Routing Items from IAB Utrecht Workshop

Geoff Huston
IAB

Address Space vs Routing Space

- One view of the rate of consumption of IPv4 prefixes from the unallocated IPv4 space appears to be declining
  - Although there are a large number of caveats surrounding this analysis of RIR allocation data
  - Note that there are a number of alternate interpretations of the allocation data that are not necessarily bounded at any level less than complete consumption

IPv4 Address Space Utilization

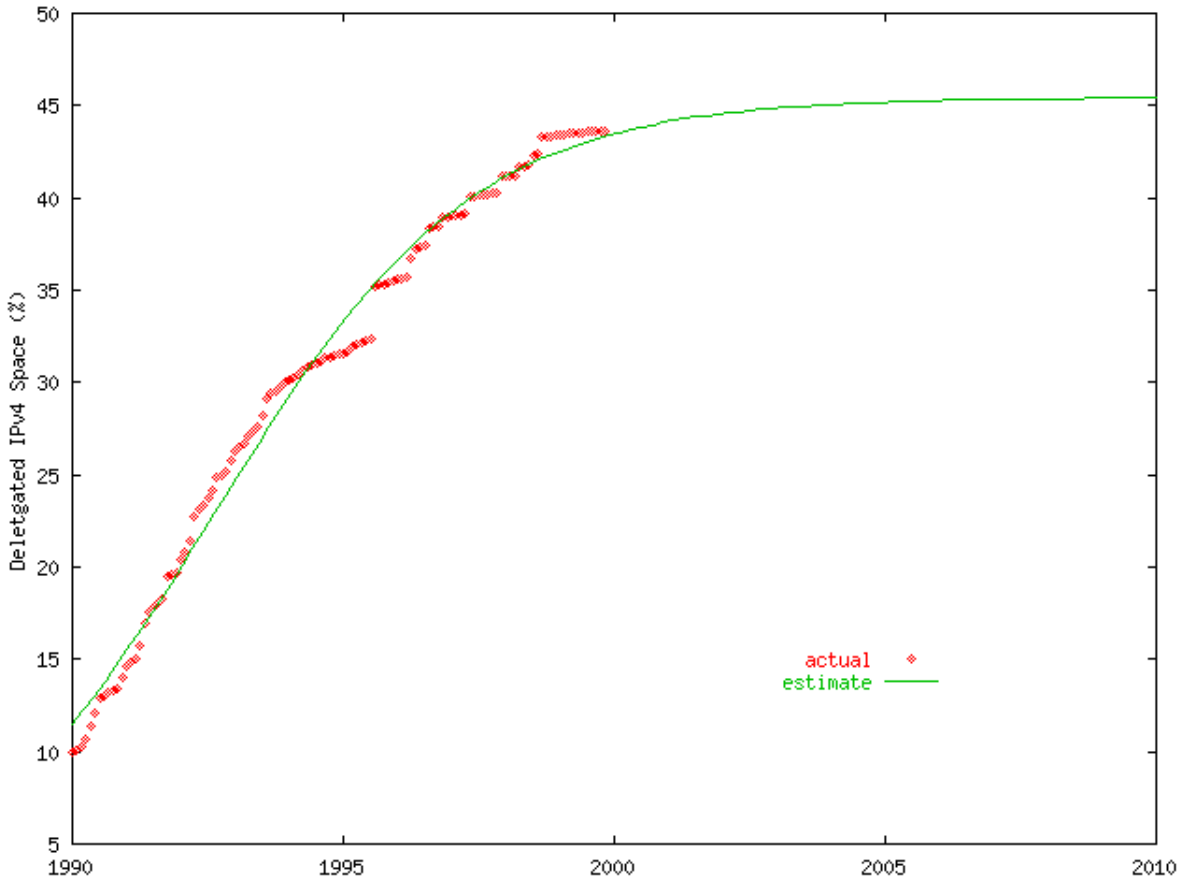

From Frank Solensky - http://ipv4space.toplayer.com/

Address Space vs Routing Space

- The rate of increase of the number of routed objects in the Internet routing tables appears to be growing
  - at a relatively constant linear rate at present

