HKIX Updates at APAN 40 in KL

Che-Hoo CHENG
CUHK/HKIX
13 Aug 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

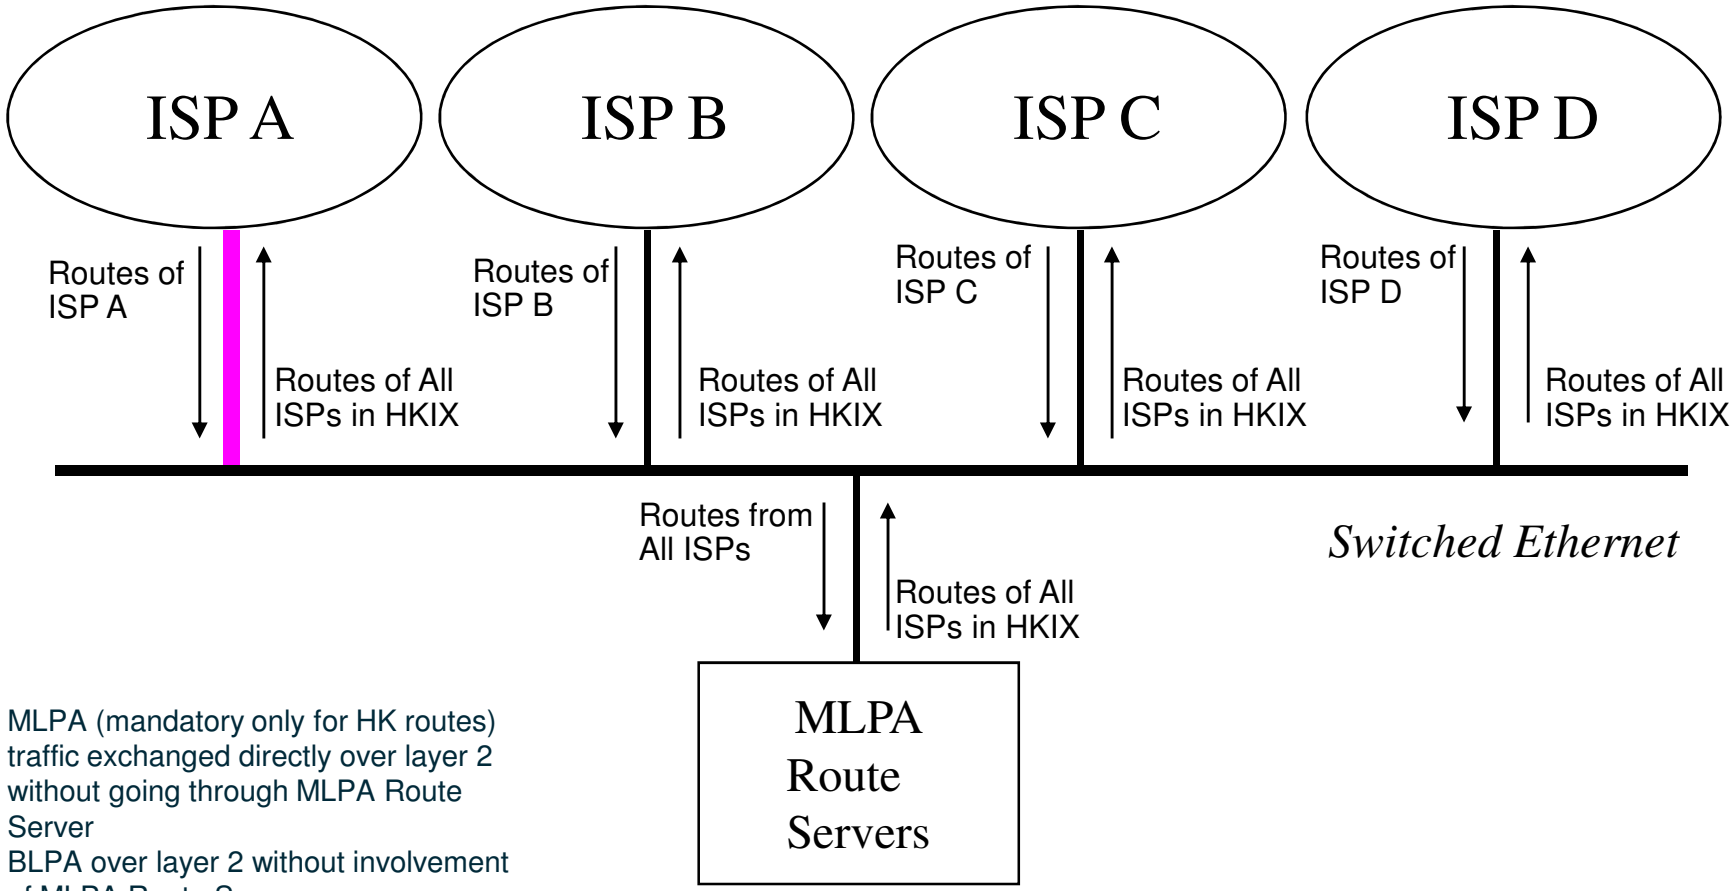

- MLPA (mandatory only for HK routes) traffic exchanged directly over layer 2 without going through MLPA Route Server
- BLPA over layer 2 without involvement of MLPA Route Server
- Supports both IPv4 and IPv6 over the same layer 2 infrastructure

