

Current Challenges in Internet Technology

with Special Focus on Routing Scalability


CNGI, Beijing, June 2007

Jari Arkko

Internet Area Director, IETF
Researcher, Ericsson Research

Outline

- Overview of the challenges
- The routing and addressing problem
- Current IETF efforts



Overview of the Challenges

- Security problems of various kinds
 - Issues related to scaling to billions of users
 - Issues related to new environments
 - Issues related to new applications
 - Deployment related issues
-

Overview of the Challenges

- Security problems of various kinds
- Issues related to scaling to billions of users
- Issues related to
- Issues related to
- Deployment relat



Dark Reading

Five Security Flaws In IPv6

Kelly Jackson Higgins 05.09.07, 10:45 AM ET

News Front Page



Africa

Americas

Asia-Pacific

Europe

Last Updated: Thursday, 1

E-mail this to a friend

Estonia hit by '

Estonia says the country's websites have been under heavy attack for the past

Ready or not, IPv6 is finally visible on the horizon... And researchers are already finding major security problems with it.




Overview of the Challenges

- Security problems of various kinds
- **Issues related to scaling to billions of users**
- Issues related to new environments
- Issues related to new applications
- Deployment related issues

In this model, IANA allocates its last IPv4 /8 to an RIR on the 18th December 2009

Overview of the Challenges

- Security problems of various kinds
- Issues related to scaling to billions of users
- **Issues related to new environments**
- Issues related to new applications
- Deployment related issues



Overview of the Challenges

- Security problems of various kinds
- Issues related to scaling to billions of users
- Issues related to new environments
- **Issues related to new applications**
- Deployment related issues

P2P



Overview of the Challenges

- Security problems of various kinds
- Issues related to scaling to billions of users
- Issues related to new environments
- Issues related to new applications
- Deployment related issues

QoS
Global PKI

End-to-end Ipsec
EAP in WLAN
BCP 38

The Routing Scaling Problem - Outline

- The concern
 - How Internet routing works
 - Pressures facing the routing system
 - The architectural principle
 - Tentative conclusions
 - Current IETF efforts
-

The Routing Scaling Problem

- The ability of the Internet routing system to cope with the growth of the Internet has been a concern during almost the entire life of the Internet
- The routing and addressing architecture has stayed very similar from the initial days
 - BGP designed in the 1980's
 - CIDR introduced in the 1990's
 - IPv6 designed in the 1990's
- Recent concern from major operators about the growth of the routing problem beyond
 - The growth of the Internet itself
 - Moore's law



Routers in the global network
of Internet Service Providers
exchange dynamic topology
information via BGP





Each router has a full global picture of the Internet - all address prefixes

Pressures Facing the Routing System

- Growth of the Internet itself
 - Organizations want independence from providers
 - Due to competition, there is a desire to be able to switch providers
 - Hard to renumber a network
 - "Provider Independent" address space
 - IPv6 had a model based solely on provider aggregated addresses
 - In reality PI is becoming available for IPv6
-

Pressures Facing the Routing System (Cont'd)

