Multihoming

or provider independent addressing

(possible usage)

János Mohácsi NIIF/HUNGARNET

Copy ...Rights

Multihoming Issues

- Many sites are multihomed in the current Internet
  - reliability
  - stability - which provider will stay in business?
  - competition
- In IPv4 we can use provider-independent addresses, or 'poke holes' in the aggregation
- But all globally aggregatable IPv6 addresses are provider-assigned!

Multihoming

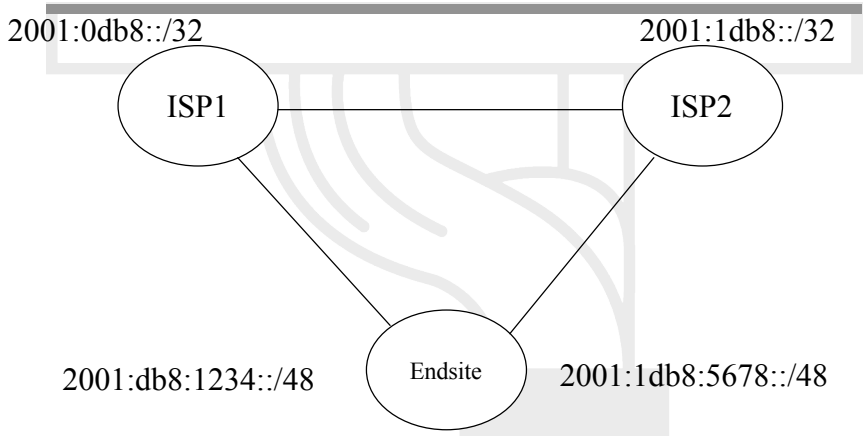

Problems With Multiple Addresses

- Host or Applications chooses from several global addresses:
  - choice should be based on the policy, not conflict with routing intentions and can break connectivity
- Address selection rules are complex and controversial: RFC 3484 - may be configurable centrally – at enterprise environment at least – draft/study exists

Problems With Provider-Independent

- Current protocols (BGP) can only control routing table growth if routes are aggregated.
- More than 10000 sites are multihomed today, but that number is constantly increasing.
- The IPv6 address space is very large
  - routing table growth could be problematical with the capability of the current hardware and protocols.

What To Do?

- IPv6 deployment on a large scale without multihoming support is rather problematical.
- It seems likely that there will be short-term fixes to allow v6 deployment, and long-term solutions.
- For now, we have some options. . .

Get PI Space

- Getting /32 (currently the PI address) is rather easy.
- But it is probably large/medium ISPs and NRENs can get.
- The IPv6 peerings should be more common among thems – but routing table will be very large!

Poking Holes – announcing more specifics

- The standard practice in IPv4 is to get addresses from one ISP, and advertise that space to all of our providers - 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.- may be changing in time
- Requires that one ISP be designated as the transit provider, and others are effectively peers – it is working very well at research communities: NRENs

