# Hong Kong Internet Exchange (HKIX) Updates

## @APRICOT-APAN 2011

## http://www.hkix.net/

HKIX

香港中文大學
The Chinese University of Hong Kong

# Introduction of HKIX *(1/2)*

- HKIX is a Settlement-Free Layer-2 Internet Exchange Point (IXP), with mandatory Multi-Lateral Peering Agreement (MLPA) for Hong Kong routes
  - ISPs can interconnect with one another and exchange inter-ISP traffic at HKIX
  - HKIX is not a Transit Provider
- HKIX supports and encourages Bi-Lateral Peering Agreement (BLPA)
- HKIX was a project initiated and funded by ITSC of CUHK in April 1995 as a community service
  - Still owned, supported and operated by ITSC of CUHK

# Introduction of HKIX *(2/2)*

- Two Main Sites for resilience:
  - HKIX1: CUHK Campus in Shatin
  - HKIX2: CITIC Tower in Central
  - Under 2 different Power Grids operated by different Companies
- Our service is basically **free of charge** as we are **not-for-profit**
  - But there will be charge for 10GE port 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 & RIRs
- Considered as Critical Internet Infrastructure in Hong Kong
- We are confident to say that because of HKIX, more than 99% of intra-HK Internet traffic is kept within HK

- More information on www.hkix.net

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

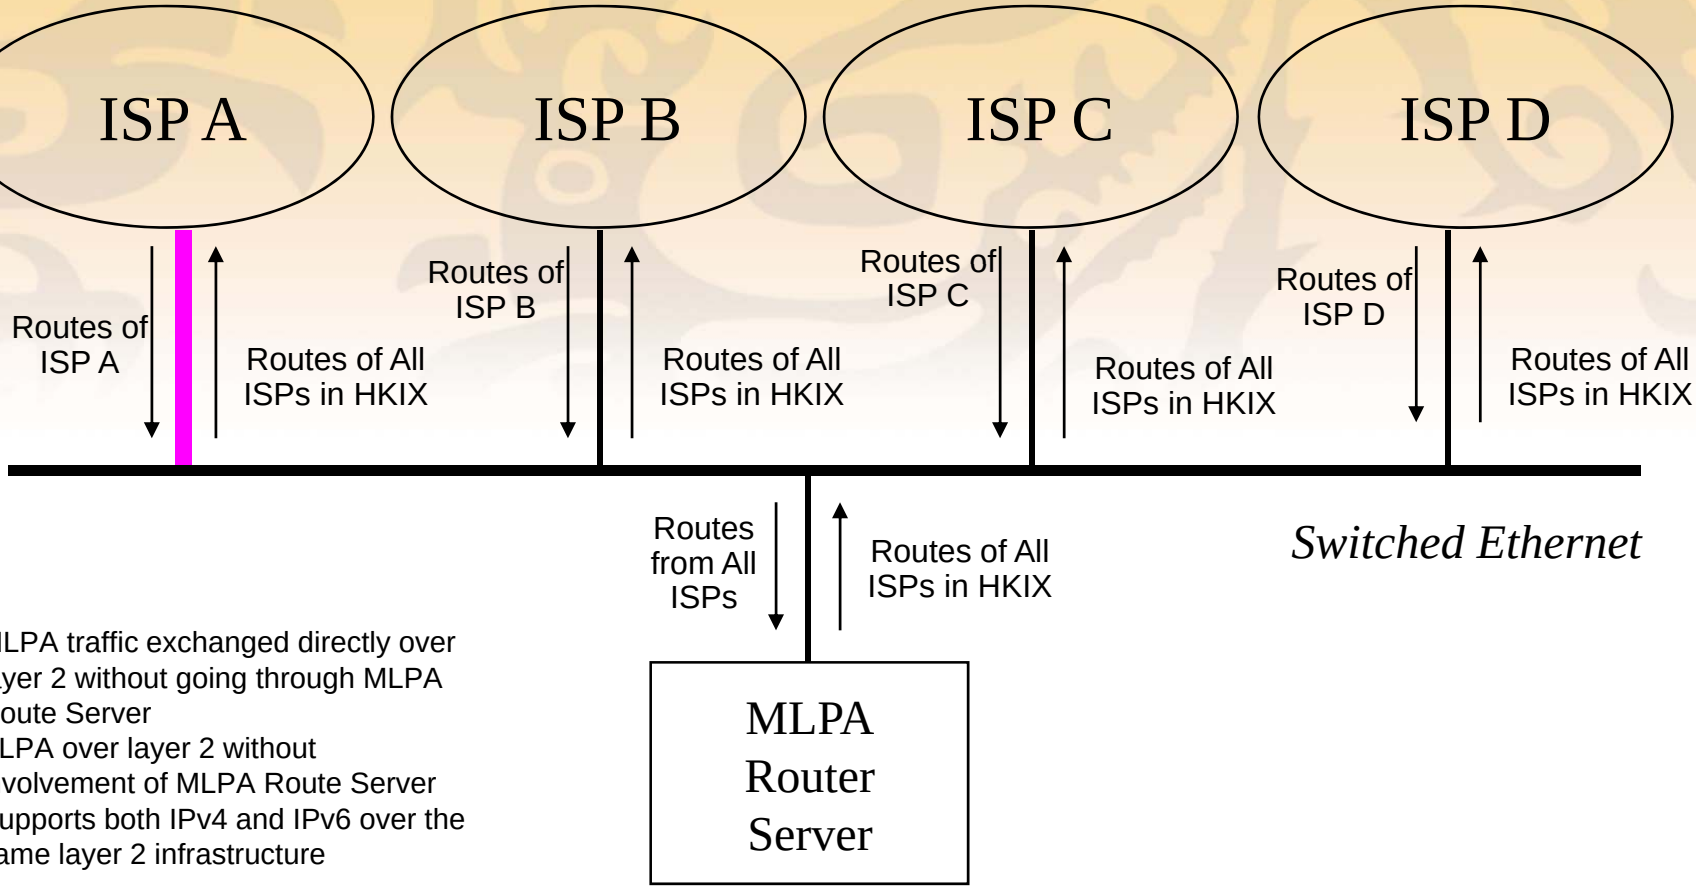


- MLPA traffic exchanged directly over layer 2 without going through MLPA Route Server