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

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

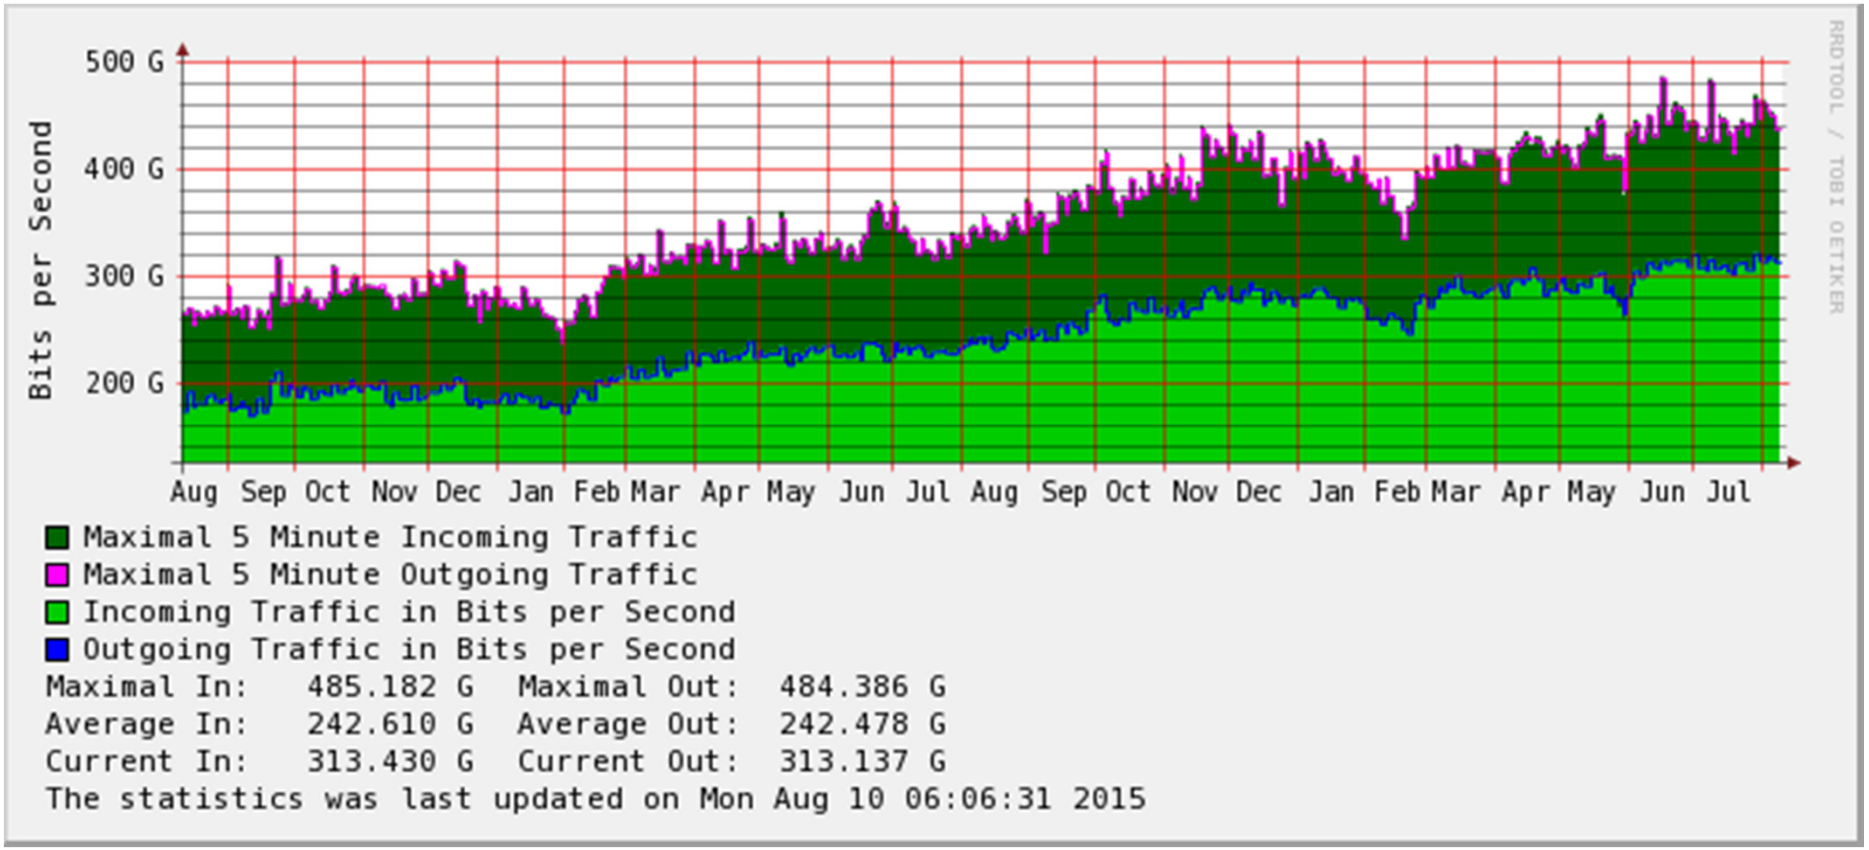

Charging Model

- 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 Charge Table (v1.2)

Standard Port Charges						NRC		MRC	
Port	Interface	Standard Interface	Availability			HKD	*USD*	HKD	*USD*
			HKIX1	HKIX1b	Satellite Sites				
GE	T	Yes	✔			Waived		**936**	*120*
	SX	Yes			✔				
	LX/LH	Yes		✔	✔				
	EX	No		✔		9,360	*1,200*		
	ZX	No		✔		15,600	*2,000*		
10GE	SR	Yes			✔	**10,140**	*1,300*	**7,800**	*1,000*
	LR	Yes	✔	✔	✔	**17,940**	*2,300*		
	ER	No	✔	✔		39,000	*5,000*		
	ZR	No	✔	✔		62,400	*8,000*		
100GE	LR4	Yes	✔	✔	Some	**117,000**	*15,000*	**46,800**	*6,000*
	ER4-Lite	No	✔	✔		468,000	*60,000*		

** Satellite Sites are to be named as HKIX2/3/4/5/6/etc which will be announced soon*
*** E/FE(10ME/100ME) connections are no longer supported*
**** There may be long lead-time for non-standard interfaces (GE-EX, GE-ZX, 10GE-ER, 10GE-ZR and 100GE-ER4-Lite)*
***** The port charges listed do **NOT** cover local circuit/loop charges, cross-connect charges, satellite-site special connection charges, or any other charges needed for making the connection*