Poke Holes

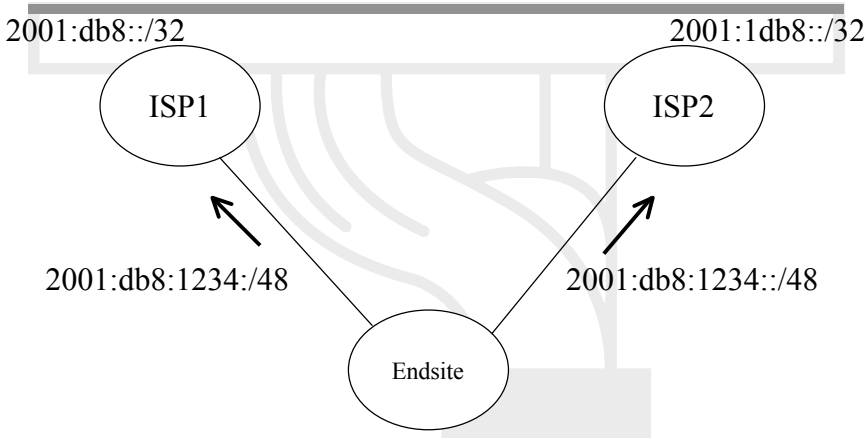

Poking holes – special cases

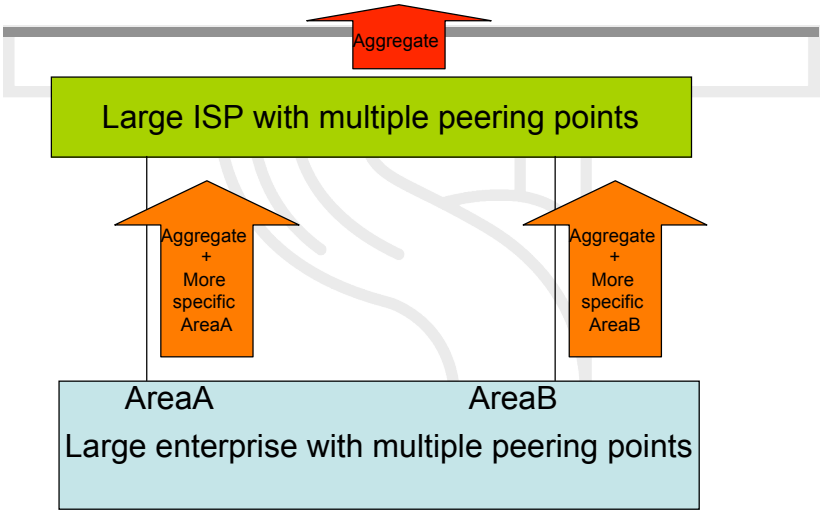

Provider-Independent Addressing?

PI Multihoming – based on geography

- One possible answer to the multihoming/multiple address problem is the use of addresses determined by geography.
- Each site uses the location of its ISP demark to determine its PI address space - put your GPS on top of your router ☺

PI Address Calculation

- Latitude/Longitude each converted to a 22-bit binary number

`40.0433N,23.2781E =`

- Two values concatenated, latitude first

`X412:1220:6cd9::/48`

- X because this scheme is not yet approved, but the expectation is that 1 will be used.

PI Address Calculation-interleaving

- Why interleave? So that as the prefix gets longer, the area included in the prefix gets smaller:

bits	degrees	nominal square	scope	sites
4	-> 90.00000	10000 km		
8	-> 22.50000	2500 km		
12	-> 5.625000	600 km	zone	
16	-> 1.406250	150 km	region	
20	-> 0.3515625	40 km	metro	16777216
24	-> 0.087890625	10 km	city	1048576
28	-> 0.02197265625	2.5 km	locality	65536
32	-> 0.0054931640625	600 m	neighborhood	4096
36	-> 0.001373291015625	150 m	block	256
40	-> 0.00034332275390625	40 m	lot	16
44	-> 0.0000858306884765625	10 m	site	1

PI Address Calculation

- If all the ISPs in an area meet at a local exchange, they may be able to aggregate PI addresses to some degree. – IX should be neutral! – regional traffic routed locally
- But using PI will inevitably mean that more prefixes are carried in the default-free zone (DFZ) at the core of the Internet.

PI Multihoming

- Proposed format: draft-hain-ipv6-pi-addr-02.txt
- Usage discussion: draft-hain-ipv6-pi-addr-use-02.txt
- Remember, this is NOT a standard yet!

PI multihoming using AS number

- Using AS number as a base to generate PI address:

  draft-savola-multi6-asn-pi-01.txt

  AS1955: 0x07a3

  After AS you might get IPv6 address automatically:

  /32 prefix: 2000:07a3::/32

  /48 prefix: 2001:0:07a3::/48

Route Selection for End-to-End Multihoming

draft-ohira-assign-select-e2e-multihome-03.txt

- Goal:
  - Small networks such as a home network or an office network with multiple upstream ISPs
  - So called ISP multi-homing is NOT addressed
- Method:
  - Hierarchical Addressing (Multi-address model)
  - Source Address Based Routing (SABR)

