Indirection
Stating the Obvious?

David Conrad
david.conrad@icann.org

Quote 1

“There is no problem in computer science that cannot be solved by an extra level of indirection.”

Why?

- Problem: Scaling
  - Everybody wants PI
    - If they know enough to choose
      - Multi-homing, renumbering, etc. all made easier
    - RIRs are responding to community demands
      - E.g., ARIN policy 2005-1, RIPE proposed policy 2006-1
  - Nobody wants to change
    - Business models (“geo-addressing is bad”)
    - Technologies and techniques (BGP, traffic engineering, etc.)
- Common proposed solution: Information hiding
  - Oh, how about adding a layer of indirection?

By Way of Example

- Indirection via ID/locator mapping
  - Map ID to locator at source edge/core boundary, remap locator to ID at destination core/edge boundary
    - How mapping/re-mapping is done is an implementation detail
      - Tunnels, address rewriting, etc.
    - How the map is propagated is an implementation detail
      - Pull (e.g., DNS), push (e.g., via routing system flooding)
- End points only know about IDs
- Routing system only knows about locators
  - No more layer violations at layer 3 and 4

Simple Multi-Homing Example

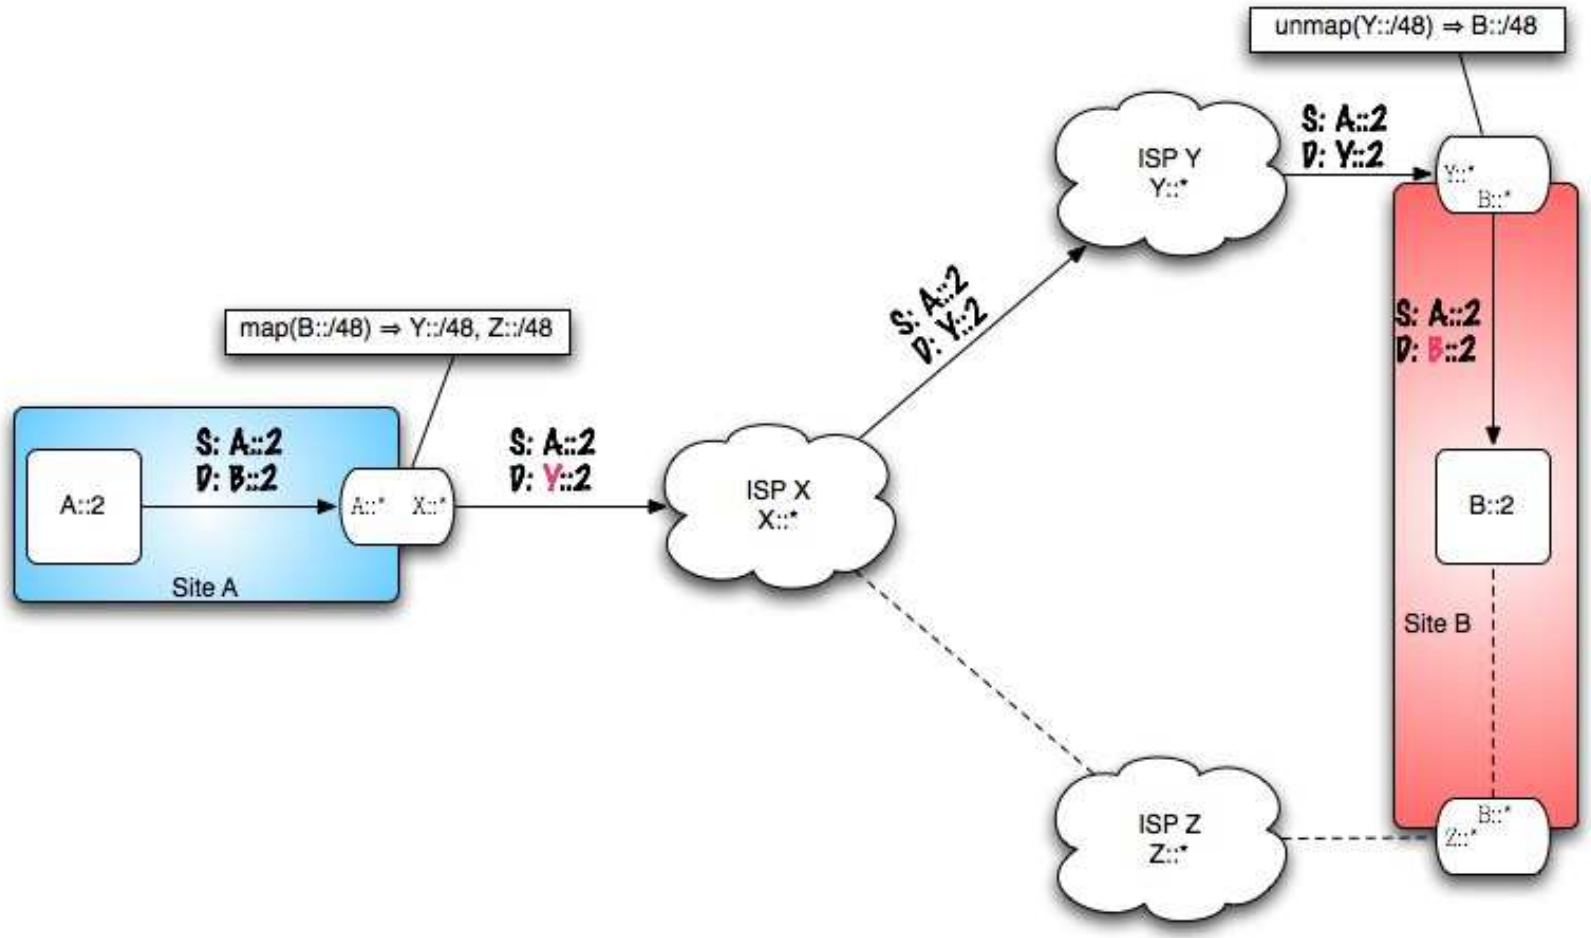

Potential Benefits

- More scalable multi-homing
  - ID maps into multiple locators
    - Selection of which locator to choose is a policy decision
      - Potentially allow multi-homed site to specify preferences
  - **Scales to number of ISPs, not number of sites**
- Renumbering/Nomadicity/Mobility
  - ID to locator map changes over time
    - “Time” depends on how long it takes to propagate the map change
      - If time is very short (e.g., planes transitioning ISP boundaries), the previous locator endpoint forwards to the next locator endpoint (“handover forwarding”)
  - **No end system changes**

More Benefits

