Emerging Routing Issues

John Scudder, David Ward

2007.03

Why is the routing community here? Tune in, turn on, no time out …

- The last 6-9 months has seen a renewed desire to visit routing and addressing architectural issues
- In fact, the routing community has been discussing this for >15 years
- This talk attempts to lay out an overview of some issues discussed by the community in various forums, consortia, working groups and design efforts
- It represents a view that we need to clearly define the problem and boundaries of the solution

  Will dictate if we enable new tools and network architectures or if we solve a smaller set and remain with the building blocks we have today

  No value judgment is given either way, just fud for discussion…

Fundamental Requirements from Routing Community

- We want the rib/fib growth to flatten or be negative
- We want the dynamics of the routing system to slow down

Let me discuss a few more goals …

Baseline Preferences (nod to Dave O. and Dino F.)

- Routing folks prefer a mechanism to:
  - Associate an ID with a set of Locator addresses
  - Forward packets using Locator addresses
    - IDs may not have to be routable
  - Maintain the reachability status of Locator addresses
    - Hosts can change, networking nodes can change

- Routing folks prefer:
  - Incremental deployability
  - Little modification of Internet infrastructure
  - Reduction of transit router state load
  - No new, specialized ID/Locator binding service
    - Much thought must be applied here
  - Provides benefits to both Sites and Providers
    - Who pays and who benefits?

Site-Based Goals

- Sites need to be multihomed
  - Connected to more than one provider
- Sites need flexibility to change providers
  - With easy or no renumbering
  - While maintaining session survivability?
- Site-supported devices need to be mobile & roam
  - While maintaining session survivability?
- Sites must be able to advertise TE/service desires
  - Enable multi-provider load balancing

Provider-Based Goals

- Providers need their routers to scale in multiple dimensions with competing requirements
  - Power || cost = f(pps, features)
- Providers need to maximize their resources to deliver cost effective connectivity
  - Including Traffic Engineering
- Provider-supported devices need to be mobile & roam
  - While achieving scalability
- Providers need to be able to prevent a bad-actor from hijacking their network paths and mapping function

… the end of the beginning

- If we are re-architecting internet routing and addressing …

  Do we want to enable a new toolset to build different network functionality?

  Should the following issues be included in or outside the bounds of the solution?

- Additional issues of concern voiced in greater routing community:

  Solutions to ***Network partitioning***

  Mobile ***A***d hoc ***NET***working

  Secure routing (paths) and forwarding between networks and sites

  Real-Time registration and resource mapping that can be used for path selection

  Service locators in topology

  Inability to have a single end-point represented in ***multiple service domains***

A tale of two databases

1. Routing Database (RIB becomes FIB)
   - No separation of provider from “customer” (aka site)
     - Provider-based addressing
     - Current site multi-homing, migration, TE and service solutions cause additional state into the routing system
       - Local operational state propagated globally

2. Mapping Database (Name to address)
   - Database (DNS) of customer name->provider address
   - Today no association in mapping database of identifier to locator

- The two primary databases (Routing and Mapping) running the internet are still in evolutionary progression from initial birth

  Issues of past ~20 years not addressed

Two databases directly related

- Architectural consideration but lack of design and development of relationship between addressing, mapping, security and routing
- How routing and mapping work together is critical to defining the problem and finding an appropriate solution
  - Both databases assume static endpoint, simple resource statements, minimal security
- Mapping requirements and Destination types result in need for at least three successive bindings:
  - name to identifier
  - identifier to locator
  - locator route/path

Operational issues (nod to Scott Shenker, Jen Rexford, Nick Feamster)

- Providers (Tier-1) have accumulated a large number of noncontiguous prefixes (site multihoming, TE, non-topological assignment policies, consolidation)
  - Effectively random numbers
  - Policy sets (based on AS) must be matched against random numbers
  - Routing policy doesn't guarantee desired results
    - Not easy to prevent erroneous announcements
- Set of transit routes == full enumeration and state maintenance of all sites and perhaps end-points
- If must announce my more specific route for TE/LB reasons, may not need operational state reflected into global internet

Summarizing the overview

- Discussion of BGP and some minor improvements later today in RTG-area

  Clarifying what toolset we have today and may be able to be done to help

- Re-chartered RRG to research routing design issues
- Need scalable router solution
  - What is the role and desired functionality of routers?
- Critical both co-dependence of routing and mapping are considered thoroughly
- Timeframe?
  - Routing community exploring short, medium and long term changes
  - No clarity in routing community of exact requirements, design or solution

End