- BLPA over layer 2 without involvement of MLPA Route Server
- Supports both IPv4 and IPv6 over the same layer 2 infrastructure

# Quick Updates *(1/3)*

- 2 x Cisco Nexus 7018 + 2 x Cisco Catalyst 6513 at HKIX1 and 1 x Cisco Catalyst 6513 at HKIX2
- Most participants connected to HKIX 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
  - Local ISPs must have proper licenses
  - Those overseas ISPs may not have Hong Kong routes…
  - Major overseas R&E networks connected since 2008

# Quick Updates *(2/3)*

- ~**145** AS'es connected with IPv4 now
  - 15 AS'es at both HKIX1 & HKIX2 for resilience
- **35** 10GE connections and **225** GE/FE connections
- >**35,000+** IPv4 routes carried by HKIX MLPA
  - More non-HK routes than HK routes
  - Serving intra-Asia traffic indeed
- Peak 5-min traffic ~**139+**Gbps now
- HKIX1 supports and encourages Link Aggregation (LACP)

- A small POP in Mega-i with 1x10GE link back to HKIX1 but it is for R&E network connections only

# Quick Updates *(3/3)*

- Basic Set-up:
  - First 2 GE ports with no colo at HKIX1 and First 2 GE ports at HKIX2: Free of charge and no formal agreement needed
- Advanced Set-up:
  - 10GE port or >2 GE ports at either site or Colo at HKIX1:
  - Formal agreement is needed and there will be colo charge and / or port charge unless aggregate traffic volume of all ports exceeds 50% (95$^{th}$ percentile)
- See http://www.hkix.net/hkix/connectguide.htm for details

