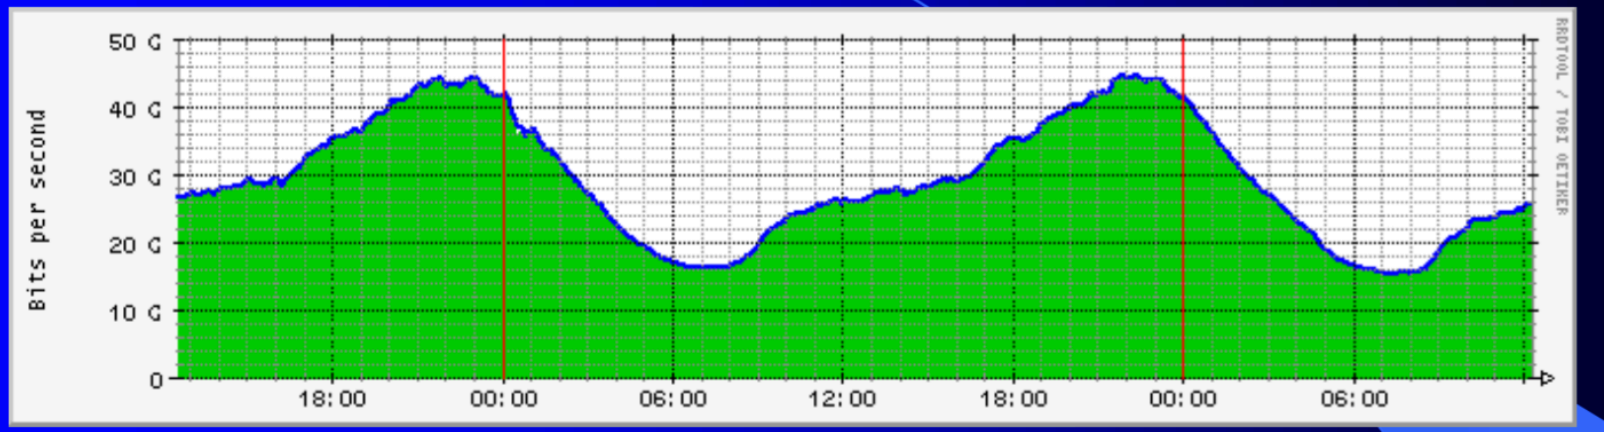

Yearly Graph (1 Day Average)