BGP Routing Table Size

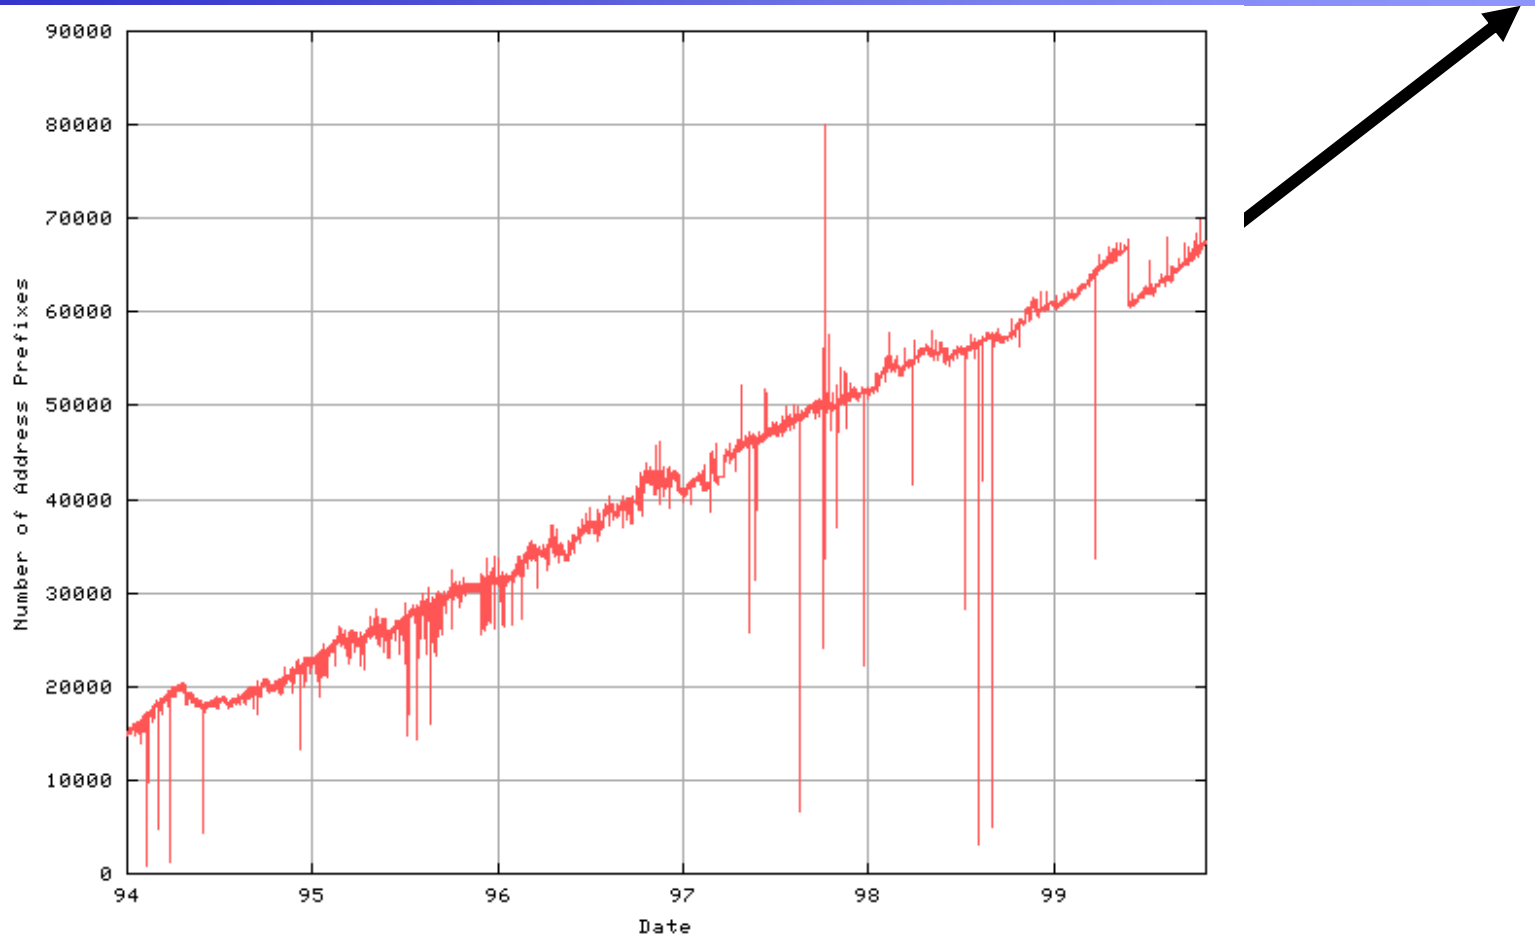

From Geoff Huston - http://www.telstra.net/ops/bgptable.html

Address Space vs Routing Space

- Routing space growth is the outcome of a number of factors, including
  - a denser mesh of interconnectivity extending to the edges of the network
  - a richer set of connection policies at the edge
  - increasing use of multi-homed AS's with distinct routing policies
  - no marginal cost for route advertisements
- Conclusion: likely **continued growth** in the size of the routing tables

Objective

- How will we scale the routing system ?
  - independent of IPv4 / IPv6
  - As V6 offers no fundamental changes to current routing issues, the routing issues are seen to be common to both environments

Routing Space Limit?

- How many IPv4 prefixes can be supported with current technology?
  - Router memory limitation? No
  - Forwarding Table lookup time? No
  - Routing Algorithm Convergence? Yes?

Routing Space Limit?

- There was speculation of a fuzzy limit of some 100,000 prefixes with current systems
  - At current linear growth rates this point may be reached in 2002
- There was also speculation that this number may be 250,000 prefixes with current systems
  - At current linear growth rates this point may be reached in 2012
- This is a very unclear area of speculation at present

Routing Convergence

- Convergence depends on many factors, including
  - number of routers that must converge to a given route
    - with route aggregation not all routes must converge everywhere
  - the number of routed objects
  - the richness of the topology mesh
  - rate of processing routing updates by a router
  - ability to dampen route flaps

Identified Routing Issues

- Computational load to calculate routing tables
- Time to distribute routing table updates across the entire network
- Routing Protocol robustness
- Vulnerabilities in terms of authenticity of injected routing information

Identified Routing Issues

- The only known scheme to produce scaleable routing systems is **topologically aggregated addressing**
- BUT such a scheme:
  - Requires renumbering to support topology changes, which is acknowledged as operationally difficult and expensive
  - Imposes a connectivity policy of simplified topologies with single-homing, which is acknowledged as counter to the current widespread connectivity practice of multi-homing

Identified Routing Issues