# HKIX Network Diagram

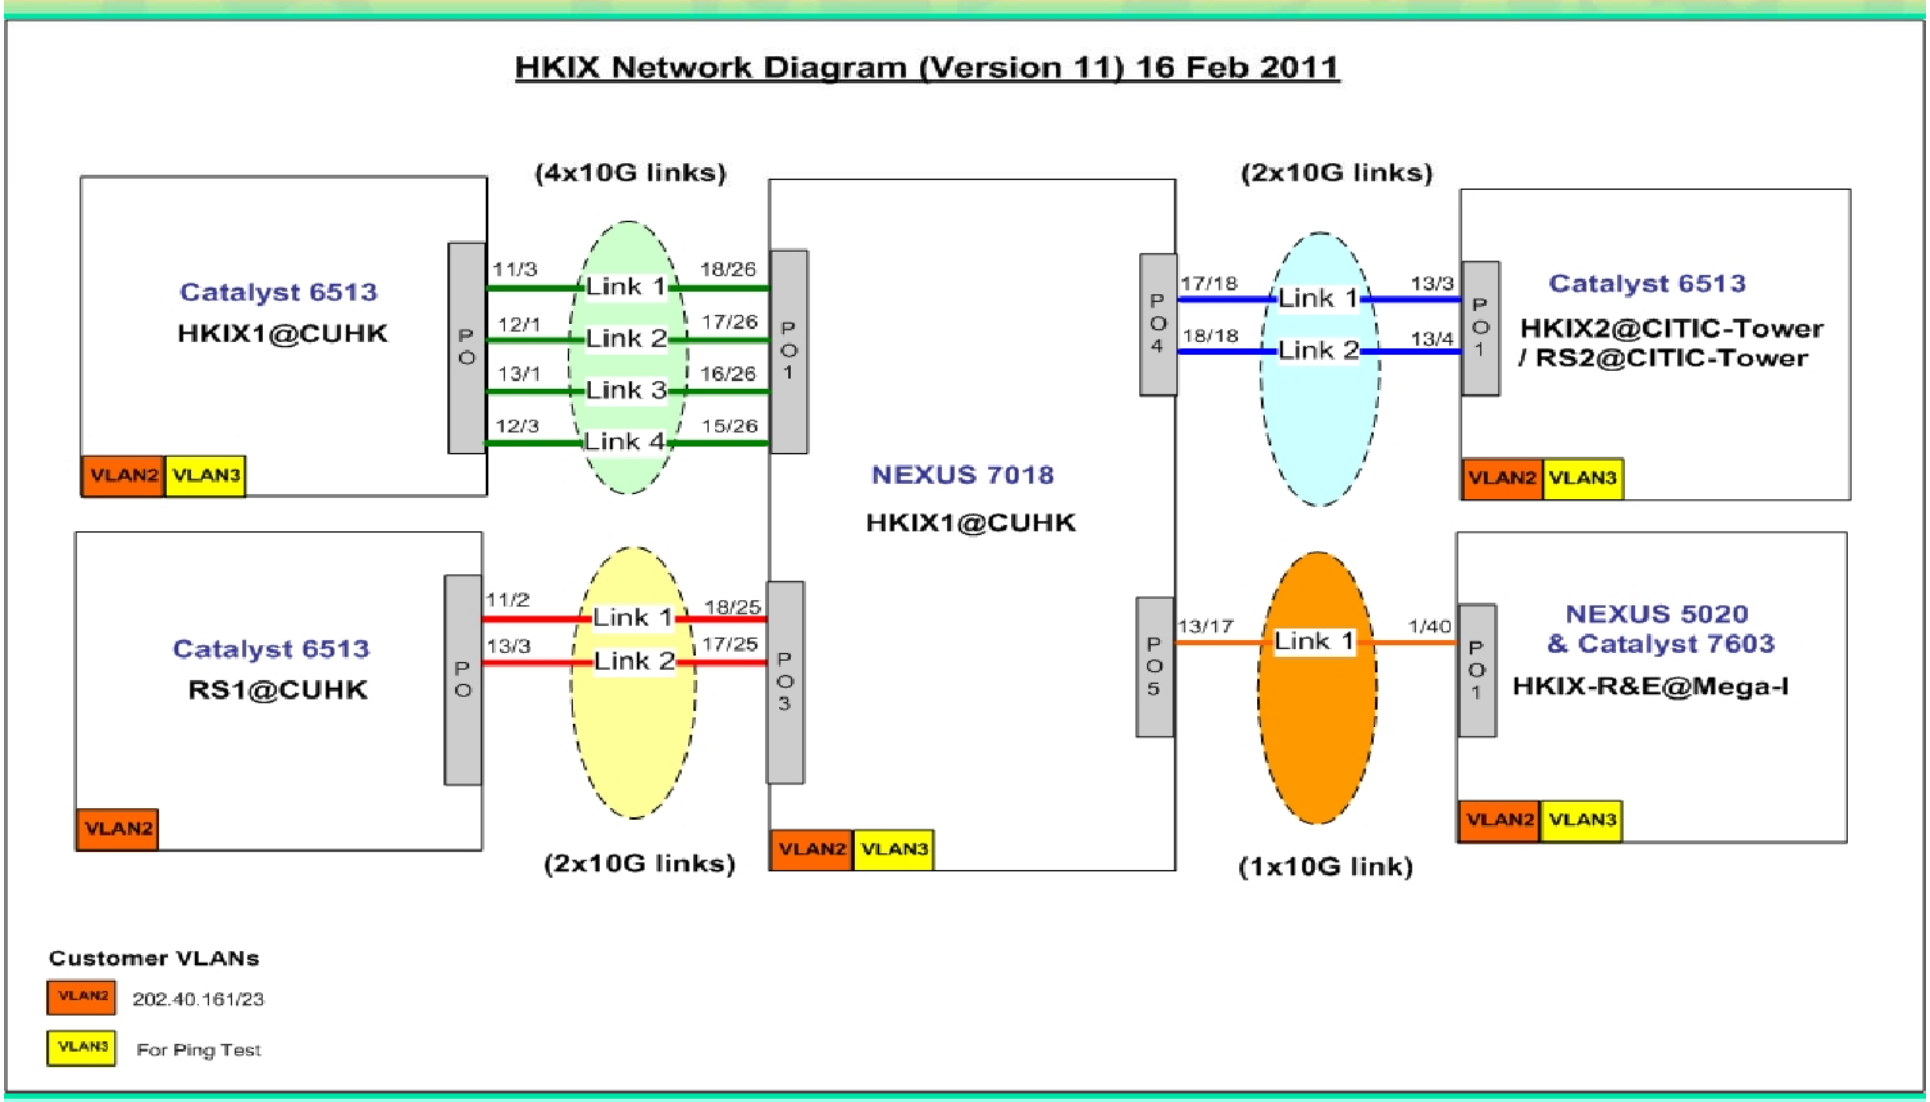

# HKIX2 Currently

- Set up in 2004 as redundant site
- IX portion managed by CUHK
- Linked up to HKIX1 by 2 x 10GE links
- Initially it is **Layer-3** connection so different broadcast domain from HKIX1
  - Same AS4635 MLPA
  - Participants cannot do BLPA across HKIX1 and HKIX2
- Extend the Layer-2 network from HKIX1 to HKIX2 has been done, participants can migrate from Layer-3 to Layer-2 now...
- But still cannot support LACP at the moment
- Will be moving to another Data Centre in 2011…