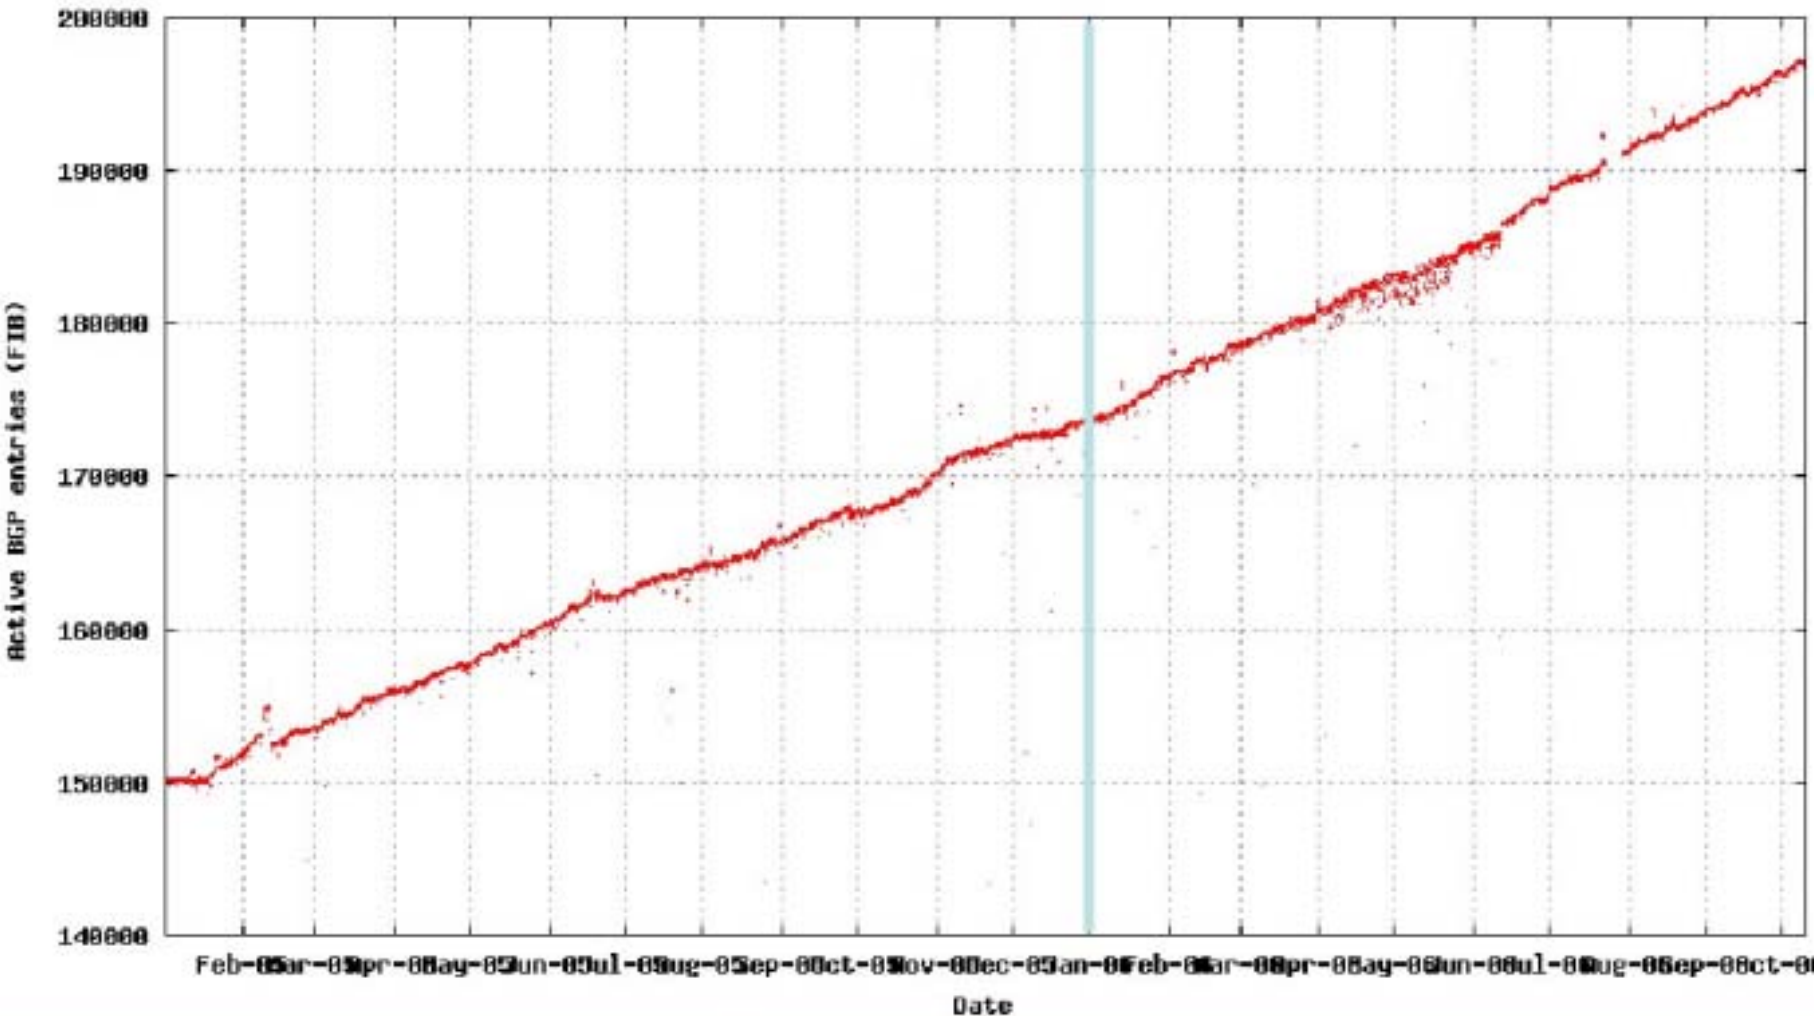
- IPv6
 - Co-existence with IPv4 implies need to hold both IPv4 and IPv6 routing tables
 - More address space implies potential for bigger routing tables
 - Multihoming
 - Solutions exist for host-based multihoming
 - But desire to work with unchanged hosts
 - Or control the multihoming behaviour from the network administrators
-

