



HKIX Updates at PHNOG

*Che-Hoo CHENG
CUHK/HKIX
15 Jun 2015*

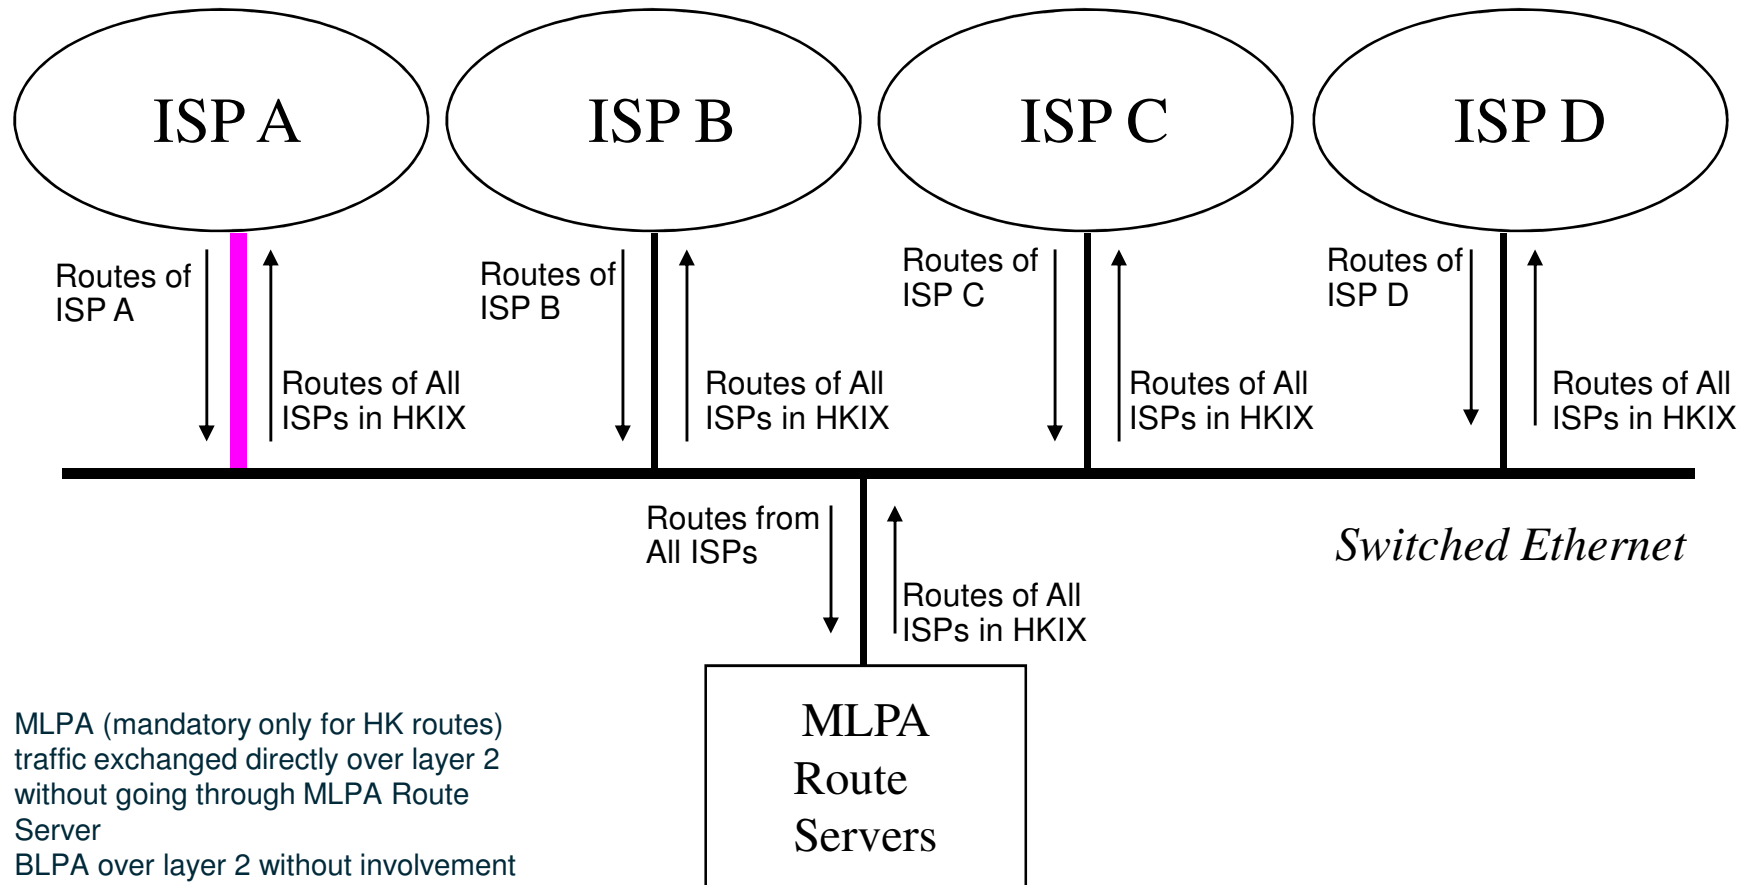


What is HKIX?

- HKIX is a public Internet Exchange Point (IXP) in Hong Kong
- HKIX is the main IXP in HK where various networks can interconnect with one another and exchange traffic
 - Not for connecting to the whole Internet
- HKIX was a project initiated by ITSC (Information Technology Services Centre) of CUHK (The Chinese University of Hong Kong) and supported by CUHK in Apr 1995 as a community service
 - Still fully supported and operated by CUHK
 - 20th Anniversary
- HKIX serves both commercial networks and R&E networks
- The original goal is to keep intra-HongKong traffic within Hong Kong



HKIX Model — MLPA over Layer 2 + BLPA



- MLPA (mandatory only for HK routes) 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



20th Anniversary of HKIX



- HKIX started with thin coaxial cables in Apr 1995
 - *Gradually changed to UTP cables / fibers with switch(es)*
 - *low-end -> high-end*
 - *One switch -> multiple switches*
- Participants had to put co-located routers at HKIX sites in order to connect
 - *Until Metro Ethernet became popular*
- It was a free service
 - *Now a fully chargeable service for long-term sustainability*



Help Keep Intra-Asia Traffic within Asia



