Metro-Based Addressing

a proposed addressing scheme for the IPv6 Internet

Steve Deering
Xerox Palo Alto Research Center
deering@parc.xerox.com
http://sandbox.xerox.com/deering

July 1995

Design Goals

- scalable routing
  - » no more than a few hundred entries per routing table
- no renumbering when changing providers (among providers serving same locale)
  - » minimize disincentive to change providers
  - » avoid need for new host technology

Metro Address Structure

FP format prefix

metro geographical region centered on major metropolitan area

site office site, private residence, campus, etc. (up to millions per metro)

note: no provider ID anywhere in address

Illustrative Example

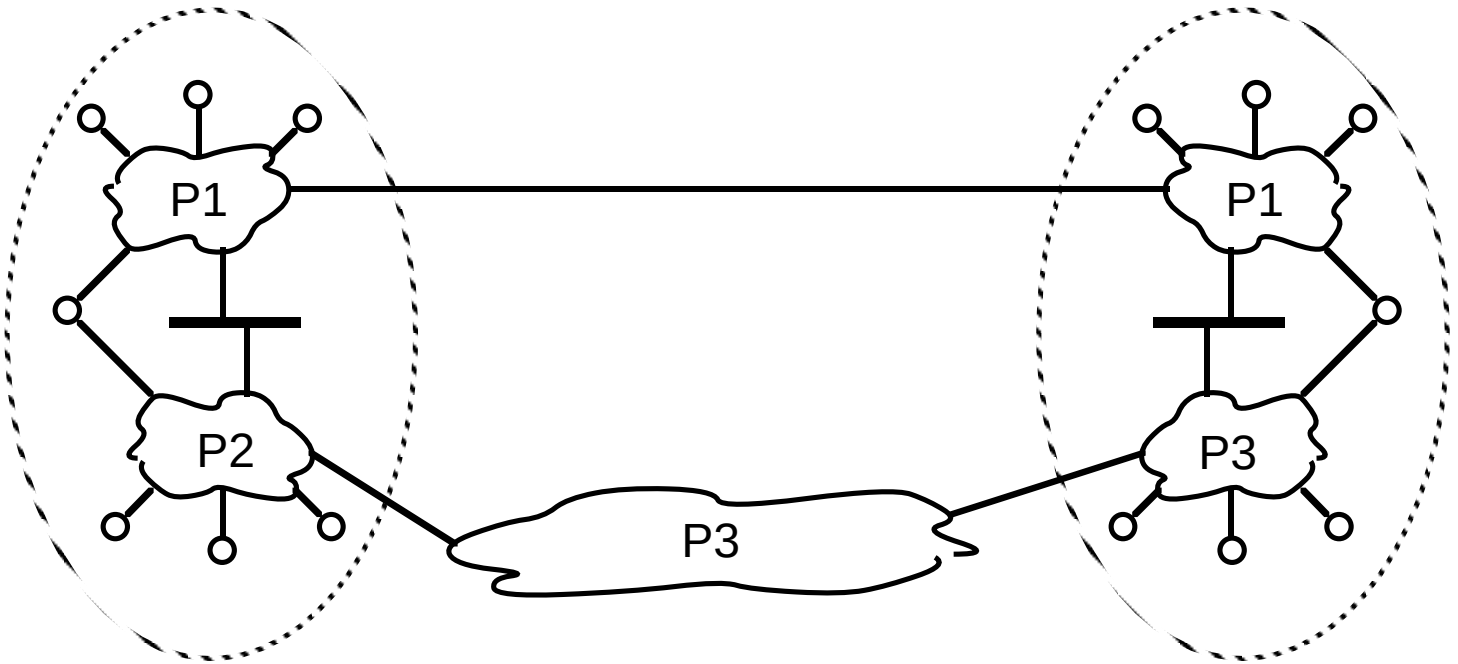

Inter-Metro Routing

- conventional longest-match on country + metro
- bilateral agreements for inter-metro carriage
- a provider may serve any number of sites in any number of metros in any number of countries

Intra-Metro Routing

- all providers serving a metro must be connected **within that metro** (either directly or indirectly)
- simplest form of interconnect is the **MIX** (metropolitan internet exchange)
  - » high-speed LAN or a layer-2 MAN service
  - » can use multiple, redundant MIXes for robustness & load-splitting
  - » operated as joint venture of providers or by a third party, with fair and equal access rules

Intra-Metro Routing (cont.)

to support up to millions of sites, a provider's intra-metro routers do 2-step route lookup:

(1) look up dest. site in customer database to learn site-connected router

(2) if found, look up site-connected router in conventional routing table, and forward

else

forward towards nearest MIX

(do normal caching to avoid lookups on every pkt)

The Customer Database

- has to be maintained by a provider anyway, for billing, maintenance, etc. purposes
- updates distributed to all intra-metro routers once a day => can change providers within 24 hours
- distributed to a provider’s local routers only

Routing at the MIX

a couple of possibilities:

- » multicast query across the MIX; cache answer
- » pre-exchange of site databases

Metro Address Assignment

- metro prefixes assigned to metros by ISOC-IANA (possibly delegated to national ISOC chapters)
- blocks of site prefixes distributed to intra-metro providers for assignment to **new** sites
- site prefix **stays with site** when it changes providers

Provider Selection (Packet-by-Packet)

outbound packets:

- by source hosts, using source routing **or** encaps
- by site boundary routers, using encaps (**plus** optional source routing)

inbound packets:

- by info placed in DNS, or
- by reply info included in packets

Stuff I Don’t Have Time to Explain Today

- rerouting around failed connections to sites, for multi-homed sites
- addressing and routing within multi-metro private networks

an internet draft is in preparation…

Open Issues

- how to achieve required intra-metro connectedness?
  - » market demand?
  - » contract between ISOC and providers, placing conditions on use of metro address prefixes?
  - » note: can co-exist with provider-based addressing
- can multiple, **competing** MIXes be supported?

The Downside

- need for provider connectedness within each metro (**eventually**)

  can phase in MIXes by having one MIX serve multiple nearby metros until density warrants separate MIXes, **but** to avoid future renumbering, must assign addresses based om metro, not MIX

- need for two-step route lookup within metros (**eventually**)

  until provider has many thousands of customers within a single metro, can use existing routing technology

The Upside

- can do scalable routing without site renumbering
- a number of secondary benefits:
  - » usually only one address per interface
  - » “non-surprising” routing (short deliver paths => low delay, possibly higher throughput)
  - » can determine rough distance (~=> delay) between addresses by examination
  - » hard to discriminate against providers
  - » simplified inter-provider coordination

a PostScript copy of these slides may be fetched from:

ftp://parcftp.xerox.com/pub/net-research/metro-addr-slides-jul95.ps