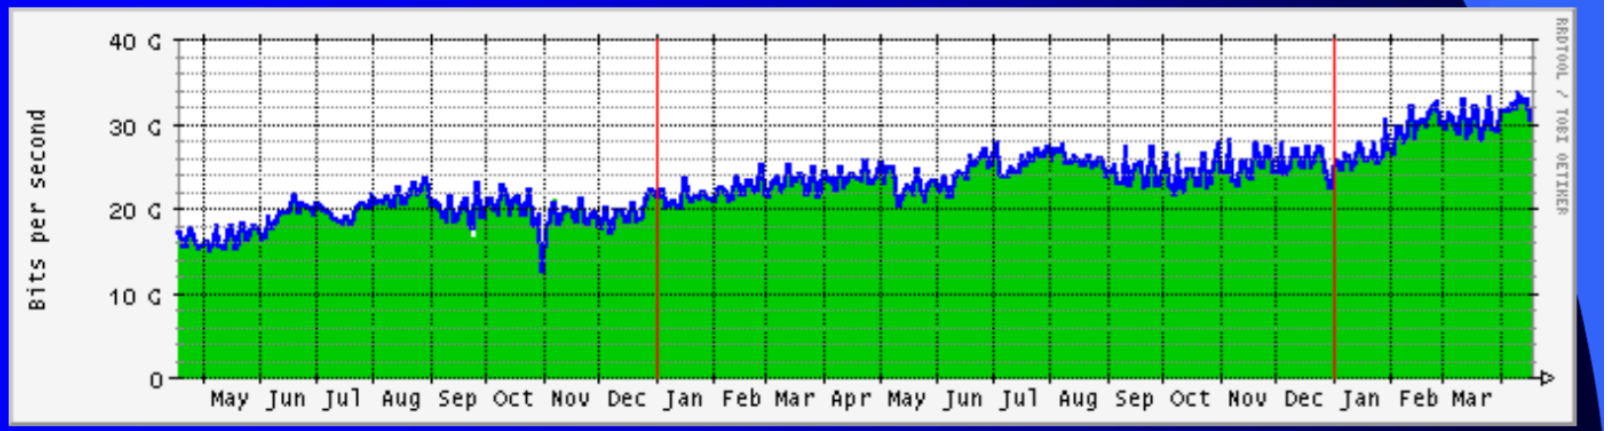

HKIX Traffic in 2010

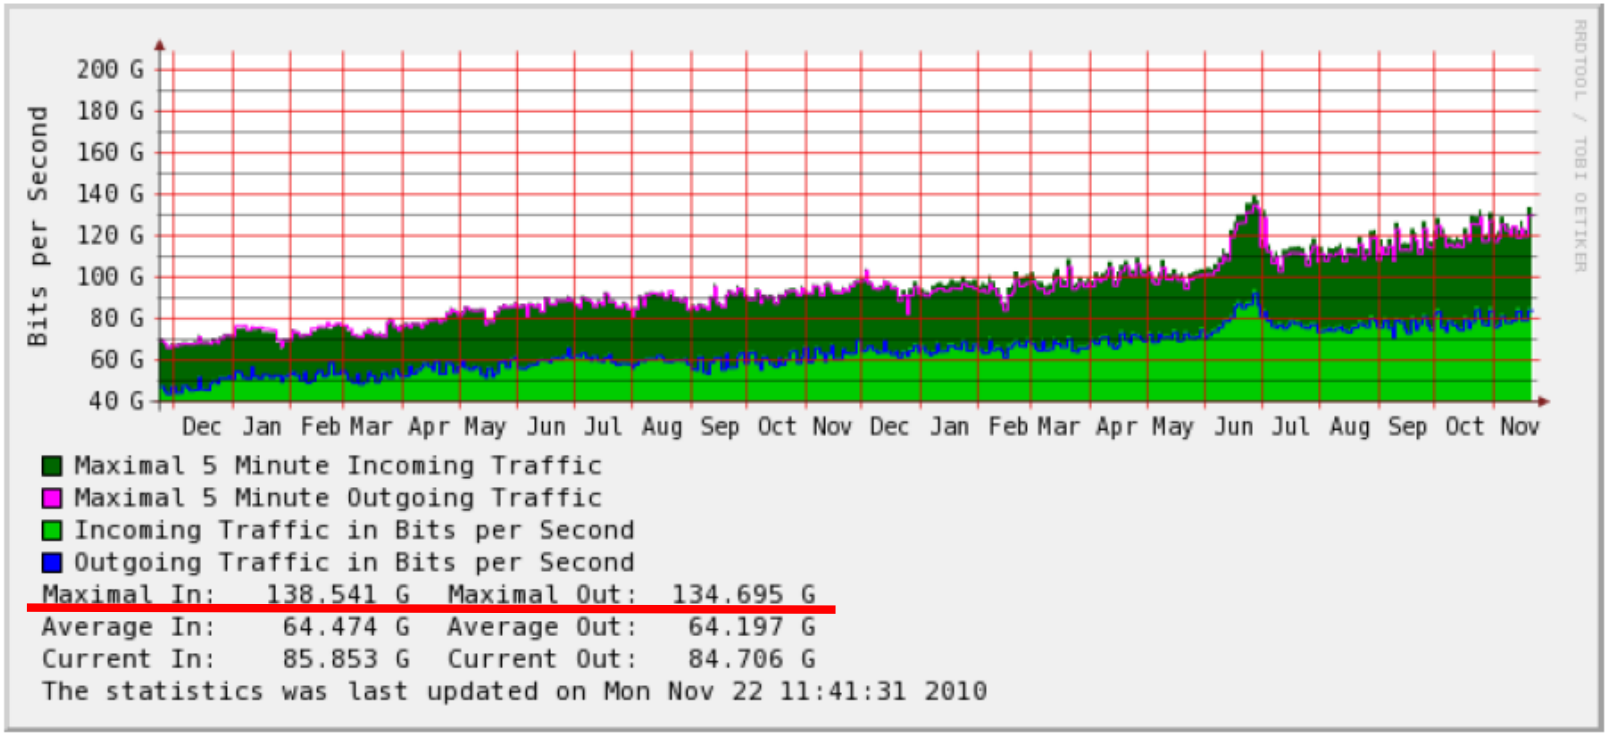

HKIX Traffic in 2013

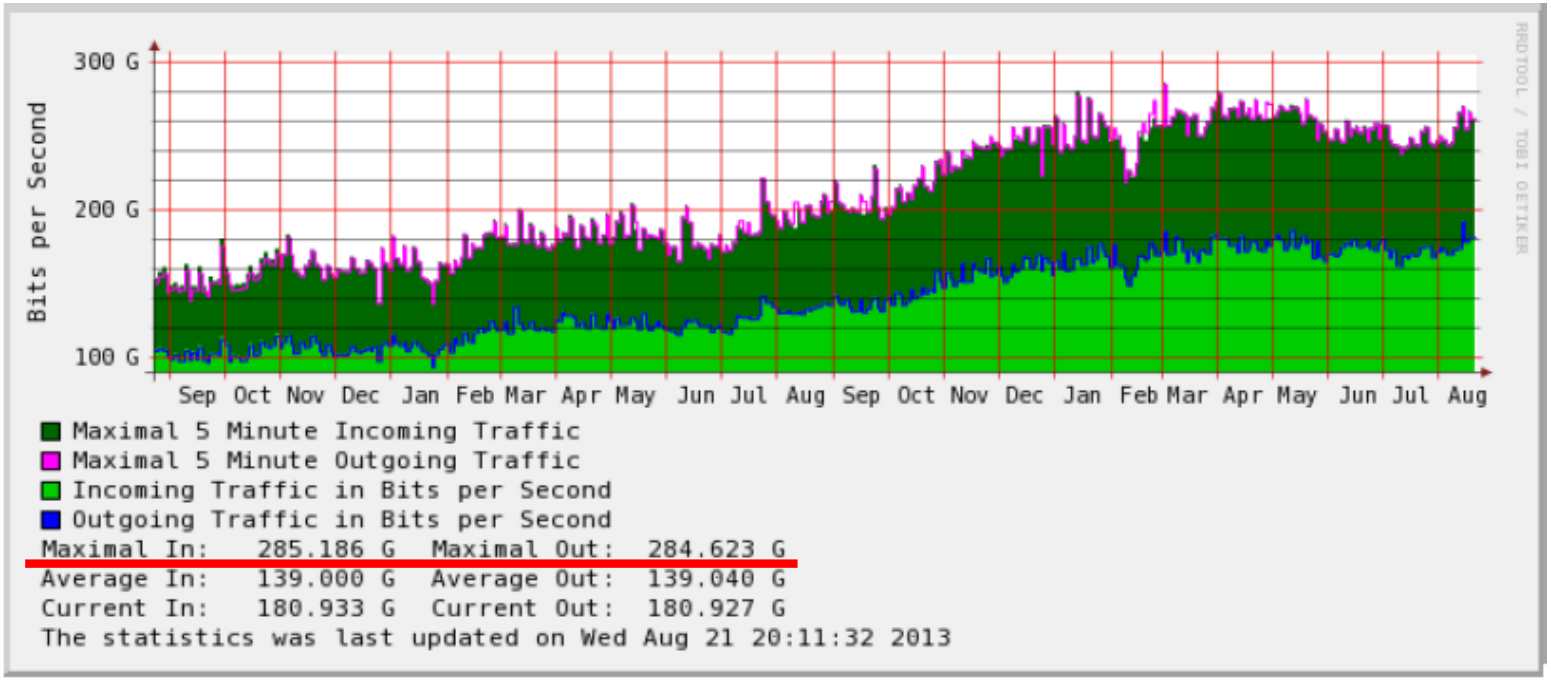

HKIX Traffic in 2016

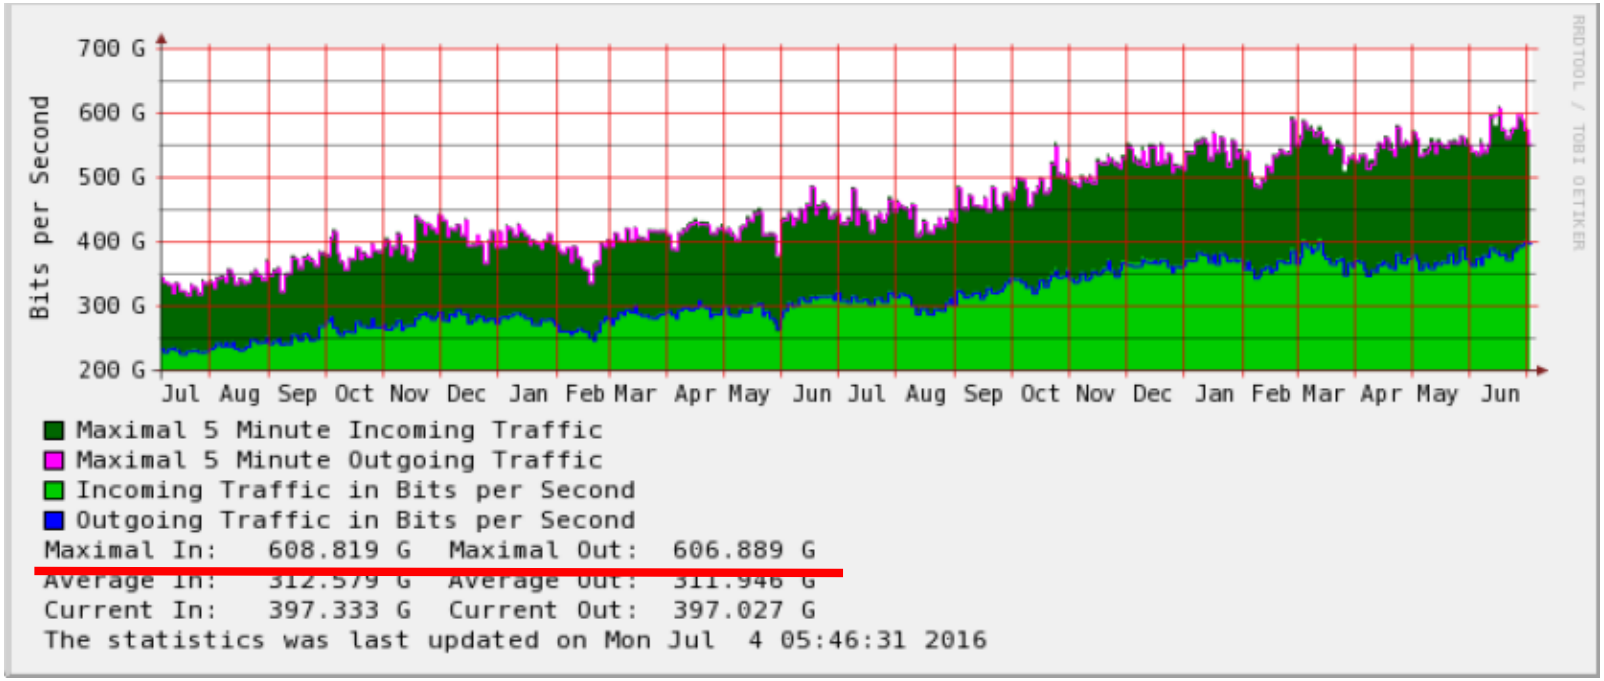

HKIX Today

- Supports both MLPA (Multilateral Peering) and BLPA (Bilateral Peering) over layer 2
- Supports IPv4/IPv6 dual-stack
- More and more non-HK participants
- 270+ autonomous systems connected
- 500+ connections in total
  - **20** 100GE, **300+** 10GE & **170+** GE
- 960+Gbps (5-min) total traffic at peak
- Annual Traffic Growth ~30%

HKIX Traffic
Daily Graph (5-min average)

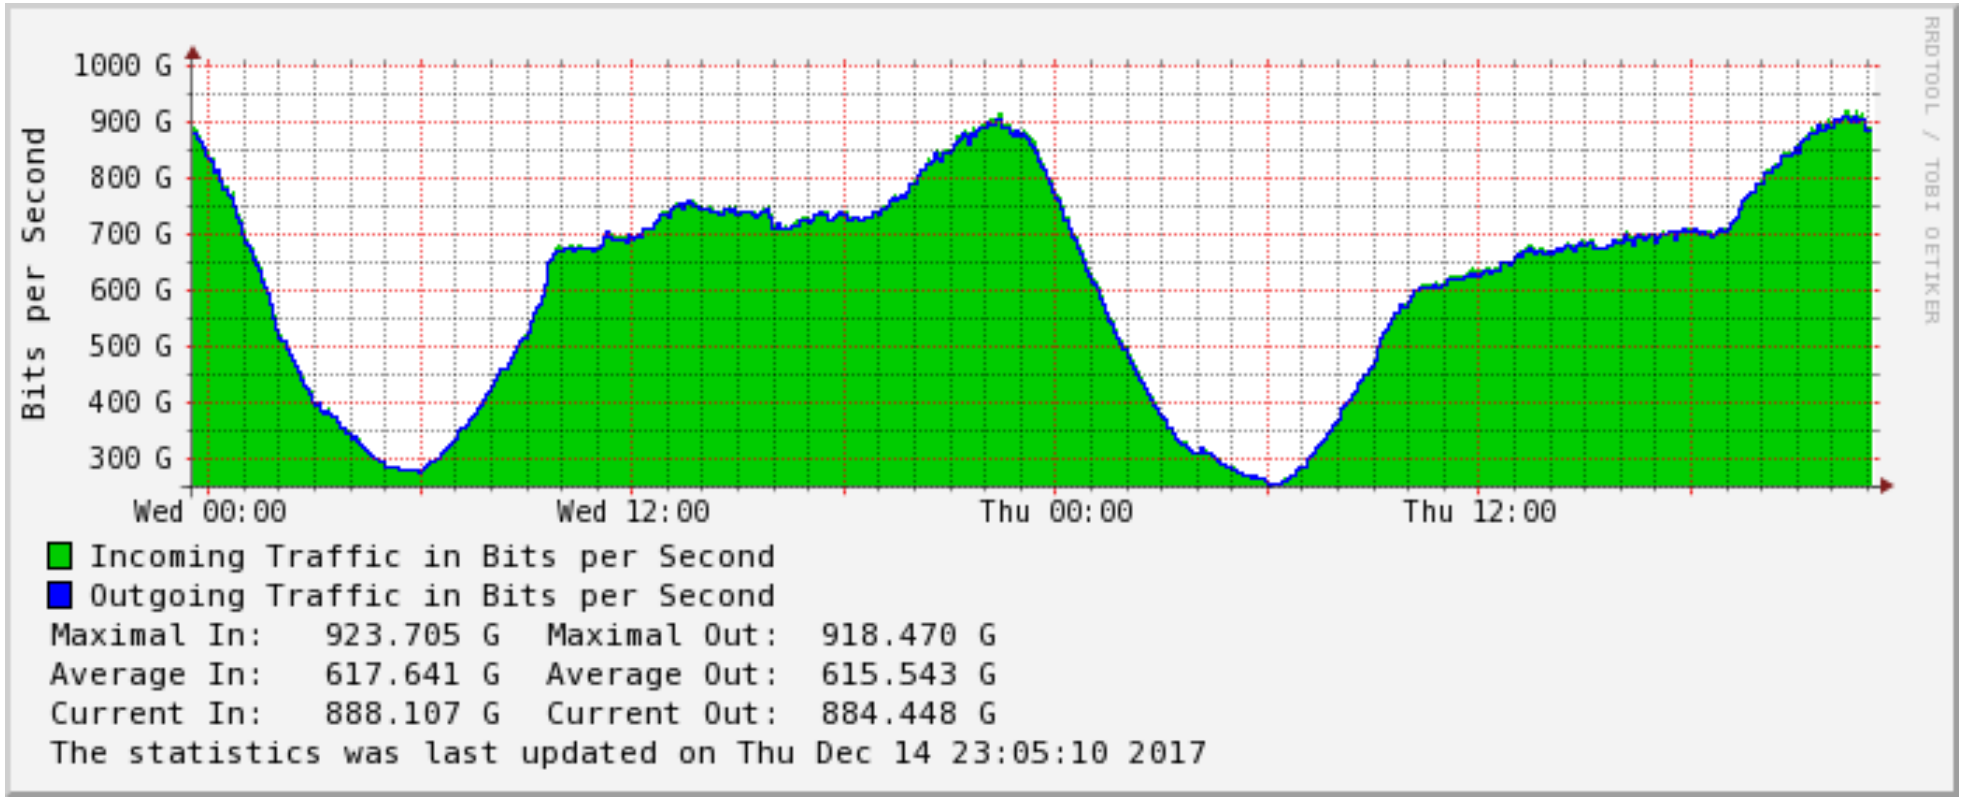

HKIX Traffic
Yearly Graph (1-day average)