# Some Statistics - Daily

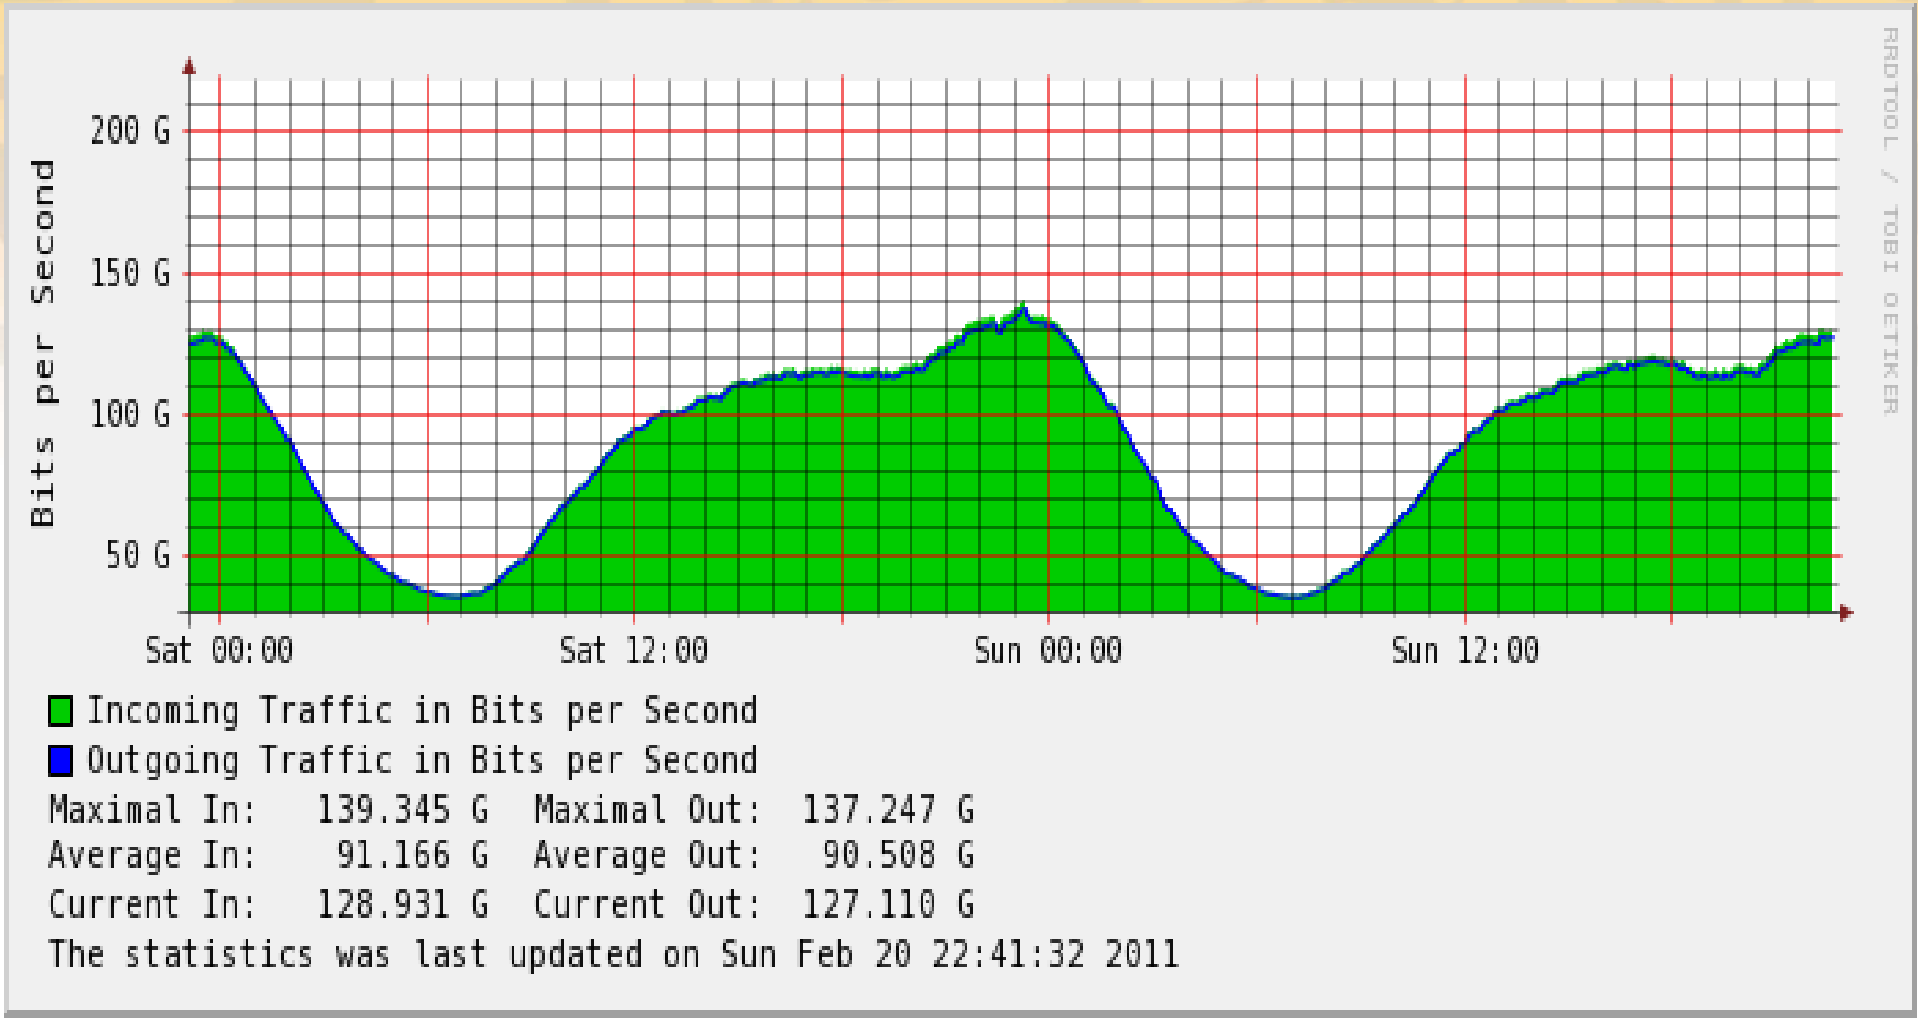

# Some Statistics - Weekly

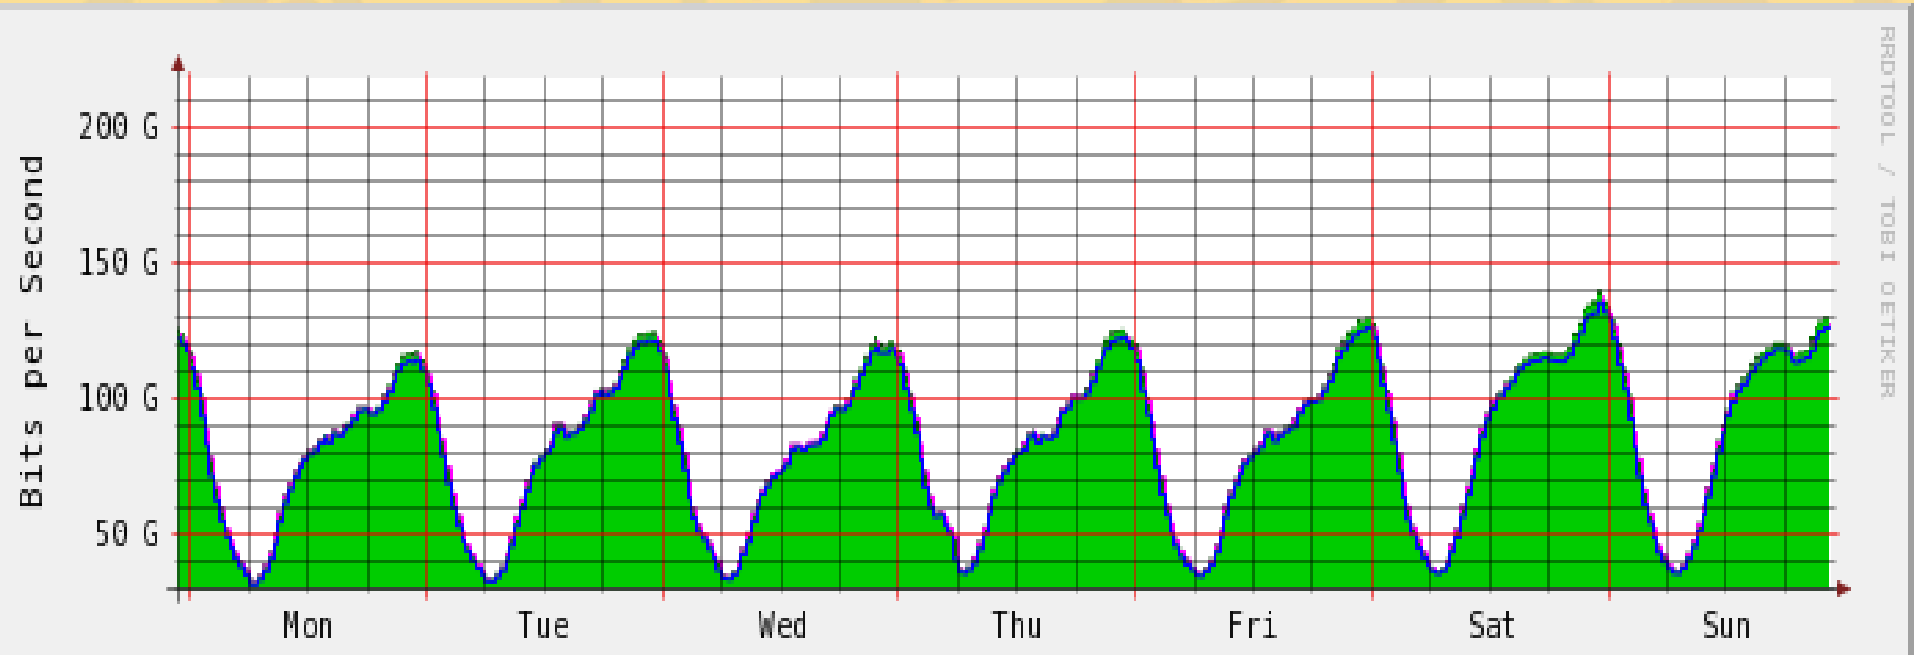


Maximal 5 Minute Incoming Traffic
Maximal 5 Minute Outgoing Traffic
Incoming Traffic in Bits per Second
Outgoing Traffic in Bits per Second

| | | | |
|---|---|---|---|
| Maximal In: | 139.345 G | Maximal Out: | 137.247 G |
| Average In: | 83.214 G | Average Out: | 82.862 G |
| Current In: | 128.192 G | Current Out: | 126.307 G |