- We have almost all the Hong Kong networks
 - We are confident to say we help keep 98% of intra-Hongkong traffic within Hong Kong
- So, we can attract participants from Mainland China, Taiwan, Korea, Japan, Singapore, Malaysia, Thailand, Indonesia, **Philippines**, Vietnam, India, Bhutan and other Asian countries
- We now have more non-HK routes than HK routes
 - On our MLPA route servers
 - Even more non-HK routes over BLPA
- We do help keep intra-Asia traffic within Asia
- In terms of network latency, Hong Kong is a good central location in Asia
 - ~50ms to Tokyo
 - ~30ms to Singapore
- So, HKIX is good for intra-Asia traffic
- HKIX does help HK maintain as one of the Internet hubs in Asia



Connectivity from Here



```
$ traceroute www.hkix.net
traceroute to www.cuhk.edu.hk (137.189.11.73), 64 hops max, 52 byte packets
 1 172.16.32.1 (172.16.32.1) 3.075 ms 9.549 ms 7.533 ms
 2 192.168.254.1 (192.168.254.1) 1.052 ms 3.014 ms 0.981 ms
 3 202.57.50.113 (202.57.50.113) 1.998 ms 1.541 ms 1.446 ms
 4 121.127.7.109 (121.127.7.109) 20.760 ms 1.837 ms 1.456 ms
 5 ip-converge.12.127.121.in-addr.arpa (121.127.12.233) 17.698 ms 16.726 ms 16.633 ms
 6 cuhk-10g.hkix.net (123.255.91.50) 18.978 ms 18.487 ms 18.991 ms
 7 137.189.192.250 (137.189.192.250) 28.506 ms 19.351 ms 18.842 ms
 8 * * *
```

