The AMS-IX switching platform

APRICOT KYOTO February 2005 Henk Steenman

Topics

- The parameters defining the AMS-IX switching platform
- The Ethernet switching platform
- Maintaining port hygiene
- Photonic switching
- Requirements for the near and long term future

Parameters defining AMS-IX

- 4 locations
- Around 322 ports connecting 210 AS numbers
 - 10 10GE ports
 - 189 GE ports
 - 93 FE ports
 - 40 E ports
- > 54 Gbit/s incoming on all customer ports

AMS-IX Ethernet switching platform

- Foundry Networks hardware
 - BI15K for edge switches
 - MG8 for core switches
- Resilient topology
 - VSRP for failover between Core switches
 - Also handles loop prevention
- Topology group with master VLAN
 - Runs Layer 2 protocols
 - Only ISL interfaces included
 - Customer ports in slave VLAN
 - follow master VLAN in case of topology change

AMS-IX switching topology



AMS-IX switching topology



- All L2 and L3 equipment allowed to connect
- BUT
 - We only want to see allowed traffic coming from L3 forwarding device MAC
- Only one MAC behind the AMS-IX switch port
 - To be functional this need to be the customer router

- Enforced by *Port security*
 - Allow only traffic from a single MAC address
 - Drop all traffic from other MAC addresses
 - Send automated e-mail in case of violation Quarantine VLAN
- New customer ports in *quarantine VLAN*
 - i.e. A unique VLAN that is not for production traffic
 - Check on proper customer router configuration
 - When OK, port will be defined in production VLAN

- Allowed Ethertypes
 - 0x0800 IPv4
 - 0x0806 ARP
 - 0x86dd IPv6
- No proxy ARP allowed
- Only Unicast
 - Exceptions
 - ARP
 - ICMPv6 Neighbor Discovery

- No link local traffic such as:
 - IRDP, ICMP redirects, IEEE802 STP
 - Vendor prop discovery protocols
 - CDP, EDP
 - IGPs
 - OSPF, ISIS, IGRP, etc
- We monitor the Exchange for broadcast and flooded traffic
 - Tools similar to "IXP watch" by LINX

- To limit the amount of ARP traffic
 - Sponge to catch ARP packets for IP addresses that are offline
 - Dedicated machine
 - Automated configuration
 - Based on number of ARPs for address
 - Automatically released when IP address is online again.

Photonic Switching

- Use Glimmerglass networks System 300 switch
 - 64 port MEMS based switch
 - Connect any port to any other port



Photonic Switch

- Main purpose Connect 10GE customers to master core switch
 - AMS-IX developed software to follow VSRP failovers
- Secondary purpose fast fiber rerouting for ISLs

AMS-IX switch requirements

- High availability edge switches
 - Aim should be 99.999% availability
 - Stable hardware
 - Failover components in switches
 - Power supplies, management blades, switch fabrics
 - Hitless software upgrades
- Much higher 10GE port density
 - More than 128 per switch required for Q3 2006
- 100GE hardware end of 2006
 - 40GE only of interest for customer connections