

Using RPSL in the New Zealand Internet Exchanges

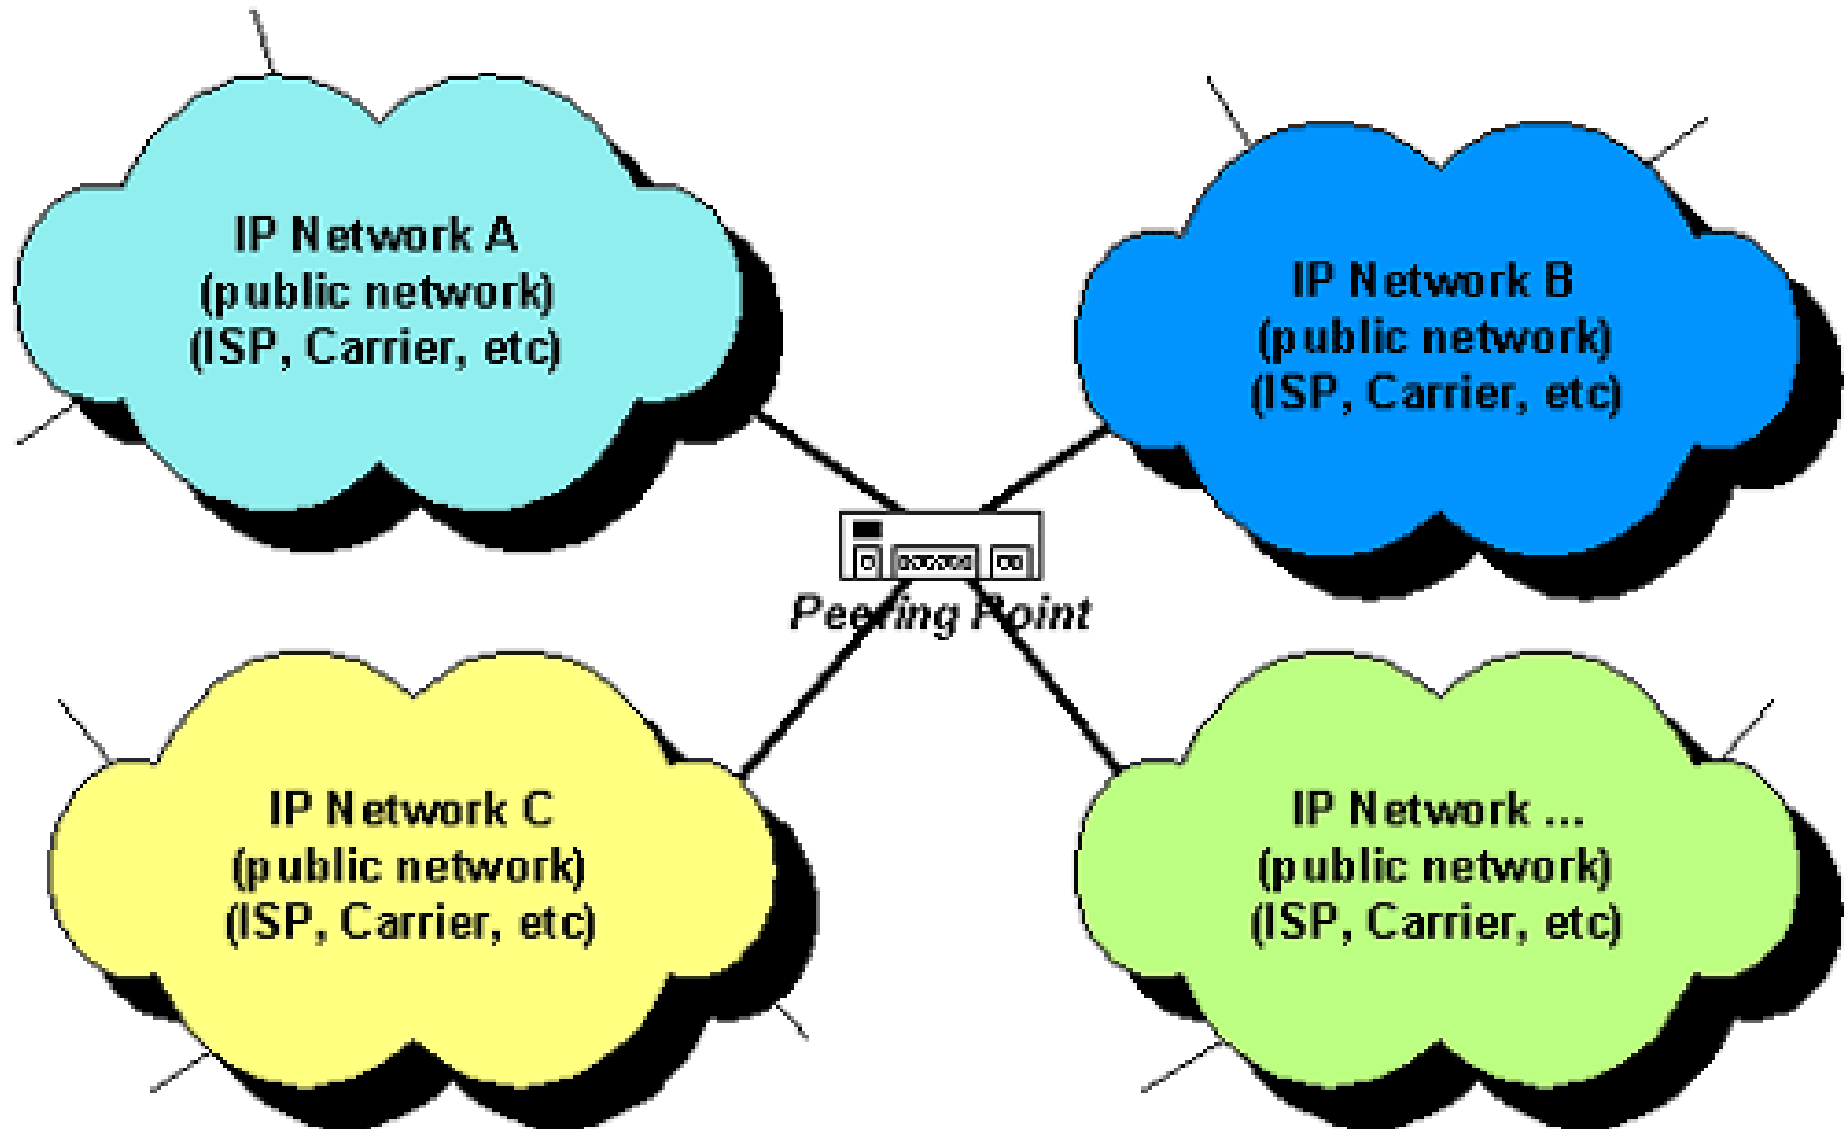
Andy Linton <asjl@citylink.co.nz>
February 2004