^C

```
$ ping www.hkix.net
```

```
PING www.cuhk.edu.hk (137.189.11.73): 56 data bytes
64 bytes from 137.189.11.73: icmp_seq=0 ttl=246 time=21.739 ms
64 bytes from 137.189.11.73: icmp_seq=1 ttl=246 time=19.531 ms
64 bytes from 137.189.11.73: icmp_seq=2 ttl=246 time=25.797 ms
64 bytes from 137.189.11.73: icmp_seq=3 ttl=246 time=19.143 ms
64 bytes from 137.189.11.73: icmp_seq=4 ttl=246 time=20.476 ms
64 bytes from 137.189.11.73: icmp_seq=5 ttl=246 time=20.447 ms
```

^C

```
--- www.cuhk.edu.hk ping statistics ---
```

```
6 packets transmitted, 6 packets received, 0.0% packet loss
round-trip min/avg/max/stddev = 19.143/21.189/25.797/2.218 ms
```




HKIX Today



- Supports both MLPA (Multilateral Peering) and BLPA (Bilateral Peering) over layer 2
- Supports IPv4/IPv6 dual-stack
- Neutral among ISPs / telcos / local loop providers / data centers / content providers / cloud services providers
- More and more non-HK participants
- >240 ASNs connected
- >430 connections in total
 - 2 x 100GE + >180 x 10GE connections
- ~450Gbps (5-min) total traffic at peak
- Annual Traffic Growth = 30% to 40%



Yearly Traffic Statistics





Charging Model



- An **evolution** from free-of-charge model adopted at the very beginning, to penalty-based charging model based on traffic volume for curbing abuse, to now simple port charge model for fairness and sustainability
- **Have started simple port charge model since 01 Jan 2013**
 - See <http://www.hkix.net/hkix/Charge/ChargeTable.htm>
- Still not for profit
 - HKIX Ltd (100% owned by CUHK) to sign agreement with participants
 - Target for fully self-sustained operations for long-term sustainability



HKIX Port Charge Table (effective 01 Jul 2014)



Standard Charges		NRC		MRC	
Port Size	Interface	HKD	USD	HKD	USD
E/FE/GE	UTP or SMF	Waived		936	120
10GE	LR	17,940	2,300	7,800	1,000
	ER	39,000	5,000		
	ZR	62,400	8,000		
100GE	LR4	117,000	15,000	46,800	6,000

* New E/FE(10M/100M) connections are no longer supported.

** Existing E/FE(10M/100M) connections have to be decommissioned or upgraded to GE or above by 30 Jun 2015

Save-IP Discount (applied ONLY if IP address is NOT needed for the port)		Reduction of MRC for each port	
Port Size	Conditions	HKD	USD
10GE	With LACP enabled; NOT applied to the 1st LACP port which needs IP	-780	-100
100GE	With LACP enabled; NOT applied to the 1st LACP port which needs IP	-4,680	-600

* No such discount for E/FE/GE connections and NRC

Volume Discount (applied under same ASN and under same contract ONLY)		Reduction of MRC for each port	
Port Size	Conditions	HKD	USD
10GE	Applied to the 5th 10GE port and onwards	-780	-100
100GE	Applied to the 3rd 100GE port and onwards	-4,680	-600

* No such discount for E/FE/GE connections and NRC

REMARKS:

NRC=Non-Recurring Charge (Non-Refundable)

MRC=Monthly Recurring Charge



Why HKIX is successful

- Neutral
 - Treat all partners equal, big or small
 - Neutral among ISPs / telcos / local loop providers / data centers / content providers / cloud services providers
- Trustable
 - Fair and consistent
 - Respect business secrets of every partner / participant
- Not for Profit
- *HKIX started very early, well before incumbent telcos started to do ISP business*



The Recent Upgrade Done in 2014



- A new highly-scalable two-tier dual-core spine-and-leaf architecture within CUHK by taking advantage of the new data center inside CUHK campus
 - HKIX1 site + HKIX1b site as Core Sites
 - Fiber distance between 2 Core Sites: <2km
 - Provide site/chassis/card resilience
 - Support 100GE connections
 - Scalable to support >6.4Tbps total traffic using 100GE backbone links primarily and FabricPath
- **Ready to support HKIX2/3/4/5/6/etc as Satellite Sites**
 - Satellite Sites have Access Switches only, which connect to Core Switches at both Core Sites



