The Map & Encap Scheme for scalable IPv4 routing with portable site prefixes

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

March 1996

Disclaimer / Acknowledgements

These are **not** my original ideas.

Basic idea has been reinvented several times:

Bob Hinden — IP-ENCAPS (**not** IPAE)

Robert Elz (kre) — big-internet posting

probably others

Not an alternative to CIDR or provider-based addressing, but an alternative to mandatory renumbering of sites when they change providers

IPv4 Routing Zones

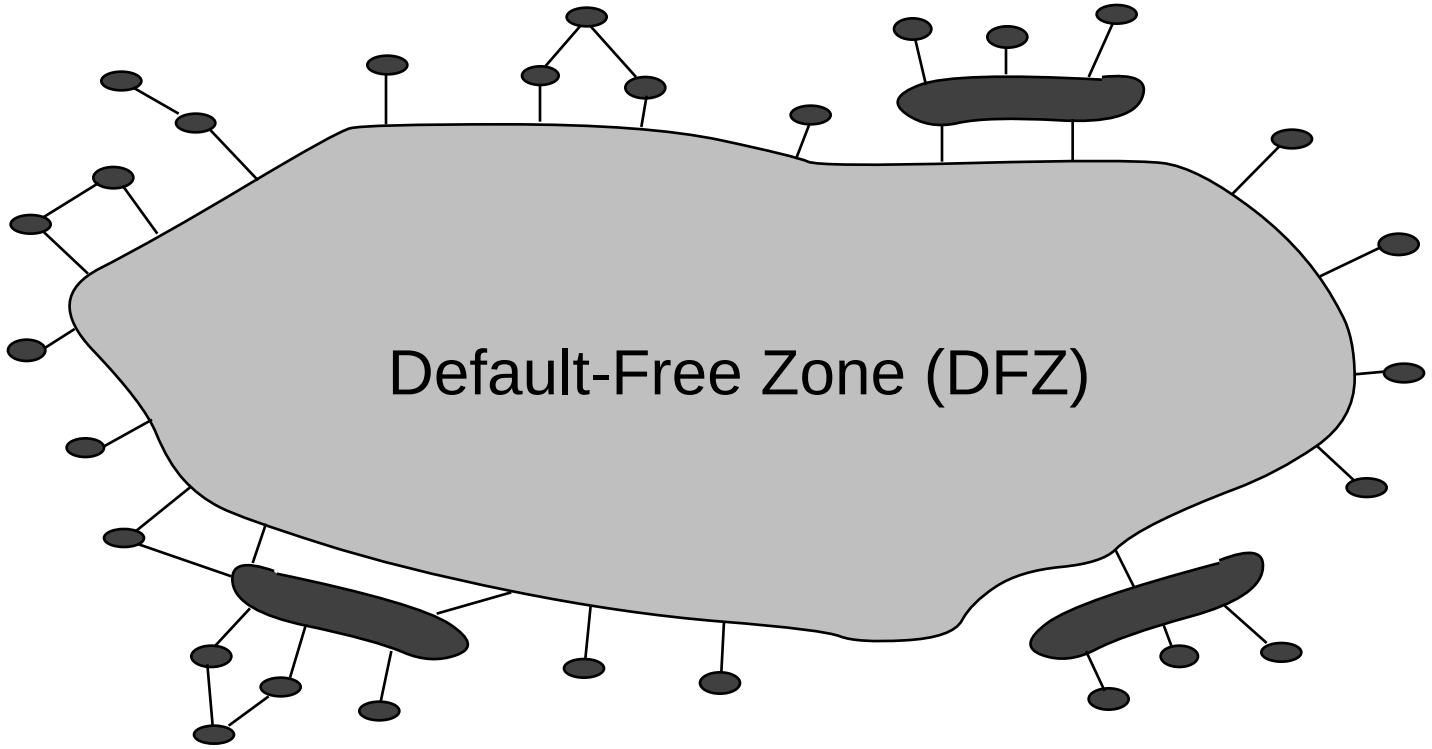

Defaulting Zones (DZs)

