Issues with IPv6

John Barlow

Advanced Communication Services Coordinator

grangenet

DNS

- New AAAA record
- Ip6.arpa versus ip6.int
- Abandoned syntax
- Dynamic DNS

DNS

- Just recently got some IPv6 addressed root name servers …
- Reverse DNS is prone to human error
  - Therefore dynamic DNS is required
- See: http://www.tldp.org/HOWTO/Linux+IPv6-HOWTO/hints-daemons-bind.html

DNS 2

Reverse entry sample:

6.a.6.3.8.b.e.f.f.f.b.5.6.0.2.0.0.1.0.0.0.0.0.
1.8.8.3.0.1.0.0.2.ip6.arpa IN PTR
jdb.aarnet.edu.au.

Forward entry sample:

jdb.aarnet.edu.au.   IN   AAAA
2001:388:1000:10:206:5bff:feb8:36a6

Hardware

- High end hardware acceleration
  - Routers (programmable ASICs)
  - Switches (ditto)
  - NICs
- 3G devices
- 64 bit buses
- 64 bit processors

IPv6 headers

- Routing header

IPv6 packet types

- ICMP required

- Neighbor Discovery defines five different ICMP packet types: A pair of Router Solicitation and Router Advertisement messages, a pair of Neighbor Solicitation and Neighbor Advertisements messages, and a Redirect message.

- Router Discovery: How hosts locate routers that reside on an attached link.
- Prefix Discovery: How hosts discover the set of address prefixes that define which destinations are on-link for an attached link. (Nodes use prefixes to distinguish destinations that reside on-link from those only reachable through a router.)
- Parameter Discovery: How a node learns such link parameters as the link MTU or such Internet parameters as the hop limit value to place in outgoing packets.

Firewalls

- Lack of v6 support
- New protocol (new problems ?)
- New IPv6 features may have no support

Tunnels

- Bad network topology
- No honouring IPv6 header options in IPv4 transit
- …

grangenet

6to4

- No method to request reverse DNS delegation
- Limited performance due to tunnels
- Lack of true header use during tunnelling
- Security issues (automatically accept all incoming tunnels …)
- Designed as a transition tool

6over4

- Standard tunnel idea, put IPv6 into IPv4 packets and run that tunnels between two pre-configured end points.
- Usually very manual process, and a good way to get IPv6 packets through a cloud of IPv4 only devices.
- This is how AARNet gets IPv6 into Australia.

grangenet

Multihoming Issues

- Many sites are multihomed in the current Internet
  - reliability
  - stability - which provider will stay in business?
  - competition
  - AUP - commodity vs. R&E
- In IPv4 we can use provider-independent addresses, or ‘poke holes’ in the aggregation
- But all IPv6 addresses are provider-assigned!

grangenet

Multihoming

- To gain redundancy you no longer route one network through two providers.
- You get network address space from each provider, and use both addresses simultaneously.
- When one provider dies your auto-configured IPv6 hosts should timeout their IPv6 address leases and stop using that address prefix …