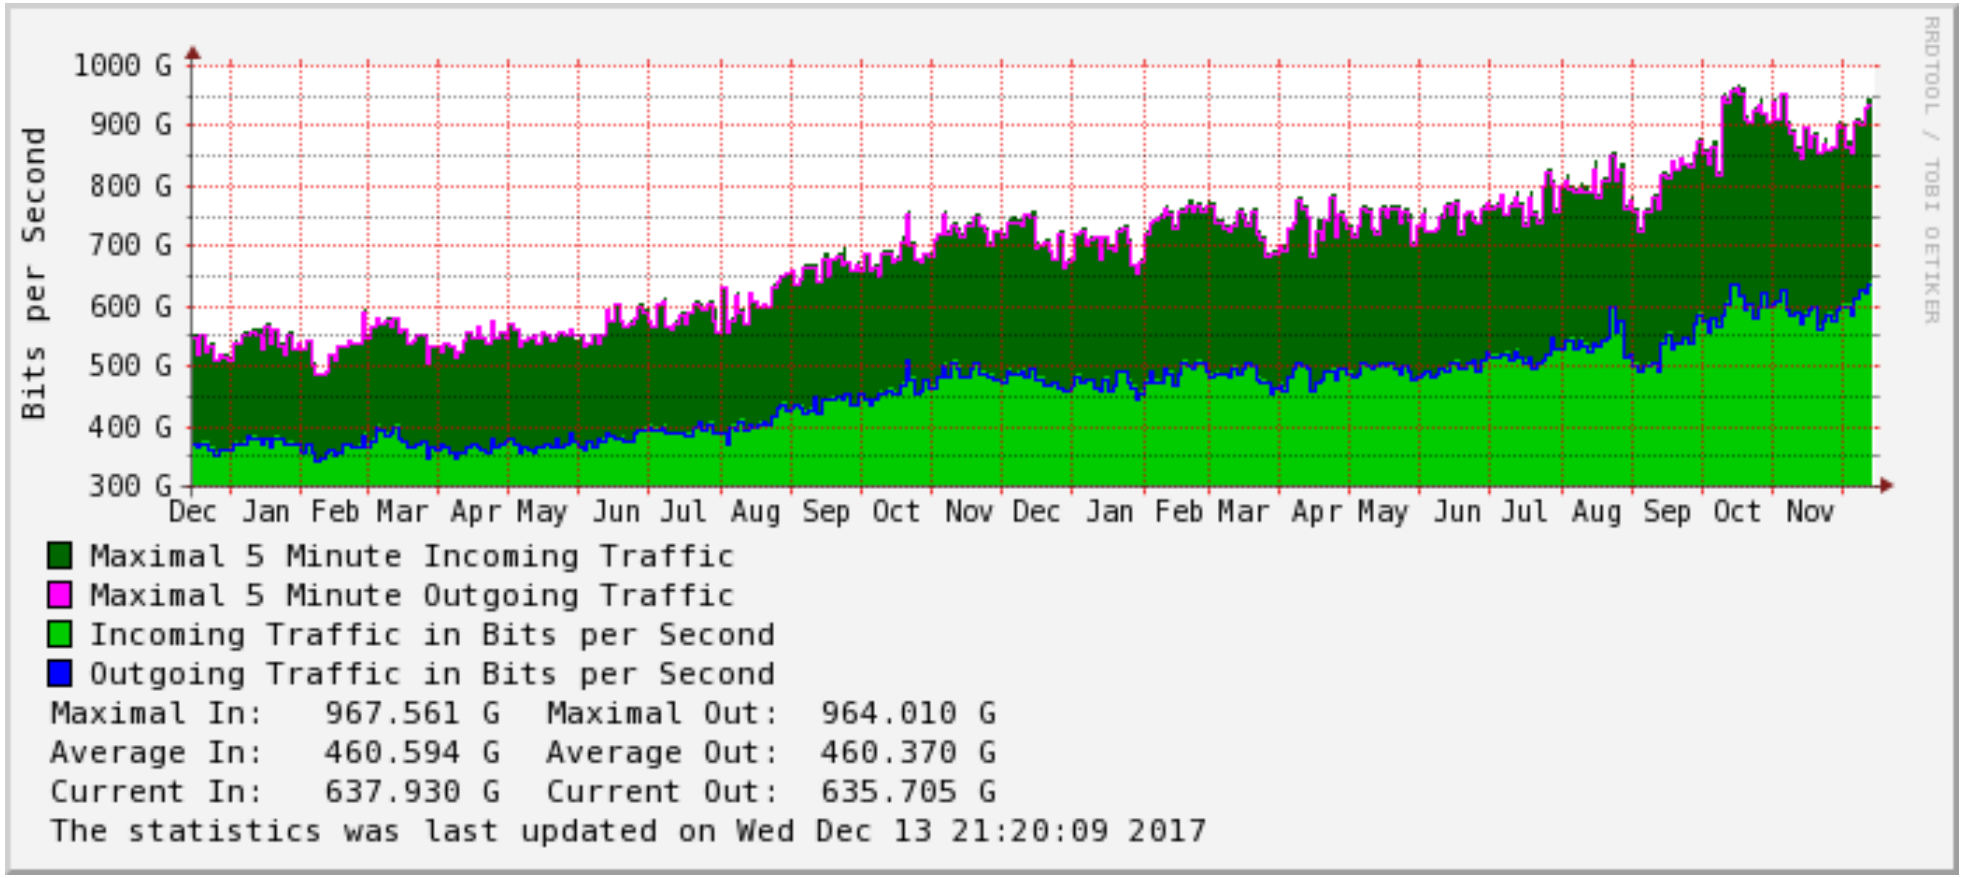

HKIX Network Diagram (Jun 2017)

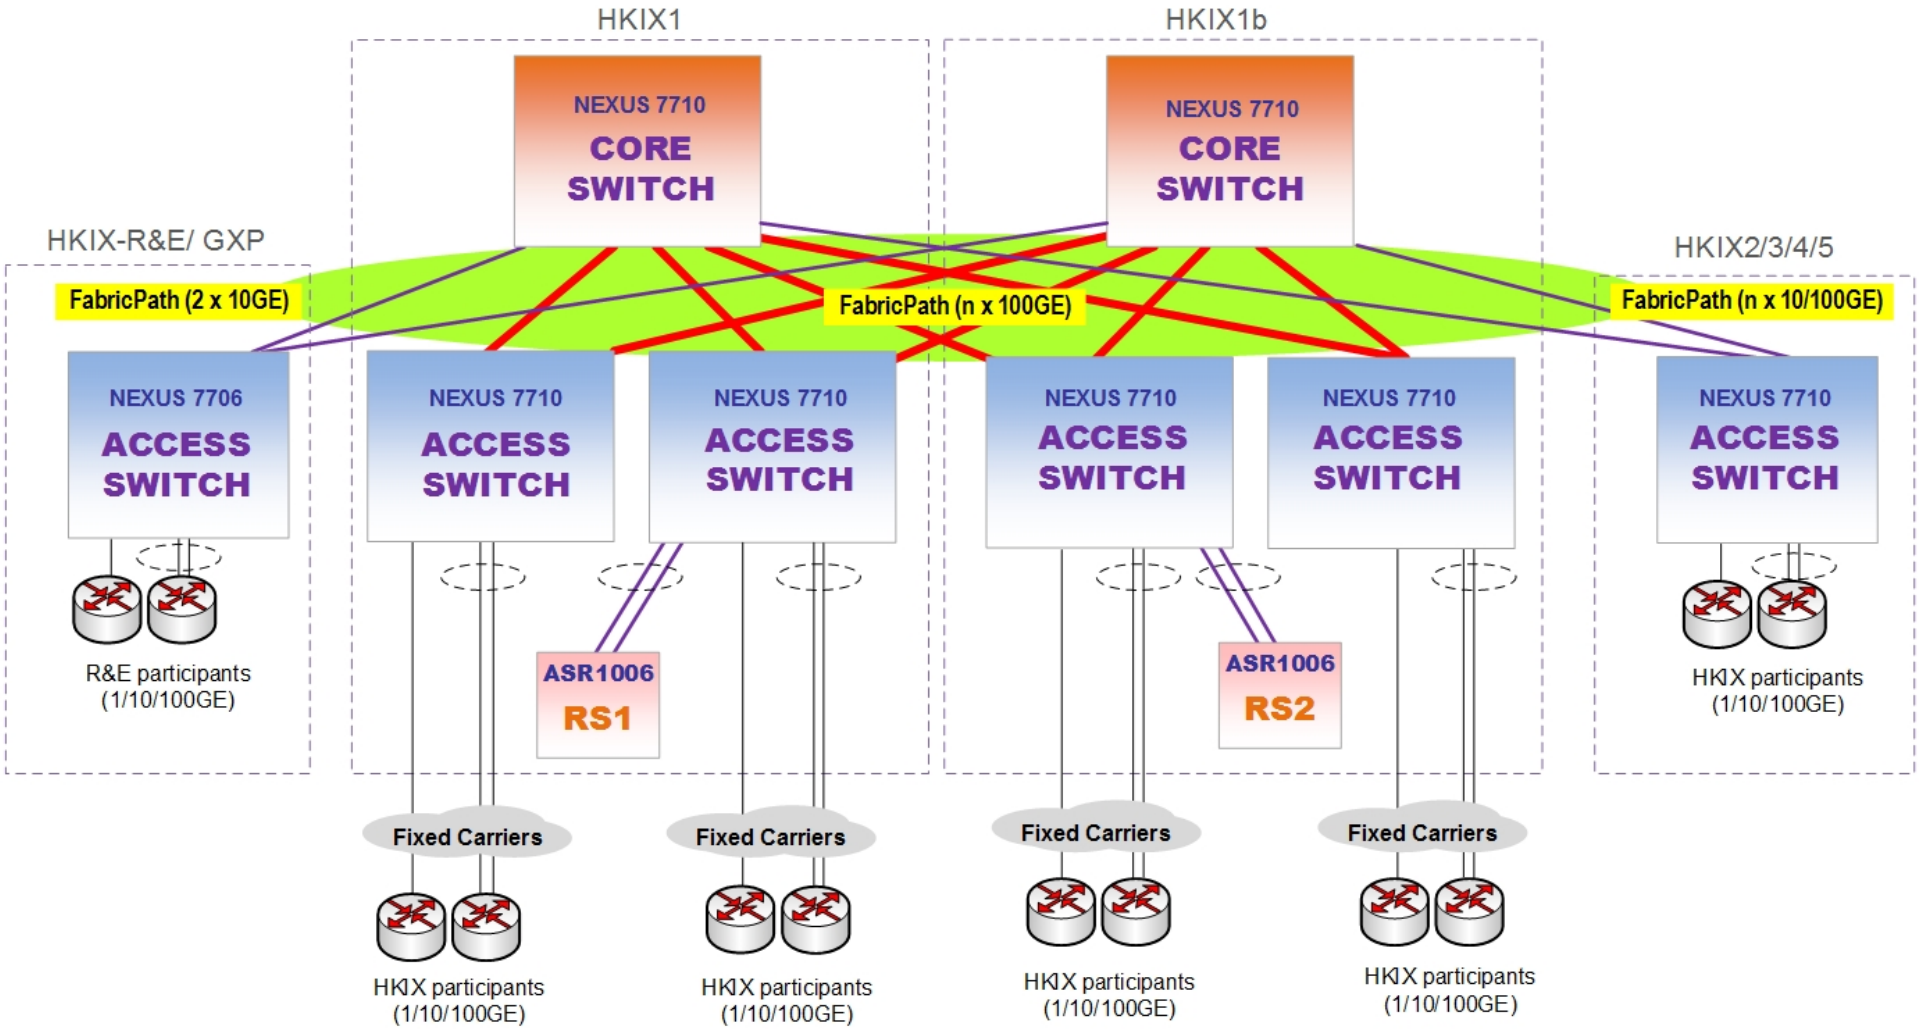

1. **HKIX1** and **HKIX1b** are the two core sites of HKIX at CUHK while **HKIX2/3/4/5** are HKIX satellite sites outside of CUHK.
2. HKIX participants are encouraged to connect to multiple sites for site resilience.