The statistics was last updated on Sun Feb 20 22:36:31 2011

# Some Statistics - Yearly

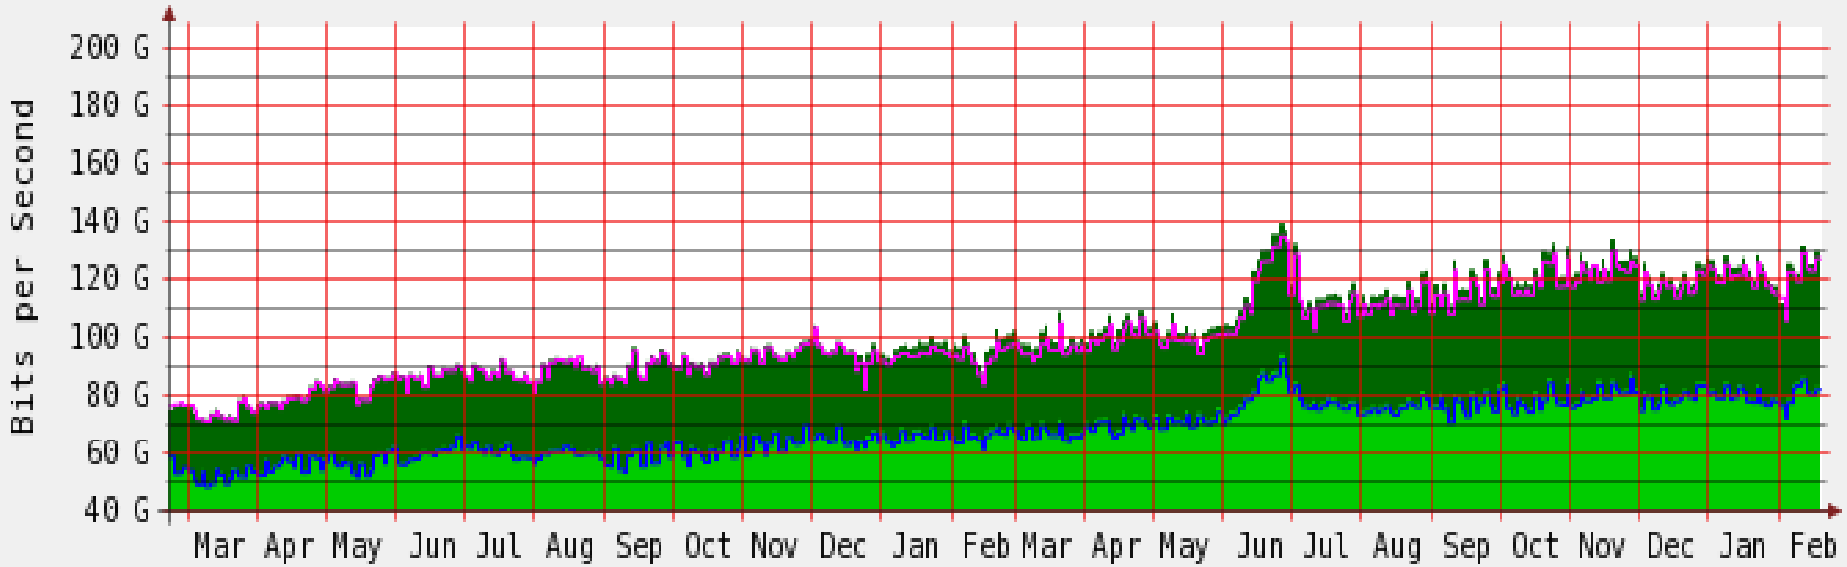


Maximal 5 Minute Incoming Traffic
Maximal 5 Minute Outgoing Traffic
Incoming Traffic in Bits per Second
Outgoing Traffic in Bits per Second

| | | | |
|---|---|---|---|
| Maximal In: | 138.541 G | Maximal Out: | 134.695 G |
| Average In: | 68.112 G | Average Out: | 67.750 G |
| Current In: | 82.230 G | Current Out: | 81.901 G |

The statistics was last updated on Sun Feb 20 07:41:31 2011

# HKIX – Statistics Summary

## **HKIX Statistics Summary** – As of February 2011

### Physical Ports

| Number of Participant Ports | 10GE | GE/FE/10M |
|---|---|---|
| HKIX Primary Site (HKIX1) | 34 | 196 |
| HKIX Secondary Site (HKIX2) | 1 | 20 |
| HKIX Research & Education Networks (HKIX-R&E) | --- | 9 |
| **Total** | **35** | **225** |

### BGP Peering

| | |
|---|---|
| Number of Peerings with RS1 | 299 |
| Number of Peerings with RS2 | 171 |
| Number of ASes connected with IPv4 | 145 |
| Number of ASes connected with IPv6 | 58 |
| Number of IPv4 Prefixes | 35,491 |
| Number of IPv6 Prefixes | 3,737 |
| Number of Licensed Participants | 91 |
| Number of Non-Licensed Participants | 44 |

# Our New Cisco Nexus 7018

- First deployed in an IX environment

- 128 10GE ports (wired speed) or
- 512 10GE ports (oversubscribed) or
- 768 GE ports
- Air-flow – side to side 8-(..

# Implementation of New High-End Switch

- To sustain growth, HKIX needed high-end switches at the core (HKIX1)
  - To support >100 10GE ports
  - To support LACP with port security over GE & 10GE ports
  - To support sFlow or equivalent
- Cisco Nexus 7018 selected
- In production since June 2009
- Migration of connections from 6513 to 7018 still in progress
  - Most 10GE connections have been migrated
  - Our 6513 will be decommissioned soon
- Added another 7018 in 2010 for resilience

# IPv6 at HKIX

- CUHK/HKIX is committed to help Internet development in HK
- IPv6 supported by HKIX since Mar 2004
- Today, **58** AS'es have their IPv6 enabled at HKIX
  - **3,736** IPv6 routes served by MLPA
  - **BLPA** encouraged
- Dual Stack recommended
  - No need to have separate equipment and connection for IPv6 so easier to justify
  - But cannot know for sure how much IPv6 traffic in total
  - Should be lower than 1% of the total traffic
  - With the new switch, we should be able to have more detailed statistics later