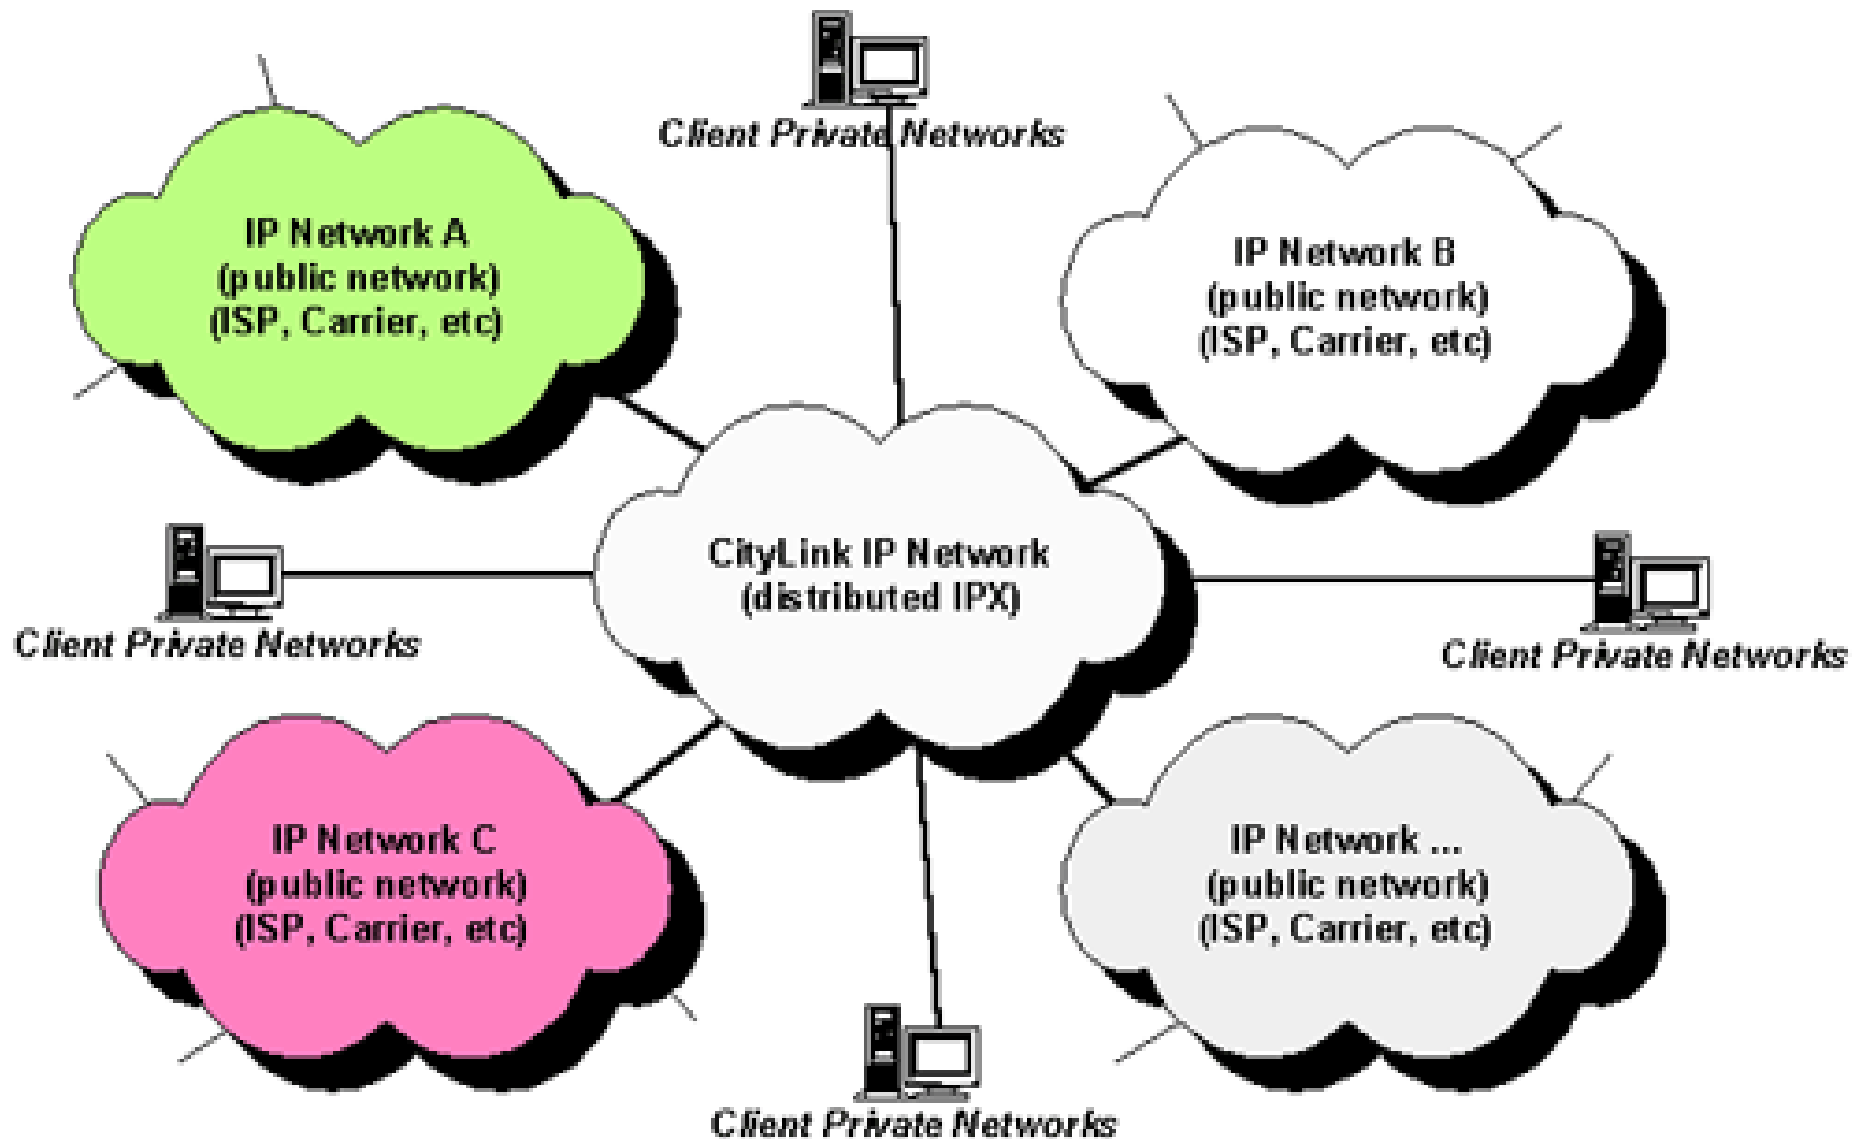
New Zealand Peering Exchanges

- Currently in Auckland and Wellington
 - Auckland Peering Exchange – APE
 - Wellington Internet Exchange – WIX
- Proposed for Hamilton, Christchurch
- Possible for Palmerston North, Dunedin, Hawkes Bay

Traditional Peering Environment



WIX Distributed Peering Environment




WIX Switches

- Highly distributed switch fabric
- 60+ Cisco switches with connections to 300+ buildings in Wellington CBD
 - 35xx, 29xx
- Connections at 10, 100, 1000 Mbps