Advantages of HKIX

- **Location**
  - Hong Kong is a good central location in Asia

    ~50ms to Tokyo and ~30ms to Singapore
- **Neutral**
  - Treat all partners equal, big or small
  - Neutral among ISPs / telcos / local loop providers/ data centers / content providers / cloud services providers
- **Trustable**
  - Treat all partners fair and consistent
  - Respect business secrets of every partner / participant
- **High Performance**
  - No internal performance bottleneck, no internal packet loss
- **Not for Profit**
  - Charging mainly for equipment upgrade and long-term sustainability, not for profit-making

100G Connections at HKIX

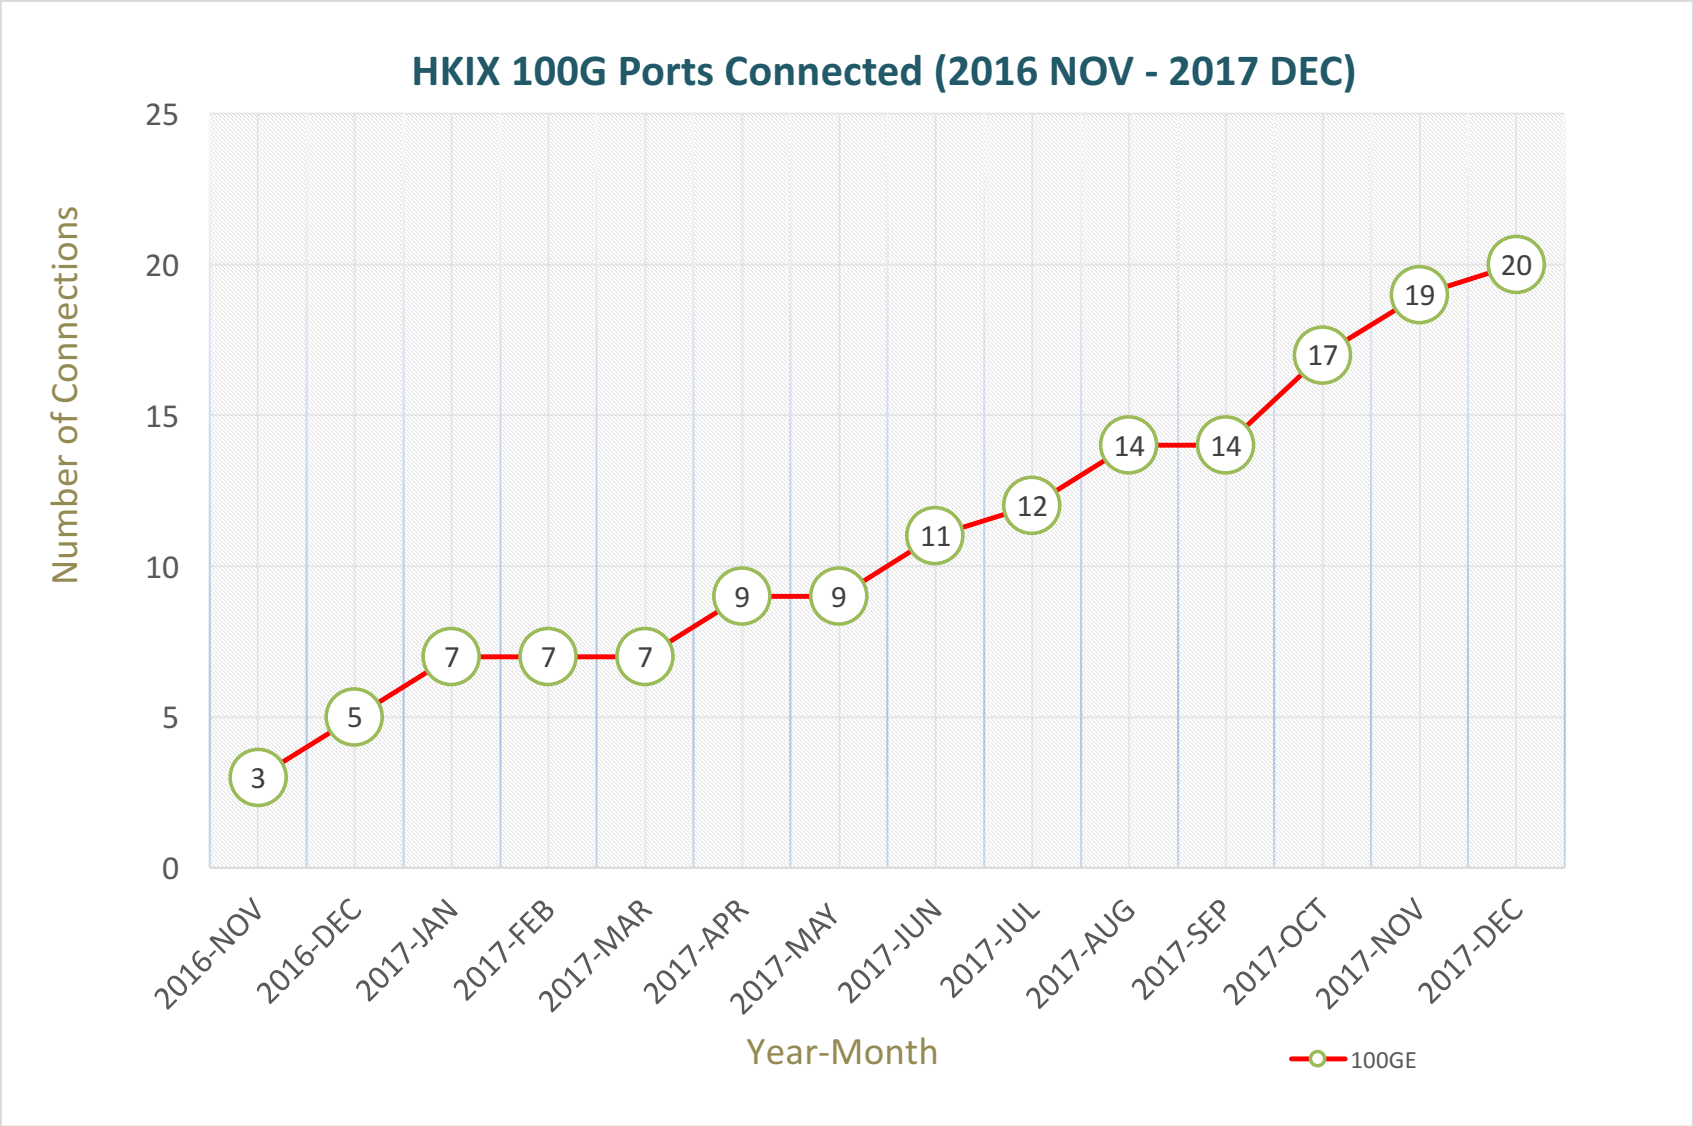

100G Participants at HKIX

- Akamai
- Amazon
- China Mobile International
- CloudFlare
- Facebook
- Google
- Hong Kong Broadband Network
- Hurricane Electric
- Tencent
- TVB
- Yahoo!

HKIX Satellite Sites

Hong Kong, 08 Feb 2017

HKIX announces that 3 new satellite sites will be established in collaboration with 3 commercial data centres which provide colocation services as well as easy connections to HKIX.

Satellite Site	Satellite Site Collaborator	District	Ports Supported	Status
HKIX2	CITIC Telecom International	Kwai Chung	GE/10GE	Ready for Service
HKIX3	SUNeVision / iAdvantage	Fo Tan	GE/10GE/100GE	Ready for Service 28 Feb 2017
HKIX4	NTT Com Asia	Tseung Kwan O	GE/10GE/100GE	Ready for Service 19 Jun 2017
HKIX5	KDDI / Telehouse / HKCOLO.net	Tseung Kwan O	GE/10GE/100GE	Ready for Service 24 Mar 2017

- For connections to HKIX at Satellite Sites, special connection charges will be charged by relevant operators, in addition to the port charges charged by HKIX.
- For HKIX participants not co-located at HKIX satellite sites, they can still connect to any of the two HKIX core sites, i.e. HKIX1 and HKIX1b sites by local loops via local loop providers.

Setup Multiple HKIX Satellite Sites

- Allow participants to **connect to HKIX more easily at lower cost** from those satellite sites in Hong Kong
- Open to commercial data centres in HK which fulfil minimum requirements so as to maintain neutrality which is the key success factor of HKIX
- Create a win-win situation with satellite site collaborators
- To be named HKIX2/3/4/5/6/etc

  Latest updates:
  - HKIX2 has been migrated from old model to HKIX Satellite Site
  - HKIX3/4/5 are new Satellite Sites and they are **Ready for Service** now
- ***HKIX1*** *and* ***HKIX1b*** *(the two HKIX core sites* ***located within CUHK Campus****) will continue to serve participants directly*

HKIX-R&E Node – Support for National R&E Networks in Hong Kong

- **HKIX helps those R&E Networks interconnect among themselves and with commercial networks without restrictions via HKIX-R&E switch at MEGA-i**

- The main purpose is to facilitate those National R&E Networks having presence in Hong Kong to do interconnections among themselves ***and*** do peering with commercial networks at HKIX more easily and at a lower cost.