The Design

- Dual-Core Two-Tier Spine-and-Leaf Design for high scalability
 - Have to sustain the growth in the next 5+ years (to support >6.4Tbps traffic level)
 - Core Switches at 2 Core Sites (HKIX1 & HKIX1b) only
 - No interconnections among core switches
 - Access Switches to serve connections from participants at HKIX1 & HKIX1b
 - Also at Satellite Sites HKIX2/3/4/5/6/etc
 - Little over-subscription between each access switch and the core switches
 - FabricPath (TRILL-like) used among the switches for resilience and load balancing
- Card/Chassis/Site Resilience
 - LACP not supported across chassis though (card resilience only)
- 100GE optics support
 - LR4 for <=10km and ER4-lite for <=25km (4Q2015)
 - Support by local loop providers is key
- Port Security still maintained (over LACP too)
 - Only allows one MAC address / one IPv4 address / one IPv6 address per port (physical or virtual)
- Have better control of Unknown-Unicast-Flooding traffic and other storm control




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



HKIX1 Core Site @CUHK

-----(<2km)-----

HKIX1b Core Site @CUHK





The Migration of Connections

Still in Progress

- New switches in production at HKIX1 site starting Mar 2014
 - Parallel run with old switches, interconnected with the new switches using multiple 100GE links
 - All new connections have been on new access switches since their production
 - While existing connections are being moved to the new access switches one by one
- By early Aug 2014, all 10GE connections had been migrated
- HKIX1b site put into production in Nov 2014 to provide true site resilience
- Remaining GE connections **MUST** be moved to new switches
 - Deadline is **30 Jun 2015**
 - The remaining old switch will be decommissioned after then
 - No E/FE support starting then
- RS1/RS2 & HKIX-R&E have been moved to the new architecture
- HKIX2 will be moved to the new architecture as new Satellite Site



The Upgrade Result in General



- The upgrade was considered successful after very much hard work
- Support 100GE connections starting Oct 2014
- Provide site / chassis / card resilience options for HKIX participants
- Support easy expansion with high scalability
- Support fast deployment for satellite sites



FabricPath



Being Used in New Architecture

- We adopt spine-and-leaf architecture for high scalability
 - Avoid connecting participant ports on core switches
- The Spanning Tree Protocol (STP) domains do not cross into the FabricPath network
- Layer 2 gateway switches, which are on the edge between the CE and the FabricPath network, must be the root for all STP domains that are connected to a FabricPath network
- Load balancing is working fine
 - Even with odd number of links
- Transparent to participants (i.e. no BGP down) when adding/removing inter-switch links



One Very Critical Point for an IXP



- An IXP must NOT be vulnerable to DDoS attack itself
 - Congestion at one port must NOT cause trouble to any other ports
- Network QoS Policy - Congestion Control Mechanisms
 - Default is “Burst optimized” which is not good for IXP because of sharing of buffer by multiple ports
 - “Mesh optimized” is more suitable for IXP



Proxy ARP Threat

- Proxy ARP **MUST be disabled** all the times on the interface connecting HKIX switches
- Participants **MUST** provide the “show ip interface” output for verification before HKIX turning up the switch ports
- Will use Dynamic ARP Inspection (DAI) to maintain static ARP list
 - But not used yet as it is manual
 - Need to input a few commands for this instead of just one command



DDoS Mitigation Trial

- Accept /32's with BGP community 4635:666 from HKIX route servers for MLPA
- Do black-holing by null route the next-hop of HKIX address 202.40.160.66 (*being changed to 123.255.90.66*) to null interface
- When under DDoS attack:
 - Signal other participants by advertising /32's under your managed address blocks with 4635:666 to do black-holing at their networks
 - Relevant sites will be sacrificed but other sites and networks will be saved
- >50 participants did the test
 - All of them had positive results
- Will declare production soon
- Will seek collaboration with similar initiatives for maximum outcome



