



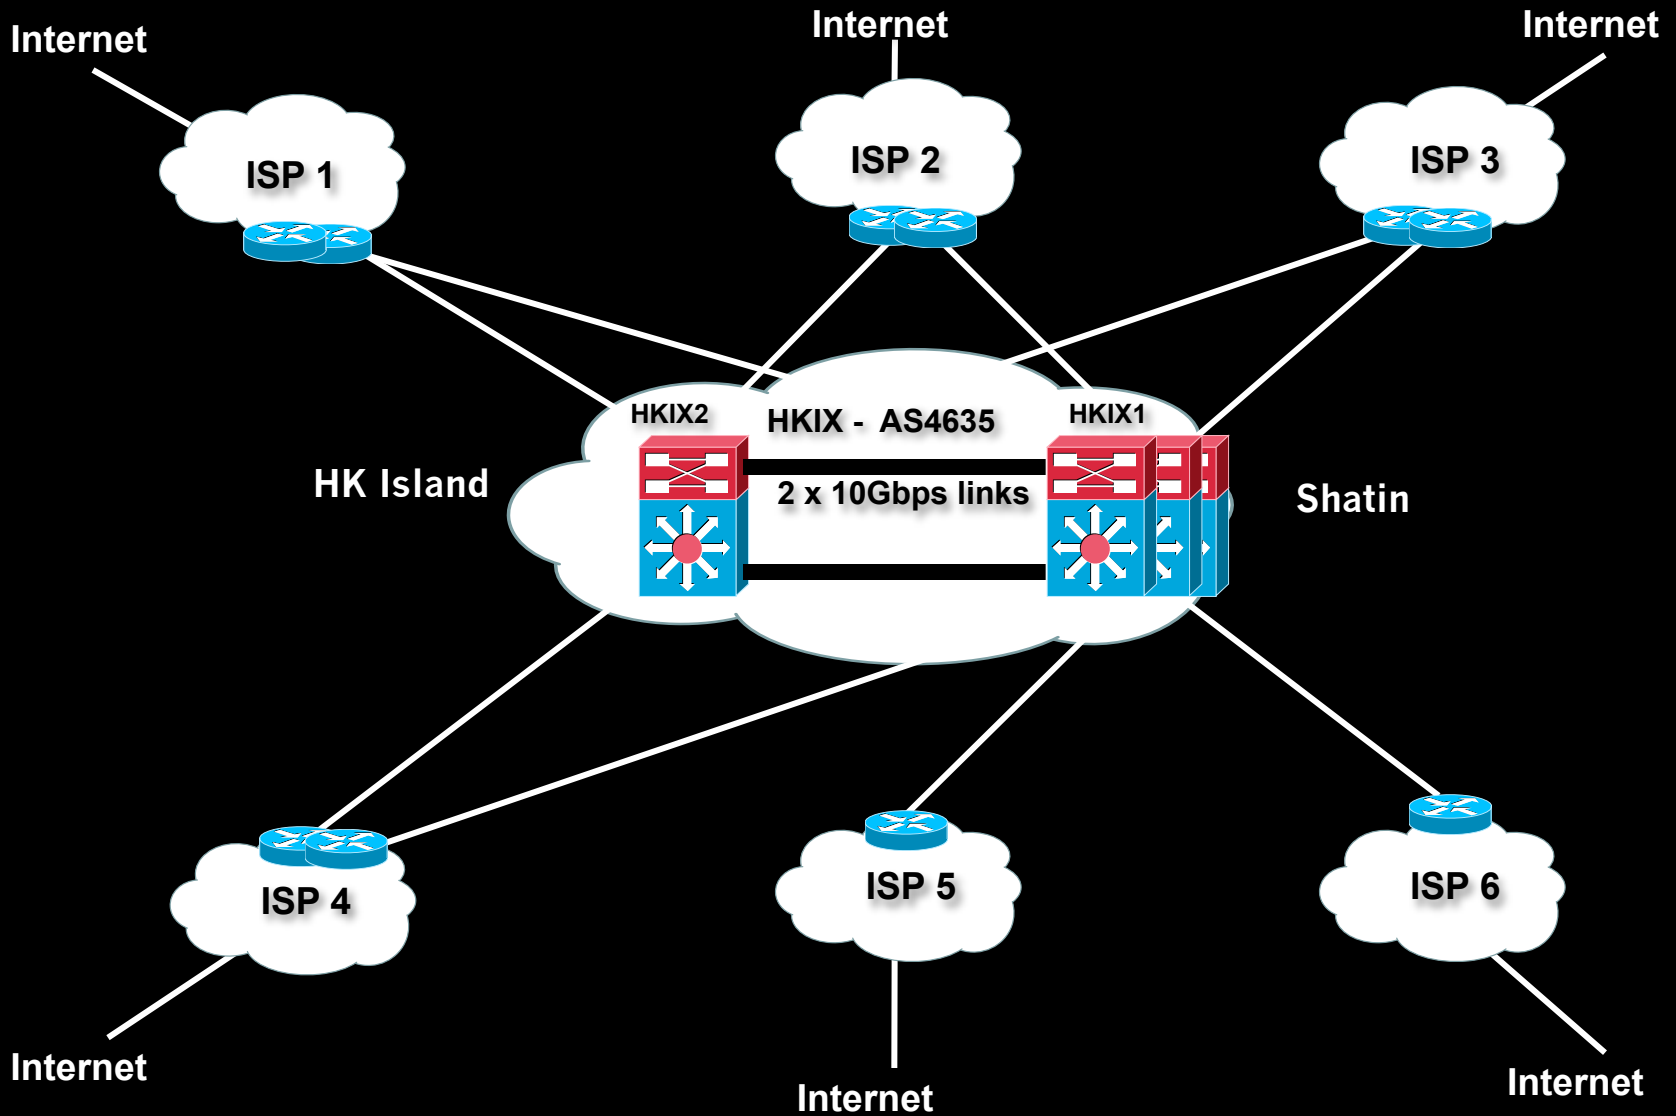
# Hong Kong Internet Exchange (HKIX)