- Started in 2008
- Located in MEGA-iAdvantage
- For Research and Education Networks (R&E) only
- Support GE/10GE/100GE Trunk Ports
- Support Point-to-point VLANs for R&E networks
  - For private interconnections among any 2 R&E networks
  - Jumbo Frame support
- Fiber Cross Connect to be provided by R&E networks
- 7x24 NOC support
- Operate by HKIX with a Nexus7700 switch at MEGA-i

HKIX-R&E Node at MEGA-i

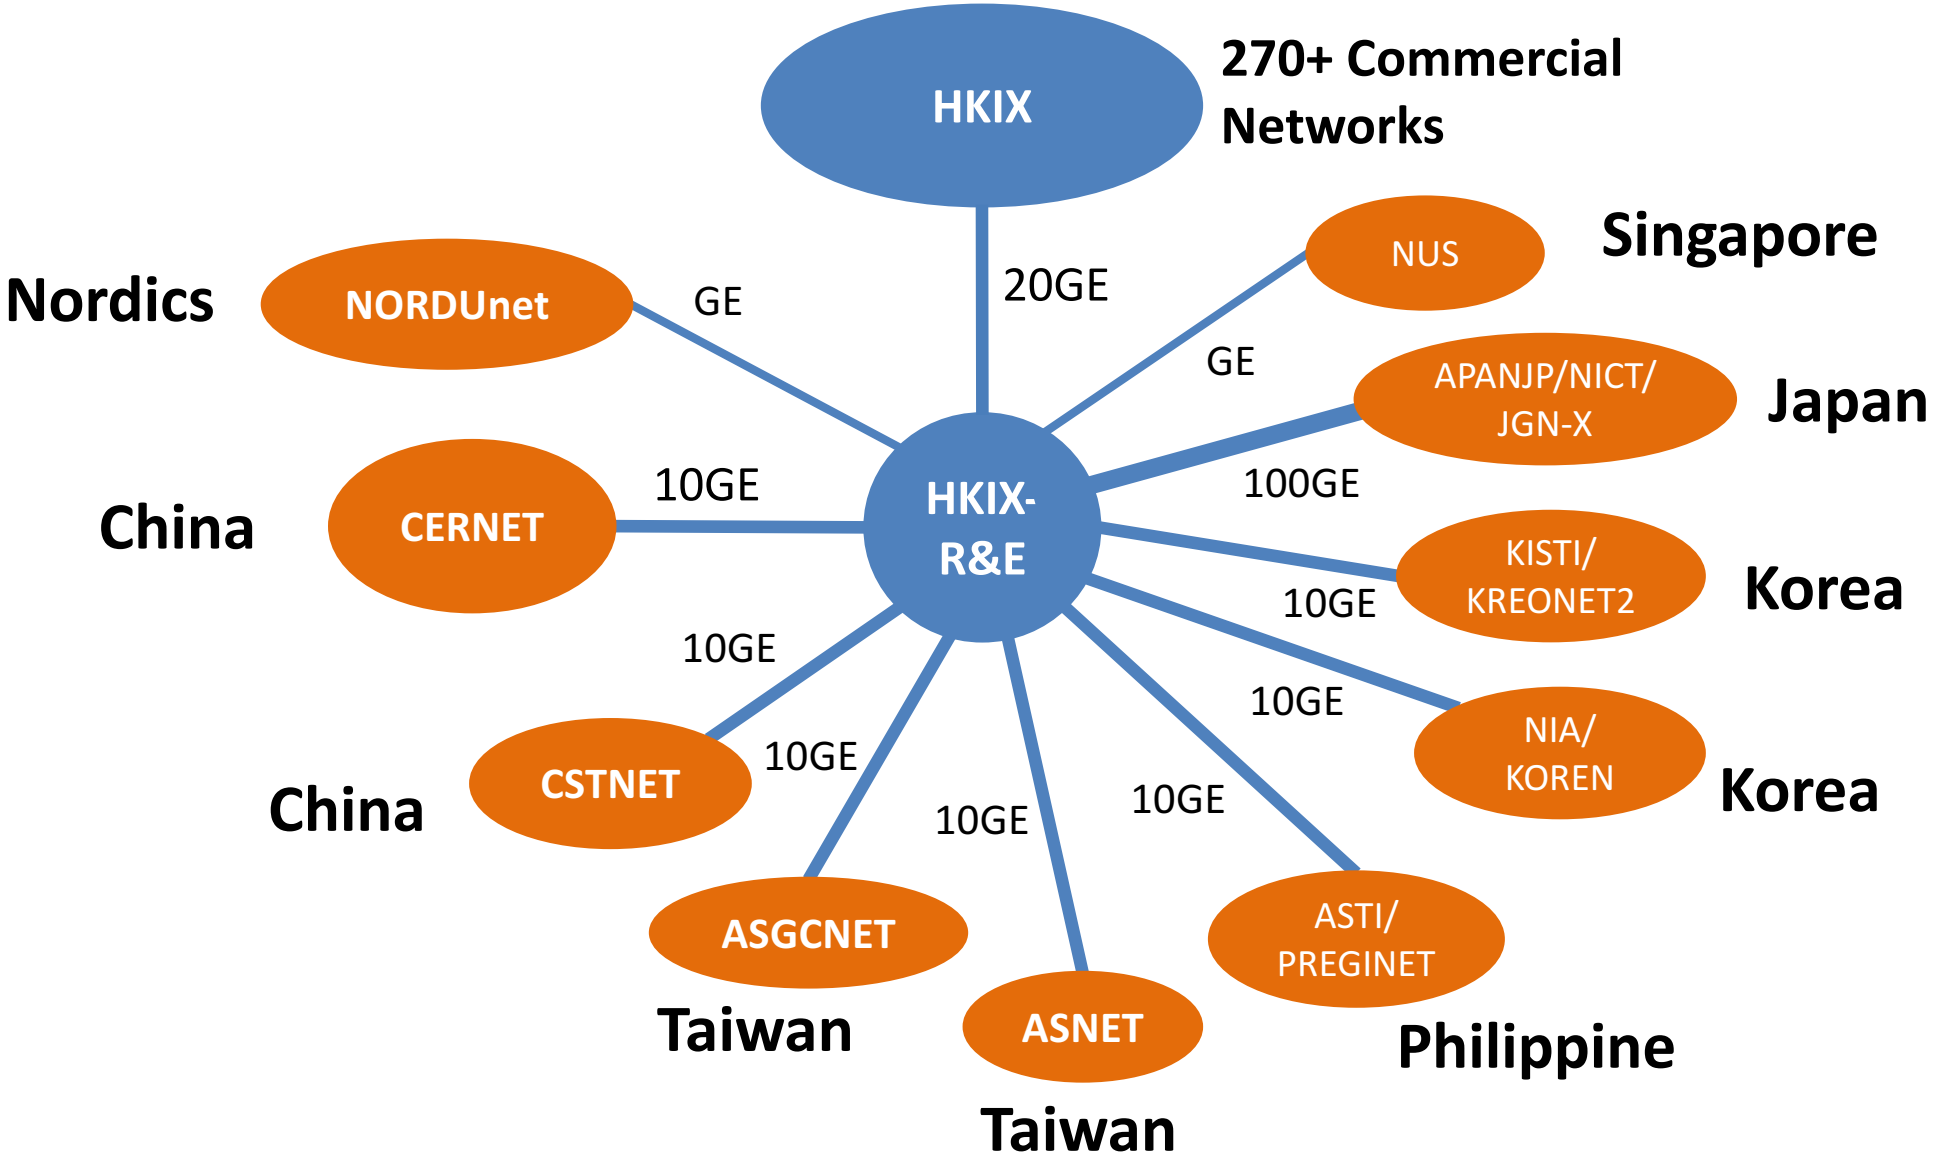

HKIX-R&E Network Diagram

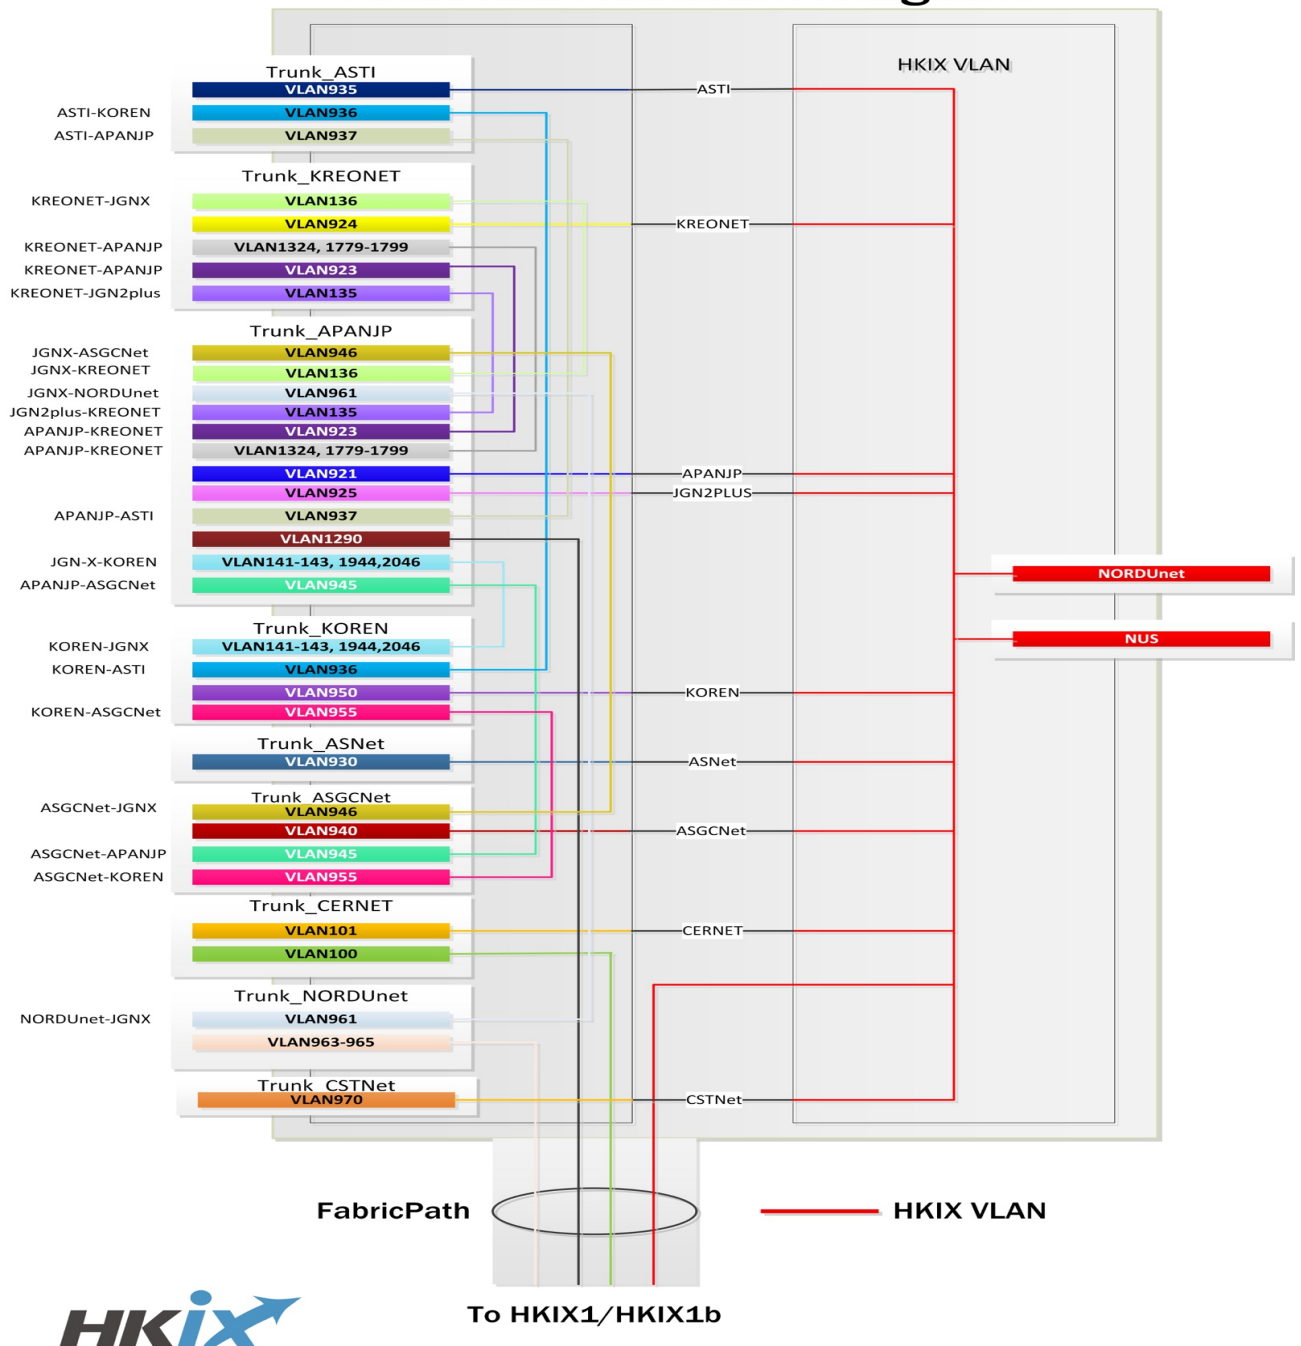

Updated on 05-Dec-2017

GNA - A Blueprint for Global R&E Network Architecture

http://gna-re.net

- The Global Network Architecture program (GNA) is an international collaboration between national research and education (R&E) networks