Pressures Facing the Routing System (Cont'd)

- Traffic engineering
 - Need to control routing on a fine granularity
 - Dynamic changes
- Errors
- Fraud and address space theft

All these imply growth in the number of address prefixes visible in the global routing table, and an increase in the rate of change

Measurements – Prefix Growth




→ Source: Geoff Huston at IAB RAWS Workshop, 2006

Table Size – Main Observation

- Table sizes grow 2x faster than real growth
- One (conservative) analysis predicts 2M entries in 10 years

Measurements – BGP Updates

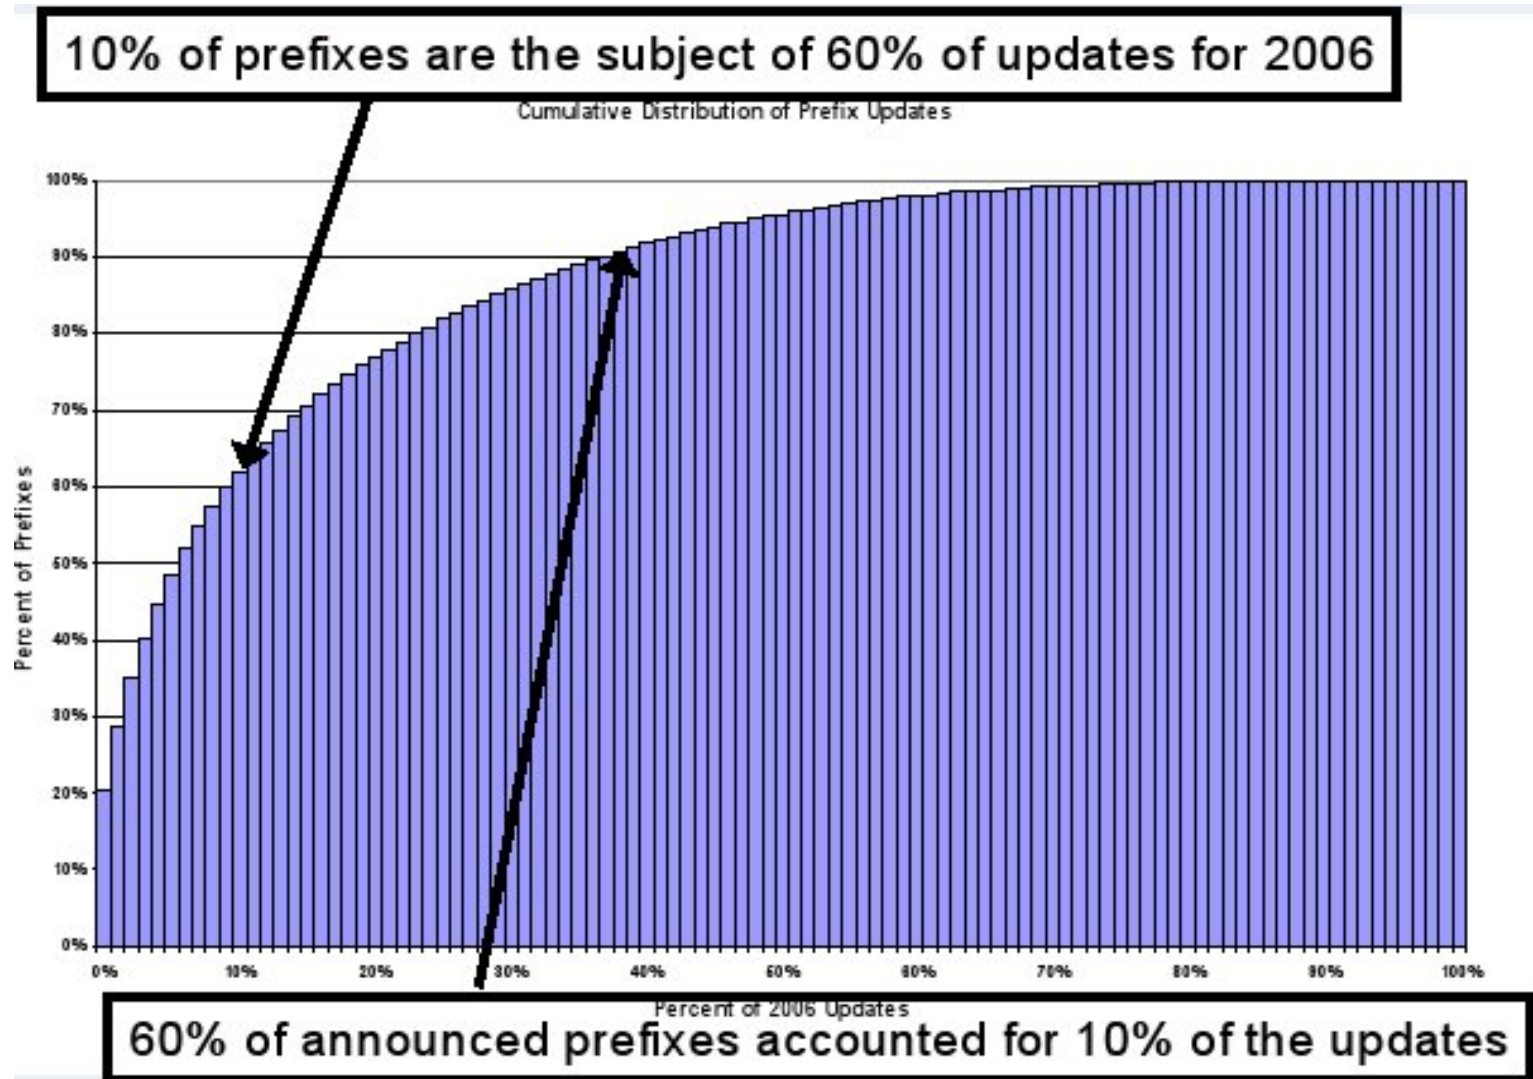


→ Source: Geoff Huston at IAB RAWS Workshop, 2006

Measurements – Update Details

- Localized both in time and source
- Main reasons why the sources are updating:
 - Traffic engineering
 - Unstable configurations
 - New routing applications
 - Address space theft
- #1 updater is an ISP with a noisy link
- Another big updater is Boeing Connexion (#13 and #20)

Distribution of Updates



Updates – Main Observation

- Most of the net is very stable
 - Parts of the net are very unstable
 - Everybody pays for the instability
 - Problem is getting worse
-

Architectural Issue

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

How Serious Is This?

- There is a concern that the growth exceeds what Moore's law can provide for routers
- Power usage, cost, functionality, investment lifecycle will suffer if this is the case
- However, there is reason for optimism
- Forwarding is a small factor in router power usage (power = line speed X per-packet actions)
- Latest generation routers use parallel processing, new memories, new generation silicon, etc
 - routing table capacity expected to last beyond useful operational lifetimes of the devices


How Serious Is This? (Cont'd)

- Punchline: The routing system is not about to fall over; lots of runway