Multihoming

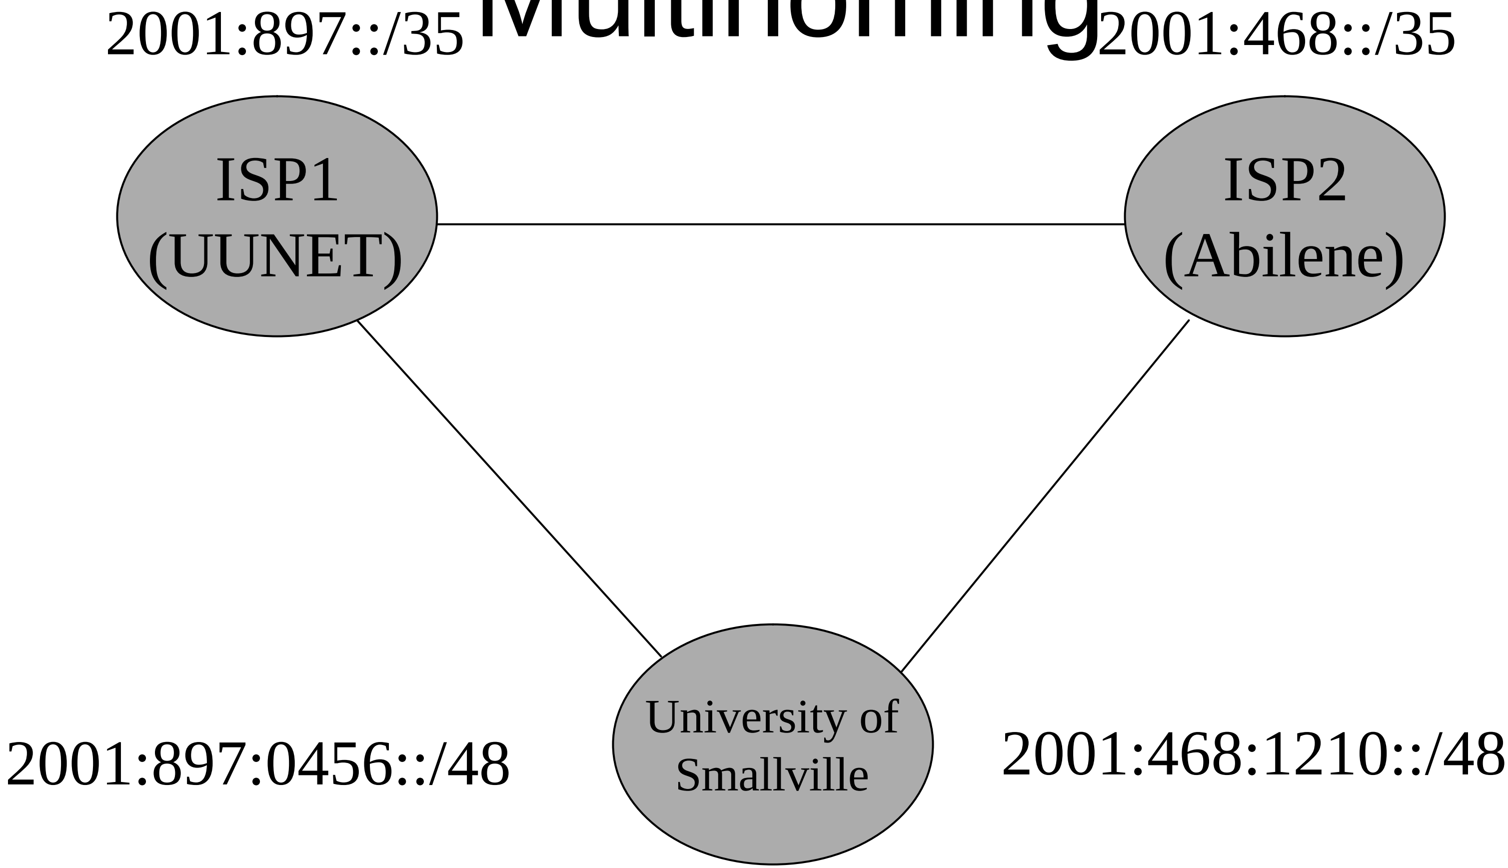

Problems With Multiple Addresses

- If the host or app chooses from several global addresses, that choice overrides policy, may conflict with routing intentions and can break connectivity
- Address selection rules are complex and controversial:

draft-ietf-ipv6-default-addr-select-09.txt

Problems With Provider-Independent

- Current protocols can only control routing table growth if routes are aggregated.
- Only about 12,000 sites are multihomed today, but that number is constantly increasing.
- The address space is so large that routing table growth could easily exceed the capability of the hardware and protocols.

What To Do?

- IPv6 can't be deployed on a large scale without multihoming support - nobody is debating this.
- It seems likely that there will be short-term fixes to allow v6 deployment, and long-term solutions.
- For now, we have some options. . .

grangenet

Get PI Space

- The RIRs have revised their rules for allocating PI space; the key is that you must plan to assign 200 /48s within 2 years.
- This isn’t as hard as it sounds, but it is probably something only gigaPoPs or large university systems can do.
- This breaks when commodity providers start offering IPv6 (unless the gigaPoP aggregates all the commodity providers as well as R&E)

Poke Holes

- The standard practice in IPv4 is to get addresses from one ISP, and advertise that space to all of our proviers - effectively making it a PI address.
- In the v6 world, most providers probably won't advertise a foreign prefix to their peers, but will carry it within their own network.
- Requires that one ISP be designated as the transit provider, and others are effectively peers.