- The discussions inside the GNA group have led to a global network architecture model that consists of a powerful intercontinental transmission substrate, consisting of:
  - Global Open Exchange Points (GXPs)
  - High-bandwidth transmission pipes (running between GXPs) for sharing

GNA – artist's impression

Credit – Mian Usman (DANTE)

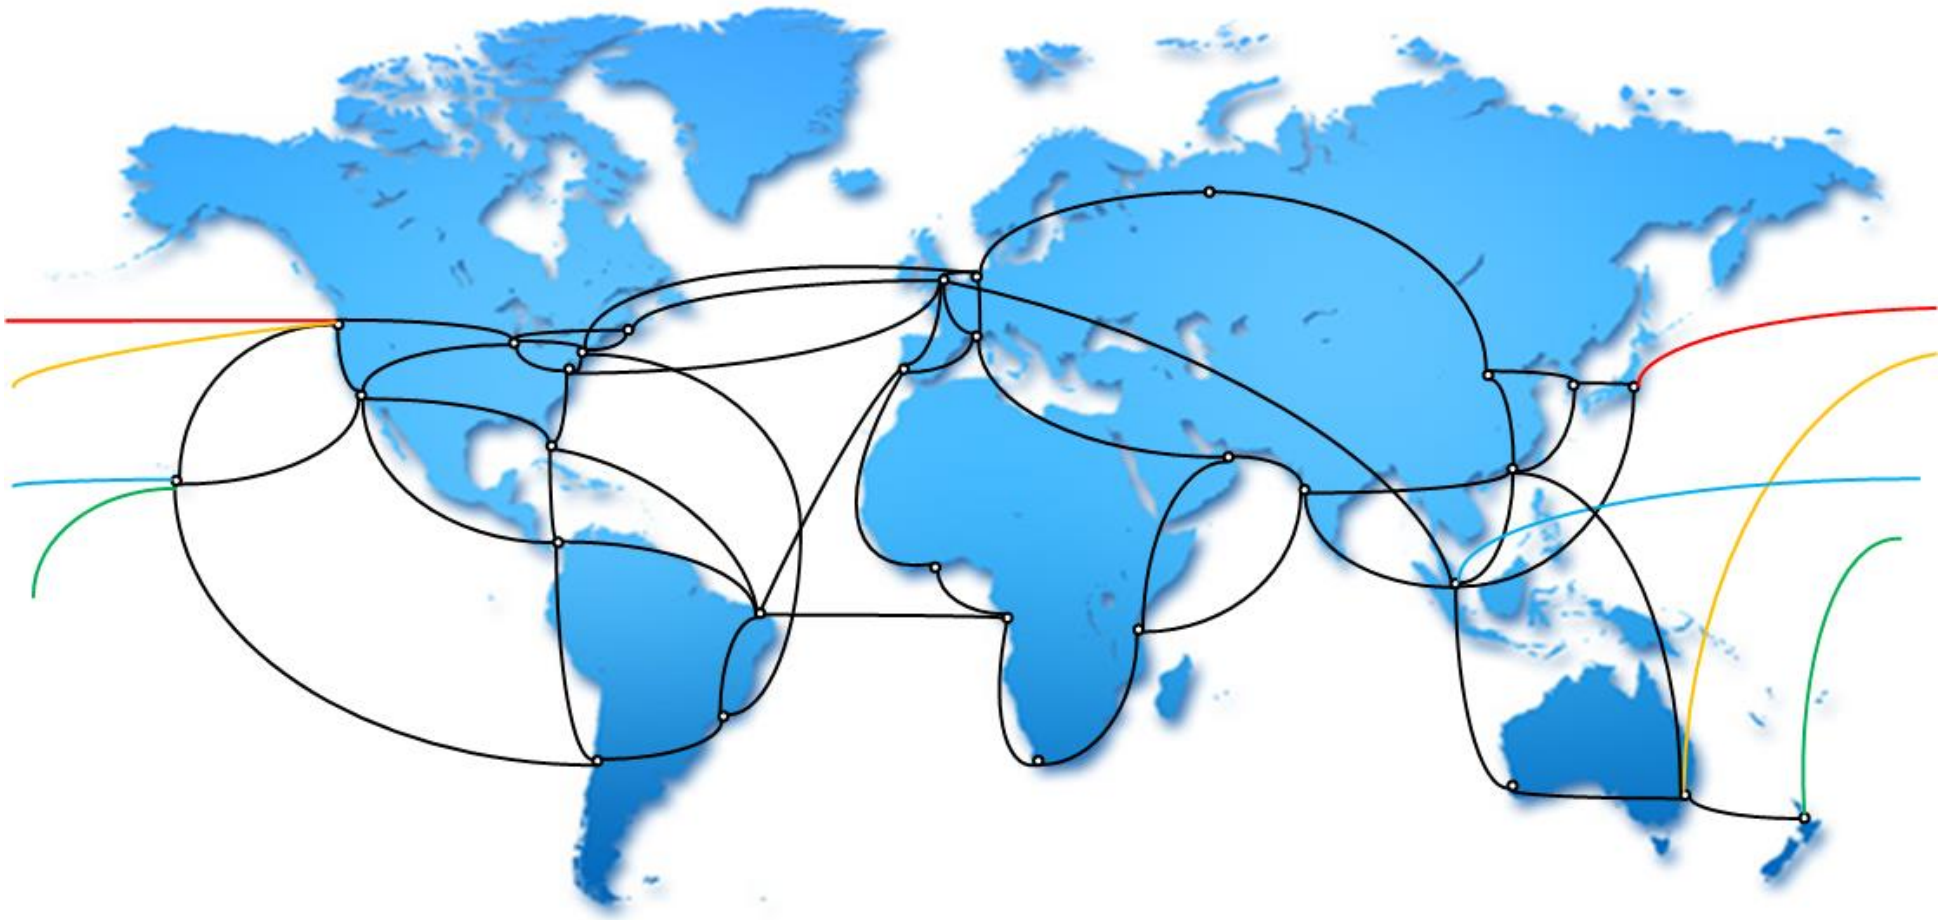

EXCHANGE POINTS IN THE GNA PHASE I

Planned Works for 2017/18

- **Improved Stability**
  - *Better Control of Proxy ARP* (DONE)
  - *L2 Control on HKIX peering LAN* (DONE)
- **Improved Services**
  - *Set up Satellite Sites in multiple commercial Data Centre* (DONE)
  - *Set up portal for HKIX participants (2018 Q1)*
  - *True 24x7 NOC* (DONE)
  - *Improve after-hour support* (DONE)
  - *More advanced Route Server features (2018 Q1)*
- **Improved Security**
  - *ISO27001 (2018 Q2)*
  - *Better support for DDoS Mitigation* (DONE)

Better Control of Proxy ARP

- Automatic Detection of Proxy ARP (implemented)
  - Based on duplicated IPv4 ARP entries learned on HKIX Route Servers
- Automatic shutdown switch port of HKIX peer causing Proxy ARP (will be implemented soon)
- Email notification to NOC of HKIX peer causing Proxy ARP (will be implemented soon)

Better Control of Proxy ARP

- Recommendation:
  - Disable Proxy ARP COMPLETELY!!
  - No restricted or unrestricted Proxy ARP
- Cisco IOS:
  - Configuration at interface:
    - no ip proxy-arp
  - Verification:
    - show ip interface | include Proxy ARP
    - “Proxy ARP is disabled”
- Juniper JUNOS:
  - Proxy ARP is not enabled by default
  - So do NOT configure restricted or unrestricted mode Proxy ARP

L2 Control for HKIX Peering LAN

- Traffic Allowed in HKIX Peering LAN:
  - Ethernet Types
    - 0x0800 - IPv4
    - 0x0806 - ARP
    - 0x86dd - IPv6
  - Unicast Only
    - No multicast or broadcast except ARP broadcast
  - Port Security Always On
    - One MAC address one port

Advanced Route Server Feature

Feature	BGP Standard Community
Send prefix to all	4635:4635
Send prefix to $Peer-AS only	4635:$Peer-AS
Do not send prefix to all	0:4635
Do not send prefix to $Peer-AS	0:$Peer-AS

- Target for Q1 of 2018
- Support 2-byte AS numbers only
- Default sending prefix to all if no BGP community is tagged

Support of Blackholing for Anti-DDoS on HKIX Route Servers

HKIX route servers support **Remote Triggered Black Hole Filtering** (RTBH) for announcement of black-hole filtering

http://www.hkix.net/hkix/anti-ddos.htm

No. of ASNs Participated : 40

How it works?

