



Experience Sharing: IPv6 at HKIX

Che-Hoo CHENG 鄭志豪

*The Chinese University of Hong Kong /
Hong Kong Internet Exchange*

27 Aug 2008



Introduction of HKIX (1/2)

- Set up by The Chinese University of Hong Kong (CUHK) in Apr 1995
- MLPA Internet Exchange over Layer 2 Infrastructure with BLPA support
- MLPA
 - Mandatory for Hong Kong routes only
 - HKIX MLPA Router Server: AS4635
 - AS4635 seen in AS Path
 - Route filters implemented strictly
 - By Prefix or by Origin AS
 - But a few trustable participants have no filters except max number of prefixes
- Support BLPA
 - One AS hop less than MLPA
 - May get more routes from your peers
 - HKIX encourages BLPA over HKIX

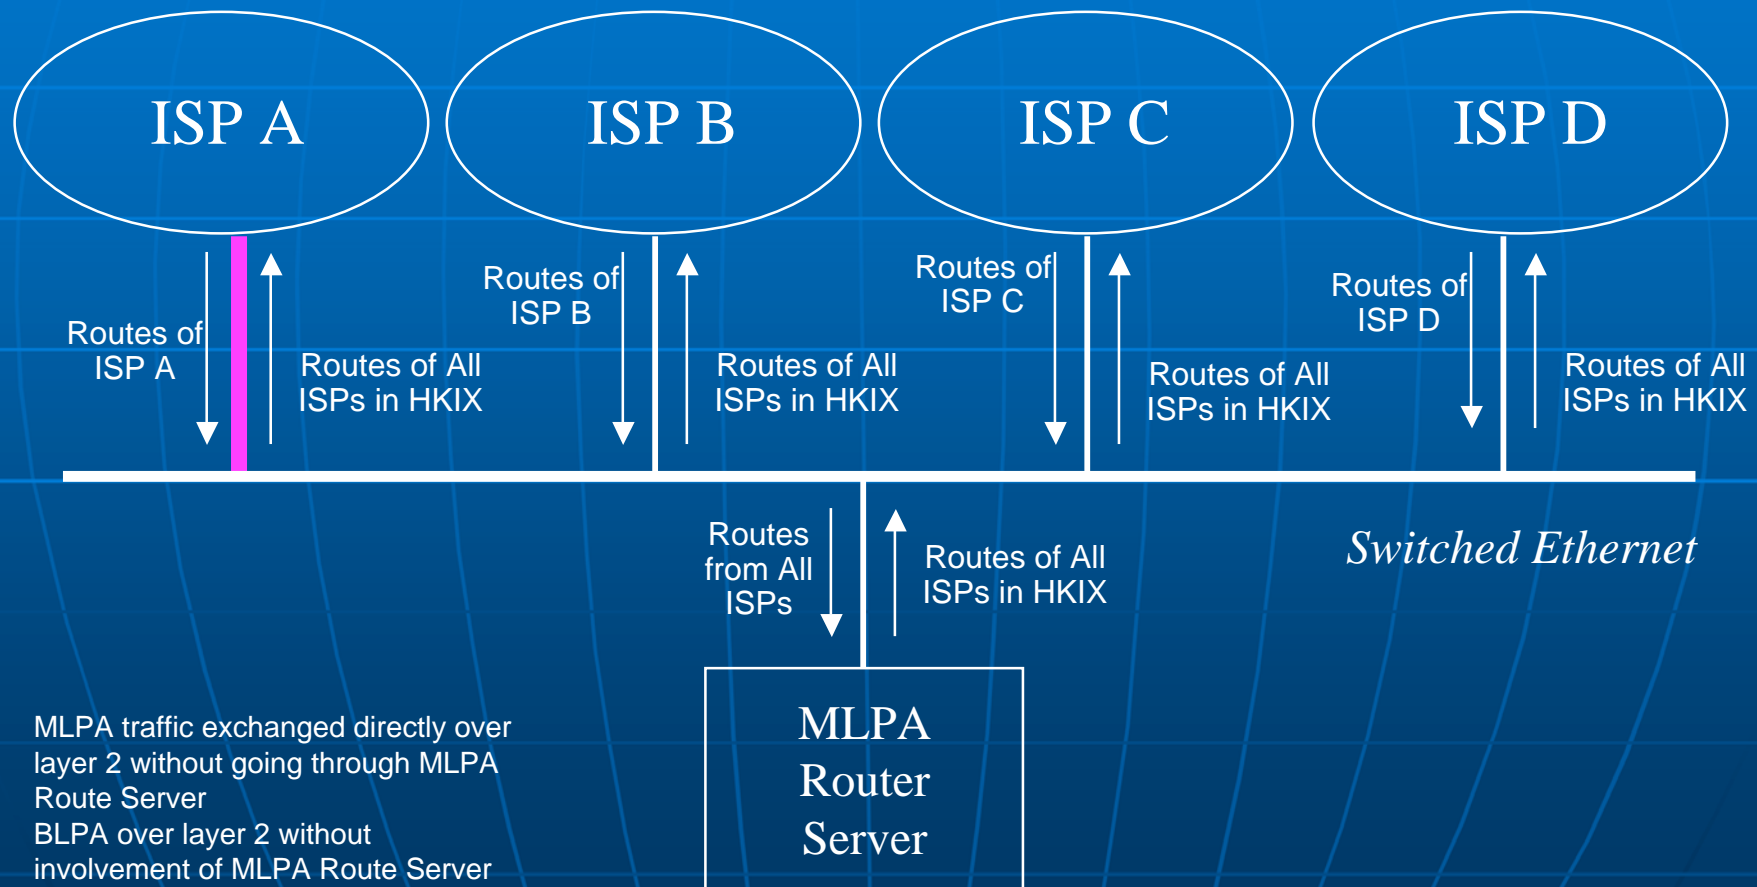


Introduction of HKIX (2/2)

- Port Security (one MAC address per switch port) implemented strictly
- Still no IX port charge as we are not-for-profit
 - But there is charge for 10GE ports or many GE ports if traffic volume is not high enough to justify the resources
- Provide colo space for strategic partners such as root / TLD DNS servers
- Still located and operated by CUHK
- Considered as Critical Internet Infrastructure in Hong Kong



HKIX Model — MLPA over Layer 2 (with BLPA support)





HKIX2

- Announced on 25 Nov 2004
- HKIX2 site in Central as redundant site of HKIX
- Linked up to HKIX by 2 x 10GE links
 - It is Layer 3 connection
 - Same AS4635 MLPA
 - Participants cannot do BLPA across HKIX and HKIX2
- Free of charge for IX ports
- IX portion managed by CUHK