IPv4 Address Renumbering and Route Servers Upgrade



Migration Date: 12-15 Jun 2015 (Fri-Mon)

IPv4 Address Renumbering

- Network mask is being changed to /21 from /23, for accommodating future growth
- ALL participants must change to **NEW 123.255.88/21**, away from *OLD 202.40.160/23*
- **Parallel run of old and new IPv4 addresses only during the 4-day migration period, having learnt from experience of other IXPs**
- MLPA: New route servers support new IPv4 addresses while existing route servers support old addresses, but IPv6 is handled separately
- BLPA: Individual participants have to coordinate with their peering partners directly
- *No change to IPv6 addresses*

Route Servers Upgrade

- The two old route servers will be decommissioned
- Two new route servers have been installed at HKIX1 and HKIX1b (the two HKIX core sites)
- More route server features will be supported later



IPv4 Address Renumbering and Route Servers Upgrade

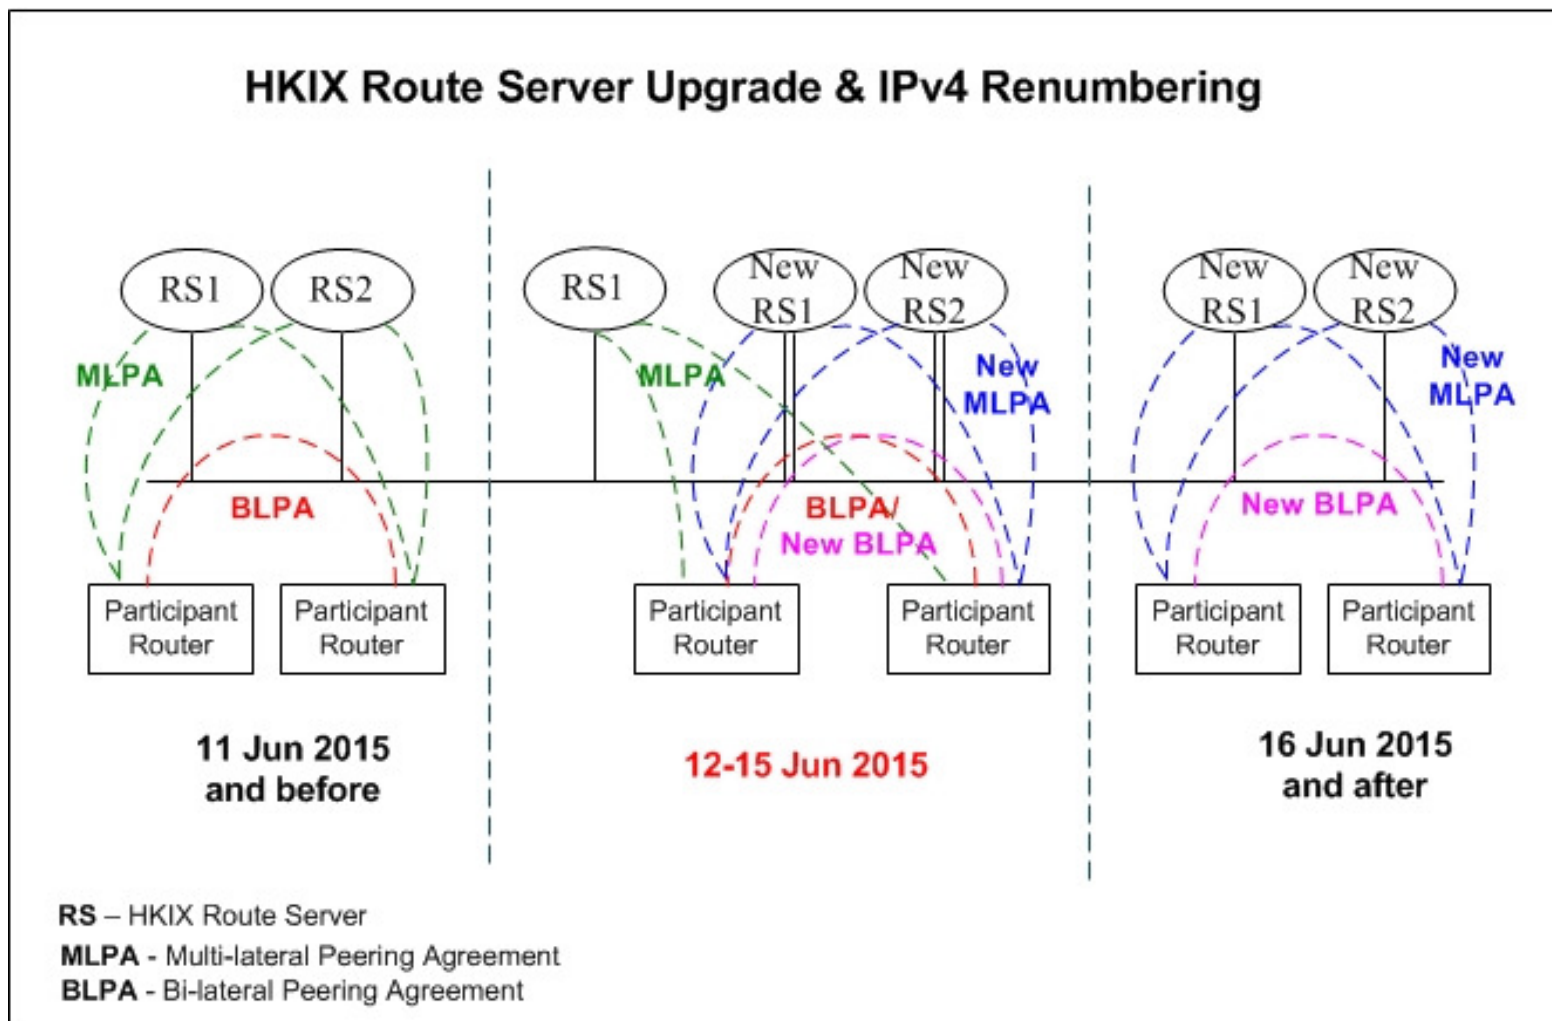


Schedule

DATE	TASKS
11 MAR 2015	HKIX announces IPv4 renumbering and route server upgrade
11 APR 2015	2nd reminders to ALL HKIX participants
04 MAY 2015	HKIX provides the information of NEW IP address block (/21) and publish the mapping of current IPs (/23) and new IPs (/21) on HKIX web site
11 MAY 2015	Final reminders to ALL HKIX participants
12 JUN 2015	HKIX setup the new route servers for MLPA peering