# HKIX – IPv6 Peering Trend

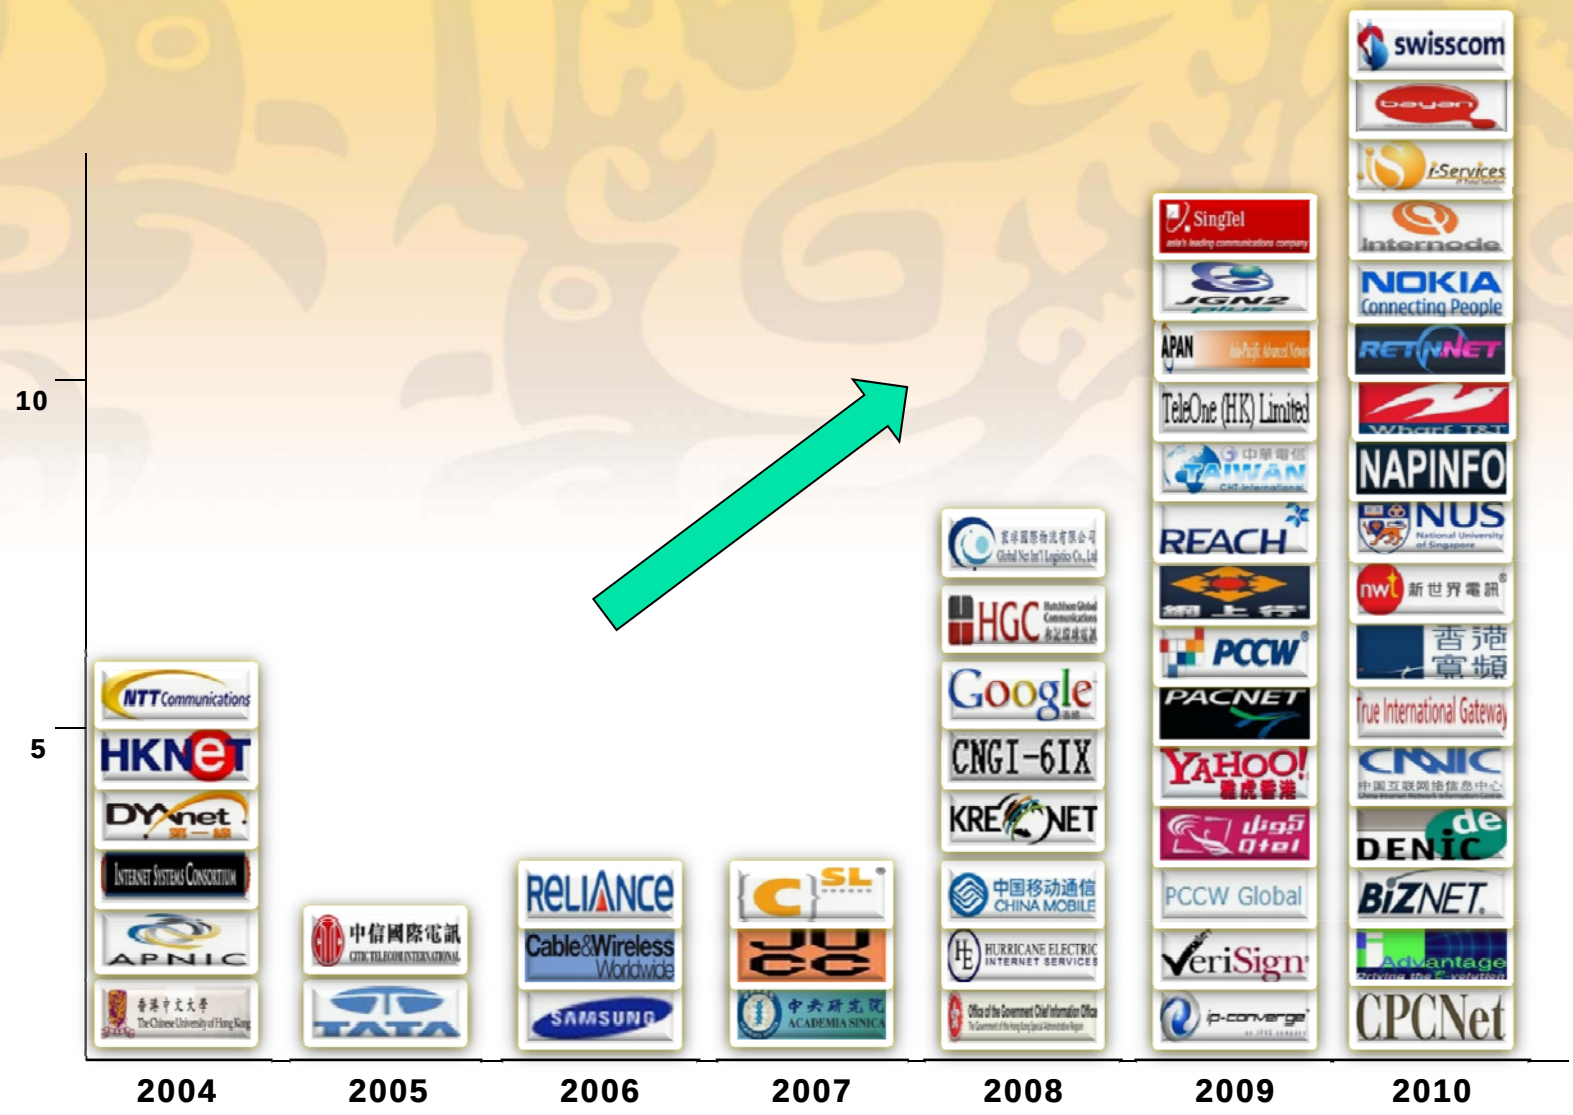

# News for IPv6 at HKIX

- HKIX can now support **IPv6-only** connections from commercial networks at MEGA-I
  - Max 1 x GE per participant
  - Must do BLPA with CUHK networks
  - This should help some participants try out IPv6 more easily
- More and more root / TLD servers on HKIX support IPv6