<http://www.hkix.net/>

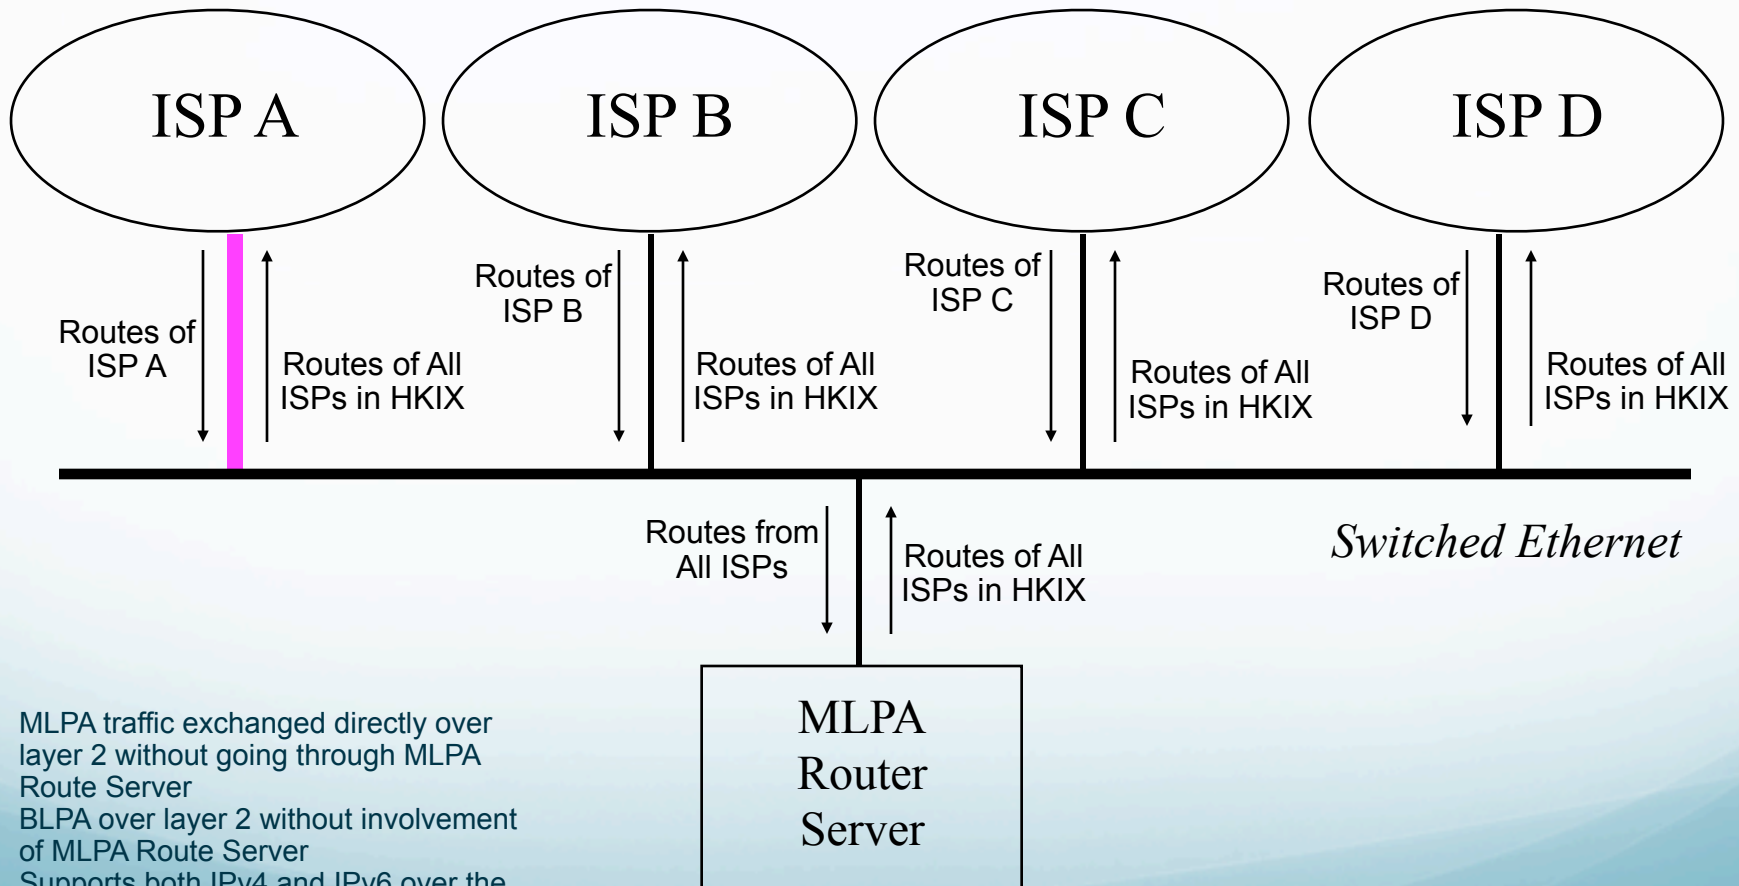
# What is HKIX?

- HKIX is a Public Internet Exchange Point (IXP) in Hong Kong – it is not a Transit Provider
- HKIX is the major domestic Interconnection point in HK where ISPs in HK can interconnect with one another and exchange inter-ISP traffic
- HKIX is a Settlement-Free Layer2 Internet Exchange Point, with mandatory Multi-Lateral Peering Agreement (MLPA) for Hong Kong routes
  - HKIX supports and encourages Bi-Lateral Peering Agreement (BLPA)
- HKIX was a project initiated and funded by ITSC of CUHK in Apr 1995 as a community service
  - Still supported and operated by ITSC of CUHK

# Current HKIX Infrastructure



# 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

# HKIX1 at ITSC of CUHK



ITSC of CUHK



# HKIX2 at CITIC Tower in Central



# HKIX History

- **Sep 91:** CUHK set up the 1<sup>st</sup> Internet link in HK to NASA Ames in US
- **Jul 92:** The HK Academic & Research Network (HARNET) IP-based Backbone was set up and JUCC/HARNET took over the management of the Internet link
- **Late 93:** 2 commercial ISPs (HK Supernet and HKIGS) were set up with their own links to US
- **94:** More ISPs were set up; ITSC of CUHK saw the needs of setting up a local exchange point and started negotiating with individual ISPs
- **April 95: ISPs started connecting to CUHK and HKIX was established**
- **Early 04:** Started supporting IPv6 and 10GE for traffic exchange and established a secondary site of HKIX (i.e. HKIX2)
- **Early 06:** International Network Services Providers and R&E networks were allowed to connect without telecom license
- **Present:** 133 AS'es connecting to HKIX; Ranked #14 in the World on Wikipedia according to traffic volume