 - But the long-term trend is wrong; if we want to change this, the time to start working is now
 - And we need an Internet which can scale to hundreds of billions of end hosts and (at least) millions of multihomed, provider independent networks
 - Also, need to be very cautious about adding new per-packet functionality to routers
-

What Can We Do?

- Faster routers – Engineering by microelectronics and router designers
- Update dynamics - BGP adjustments, better operational practices
- Bigger network, traffic engineering, multihoming, e2e transparency, and mobility would benefit from architectural changes

- Identifier/locator separation and/or multilevel locators form a hopeful approach



Some of the Hard Parts...

- Incentives for deployment must be in the right places
 - Costs and benefits in the same organization
 - Incrementally deployable
 - Avoid application changes, losing connectivity to others, ...
 - Securing a mapping
-

Activities on the Routing Problem

- GROW WG – Measurements, operational advice
 - IDR WG – Modest BGP improvements
 - RRG RG – New routing architectures (research)

 - Vendors, silicon makers – Faster, better implementations
-

Summary

- A lot of interesting problems in the Internet...
- The routing and addressing problem is being looked at actively at the IETF -- participate!
- Many of the problems are hard – no easy solutions, have to make tradeoffs
- Need to focus on motives of the community to take new technology into use



Reading Material

<http://www.ietf.org/internet-drafts/draft-iab-raws-report-02.txt>

http://www.ripe.net/ripe/meetings/ripe-54/presentations/Router_Scaling_Trends.pdf

http://www.ripe.net/ripe/meetings/ripe-54/presentations/Moderate_BGP.pdf

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

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