Inside the DFZ

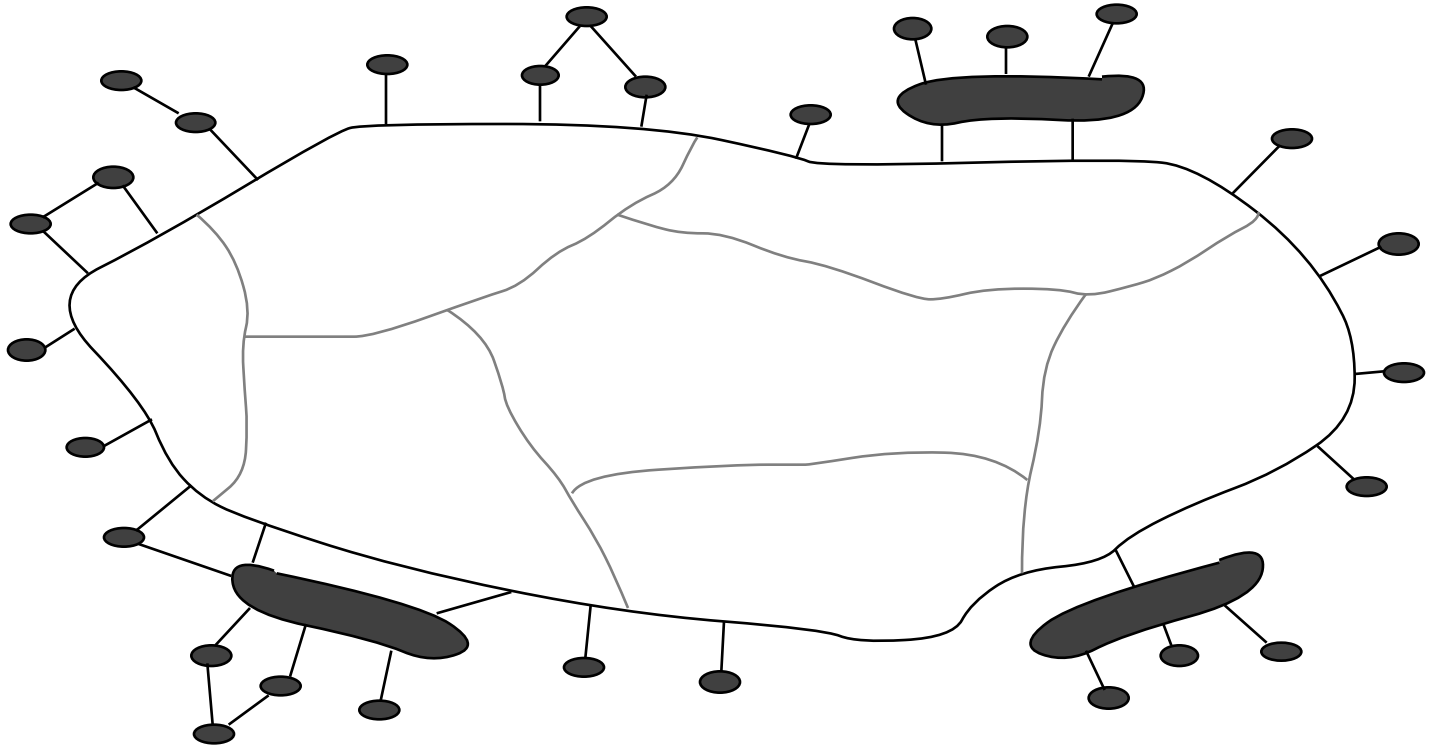

The Basic Idea

- DZs do normal routing on site prefixes (CIDR allocation, but portable across providers); unknown prefixes are routed to DFZ
- when packet hits fist DFZ router, destination site prefix is mapped to DFZ egress router address (or anycast address of set of routers)
- original packet encapsulated with an IPv4 header destined to egress router
- DFZ maintains routes for intra-DFZ topology only
- egress router removes encapsulating header

Intra-DFZ Addresses

- new allocation from IPv4 address space, e.g., one or more Class A's
- CIDR-aggregated according to DFZ provider topology
- when provider topology changes, DFZ routers renumbered as necessary to keep DFZ routing tables small (and mappings updated accordingly)

The Mapping Tables

- each DFZ provider maintains mapping from own customer site prefixes to DFZ addresses of the routers to which they attach
- each DFZ provider distributes own mappings to all DFZ-border routers, e.g., by scope-limited, low-frequency multicast
- resulting mapping tables will have more entries than current DFZ routing tables, but smaller entries and less volatile
  (simple indexed table is max 64 MB)