# HKIX Policies for Joining

- **Membership requirements:**
  - Local ISPs with proper licenses (SBO, PNETS or FTNS)
    - Research & Education Networks
    - International Network Services Providers
      - Must warrant not to conduct ISP business in Hong Kong (otherwise they need to have PNETS license)
  - Have global Internet connectivity independent of HKIX facilities
  - Provide its own local circuit to HKIX
  - Must agree to do MLPA for Hong Kong routes




# HKIX Charging Model

- HKIX provides 2 GE ports at each HKIX site for each member free of charge as Basic Setup
  - No formal agreement is needed for Basic Setup
- Requesting for 10GE ports or additional GE ports involves formal agreement
- If port utilization is lower than 50%, there will be charges
  - If higher, no charges
  - This is to curb abuse
- Co-location service is chargeable now
- Not really for profit
  - Target for self-sustained

# HKIX2

- Announced on 25 Nov 2004
- HKIX2 site in CITIC Tower, Central as redundant site of HKIX
- Linked up to HKIX1 by 2 x 10GE links
  - It is Layer 2 connection now
  - Same MLPA domain as HKIX
  - Members can do BLPA across HKIX1 and HKIX2
- IX portion managed by ITSC of CUHK
- Same policies same charging model as HKIX1

# Some Statistics - Daily



■ Incoming Traffic in Bits per Second

■ Outgoing Traffic in Bits per Second


Maximal In: 134.968 G Maximal Out: 131.462 G

Average In: 88.151 G Average Out: 87.354 G


Current In: 45.732 G Current Out: 45.695 G

The statistics was last updated on Sat Jun 26 08:01:31 2010


# Some Statistics - Weekly




# Some Statistics - Monthly



# Some Statistics - Yearly



# Some Statistics - Number of Routes on MLPA



# HKIX Members – Beyond Asia





# Help Keep Intra-Asia Traffic within Asia

- We have members from Mainland China, Taiwan, Korea, Japan, Singapore, Malaysia, Thailand, Indonesia, Philippines, Bhutan, Qatar and other Asian countries
- Ten members are announcing more than 1,000 routes to MLPA so we have more non-Hong Kong routes than Hong Kong routes
- BLPA over HKIX facilitates 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



# 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 in HK

- As important as Root Servers
- Anycast is getting more and more popular at TLD level
- During the disaster, we had Root Servers 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 there was 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 neighboring economies
- The authoritative servers of the following TLDs are connecting to HKIX directly:
  - .com, .net, .org, .asia, .info, .hk, .mo, .\*tw, .sg, .my and many others

# IPv6 at HKIX

- CUHK/HKIX is committed to help Internet development in HK
- IPv6 supported by HKIX since Mar 2004
  - Dual stack
- Today, 48 AS'es have been assigned addresses at HKIX and have joined MLPA
  - BLPA encouraged
- Root server instance F supports IPv6 transport at HKIX
- Dual stack so cannot know for sure how much IPv6 traffic in total
  - Should be lower than 1% of the total traffic
  - With the new switch installed, we should be able to have more detailed statistics later

# HKIX – Member of IILG

- Considered as Critical Internet Infrastructure in HK
- Internet Infrastructure Liaison Group (IILG)
  - Coordinated by OGCIO of HKSARG
  - Members
    - OGCIO
    - OFTA
    - Hong Kong Police
    - HK Computer Emergency Response Team (HKCERT)
    - Major FTNS operators / ISPs
    - HKDNR
    - **HKIX**



# Technical Updates (1/3)



- HKIX-R&E in Mega-i with 2 x GE links back to HKIX1 but it is for R&E network connections only
- 1 x Cisco Nexus 7018 + 2 x Cisco Catalyst 6513 at HKIX1 and 1 x Cisco Catalyst 6513 at HKIX2 plus 1 x Cisco 7603 at HKIX-R&E
- 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
  - Local ISPs must have proper licenses
  - Those overseas ISPs may not have Hong Kong routes...
  - Major overseas R&E networks connected since 2008