Conditions of a Target Site

- Small site as a home network
- A /48 address space for a site
  - assemble a network flexibly
- Multi links & multi exit routers
- Lower 80 bits are set up in advance

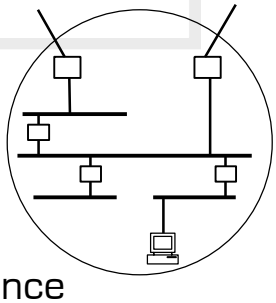

← upstream independent →

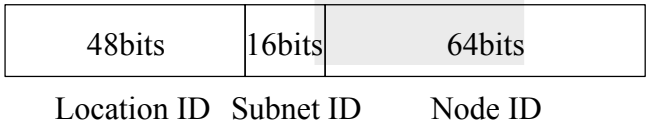

Route/Address Information Management Mechanism (1/2)

- Kinds of information
  - from site external (address delegating)
    - Delegated PA address prefix
    - Proper exit router for each PA address prefix
  - site internal
    - State of links in a site
    - State of links which site exit routers have

Route/Address Information Management Mechanism (2/2)

- Candidate methods to carry the information
  - from site external (address delegating)
    - manual configuration
    - DHCP with prefix option (an I-D is proposed by dhc wg)
      - server: some node in upstream ISP side
      - client: site exit routers
  - site internal
    - manual configuration
    - IGPs with SABR extension

Setup of SABR

- FreeBSD/NetBSD/OpenBSD
  - pf (packet filter)
    - pass out quick route-to (dc0 fe80::1) from 2001:db8:7000:f00::/64 to any
    - pass out quick route-to (dc1 fe80::1) from 2001:1db8:190:f00::/64 to any
- NetBSD (1.6.1)
  - ICMP Extension & ipfilter (need some modifications)
    - route add default fe80::1%fxp0
    - route add default fe80::2%fxp0 -sabrnet 2001:db8:190:f80:: -sabrmasklen 64
- Cisco (after ~~IOS 12.3(7)T~~)
  - working

Intention to link this with DHCP/RA.

Source Address Based Routing (SABR)

- Select an external connection from multiple entries according to a source address
- Pros:
  - No route information from outside
  - No tunnels
  - No servers to mapping between src/dst address
  - No labels nor extensible headers
- Con:
  - Most of intermediate routers and interior gateway routing protocols in a site must be modified

Multihoming with tunnels

RFC 3178 (Informational)

RFC3178 context

- Very little assumption on ISP
- No changes in Router/Hosts
- Copes with p2p link ISP
  - Reduce downtime
- May require ISP cooperation

- Simple elegant solution

RFC3178 proposal

- Configuration of secondary links
- Announce lower preference router over secondary links

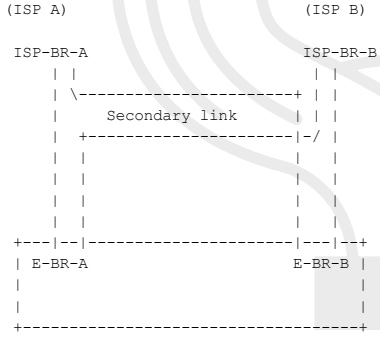

RFC 3178 - initial setup

- Get Address from multiple ISP – route them locally
- IPv6: End host can get multiple address or, single address

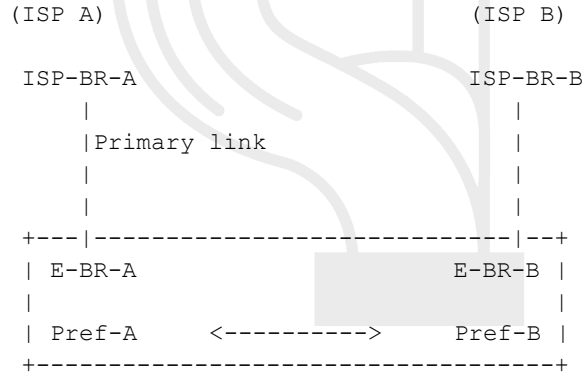

RFC 3178 – link failure

- Link to ISP-A is down, secondary link is used, reachability guaranteed, convergence depends on the routing protocol used