- The victim's address must be included in the participant filter on the HKIX route servers for BGP announcement
- Participant tag the /32 prefix with 4635:666 for its customer
- HKIX route servers set the prefix with next hop 123.255.90.66
- Other HKIX participants accept the /32 prefix and set the next hop address for 123.255.90.66 to null

Expected Results:

- Only the victim (/32) will be unreachable via HKIX network while saving the others
- The DDoS traffic will be black-holed at the side of the participating routers which are closer to the DDoS traffic sources

Support of Blackholing for Anti-DDoS on HKIX Route Servers (BEFORE)

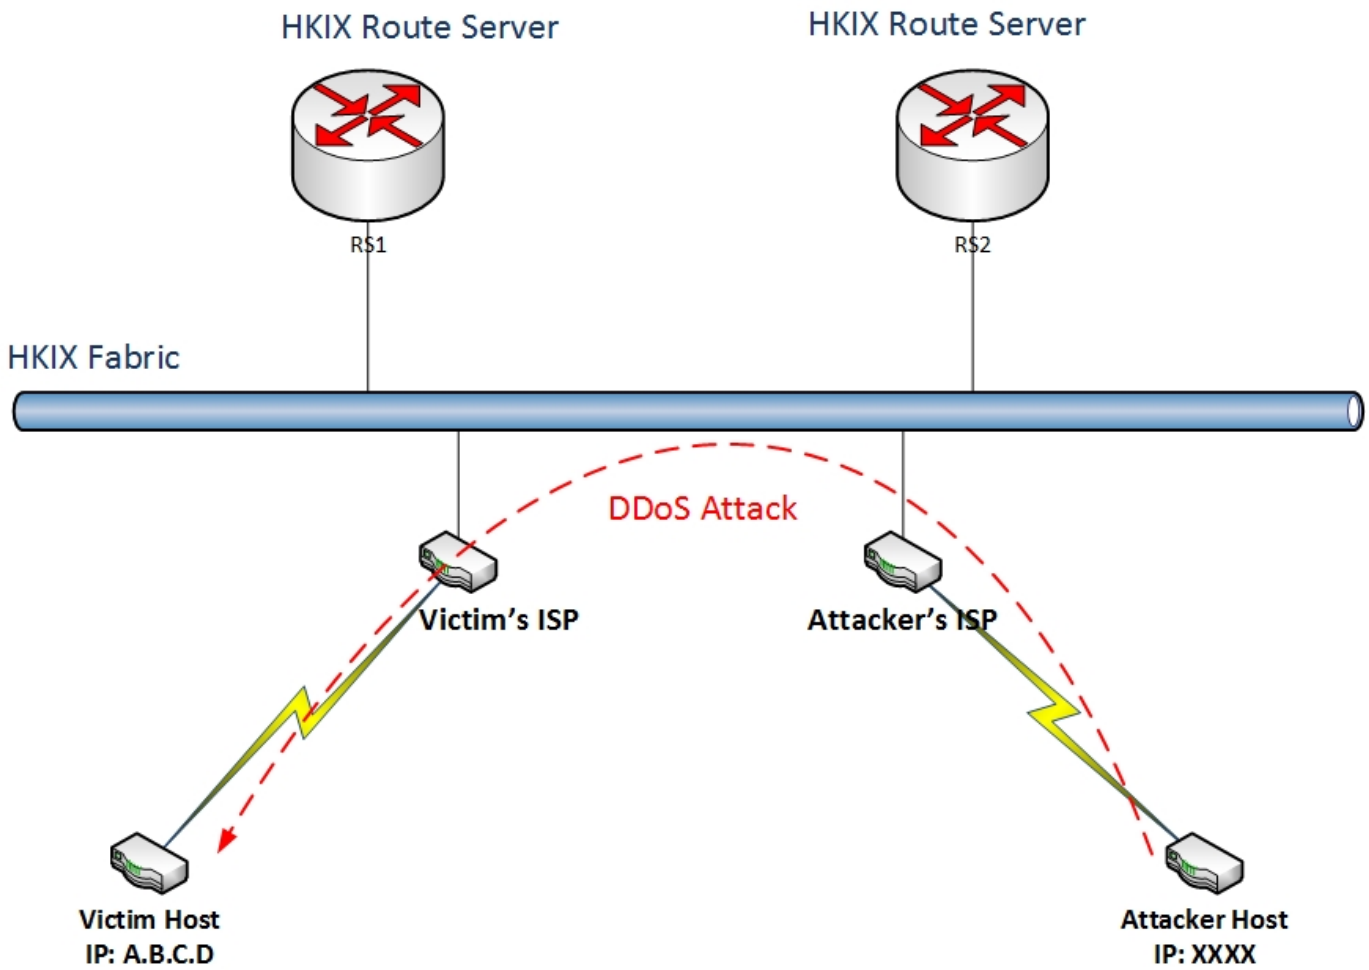

Support of Blackholing for Anti-DDoS on HKIX Route Servers (AFTER)

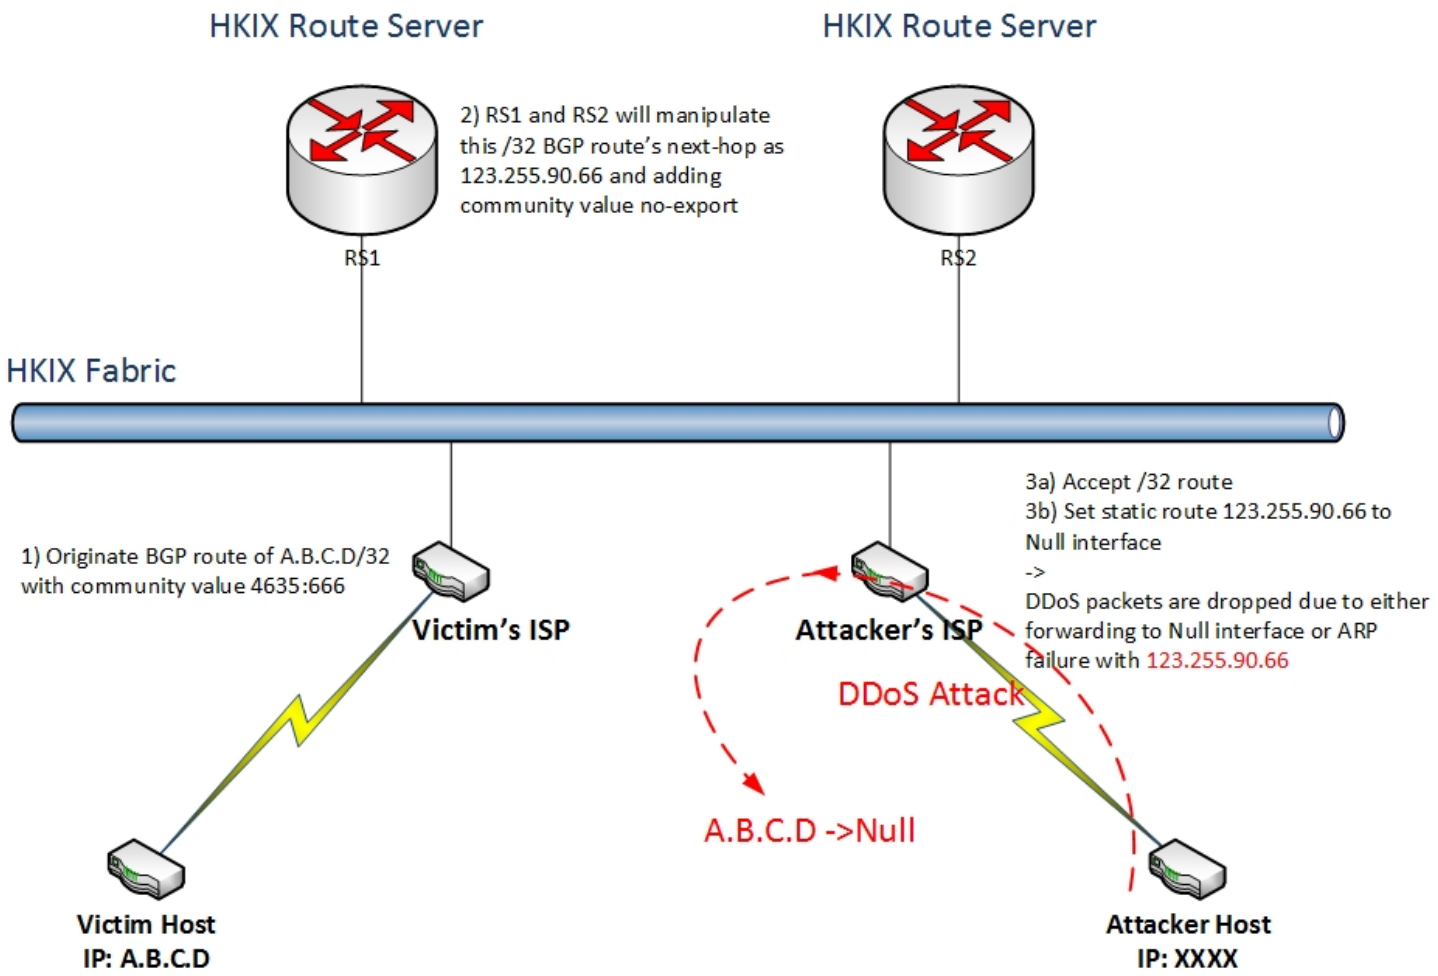

Support of Hiding AS4635 from HKIX Route Servers

- Hiding AS4635 (ASN of HKIX) on the AS Path in the BGP routes received from HKIX route servers
- Support both IPv4 and/or IPv6

HKIX Participant should proceed the following steps:

1. Disable BGP Enforce the First Autonomous System Path on your HKIX peering router
   - Sample configuration for Cisco routers:
     Router(config)# router bgp <Your-ASN>
     Router(config-router)# no bgp enforce-first-as
2. Notify HKIX for hiding AS4635 in the BGP routes
3. HKIX will hide the AS4635 on the AS Path for the IPv4 and/or IPv6 routes sending from HKIX route servers to your HKIX peering

Portal for HKIX Participants

- https://portal.hkix.net
- Functions:
  - Change Port Security
  - MRTG Statistics
    - Physical port
    - LAG port
    - Aggregated per Customer
  - Schedule Maintenance Window
- Contact HKIX Team at provision@hkix.net for pilot testing of HKIX Portal

Portal for HKIX Participants

- Login Page (URL: https://portal.hkix.net/)