WIX Peers

- Currently over 120 peering sessions with the WIX route servers
 - Smallest prefixes announced - /29
 - 41 peers have public AS numbers
 - 85 peers have private AS numbers
- We operate our own route registry
 - <http://www.wix.net.nz/cgi-bin/whois.pl>
- We use RPSL to manage all IPv4 route servers
- Looking glass available
 - <http://www.wix.net.nz/cgi-bin/mrlg-wix.cgi>

Route Servers



- 266 Mhz 586 class processor from www.soekris.com
- Uses LEAF code - leaf.sourceforge.org
- Quagga Routing Suite - www.quagga.net

BGP configuration

- Two route servers which need to be identical
 - Config is currently 131471bytes and 4931 lines long
- We need tight control on what each peer can announce to the servers
- If a server fails we need to be able to build a replacement quickly
- All this needs to be open and transparent

BGP policy

- Each peer can announce only their networks to the route servers
- Route servers manipulate data in several ways
 - remove private AS details
 - filter bogons and default (just in case)
 - advertise all routes learnt with a *no_export* community
- *Policy is made visible to community*
 - *whois -r -h whois.wix.net.nz as9439*

AS9439 aut-num object

```
aut-num:    AS9439
import:     from AS9439:AS-PRIVATE
            accept AS9439:RS-ROUTES:PeerAS
import:     from AS9439:AS-PUBLIC
            accept AS9439:RS-ROUTES:PeerAS
export:     to AS9439:AS-PRIVATE
            action community = { no_export };
            announce ANY AND NOT {0.0.0.0/0}
            AND NOT fltr-bogons
export:     to AS9439:AS-PUBLIC
            action community = { no_export };
            announce ANY AND NOT {0.0.0.0/0}
            AND NOT fltr-bogons
```

AS9439:AS-PUBLIC

as-set: AS9439:AS-PUBLIC
descr: Public ASes for NZRR
members: AS42
members: AS24005
members: AS23959
members: AS23905
members: AS23977
members: AS2687
members: AS23755
members: AS23904
members: AS9727
members: AS681

.....



AS9439:RS-ROUTES:AS23754

route-set: AS9439:RS-ROUTES:AS23754

descr: Route set for Citylink to NZRR

members: 203.97.231.224/28,

210.48.103.0/28,

210.48.103.144/28,

210.48.103.136/29,

203.118.144.0/24^24-29,

202.37.19.0/24,

198.48.0.0/22,

202.7.4.0/24,

202.8.44.0/22

source: NZRR



fltr-bogons

- Formed by combining two other filter-sets
 - fltr-unallocated OR fltr-martian
- Based on the work of Rob Thomas. See:
 - <http://www.cymru.com/Documents/bogon-list.html>

fltr-martians

filter-set: fltr-martian

```
filter: {  
    0.0.0.0/8^+ ,  
    10.0.0.0/8^+ ,  
    127.0.0.0/8^+ ,  
    169.254.0.0/16^+ ,  
    172.16.0.0/12^+ ,  
    192.0.2.0/24^+ ,  
    192.168.0.0/16^+ ,  
    198.18.0.0/15^+ ,  
    224.0.0.0/3^+  
}
```

fltr-unallocated

filter-set: fltr-unallocated

filter: {1.0.0.0/8^+,
2.0.0.0/8^+,
5.0.0.0/8^+,
7.0.0.0/8^+,
23.0.0.0/8^+,
27.0.0.0/8^+,
31.0.0.0/8^+,
36.0.0.0/8^+,
37.0.0.0/8^+,
39.0.0.0/8^+,
41.0.0.0/8^+,
42.0.0.0/8^+,

.....



How does all this help?

- We have software that reads a list of peers, extracts the relevant data from the Route Registry and generates a BGP config file
- Data file looks like:
 - 202.7.0.1:23754:Citylink:passive
 - 202.7.0.7:23754:Citylink:passive
 - 202.7.0.8:23755:Citybridge:passive
 - 202.7.0.48:64585:NRG:passive
 - 202.7.0.49:2687:ATT:passive
 - 202.7.0.50:17412:WalkerWireless:



Software details

- A set of ***shell*** and **perl** scripts that call RtConfig to construct a Cisco/Zebra/Quagga bgpd.conf file
- Can also produce Juniper code
- Managed using **make** and **RCS** to provide **backups and consistency**
- **Not a polished software distribution – site specific but available on request**

What's next?

- IPv6 peering also available in both locations
 - Small number of peers so far so configs done by hand
 - RPSLNg is close to standardisation and Rtconfig is available to manage IPv6 policy
 - When manual load gets too high, we'll do this for IPv6 as well

Questions?