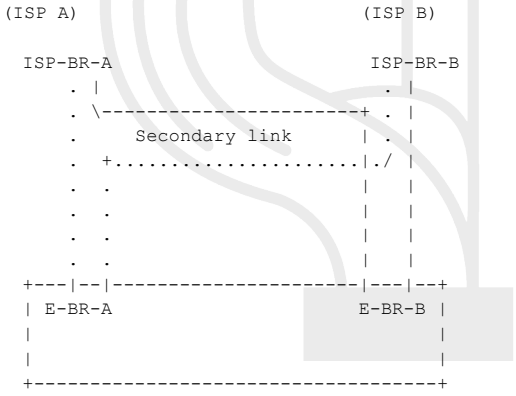

Not quite multihoming – ULA (Unique Local Addresses)

János Mohácsi NIIF/HUNGARNET

ULA Features

- Globally unique prefix.
- Well known prefix to allow for easy filtering at site boundaries.
- Allows sites to be combined or privately interconnected without creating any address conflicts or require renumbering of interfaces using these prefixes.
- Internet Service Provider independent and can be used for communications inside of a site without having any permanent or intermittent Internet connectivity.
- If accidentally leaked outside of a site via routing or DNS, there is no conflict with any other addresses.
- In practice, applications may treat these address like global scoped addresses.
- These addresses are also candidates for end-to-end use in some classes of multihoming solutions.

Format

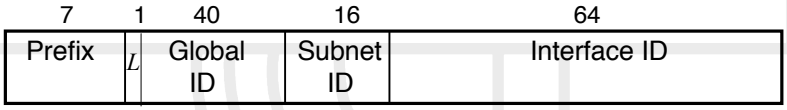

Prefix	7-bit Prefix to identify Local IPv6 unicast addresses (FC00::/7 assumed)
L	Local/Global assignments
Global ID	40-bit Global identifier used to create a global unique prefix (1.1 trillion assignments)
Subnet ID	16-bit subnet ID is an identifier of a subnet within the site
Interface ID	64-bit Interface ID

Global ID

- Generated with a SHA1 based pseudo-random algorithm (specified in draft)
- Two allocations approaches
  - FC00::/8 Centrally assigned
  - FD00::/8 Locally assigned
- Centrally assigned
  - Allows for higher likelihood of uniqueness
  - Escrowed to allow for resolution of duplicate assignment conflicts
- Locally Assigned
  - Generated locally without any central coordination

Centrally assigned

- Single allocation authority to ensure uniqueness and allow for conflict resolution
- Requirements
  - Available to anyone in an unbiased manner
  - Permanent with no periodic fees
  - One time non-refundable allocation fee very low cost per allocation
  - The ownership of each individual allocation should be private, but should be escrowed
- Public Internet Registry (PIR) used as example of allocation authority
  - IANA to establish

Locally assigned

- Locally generated Global ID with pseudo-random algorithm
  - Reasonable likelihood of uniqueness
- No need to contact a assignment authority or ISP

ULA-Review

- Simple - no registration or approval required
  - Local and Central allocation
- Stable addresses
  - Yes, permanent allocations independent of an ISP or ISP connectivity state
- Private
  - Yes, easy to filter on FC00::/7 prefix
- Multiple link operation
  - Yes, 16-bit subnet field
  - Compatible with RFC3177

ULA - Review/2

- Compatible with any site naming system
  - Yes, unique prefix and resulting addresses
- Unambiguous prefixes
  - Yes, pseudo-random generated with local and centralized allocation
- Compatible with VPN
  - Yes, unique prefixes all for inter-site communications and restricted routing

ULA-Review/3

- Makes renumbering easier
  - Internal communication stable ULA
  - External communication – Global address based on names
  - VPNs are problematical
- Proper RFC 3484 implementation is a MUST!
- Proper ICMPv6 error handling is necessary – blackhole has bad side effects for TCP
- May break IPv6 multicasting – ULA is global address
- See more on Network Architecture Protection