21st APNIC Open Policy Meeting
27 February - 3 March 2006 Perth - AUSTRALIA

APRICOT PERTH 2006

Exchange point operational experiences

Stephen Baxter
Co-Founder/Technical Director
PIPE Networks

About PIPE Networks

Founded December 2001, trading May 2002, founded by ex ISP industry people

Australia's Largest Peering Point Operator, professional staff, SLA available

Profitable, ASX listed company, sound proven business model

Distributed, multi-site peering point, no one point of failure

Over 80 ISPs/carriers peering with over 200 connections at 14 sites in 6 states

Major sources of Australian content on-net – ABC, Hostworks, Webcentral

All 10G metro networks – Australia's fastest, largest Metro Ethernet Network

Exchange point operational experiences

PIPE Networks peering description

1. MLPA
2. Single peering VLAN (with some ATM exceptions)
3. Route Server Based Routing Updates
4. Central web based routing policy/database tools
5. Automated filter list generation

Exchange point operational experiences

Disclaimer

All stories here are not based in fact, any actions recounted in this presentation could not be those of our customers.

All names have been changed to protect

the guilty^H^H^H^H^H^H innocent.

Exchange point operational experiences

Contents for this presentation

What happens at our peering points

Routers

Switches

Routing

Other Experiences

Suggestions from the operator to potential peers

Exchange point operational experiences

What happens at our peering points

What our customer use the PIPE Networks peering points for:

(The following are all anecdotal as recounted from conversations with our customers)

Least cost routing (go figure!)

Exchange point operational experiences

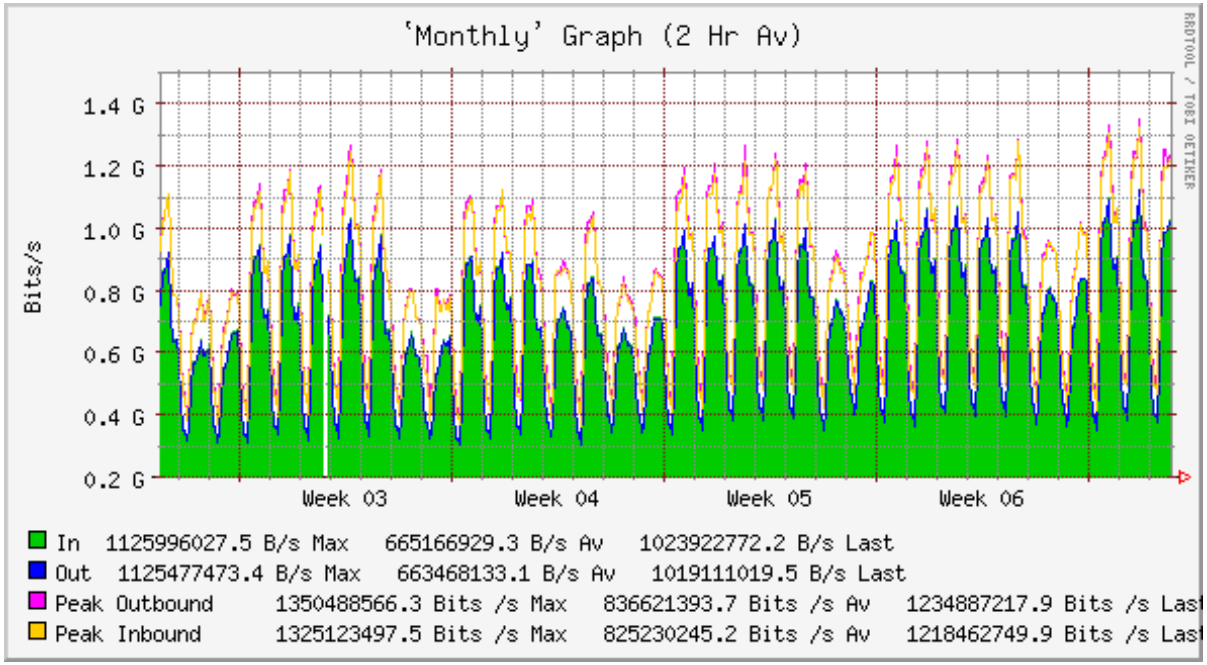

Graph showing closed circuit traffic via PIPE Networks

(Note there is little 'PIPE Only' content in our traffic)

Exchange point operational experiences

What happens at our peering points

What our customer use the PIPE Networks peering points for.

(The following are all anecdotal as recounted from conversations with our customers)

Least cost routing (go figure!)

Transit supply to and from other peers

Exchange point operational experiences

What happens at our peering points

What our customer use the PIPE Networks peering points for.

(The following are all anecdotal as recounted from conversations with our customers)

Least cost routing (go figure!)

Transit supply to and from other peers

Layer 2 forwarding resale

Exchange point operational experiences

What happens at our peering points

What our customer use the PIPE Networks peering points for.

(The following are all anecdotal as recounted from conversations with our customers)

Least cost routing (go figure!)

Transit supply to and from other peers

Layer 2 forwarding resale

Intra-city high speed, low cost transmission

Exchange point operational experiences

What happens at our peering points

What our customer use the PIPE Networks peering points for.

(The following are all anecdotal as recounted from conversations with our customers)

Least cost routing (go figure!)

Transit supply to and from other peers

Layer 2 forwarding resale

Intra-city high speed, low cost transmission

Access to USENET feed

Exchange point operational experiences

Routers

Cisco IOS based routers – hard to pass up. Nearly all of our customers use the platform ensuring good compatibility. We do not route with our routers merely use them as route servers.

7000 series for main route servers.

Linux and Zebra/Quagga also used for free content area routing (high volume, low monetary value content)

Routing table size 4% of global table – CRS-1 is not required

Routing table very stable

Flap controls not used due to table size and stability

Multicast support using MSDP and MBGP in IOS

Exchange point operational experiences

Switches

What did we use – Cisco 3500 series desktop switches (desktop switches have issues). These have been replaced for well over 1 year.

What did we look for in new switches :

10G backhaul required – diverse backhaul without 30 sec STP closure

Fine grained port control – access lists, STP control, broadcast storm control

Broadcast storm control, dual PSU, hot swap

Foundry Networks Super-X and FES-X series

Exchange point operational experiences

Routing

Customer must enter routes into a web based tool called PIPEdb.

ACL are built daily from this tool. We only allow routes that our customer inform us of in advance.

Monitor customer route numbers – faults usually occur with an increase or decrease in route numbers

Place limits on the number of routes customers can send

Filter routes in and out – routes and ASN associations

No mandatory RADB or other IRR

Exchange point operational experiences

Other experiences

MD5 passwords

Interesting routing and other behavior :

- Default Route
- Global Tables
- Most IGP
- All portions of RFC 1918

There has never been a customer impacting leak of routes at a PIPE Networks peering point.

Exchange point operational experiences

Suggestions from the operator to potential peers

Do not use "*ip verify unicast reverse-path*" etc unless you really know what you are doing.

Symmetrical routing solves many issues

BGP decision algorithm usually makes poor commercial choices

TURN OFF SPANNING TREE ON YOUR IX INTERFACE

Use ACL's and communities

Use common aggregation policies on your peering and transit

Don't bridge your management network to the public peering fabric

Don't enable an IGP on your peering interface

TURN OFF SPANNING TREE ON YOUR IX INTERFACE

Exchange point operational experiences

Thanks

steve@pipenetworks.com