- Change only at the edge/core boundary
  - No change to end user deployed base
    - No application or end device IP stack changes required
  - No change to ISP core infrastructure
  - No change to ISP business models
    - No forced peering/settlements necessary
- Can be used for IPv4 to IPv6 migration
  - Or rather, IPv4 sites communicating with other IPv4 sites over an IPv6 infrastructure

Yet More Benefits(?)

- No change to most ISP routing technology or techniques
  - Still BGP/CIDR within the routing system
- IDs do **not** need to be allocated hierarchically
  - No hierarchy waste
  - Geo-political address allocation would be fine
    - Get the ITU folks off our back
- Locators can continue to be allocated by the existing players (RIRs/LIRs)
  - Geoff's job is safe!

Mapping

- Common question: How do you do the mapping between locator and ID?
  - Plenty of ways, all have cost/benefit tradeoffs, e.g.:
- Pull (e.g., DNS)
  - "Indirector" looks up ID gets back "LOCATOR" RRs, caches them
    - No new protocol need
    - First indirection at edge takes longer than subsequent but only useful data is fetched
- Push (e.g., Routing system-like flood)
  - "Indirector" receives updates propagated like any other routing update
    - Probably need new protocol
    - No delays for first indirection, but more memory and bandwidth required, even though most data not used

Quote 2

“Any performance problem can be solved by removing a level of indirection"

Drawbacks

- Loss of information
  - To scale, information is hidden
    - Some of this information can be valuable
- Performance
  - Resolving the indirection takes time
- Increased complexity
  - How much depends on how the indirection is implemented
  - Additional network element (the “indirector”)
- Security?
  - Does it make spoofing easier?
    - Does that matter?
  - What do you filter on?