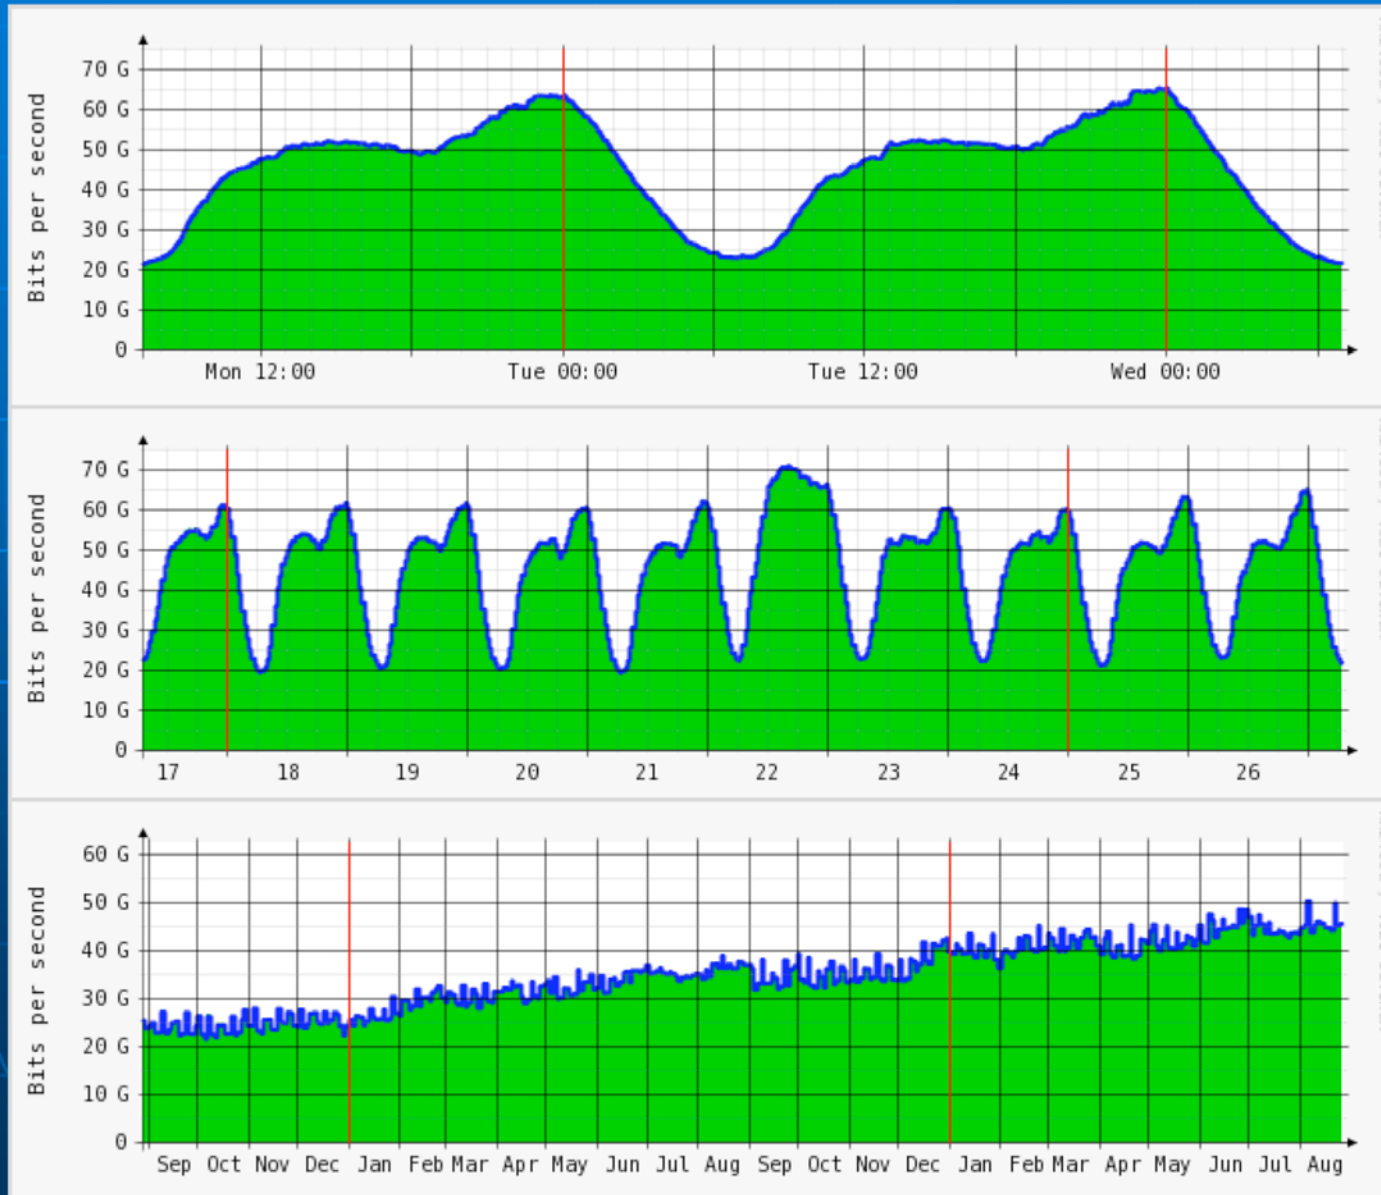


Quick Updates

- 2 x Cisco Catalyst 6513 at HKIX and 1 x Cisco Catalyst 6513 at HKIX2
- Most connected to HKIX switches without co-located routers
 - Cross-border layer 2 Ethernet connections to HKIX possible
 - Ethernet over MPLS or Ethernet over SDH
- Officially allow overseas ISPs to connect now
 - No co-located router needed in Hong Kong
 - Those overseas ISPs may not have Hong Kong routes...
- > 90 HKIX participants now (>10 on HKIX2 + HKIX)
 - Some have multiple AS'es
- 18,000-26,000 IPv4 prefixes carried by HKIX MLPA router server now
- Peak 5-min traffic >65 Gbps now



Some Statistics





Help Keep Intra-Asia Traffic within Asia

- We have participants from Mainland China, Taiwan, Korea, Indonesia and Philippines
- MLPA of HKIX
 - 6 AS'es are announcing more than 1,000 prefixes
 - More non-Hong Kong routes than Hong Kong routes
- BLPA over HKIX
 - Even more non-Hong Kong routes
- So, 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
- HKIX is good for intra-Asia traffic
- Many overseas and mainland China academic networks have presence in Hong Kong



Plan for 2008

- Replace one Cisco Catalyst 6513 with a higher-end layer 2 switch
 - To support more 10GE ports
 - To support Link Aggregation with port security
 - Remote participants have to check whether their tail providers can support the configuration with enough transparency
 - sFlow
- MLPA: Support daily automatic route filter updates from routing registry database
- MLPA: Support BGP community for easier traffic load balancing
- We continue to encourage BLPA ☺
- Portal for Participants
- Suggestions are welcome



Submarine Cable Disaster in Dec 2006

- Due to Earthquake in South of Taiwan (Luzon Strait) on 26 Dec 2006
- Most cable systems going through Luzon Strait were cut then
- HK was almost isolated from Global Internet
