CS 4/55231
Internet Engineering

Kent State University
Dept. of Computer Science

LECT-8B

Autonomous Systems Networking

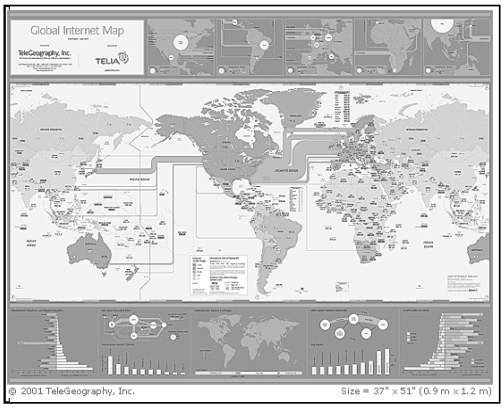

INTERNET ENGINEERING

Concepts (POP, ISP, NP)

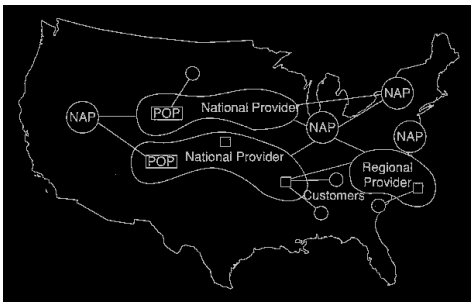

INTERNET ENGINEERING

Basic Concepts (2)

- National Access Point (NAP)
  - The facility where various NPs networks can interconnect.
  - Formerly these were organized as Federal Internet Exchange (FIX) & Commercial Internet Exchange (CIX). FIX/CIX model did not scale well.
  - It is physically a high-speed network switch or network to which a number of routers can be connected for the purpose of traffic exchange (example: FDDI or ATM switch).

INTERNET ENGINEERING

NAP

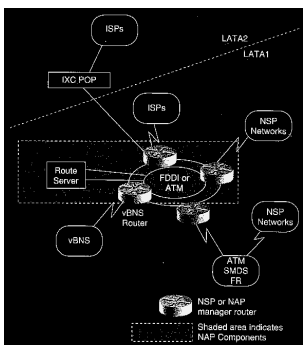

INTERNET ENGINEERING

Basic Concepts (3)

- Route Server
  - Route server exchanges routing information and policy with the service provider routers attached to the NAP.
  - It does not perform any traffic forwarding.
  - A group of servers facilitates interconnections between ISPs by gathering routing information from each ISP applying ISPs predefined set of rules, policies, and then redistributing the processed information to each ISP.
  - It saves routers of each individual ISPs to peer with all other routers, thus cutting down the number of peers from (n01) to 1.

