



HKIX IPv4 Address Renumbering from /23 to /21 – Experience Sharing

Che-Hoo CHENG

CUHK/HKIX

08 Sep 2015



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*



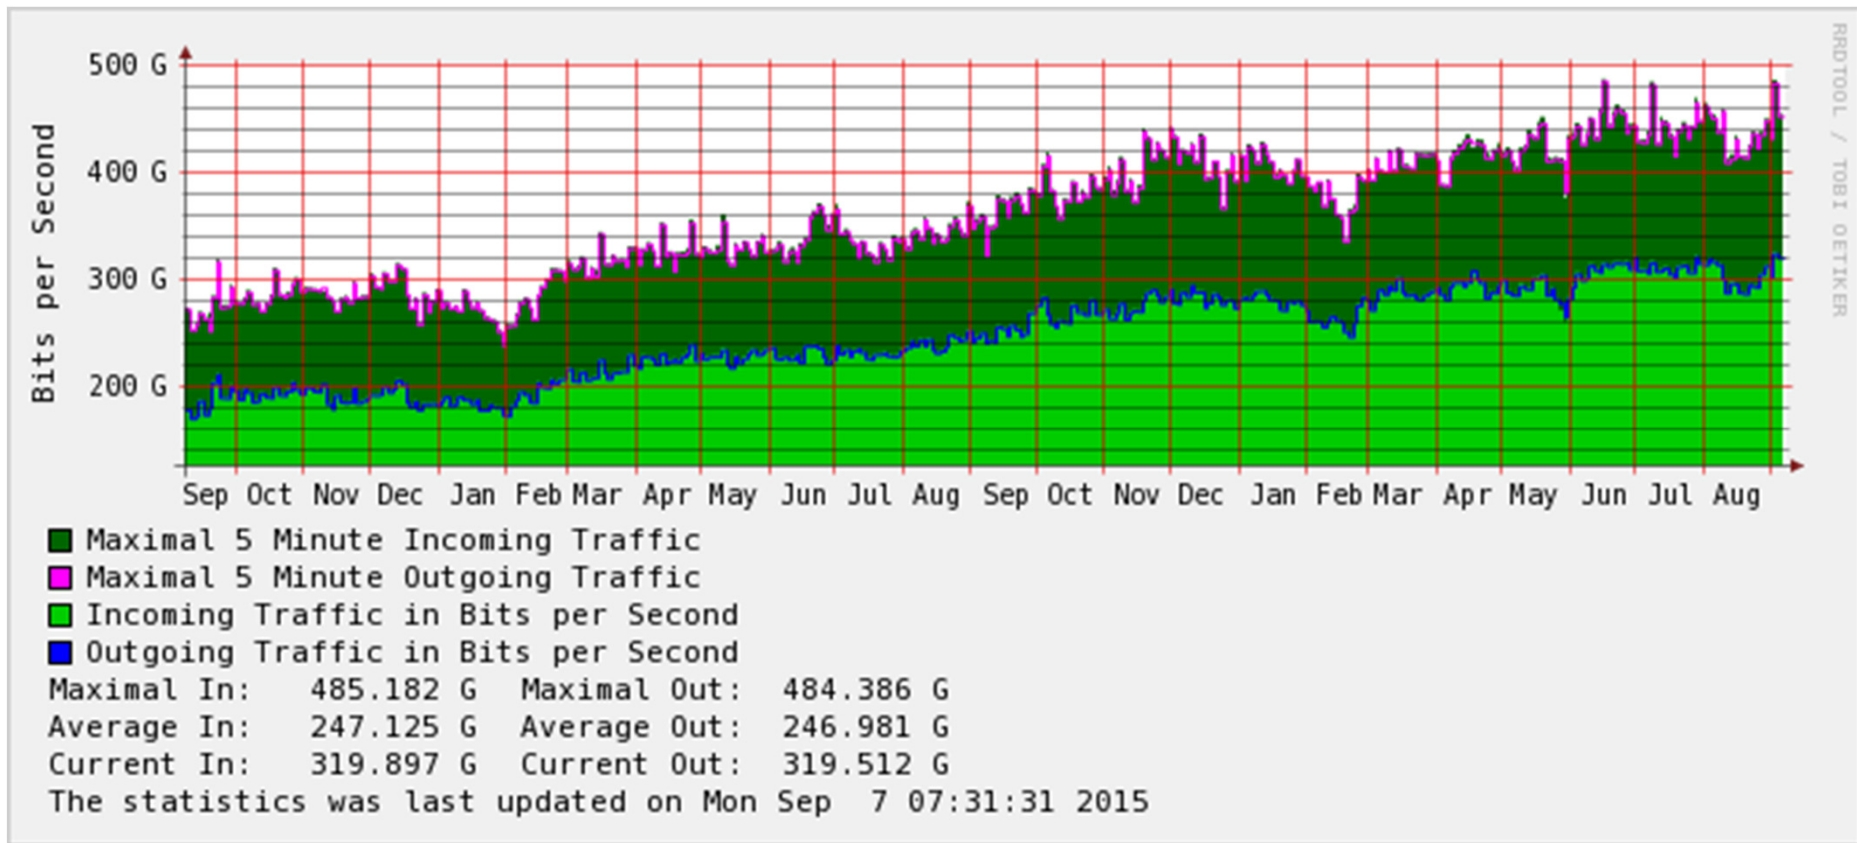
HKIX Today



- 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
- >230 AS'es connected
- >420 connections in total
 - 2 x 100GE + >**190 x 10GE** + >220 x GE
- ~485Gbps (5-min) total traffic at peak
- Annual Traffic Growth = 30% to 40%