Simple Multi-Homing Example

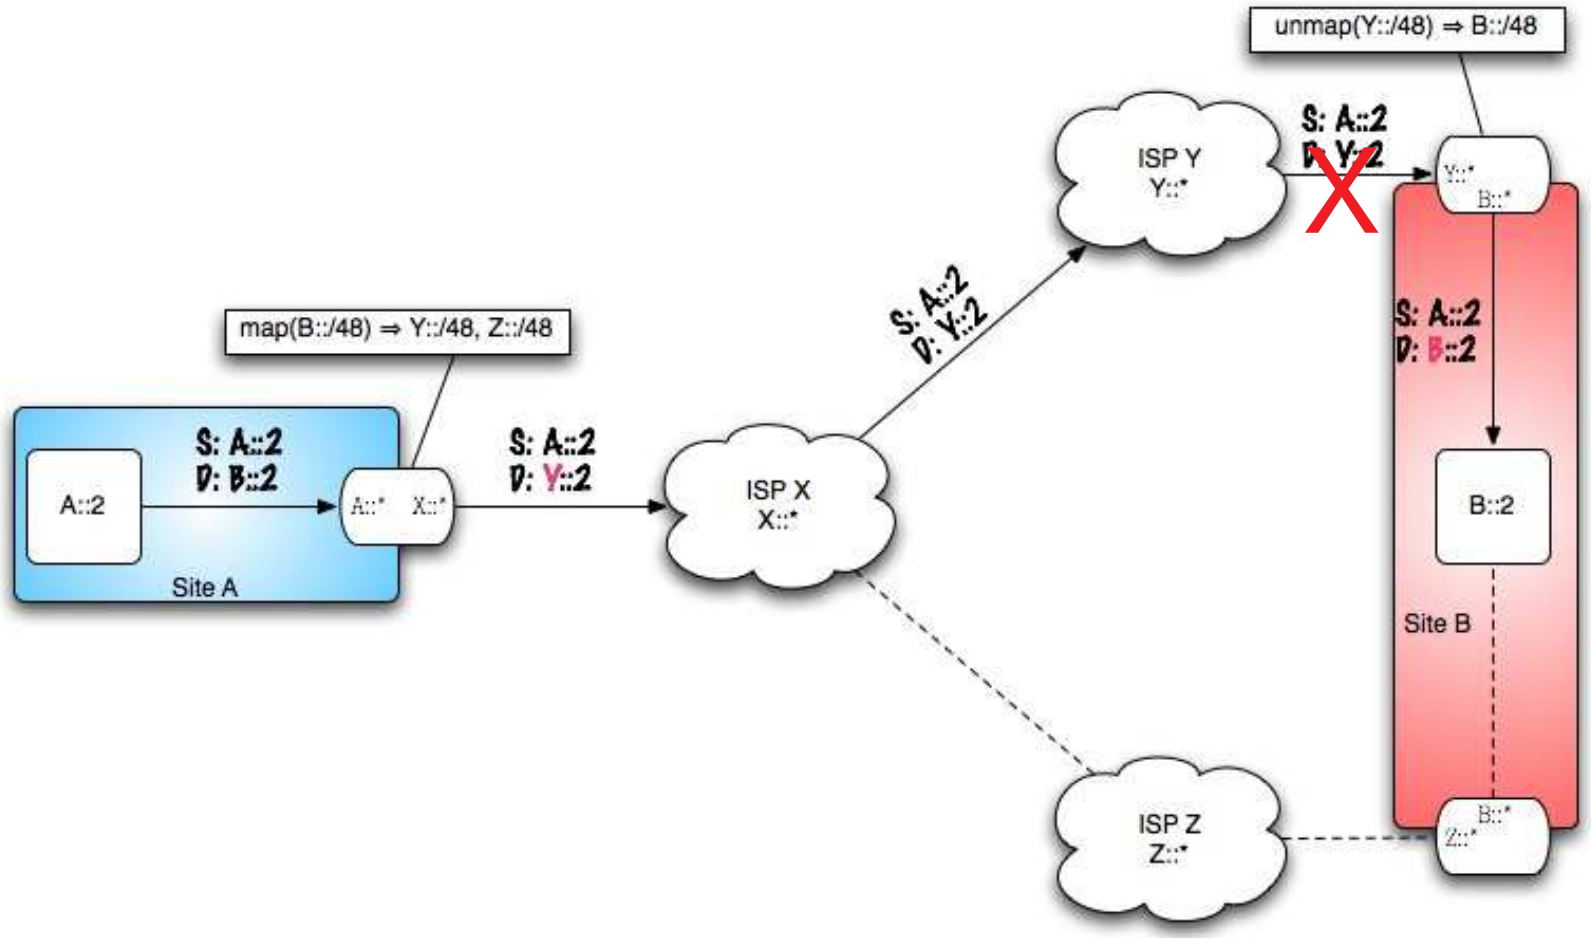

More Drawbacks

- No change to ISP routing technology or techniques
  - Still BGP/CIDR
- Map maintenance
  - Reachability notification
    - When should a mapping change?
    - How is the change propagated?
- Increased address space usage as compared to PI
  - In the example, each homing requires separate locator prefix to be mapped to a single ID prefix
    - 2x address space consumption at a minimum
  - If this matters in IPv6 (probably not)

Observations

- Indirection isn't new
  - RFC 1955 (IP Encaps), CRIO (Paul Francis, et al), among many, many others
    - All indirection of one form or another
  - We already have locator/ID splits
    - Domain names → IP addresses (when HTTP is the new IP)
    - Network part → host part
- **Any** solution is a cost/benefit tradeoff
  - What are customers/ISPs willing to give up vs. what are they willing to pay for.
- The solution **will** involve information hiding
  - What do you want to lose today?