INTERNET ENGINEERING

Route Server

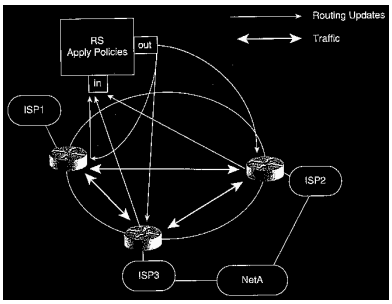

INTERNET ENGINEERING

Internet Exchange /NAP

Internet Exchange

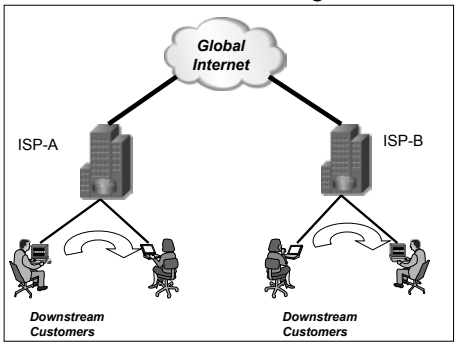

INTERNET ENGINEERING

Internet Exchange

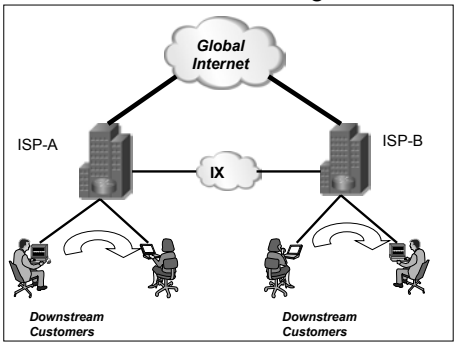

INTERNET ENGINEERING

Some Internet Exchanges

- USA – Major NAPs
  - MAE-West California, MAE-East Wash. DC operated by WCOM
  - Chicago NAP operated by Ameritech
  - New York NAP operated by Sprint
  - Nap of the Americas – operated by Terremark
- China - TerreNAP (Beijing), ShangHai IX (SHIX)
- UK - MaNAP, LINX, LoNAP, ScotIX…
- Japan - JPIX, Media Exchange (TTNet), NSPIXP, NSPIX2, NSPIX3
- Korea - KINX, KIX, KTIX
- Taiwan - TWIX
- Singapore - SingTel IX
- HKSAR – HKIX, ReachIX, Pilhana

INTERNET ENGINEERING

IX Technology

- The high speed interconnect was typically provided with Ethernet and FDDI.
- The speed requirement introduced ATM and Frame Relay.
- More recent trend is to combine ATM, GigE and Frame Relay services.

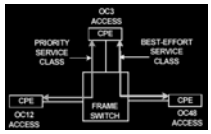

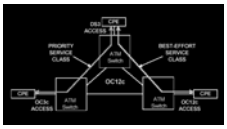

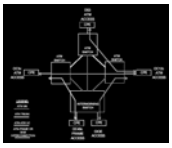

Sample IX: Hong Kong Internet Exchange

- Hong Kong Internet eXchange (HKIX) is initiated and coordinated by Information Technology Services Centre (ITSC) of the Chinese University of Hong Kong (CUHK).
- HKIX is a layer-two settlement-free multilateral exchange point mainly for routing of intra-Hongkong Internet traffic.
- HKIX can also be used for routing of traffic between the networks in Hong Kong and the peer or downstream networks of HKIX participants in other countries. The peering model of HKIX is a SKA (Sender Keep All) peering model.

Sample IX: HKIX Requirements

- Participants should be an ISP with proper licenses (such as issued by the Office of Telecommunication Authority of Hong Kong),
- Each participant must have Autonomous System (AS) number issued by InterNIC, RIPE-NCC, APNIC or their sub-registries. If they are single-homed, their upstream providers need only send a BGP default route to them. The source addresses of the data traversing the HKIX must be officially assigned by InterNIC, RIPE-NCC, APNIC or their sub-registries.
- Hong Kong Academic and Research Network (HARNET) members can also connect.
- Participants should arrange a link to the HKIX.
- Participants should be able to exchange routing tables (or peer) with the HKIX route server and (if any) peer with their upstream/peer/downstream providers using Border Gateway Protocol (BGP4).
- Participants should have primary global Internet connectivity independent of HKIX. The HKIX must not be used as the primary connection to the global Internet. Participants must be self-sufficient. They should have their own primary DNS, Email, WWW and News servers.

Example IX: HKIX Guidelines

- All participants of HKIX are exchanging data via an Ethernet. HKIX has set up route servers for peering;
- A data link- at least T1 to be provided by the participants.
- Participants must peer with the HKIX route server using BGP4.
- Participants can also peer with other participants directly as long as the arrangement does not violate any of the policies and guidelines stated here.
- HKIX is a settlement-free interconnection point. No settlement needs to be paid by the peering participants for the incoming and outgoing traffic.
- All participants of HKIX are considered equal. All HKIX participants must peer with one another via the HKIX route servers, should not filter traffic or routing table entries to or from any other participants unless it is justifiable.
- The main purpose of the HKIX at CUHK is for routing of intra-Hong Kong traffic but it is acceptable if participants allow others to exchange traffic with their peer or downstream network(s) at other countries free of charge.
- For colocation, HKIX will provide rack space, electricity, air-conditioning for their routers. HKIX will not provide any operational or management support.