- Restoration was done slowly and gradually
- Cable repair finally complete in late Jan 2007
- Lessons learnt:
 - Cable route diversity must be observed
 - Should not rely totally on cables of East routing which all go through Luzon Strait
 - Should be prepared to pay more for cables of West/North/South routing for better reliability
 - **DNS infrastructure in HK must be improved**
 - **.com, .net and .org TLD servers could not be found on HKIX MLPA route server**
 - HKIX (layer 2 part) could be used for acquiring temporary IP transit services during emergency period



Authoritative TLD Servers at HKIX

- As important as Root Servers
- Anycast is getting more and more popular at TLD level
- During the disaster, we had Root Server instances F & I connected to HKIX so .hk, .mo and .cn are fine
 - .com/.net/.org were half dead even though IP connectivity among HK, Macau and Mainland China was fine
 - Although UltraDNS had anycast servers in HK serving .org and others, they did not have connectivity to HKIX MLPA so could not help the situation!
- Key players at HKIX directly or indirectly now or soon:
 - Afilias, APNIC, CommunityDNS, PCH & Verisign
- Root Server instance J added to HKIX in Feb 2008
- We will continue to spend effort to encourage set-up of DNS server instances of major TLDs in Hong Kong with connection to HKIX MLPA (plus BLPA over HKIX) to improve DNS performance for the whole Hong Kong and neighbouring economies



IPv6 in Hong Kong (1/2)

- HK is lagging behind very much
 - We need to catch up
- Only a handful of backbone ISPs provide IPv6 transit service in HK
 - Such as NTT Com and Reliance Globalcom/FLAG
- Only very few retail ISPs provide IPv6 access service but they are not active at all
 - Such as NTT-HKNET, Diyixian and CITIC1616 for business customers only
- No residential broadband ISPs provide IPv6 service



IPv6 in Hong Kong (2/2)

