

Routing and Addressing: where we are today

Prague, IETF 68

March 2007

Agenda

- Recent activity in context
- This week's activity
- Analysis since the Amsterdam workshop
- Where to search for solutions
- Summary
- Appendix: further off-line reading

Disclaimer: there is a wide range of views on these questions; these slides give only one view.

- Open microphone

Context for Recent Activity

- Historical timeline
 - Packet switching invented (1962)
 - Internet concept invented (1974)
 - IP designed (~1978)
 - BGP designed (~1988)
 - CIDR designed (1992)
 - IPv6 designed (1995)
- Growing concern about scaling, transparency, multihoming, renumbering, provider independence, traffic engineering, IPv6 impact (1995-2006)
- IAB Routing & Addressing workshop (2006)

Context (2) - Architecture

Architectural principle to uphold:

- A network should be able to implement reasonable internetworking choices without unduly impacting another network's operation

The issue, at an architectural level:

- Some of today's internetworking seem only implementable in ways that threaten this principle. This is the root cause of ISP problems and end site dissatisfaction. What can be done to harmonize the network to that architectural principle?
 - Tragedy of the commons from PI address usage

Context (3) - Technical goals

- Solve end user problems
 - Connect to multiple ISPs while maintaining support for current TE capabilities
 - Change ISPs without major switching costs
 - Support e2e session transparency, ...
- Solve ISP problems
 - Keep BGP table size and dynamics within router operational capability
 - Provide ISPs with ability to engineer traffic flows to match business needs

Context (4) - Scaling

- Today, 200k Internet BGP routes and several times *more* customer and VPN routes is common
- A goal for 2050 is 10 billion end nodes with 10 million multihomed customers
- Can we get there from here at reasonable costs for vendors, ISPs and user sites?
- What should be the 5 year goal?

Recent Activity (1) - general

- Refining the IAB workshop report
- Analyzing the concerns it raised
- Looking at solutions
- Meetings, including
 - NANOG BOF February 5
 - DARPA R&A workshop February 21/23
 - APRICOT session February 27
 - TERENA workshop* February 22
 - NSF/OECD workshop* January 31

*More general meetings, but touching on R&A

Recent Activity (2) - IETF, IAB, IRTF

- R&A Directorate established
 - to expand after the IETF, review in 6 months
 - role is coordination and communication
- Routing Research Group recharter
- R&A discussion list active (ram@iab.org)
- Internet and Routing ADs have been preparing for this meeting

This week's activity

- This report
- Internet Area (Thursday, 09:00--11:30)
 - ROAP (ROuting & Addressing Problem) discussion
 - BOF-style focus on future IETF work on identifier-locator separation and multi-level locator designs
- Routing Area (Thursday 13:00--15:00)
 - focus on BGP table growth and dynamics
 - early thoughts about BGP extensions/practices that might help
- Routing Research Group
 - met on Saturday and is very active

Analysis since IAB R&A workshop

- BGP4 used since 1994 with little change
- Mounting concern at growth in BGP table
 - size & update rate, impact of multihoming
 - largely orthogonal to IPv4 v IPv6
- IAB Routing & Addressing workshop 10/06
 - concluded that there is a problem, but more work needed to define it
 - see draft-iab-raws-report
- Multiple discussions since then, sometimes confused and confusing
- Need to clarify what we know, what we don't know, and how to proceed systematically

Is there a "hot box" problem?

- The workshop looked at hardware trends that raise economic and engineering concerns about the scaling of the FIB (Forwarding Information Base)
 - Unless the FIB scales sub-linearly with the number of end sites, it breaks current multihoming and PI practice
- No agreement in the community...
- In our own opinion, recent analysis (and new architectures that people have used to build routers) indicate that for at least for two more generations of microelectronics (45nm, 32nm) this isn't a problem
 - looks like 5 to 10 years of growth; core router scaling seems to be dominated by line speed multiplied by functional complexity (in the forwarding path), not RIB/FIB size.
- *No need for panic*

Is there a "hot box" problem? (Continued)

- But... we are still not following the architectural principle of avoiding global impacts of local decisions
- Is there a way to improve?
- Observation: we have spent a lot of time arguing about whether there is a problem, crystal ball predictions, hardware architectures, etc.
- We would like to put this debate behind us – focus on whether there is some useful work that addresses the architectural issue (without breaking backwards compatibility, etc)

Is there a "hot wire" problem?

- It's been said that even if we contain the RIB size, BGP4 dynamics (update messages) will saturate... something.
- There is experimental evidence of a lot of update traffic that is potentially redundant (i.e. wasting a lot of energy on transient connectivity glitches)
 - but we seem to have no analytical model for the impact of this as the network continues to scale
 - if problematic, it means the update rate has to scale sub-linearly with the number of end sites
- *This problem needs continued investigation - see Routing Area meeting tomorrow*

The transparency problem

- Since 1981, upper layers have assumed that a *Thing That Looks Like An Address* **is** an address.
 - Application programmers often assume that an IP address is a valid end system identifier that can also be passed on to third parties
- In consequence, the fact that addresses are sometimes merely locators is a problem (which NAT, STUN etc. deal with in their own ways)
- The historical reliance on address transparency creates specific difficulty for multihoming and traffic engineering
- *This problem seems to need a solution*

Solution directions*

- RIB/FIB scaling - engineering by microelectronics and router designers
- Update dynamics - BGP adjustments, better operational practices
- Traffic engineering, multihoming, e2e transparency, and mobility would benefit from architectural changes
 - identifier/locator separation and/or multilevel locators form a hopeful approach
- All these are orthogonal to both IPv6 deployment and application level namespace issues

* See Appendix material for further thoughts

Summary 1: The IETF role

- We can provide a forum for open problem analysis
 - vendors, users and operators together
- We can provide a forum for open development of solutions
 - vendors, users and operators together
- We can't do research (the IRTF can)
- We can't control economic or business behaviors

Summary 2: Technical

- In the short term, routing table growth is "only" an engineering issue
- Routing dynamics needs to be better understood, but is likely also an engineering issue
 - to be addressed by stronger pushback in the ISP community, implementation improvements, protocol improvements
- Thus, there is reason to believe we do not have a short term technology problem
 - But hard work and business issues are ahead.
- But there are architectural problems
- IETF can help in short term protocol work
 - Such as tuning BGP better for today's challenges
- IETF can also help by looking at architectural changes
 - Such as identifier/locator separation & multi-level locators

Summary 3 - Overall Plan

Divide and conquer:

- Continue dialog with the operator community
- Pursue implementation improvements
- Evaluate incremental BGP improvements
- Encourage Id-Loc separation and multilevel locator research and experimentation
- Define an IETF Id-Loc separation or multilevel locator solution

Appendix – Further Reading

<http://www.ietf.org/internet-drafts/draft-iab-raws-report-01.txt>

http://submission.apricot.net/chatter07/slides/future_of_routing/apia-future-routing-john-scudder.pdf

http://submission.apricot.net/chatter07/slides/future_of_routing/apia-future-routing-jari-arkko.pdf

And more material in

<http://www3.ietf.org/proceedings/07mar/agenda/intarea.txt>

<http://www3.ietf.org/proceedings/07mar/agenda/rtgarea.txt>

<http://www1.tools.ietf.org/group/irtf/trac/wiki/RRG>

Directorate:

<http://www.ietf.org/IESG/content/radir.html>