Schematic Diagram of HKIX (Phase I)

Apr. 95

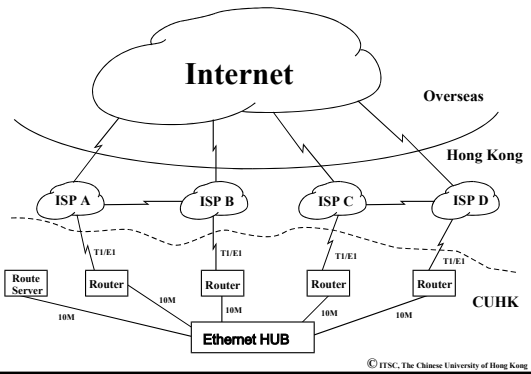

Schematic Diagram of HKIX (Phase IV)

November 2000

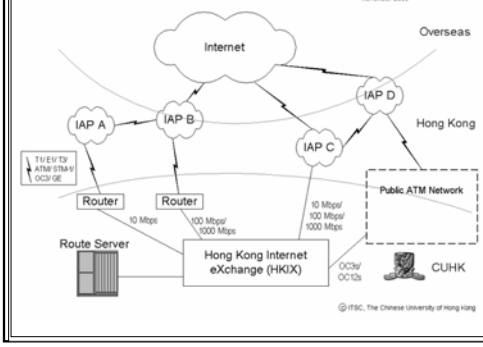

HKIX: Switching Traffic Trend

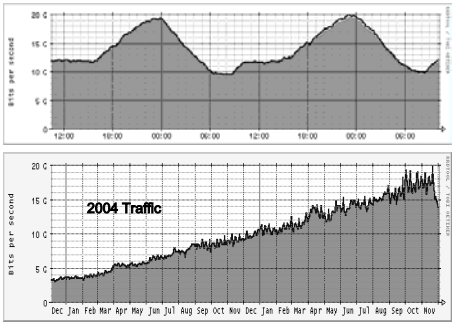

Conclusion: Research Issues

- BGP's rich feature set of tunable knobs and complex cross protocol interactions make it highly subject to a variety of problems, including mis-configuration, oscillations, and protocol divergence. The challenge of supporting many different complex policies in BGP without significantly complicating the protocol or degrading its performance has led to much research activity.
- **Configuration checking:**
  - Configuration checking tools can avoid mis-configurations by verifying certain consistency criteria hold, and modeling tools can predict side-effects of configuration changes on routers within an ISP. Across ISPs, uncoordinated routing policy can worsen route convergence and stability. The Routing Arbiter project introduced a distributed architecture for publishing and coordinating routing policies so as to avoid these problems, but was not widely deployed.
- **Language design:**
  - Routing Policy Specification Language (RPSL) is a vendor-neutral language proposed to describe an ISP's policy. It was envisioned these descriptions could be bound together in a database and checked for consistency. RPSL, though mature, is somewhat low-level and mechanism oriented. It may be possible to substantially improve upon RPSL by designing router configuration languages with higher level constructs that allow diverse policies while precluding certain misconfigurations, enforcing certain consistency properties to hold, simplifying configuration of certain common design patterns [22], however the design of such a language remains an open problem.
- **New architectures:**
  - HLP [23] is a proposed replacement for eBGP. The design philosophy of HLP is to expose common policies that can typically be inferred in BGP today and optimize the routing protocol based on the resulting structure, with the aim to improve scalability and convergence of interdomain routes.
  - Routing Control Platform (RCP) [24] is a logically centralized system that computes and distributes routes to routers inside an ISP. The centralization allows policies to be applied at the AS level, and the RCP applies the policies and its own decision process to select the best BGP route for each destination prefix on behalf of each router. This simplifies the configuration and application of policies and avoids misconfiguration.

INTERNET ENGINEERING

End