12-15 JUN 2015	HKIX participants change to new IPs and setup MLPA peering with new route servers. Re-establish the BLPA peering (if any) with individual participant(s) with the corresponding IPs in new /21
16 JUN 2015	HKIX decommission the old route servers and IP address block 202.40.160/23 ceased




IPv4 Address Renumbering and Route Servers Upgrade





Traffic During Migration



- For details including the status of the migration, please check:
<http://www.hkix.net/hkix/renumbering.htm>



Setting up Multiple HKIX Satellite Sites



- To be named HKIX2/3/4/5/6/etc
- Allow participants to **connect to HKIX more easily at lower cost** from those satellite sites
- Open to all commercial data centres in HK which fulfil minimum requirements so as to maintain neutrality which is the key success factor of HKIX
 - ISO27001 requirement
 - Minimum size requirements
 - Requirements on circuits connecting back to the two HKIX core sites
 - Non-exclusive
- Intend to create win-win situation with satellite site partners
- **Contact us if you are interested in this**
- *NOTE: HKIX1 and HKIX1b (the two HKIX core sites) will continue to serve participants directly*



Planned Work in 2015-2016



- Introduce advanced Route Server functions
- Better Control Proxy ARP
- Better support for DDoS Mitigation
- More L2 ACL on HKIX peering LAN
- Portal for HKIX participants
 - Port info and traffic statistics
 - Self-service port security update
 - Network maintenance schedule
- Improve after-hour support
- ISO27001



Thank you!