Yearly Traffic Statistics





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



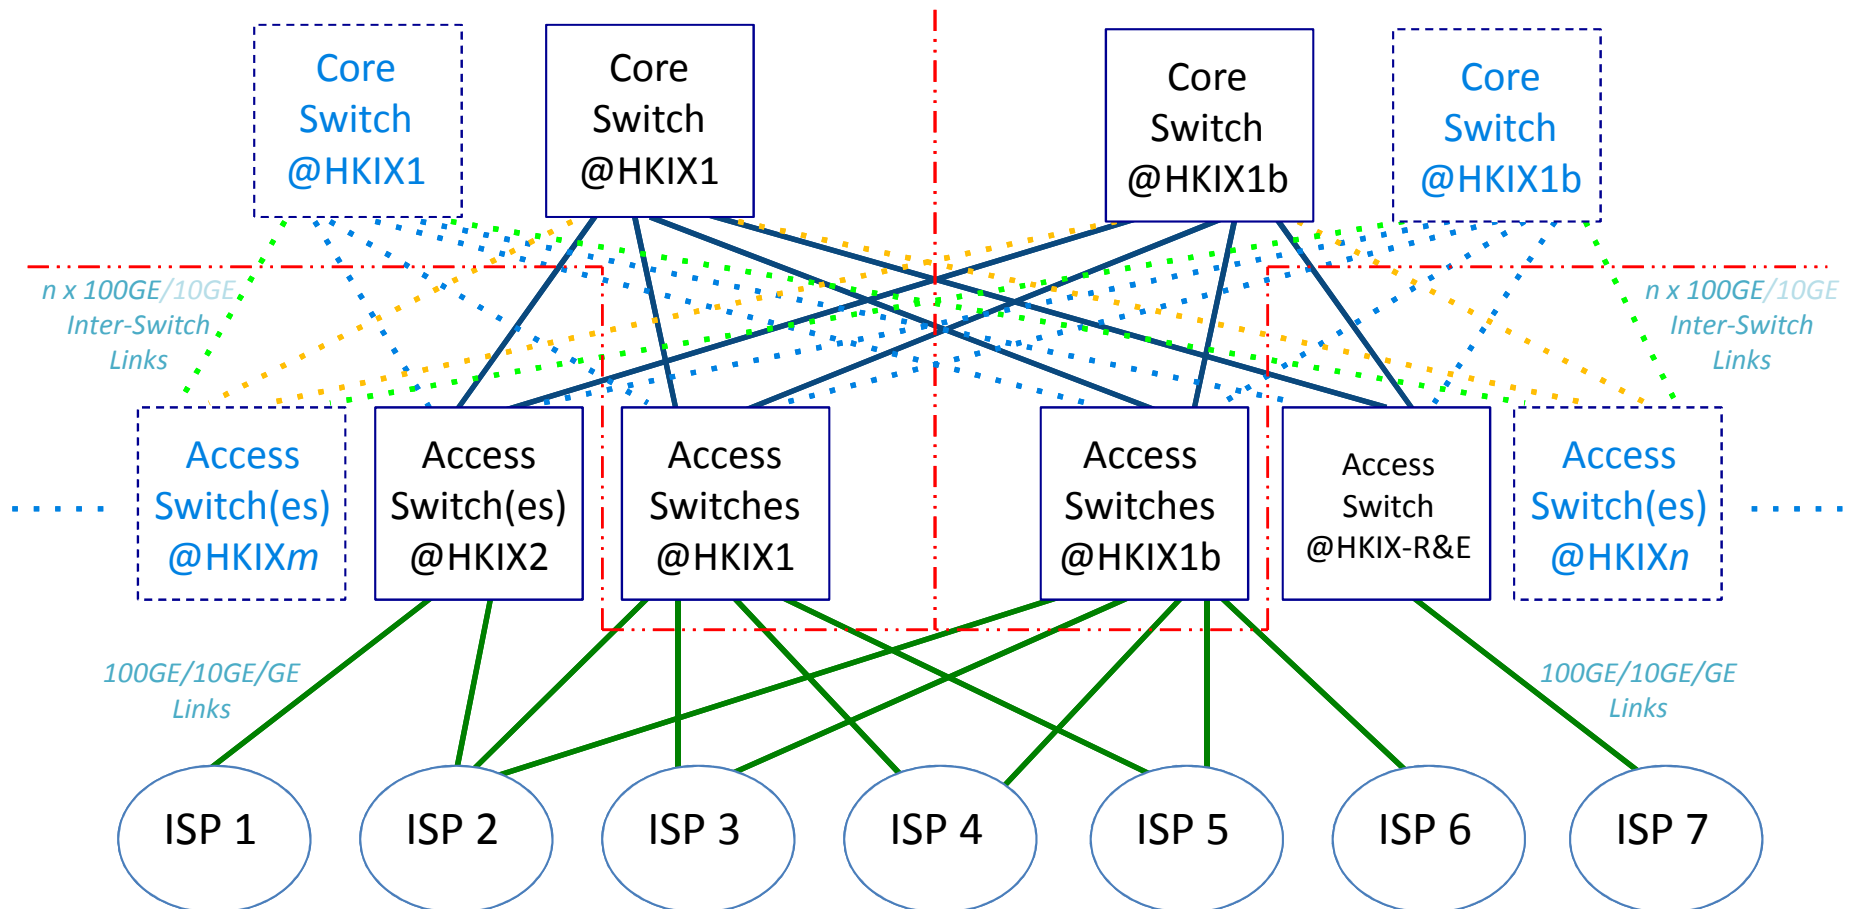
New HKIX Dual-Core Two-Tier Spine-and-Leaf Architecture For 2014 and Beyond



HKIX1 Core Site @CUHK

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

HKIX1b Core Site @CUHK





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



IPv4 Address Renumbering and Route Servers Upgrade



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

IPv4 Address Renumbering

- Network mask was changed to /21 from /23, for accommodating future growth
- ALL participants had to 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 old route servers supported old addresses, but IPv6 was handled separately
- BLPA: Individual participants had to coordinate with their peering partners directly
- *No change to IPv6 addresses*

Route Servers Upgrade

- The two old route servers were decommissioned at the end of the period
- Two new route servers had 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



Considerations beforehand:

- Peak traffic level: ~450Gbps (5-min average)
- # of prefixes on MLPA route servers: ~80K IPv4 prefixes & ~12K IPv6 prefixes
- Complexity of migration: participants from many different time zones & 330+ BGP sessions
- Have to minimize topology changes and configuration changes to participants
- Also need to care about bilateral peering – both old and new networks on the same VLAN
- Have to take care of capacity requirements and routing performance if transit is to be provided between old and new networks

Three options had been looked into:

- **Big Bang Approach** – Pros: Minimum effort to HKIX / Cons: Need coordination with ALL participants for aligning the maintenance window which is extremely difficult
- **Parallel Run with Transit** – Pros: Easier for participants / Cons: Transit routers would need to handle huge traffic of up to 300Gbps and would not be able to support BLPA across old and new networks
- **Parallel Run with Secondary Address** – Pros: Flexible changing time as secondary address can be configured before migration / Cons: Participants need to configure 2nd address on all the router interfaces connecting to HKIX



IPv4 Address Renumbering and Route Servers Upgrade



After careful studies and making reference to other IXPs around the world, we finally decided to take the approach of **Parallel Run with Secondary Address + Transit Router** (for backup and contingency) and do the renumbering within **4-day period** (Fri to Mon)



IPv4 Address Renumbering and Route Servers Upgrade



Communication Part:

Before Migration

- 3-4 months – Announced the address renumbering at APRICOT-APAN 2015 and then HKN0G 1.1
- 3 months – Made the announcement by emails to all HKIX participants without detailed info and requested them to provide their contact points for the IPv4 renumbering tasks
- 2-3 months – Replied acknowledged participants and let them know that a migration webpage had been established and the latest information would be published there
- 2 months – Sent reminders to participants who had not respond through all contact points (i.e. contractual / billing / technical contacts) as their commitment to the address renumbering would be very important to the whole project
- 5 weeks – Provided the information of new IP addresses and published the mapping of old address to new address on the migration webpage
- 4 weeks – Published final schedule, migration details, sample configurations and FAQs
- 3 weeks – Sent individual emails to participants and asked them to confirm and specify the intended migration time within the 4-day period
- 1-2 weeks – Followed up again if reply was still not received from the participants
- 1 week – Set up the Command and Control Center (CCC) and ensure that all email templates were ready in place
- 1 day – CCC in operations, 24-hour technical team standby



IPv4 Address Renumbering and Route Servers Upgrade



Communication Part:

During Migration

- Closely monitored the migration progress and escalated the cases to technical team in case problem reported by HKIX participants
- Provided the latest renumbering status on migration webpage and let participants know the up-to-date progress

After Migration

- Followed up with participants and requested them to remove the old addresses from their router interfaces



IPv4 Address Renumbering and Route Servers Upgrade



Technical Part:

Before Migration

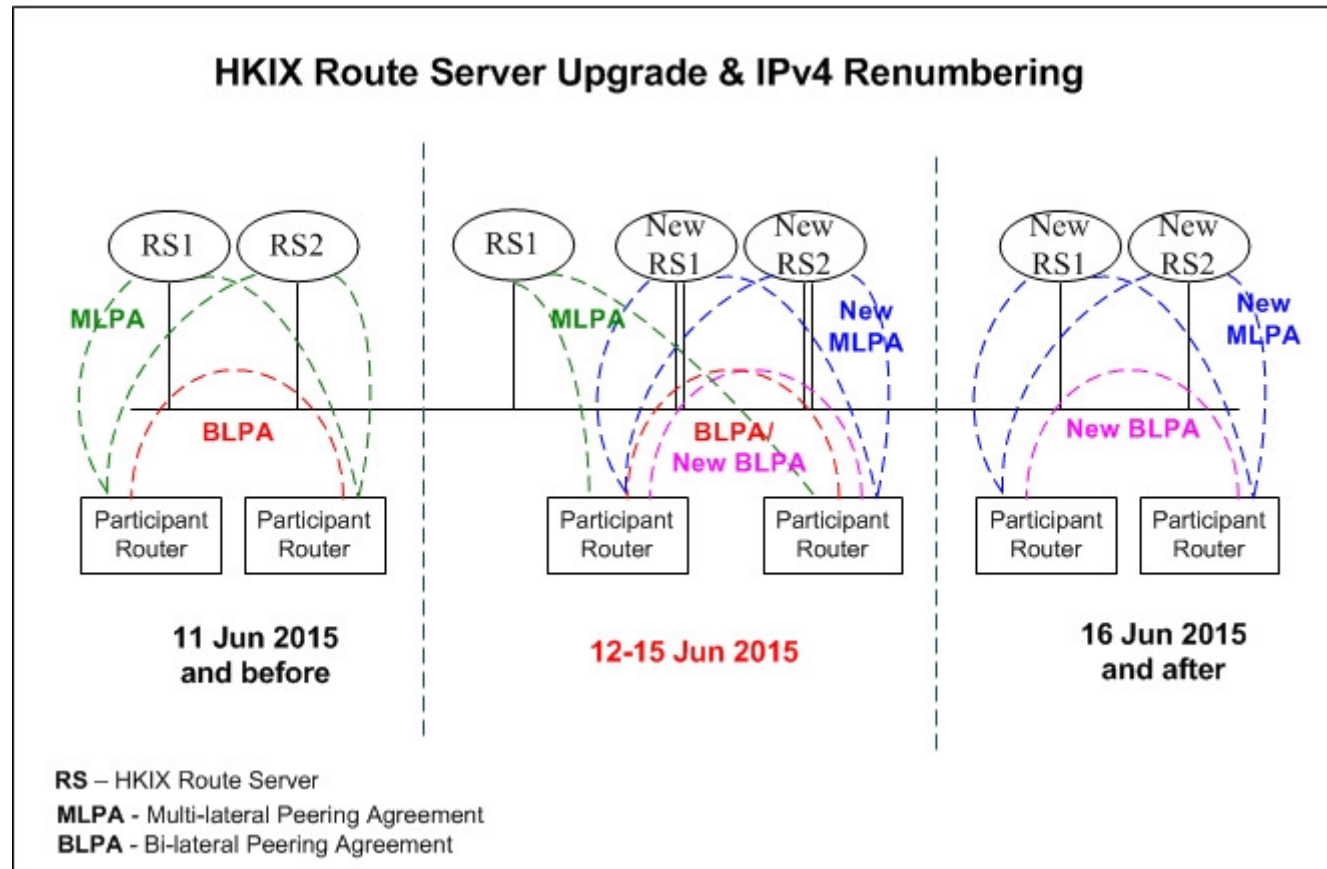
- 2 months – Tested the equipment in lab and did simulation with different scenarios
- 1 month – Equipment trial run and final acceptance test
- 3 weeks – Installed the new route servers & backup transit router
- 2 weeks – Replaced RS2 with new route server using OLD address
- 1 week – Deployed new RS1 and invited some participants for pilot testing

During Migration

- Start of Day 1 – Re-configured RS2 to use NEW address; Set up new RS1 with NEW address; old RS1 still in production
- Day 1-4 – Set up BGP sessions with participants on new RS1 & RS2; Parallel run of new and old Route Servers
- Day 1-4 – Monitored the traffic and the overall progress with the migration schedule provided by participants
- Day 1-4 – Provided instant technical assistance (including trouble-shooting) to participants in case they had difficulties in setting up the BGP sessions
- Day 1-4 – No observable traffic drop during the period
- End of Day 4 – Shut down and decommissioned old RS1 & RS2



IPv4 Address Renumbering and Route Servers Upgrade





IPv4 Address Renumbering and Route Servers Upgrade



Lessons Learnt:

- The Key to Success is Good Planning and Good Communication
- Parallel run for migration is a must but there is no need to do parallel run for too long
- Making contact with all the participants is most time-consuming but is also most important

Many thanks to the whole HKIX Team and all the HKIX participants involved



Setting up Multiple HKIX Satellite Sites



- Allow participants to connect to HKIX more easily at lower cost from those satellite sites in Hong Kong
- 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 collaborators
- To be named HKIX2/3/4/5/6/etc
- **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 of 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!