# Authoritative TLD Servers in HK

- As important as Root Servers
- Anycast is getting popular at TLD level
- During the disaster in 2006, we had Root Servers F & I connected to HKIX so .hk, .mo and .cn were fine
  - .com/.net/.org were half dead even though IP connectivity among HK, Macau and Mainland China was fine
  - Although there were anycast servers in HK serving .org and others, they did not have connectivity to HKIX MLPA so could not help the situation!
- We 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
- The authoritative servers of the following TLDs are connecting to HKIX directly now:
  - .com, .net, .org, .asia, .info, .hk, .mo, .*.tw, .sg, .my, .cn, .de and many many others

# MLPA at HKIX

- Mandatory for Hong Kong routes only
- Our MLPA route servers do not have full routes
- We do monitor the BGP sessions closely
- ASN of Router Server: AS4635
  - AS4635 seen in AS Path
- IPv4 route filters implemented strictly
  - By Prefix or by Origin AS
  - But a few trustable participants have no filters except max number of prefixes and bogus routes filter
  - Accept /24 or shorter prefixes
- IPv6 route filter not implemented in order to allow easier interconnections
  - But have max number of prefixes and bogus routes filter
  - Accept /64 or shorter prefixes
- See http://www.hkix.net/hkix/route-server.htm for details

# Bilateral Peering (BLPA) over HKIX

- **HKIX does support and encourage BLPA** as HKIX is basically a layer-2 IXP
- With BLPA, your can have better routes and connectivity
  - One AS hop less than MLPA
  - May get more routes from your BLPA peers than MLPA
- Do not blindly prefer routes learnt from HKIX's MLPA by using higher LocalPref
  - Doing more BLPA recommended
- Set up a record of your AS on www.peeringdb.com and tell everyone that you are on HKIX and willing to do BLPA
  - Also use it to find your potential BLPA peers
- Most content providers are willing to do bilateral peering
- Do set up bilateral peering with root / TLD DNS servers on HKIX to enjoy faster DNS queries

# Participants from Other Asian Economies

- The number is increasing
- Those are among the top 5 ISPs in their corresponding economies and they are not really regional players so they do interconnections only in HK
- From Australia, Bhutan, India, Indonesia, Korea, Malaysia, Philippines, Qatar, Taiwan, Thailand, Russia and so on
- They seek for better interconnections and better connectivity
- They may be willing to do BLPA at HKIX so contact them for BLPA

- HKIX is indeed serving as an Asian IXP

# Port Security

- Port Security implemented strictly
  - Also for LACP connections
- One MAC address / one IPv4 address / one IPv6 address per port (or LACP port channel)
- Some participants are unaware of this and do change of router / interface without notifying us

# Link Aggregation (LACP)

- Having many connections to HKIX increases difficulties of traffic engineering
- May not be able to support many connections if you only have a few routers
  - Each router can only have one interface connecting to HKIX
- LACP is a solution to solve these issues when your traffic grows
- Now, 7018 at HKIX1 can support LACP
- However, please do check whether your circuit providers can provide clear channel Ethernet circuits to HKIX1 with enough transparency before you place orders
- Please also check whether your routers can support LACP

# Other Operational Tips

- HKIX cannot help blackhole traffic because HKIX is basically a layer-2 infrastructure
- If there is scheduled maintenance, please notify hkix-noc@cuhk.edu.hk in advance so that we will not treat your BGP down message as failure
- Make sure proxy ARP is disabled on your router interface towards HKIX
- Do monitor the growth of number of routes from our route server and adjust your max prefix settings accordingly
- Do monitor the utilization of your links closely and do upgrade before they are full
- When your link / BGP session is down, do also check with your circuit providers at the same time
- Do your own route / route6 / as-set objects on IRRDB and keep them up-to-date

# Other Plans for 2011

- Start assigning 202.40.160/23 for IPv4 connections in **March 2011** (originally 202.40.161/24)
- 218.100.16/24 in HKIX2 will be replaced by 202.40.160/23
- MLPA:
  - Support daily automatic route filter updates from routing registry database (IRRDB)
  - Support more BGP community for easier traffic engineering
- Portal for Participants
  - Traffic statistics with data from Layer-2 Netflow
- Improve after-hour support
- Moving of HKIX-2 to a new Data Centre
- Starts supporting LACP at new HKIX-2?

- **Suggestions are welcome**

# Challenges of HKIX

- Space in CUHK needed for co-location requirements
  - Mainly serve Root / TLD servers, RIRs (such as APNIC) and a few other strategic partners only
  - A lot of requests from time to time
- Presence in other Data Centres?
- Better Redundancy – equipment and locations
- Peer-to-Peer Traffic and Video Traffic Growth
- DDoS Attacks
- 40G & 100G support

# Thank You

# Questions?