- Two mobile phone service providers are testing out IPv6
 - CSL/Telstra and China Mobile-Peoples
- No IPv6 Tunnel Broker in HK
- Content providers are not ready for IPv6
- Government networks do not run IPv6 yet
- HARNET relied on ABILENE mainly so did not have full routes (< 40%) and most overseas IPv6 traffic had to be routed through US



IPv6 at HKIX

- CUHK/HKIX is committed to help Internet development in HK
- IPv6 supported by HKIX since Mar 2004
 - Dual stack
- Today, 16 different AS'es have been assigned addresses at HKIX/HKIX2 and have joined MLPA
 - BLPA allowed
- Root server instance F supports IPv6 transport at HKIX
- Dual stack so can't know for sure how much IPv6 traffic in total
 - Hopefully with the new switch, we can have more detailed statistics



IPv6 Participants at HKIX (1/2)

- APNIC (AS18366)
- ASCC-ASNET (AS9264)
- China Mobile-Peoples (AS9231; at HKIX2)
- CITIC1616 (AS17554; at HKIX2)
- CNGI-6IX/CERNET/CERNET2
(AS23911/AS4538/AS23910; IPv6 only)
- CUHK (AS3661 & AS4641)
- Diyixian (AS9584)
- ISC (AS23709)
- JUCC-HARNET (AS3662)



IPv6 Participants at HKIX (2/2)

- KREONET2 (AS17579)
- NTT Com (AS2914)
- Reliance Globalcom / FLAG (AS15412)
- Samsung (AS6619)
- Telstra-CSL (AS38819)
- TIC (AS1836)

- Good mix of academic networks and commercial networks
- Can set up BLPA with them over HKIX



Recent IPv6 Work at HKIX

- Remove route filters for IPv6 at MLPA route server
 - Still provides minimal protection such as bogus routes
 - Essentially a transit exchange for IPv6
 - Total number of routes close to 1,100 now
- Add BGP community tagging to distinguish upstream routes (for transit purpose) from downstream routes (for peering purpose) soon
- 3-month free IPv6 full transit service provided by NTT Com (AS2914) for HKIX participants
 - Announced on 18 Aug 2008:
<http://www.hkix.net/hkix/announce.htm>
 - Available until 31 Dec 2008



Observations (1/4)

- Dual Stack seems to be the norm for IX
 - No need to have separate equipment and connection for IPv6 so easier to justify
 - Same AS# for both IPv4 and IPv6 seems to be the norm
 - One exception is Pacnet (AS10026 for IPv4 and AS18084 for IPv6)
 - sFlow should help give ideas of IPv4 traffic volume versus IPv6 traffic volume
- Using tunneling as IPv6 IX seems not acceptable by community anymore



Observations (2/4)

- The first day HKIX removed the route filters, someone from Germany noticed that immediately and warned us
 - Some people do care about IPv6
- Commercial providers care very much about routing and operations
 - They treat IPv4 and IPv6 the same as much as possible
 - Transit versus Peering
 - Customer versus Peer



Observations (3/4)

- HKIX now has /64 in total for HKIX & HKIX2 and use /120 for each
 - Will get /48 and use /64 for each
- Should not blindly prefer routes learnt from HKIX MLPA
 - Should add more BGP community tagging control
- Cannot help black-holing traffic because HKIX is basically a layer 2 infrastructure
- /32 of HARNET split into /35's and/or /48's because of various requirements
 - Different sets of routes are being announced to different upstream networks and peer networks
 - Routing is messy



Observations (4/4)

- MyAPNIC still has not provided the same functionality for IPv6 as IPv4, such as:
 - Add reverse DNS domain object
 - Add route6 object
- Not many people are aware of APNIC's policy change of IPv6 initial allocation criteria
- Need TLD servers to support AAAA glue records and run IPv6 transport in order to have full IPv6 experience
 - .HK does not AAAA glue records yet!!!
- Still far from universal development in HK, especially in commercial networks
 - Nobody is pushing and no demand
 - IPv6 knowledge is bad
- iPhone 2.0 S/W does not support IPv6 yet!!!



Thank You