HKIX Portal – Port Security

- Change port security

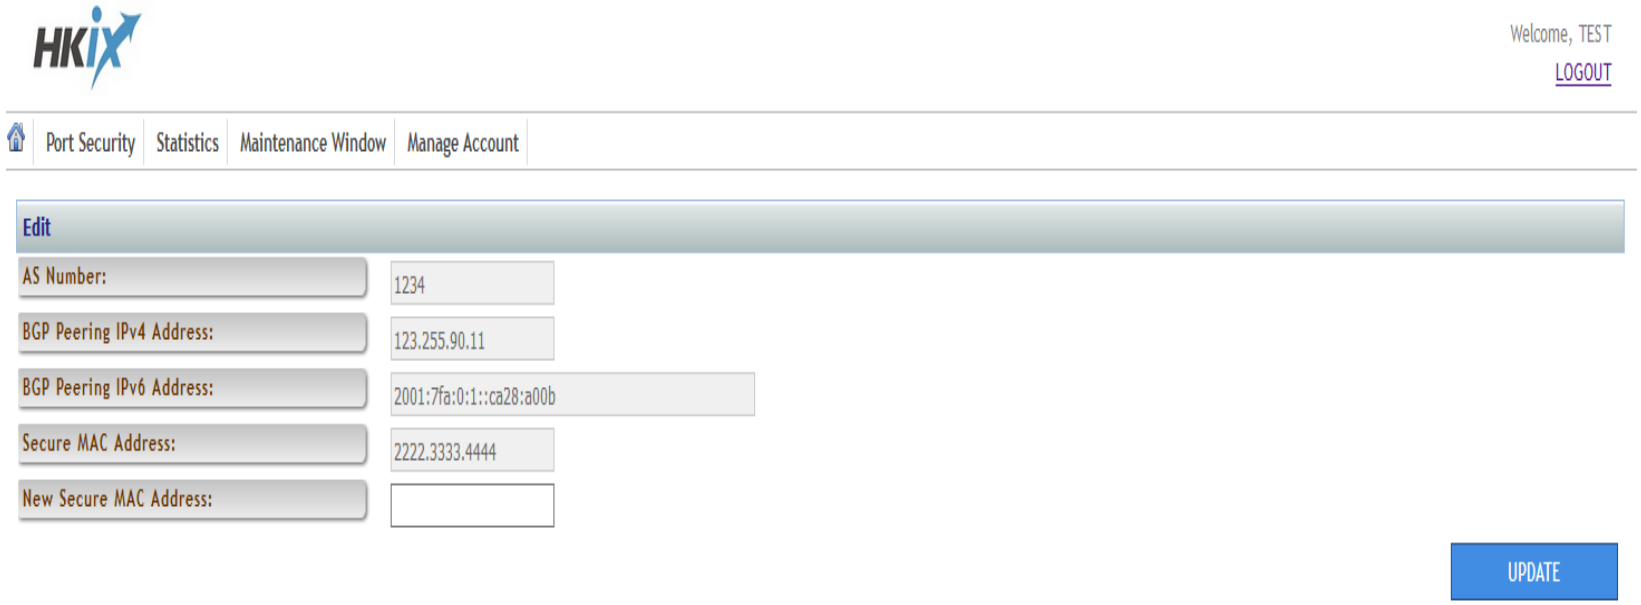

HKIX Portal – MRTG Statistics

- Review an individual statistics / HKIX total statistics

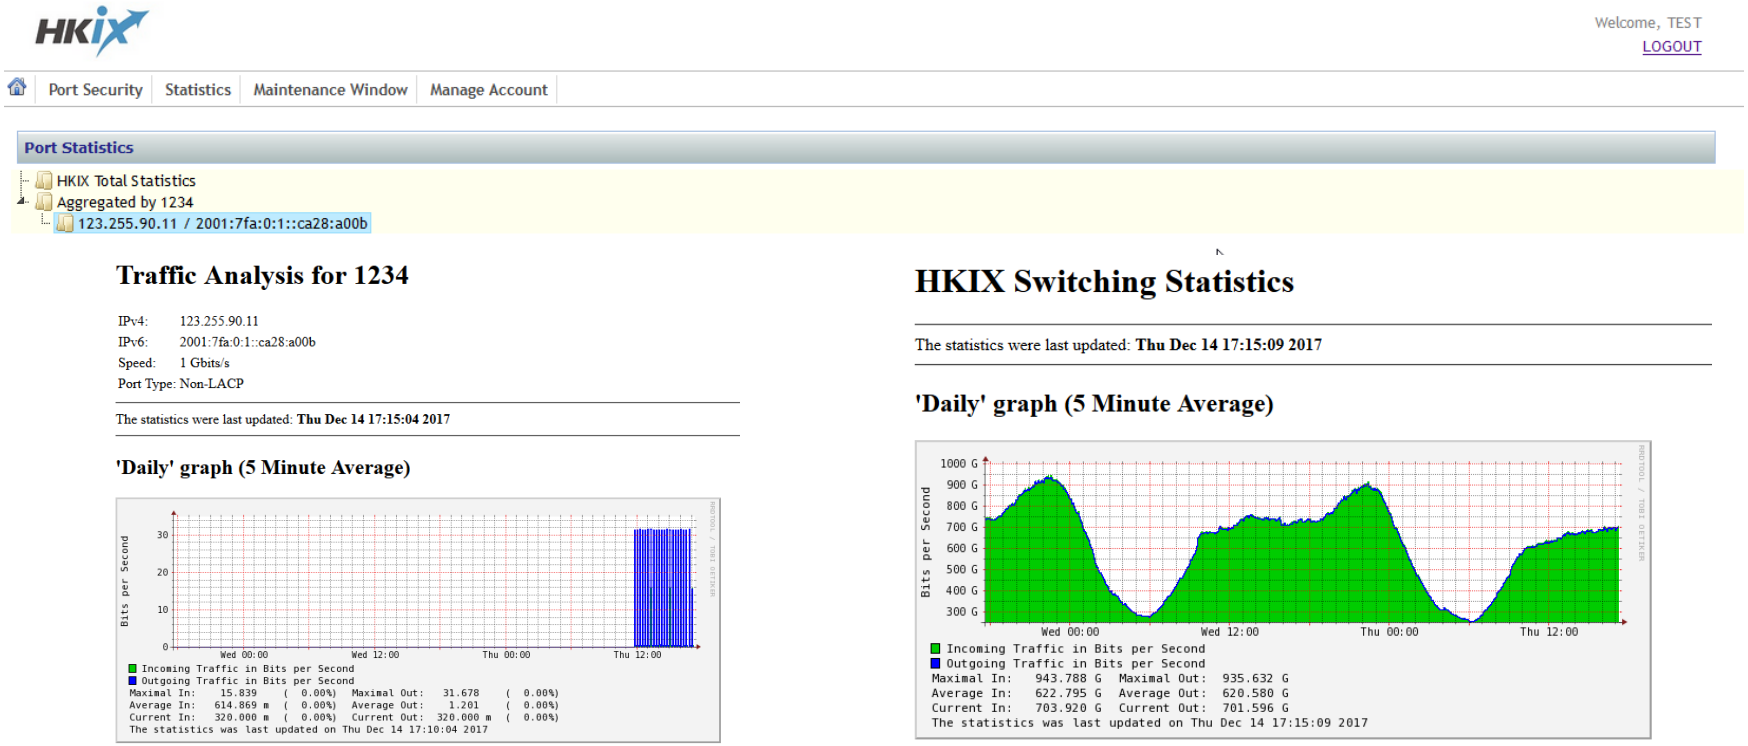

HKIX Portal - Maintenance Window

- Schedule Maintenance Window

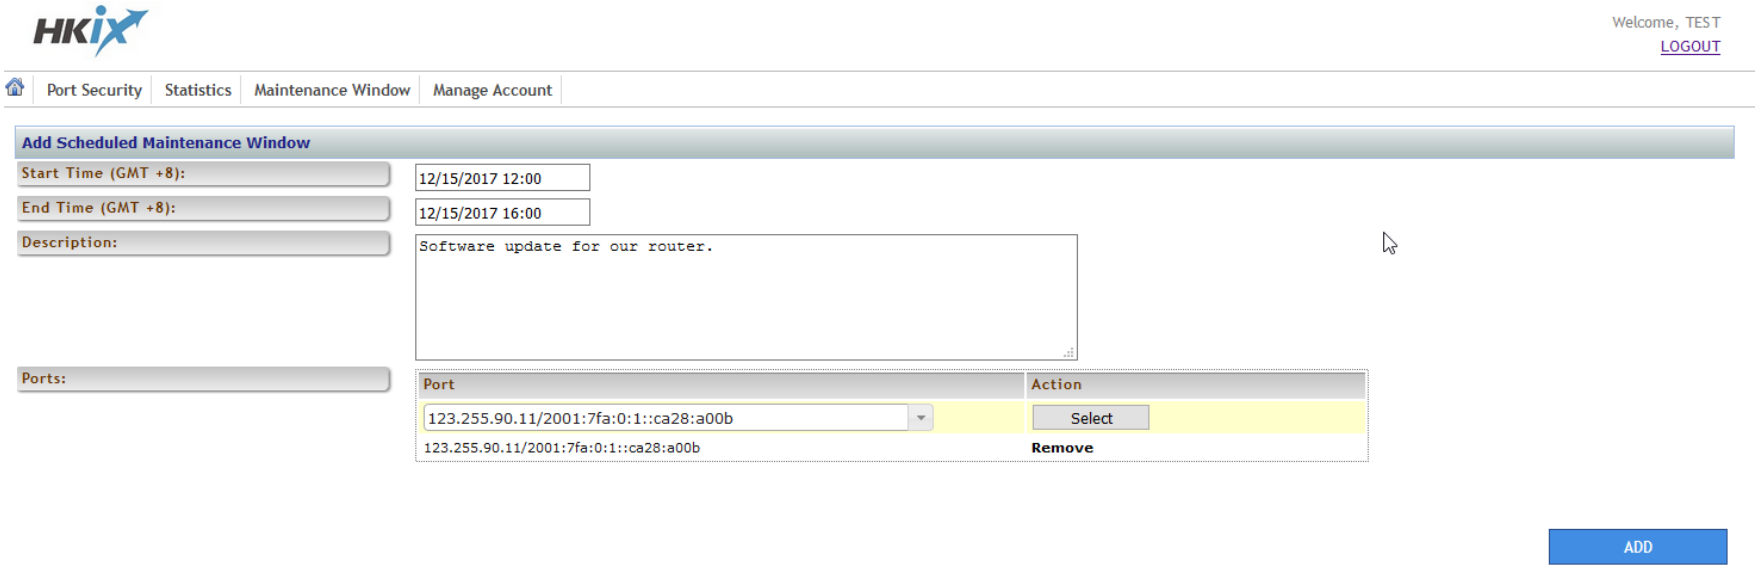

24x7 HKIX NOC

- Full operation starting Q1 of 2017
- Contact us at noc@hkix.net for security or operational related matters
- Keep your contact point at HKIX updated for security incident reporting

Other Operational Tips

HKIX Participants SHOULD NOT:

- Perform testing or looping on HKIX networks
- Announce full/default route to HKIX route servers
- Advertise HKIX peering LAN to other networks
- Forward link-local protocols to HKIX Peering LAN
  - IRDP
  - ICMP redirects
  - IEEE 802 Spanning Tree
  - Vendor proprietary protocols such as discovery protocols: CDP, EDP
  - VLAN/ Trunk protocols: VTP, DTP
  - Interior routing protocol broadcasts (e.g. OSPF, ISIS, IGRP, EIGRP)
  - BOOTP/DHCP
  - PIM-SM
  - PIM-DM
  - DVMRP
  - ICMPv6 ND-RA
  - UDLD
  - L2 Keepalives

Other Operational Tips

HKIX Participants SHOULD DO:

- Make sure proxy ARP is disabled
- Establish BGP MLPA peering with BOTH HKIX route servers
- Notify HKIX NOC for schedule maintenance in advance so that we will not treat your BGP session down as failure
- Monitor the growth of number of prefixes from our route servers and adjust your max prefix setting accordingly
- Monitor the utilization of your links closely and do upgrade before they are full
- Do your own route / route6 / as-set objects on IRRDB and keep them up-to-date
- Do update your contact and peering info in PeeringDB

Thank You!

For enquiries, please contact us at

info@hkix.net