# Technical Updates (2/3)



- 133 AS'es connected with IPv4 and 48 AS'es with IPv6
  - 17 AS'es at multiple HKIX sites for resilience
- 26 10GE connections and 211 E/FE/GE connections
  - 25 + 182 @HKIX1
  - 1 + 19 @HKIX2
  - 0 + 10 @HKIX-R&E
- >31,000 IPv4 routes and >2,400 IPv6 routes carried by HKIX MLPA
  - More non-HK routes than HK routes
  - Serving intra-Asia traffic indeed
- Peak 5-min traffic >130Gbps
- HKIX1 supports and encourages Link Aggregation (LACP)





# Technical 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
- Advanced Set-up:
  - 10GE port / >2 GE ports at either site / Colo at HKIX1: Formal agreement is needed and there will be colo charge and a small port charge unless aggregate traffic volume of all ports exceeds 50% (95<sup>th</sup> percentile)
- See <http://www.hkix.net/hkix/connectguide.htm> for details



# Implementation of New High-End Switch



- To sustain growth, HKIX needed a brand new high-end switch 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 after extensive pre-tender POC tests and complicated tendering
- In production since 15 June 2009
- Migration of connections from 6513 to 7018 still in progress
  - Most 10GE connections have been migrated
- Have ordered another 7018 chassis for resilience



# Our New 7018





# 7018 Preparation (Before 15 Jun 2009)



- Non-standard equipment rack needed:
  - Delivery issue, installation issue and high price
- Chassis failure: fast replacement
- Port Security problem
  - Had to wait for NX-OS 4.2(1) with major fix on Port Security
- SFP+ contact problem: unplug->plug to solve
- ISSU seems working fine
- First IX customer so had good support from Cisco



# Migration Issues (After 15 Jun 2009)



- 7018 in production since 15 Jun 2009
- Large participants' migration to new switch is a big issue
  - Layer 2 Netflow would help but we do not have it yet
- 6513 as central hub -> 7018 as central hub
- Inter-switch links 2x10GE -> 4x10GE
  - But we did not have enough 10GE ports on 6513's
  - 7018 does not support ER/ZR yet
- Xenpak changed to SFP+
  - Providing upgrade options to 10GE participants
  - Cabling patching done by fixed networks
- Concerns on migration by individual participants



# 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 over HKIX



- **HKIX does support and encourage BLPA** as HKIX is basically a layer-2 IXP
- With BLPA, you 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](http://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



# 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)
- UFB (Unicast Flood Blocking) feature is important
- 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



- Must disable Proxy ARP
- 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](mailto:hkix-noc@cuhk.edu.hk) in advance so that we will not treat your BGP down message as failure
- 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
  - APNIC RRDB is free if you are a member



# To Be Done By End of 2010

- HKIX1 broadcast domain / VLAN has been extended to HKIX2
  - To move all HKIX2 participants to HKIX1 VLAN which will involve change of IP addresses
- All IPv4 connections to migrate to 202.40.160/23 from 202.40.161/24 (*and 218.100.16/24*):
  - Change of network mask only
- All IPv6 connections to migrate to 2001:7FA:0:1::/64 from 2001:7FA:0:1::CA28:A100/120 (*and 2001:7FA:0:1::DA64:1000/120*):
  - Change of network mask only
- Support MLPA route server redundancy:
  - 202.40.161.1 (rs1.hkix.net) & 202.40.161.2 (rs2.hkix.net)
- Support 4-byte ASN



# Our Goals

- To have one single HKIX broadcast domain to better support BLPA
- To have better resilience
- To sustain future growth
- To reduce confusion



# Other Plans for 2010



- MLPA: Support daily automatic route filter updates from routing registry database (IRRDB)
- MLPA: Support more BGP community for easier traffic engineering
- Portal for Participants
  - Traffic statistics with data from Layer-2 Netflow
- Improve after-hour support
- **Suggestions are welcome**

# APRICOT-APAN.ASIA

2011 <http://apricot-apan.asia>  
HONG KONG

Hosts:



Co-hosts:



# WELCOME to Hong Kong!



**Main Conference: 21 – 25 Feb, 2011**

Workshop: 15 – 19 Feb, 2011 (@Cyberport)

**Hong Kong Convention and Exhibition Centre**

HOSTS:



## Why APRICOT-APAN.Asia/2011?

- The first ever joint event of APRICOT and APAN making it the biggest Internet conference in Asia Pacific
- The 2nd APRICOT in Hong Kong (last one in 1997) and the 1st APAN meeting in Hong Kong
- Targeting 1,000 participants from all over Asia Pacific region

HOSTS:







Questions?