- Convergence time for routing may increase
  - current convergence times are estimated to be of the order of 30 seconds
    - according to one set of observations, although there are other observations of different behavior of the routing system
  - with larger BGP tables and a larger interconnection mesh, the time to reach full convergence is likely to increase
- **Further research required on routing mechanisms to alleviate routing table entropy and speed up routing convergence**

Identified Routing Issues

- What would be the impact of widespread use of RSIP models on the routing system?
- If there was no distinguished global address space, but multiple address realms, can routing be made to work coherently?

- No ideas as to how to make such a system work in a robust fashion

Identified Routing Issues

- How will policy-qualified Route Objects impact the route space?
  - Current routing convergence is based on convergence to single ‘best path’.
  - If the path objects are qualified by service quality attributes, for example, will this be supportable within the current routing mechanisms?
    - Attribute tagging and selection mechanisms are feasible
    - impact on route table size could be very dramatic if adopted for Inter-Provider QoS routing

Identified Routing Issues

- IPv4 tightly associates node identity and the routing path to the node
  - **desireable to research mechanisms that delineate node identity from node routing path (or ‘node address’)**

Routing and NAT

"I'd like to believe that native IP routing will almost always be the path of least resistance, but with increasing penetration of NATs I have a difficult time convincing myself of that."

Keith Moore

Routing and Scale

"In any case, in the long run, the only thing that *will* keep the routing running as the network gets larger is topologically aggregated addresses.

We can change routing architectures (and algorithms) until we are blue in the face, and it won't change that.

There is **no** magic bullet."

Noel Chiappa

Routing Recommendations

"In the absence of a new addressing model to replace topological aggregation, and of a clear and substantial demand from the user community for a new routing architecture there is no reason to start work on standards for a "next generation" routing system in the IETF.

We recommend that work should continue in the IRTF Routing Research Group."

Draft Workshop Report