Poke Holes

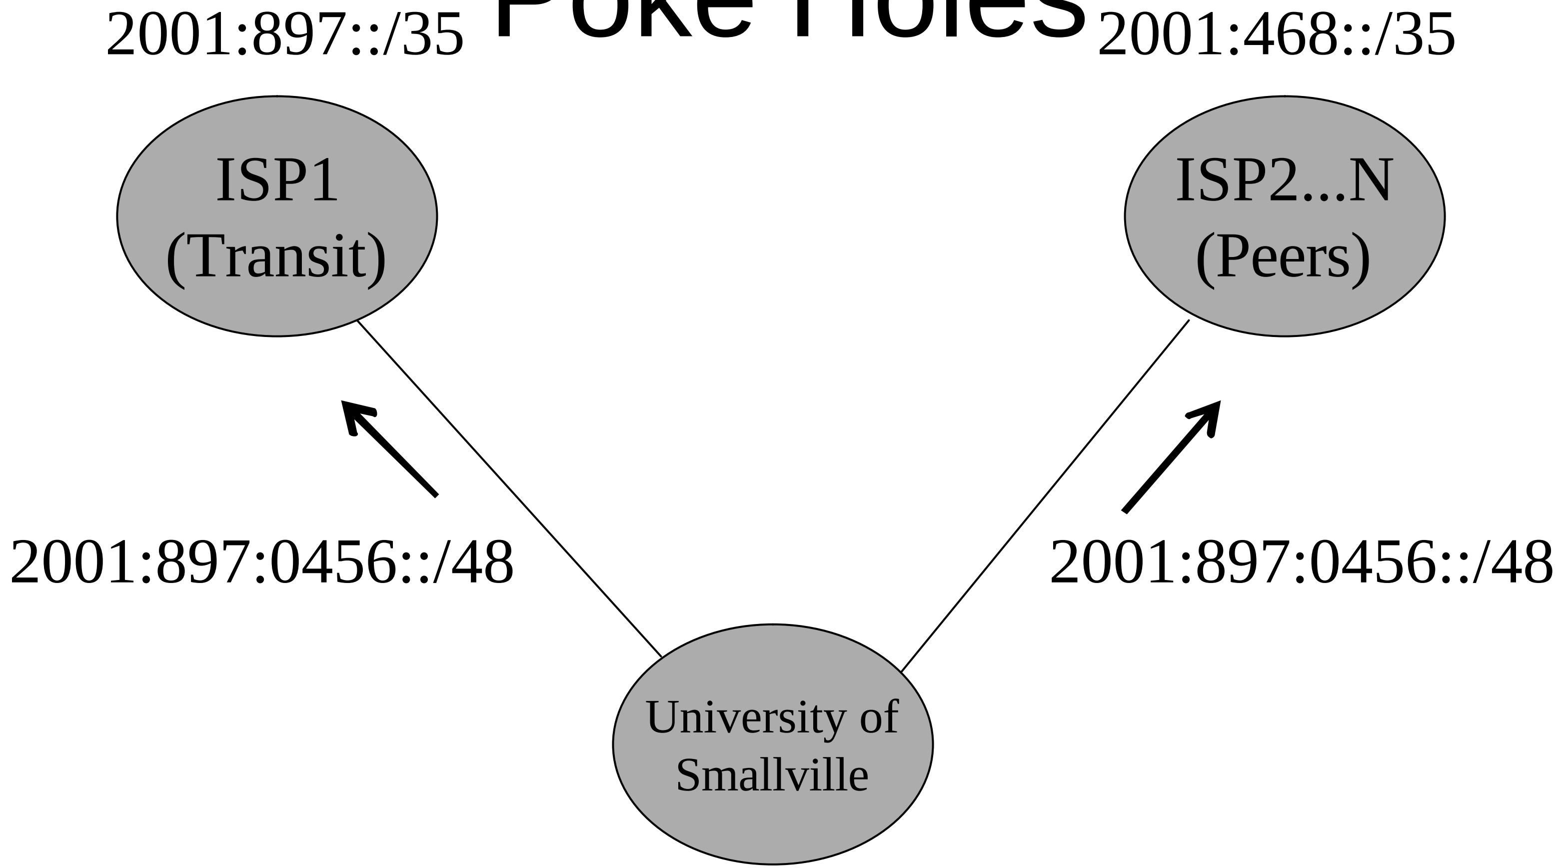

2001:897::/35
2001:468::/35
ISP1
(Transit)
ISP2...N
(Peers)
2001:897:0456::/48
2001:897:0456::/48
University of
Smallville

Migration Plans

- Tunnel individual hosts
- Tunnel to an IPv6 router on your LAN
- Upgrade to get native IPv6 (on your router and/or from your ISP)

Native IPv6 Connection

- Would be really nice, dependant on router support (hardware acceleration and software options).
- Works fine over most layer 2 devices (including wireless).

LAN Issues

- Most Layer 2 devices are fine for IPv6
  - Caveat on the above for IPv6 multicast, which has not been finalised – the issue is the equivalent function of IPv4 IGMP snooping
- Layer 3 devices require software upgrade to handle IPv6
- Hardware accelerated layer 3 devices probably need replacement to accelerate IPv6 (put this requirement on all future purchases)

LAN Issues …

- Can phase IPv6 in gradually using dedicated boxes on each layer 2 segment (in addition to your current IPv4 layer 3 routers)
- Need to rethink the basics
  - Address allocation (Phones, building control, new IP devices)
  - Auto-configuration (compared to DHCP)
  - Multicast services (DNS ? NTP ?)

NAT-PT

- IPv6 “nat” to IPv4 (and back again)
  - Requires DNS server hack
  - As per NAT, every protocol needs to be handled independently
- Allows IPv6 only host to use the (IPv4 and IPv6) Internet