Save-IP Discount (applied **ONLY** if IP address is **NOT** needed for the port)		Reduction of MRC for each port entitled	
Port	Conditions	HKD	*USD*
10GE	With Link Aggregation over multiple ports; **NOT** applied to the 1st port which needs IP address	**-780**	*-100*
100GE	With Link Aggregation over multiple ports; **NOT** applied to the 1st port which needs IP address	**-4,680**	*-600*

** No such discount for GE connections and NRC*
*** With Link Aggregation over multiple ports, only card resilience can be provided but not chassis resilience and site resilience*

Volume Discount (applied under the same AS and the same contract **ONLY**)		Reduction of MRC for each port entitled	
Port	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 GE connections and NRC*

REMARKS:
***NRC** = Non-Recurring Charge (**NON**-refundable & **NON**-transferrable to other AS or other company under different name)*
***MRC** = Monthly Recurring Charge*

Updated on 30 Jun 2015

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

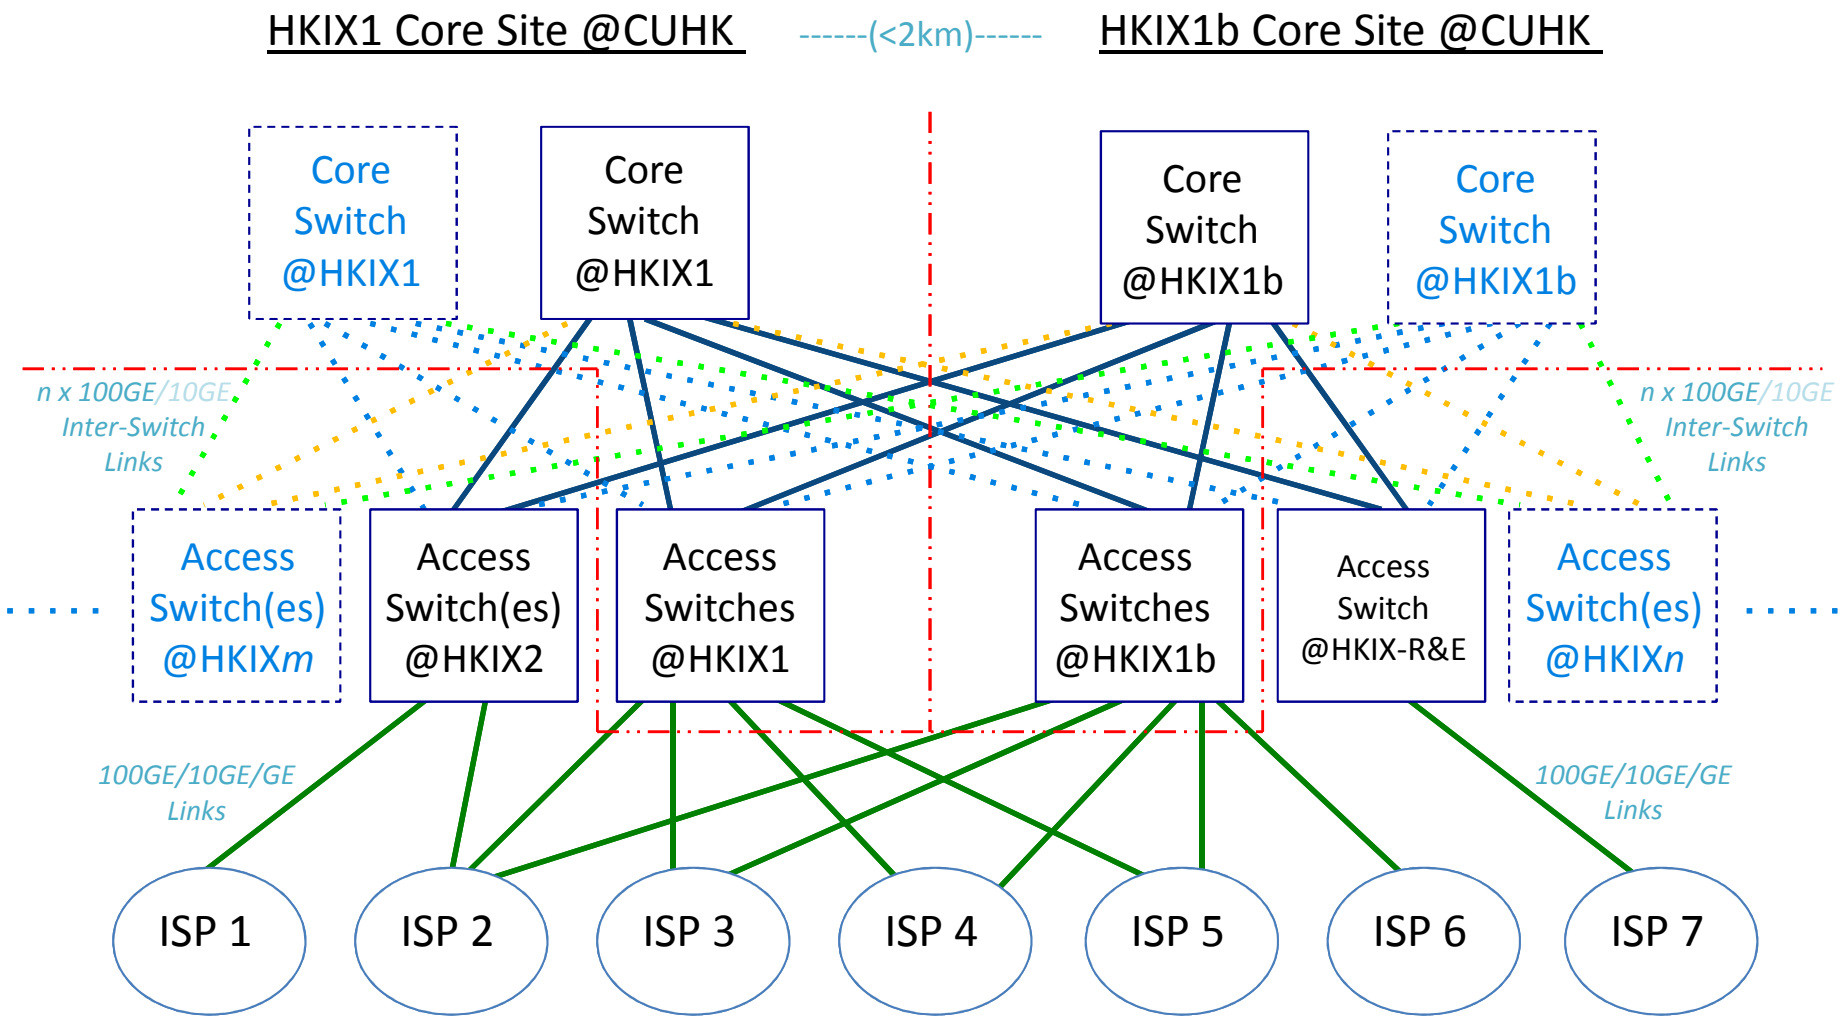

FabricPath Being Used in New Architecture

- 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 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

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
- *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

Special Support for R&E Networks Co-located in HK

- Support Trunk Ports only for R&E Networks
- Support special VLANs only for R&E Networks
  - For private interconnections among any 2 R&E networks
  - *One special R&E IX-VLAN for interconnections among R&E networks with no commercial networks*
  - Jumbo Frame support
- *Limited colo at new HKIX1b site at CUHK Campus*
  - *Up to 2 rack per R&E network*
    - *Discounted MRC*
  - *No MRC for fiber cross-connects*

R&E Networks Already Connected to HKIX

- APAN-JP/JGN-X
- ASCC-ASNET
- ASGC
- *ASTI/PREGINET (coming)*
- CERNET/CNGI-6IX
- CSTNET
- CUHK
- HARNET
- KISTI-KREONET2
- NUS

- **CUHK/HKIX is more than willing to deploy changes in order to serve R&E community better**
  - **Suggestions are welcome**

Thank you!