

Internet Research Task Force  
Internet-Draft  
Intended status: Informational  
Expires: January 9, 2008

T. Li, Ed.  
Cisco Systems, Inc.  
July 8, 2007

**Design Goals for Scalable Internet Routing**  
**draft-irtf-rrg-design-goals-01**

Status of this Memo

By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she becomes aware will be disclosed, in accordance with [Section 6 of BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

This Internet-Draft will expire on January 9, 2008.

Copyright Notice

Copyright (C) The IETF Trust (2007).

Abstract

It is commonly recognized that the Internet routing and addressing architecture is facing challenges in scalability, mobility, multi-homing, and inter-domain traffic engineering. The RRG is designing an alternate architecture to meet these challenges. This document consists of a prioritized list of design goals for the architecture.

## Table of Contents

<a href="#">1.</a>	<a href="#">Introduction . . . . .</a>	<a href="#">3</a>
<a href="#">1.1.</a>	<a href="#">Requirements Language . . . . .</a>	<a href="#">3</a>
<a href="#">1.2.</a>	<a href="#">Priorities . . . . .</a>	<a href="#">3</a>
<a href="#">2.</a>	<a href="#">General Design Goals Collected from Past . . . . .</a>	<a href="#">3</a>
<a href="#">3.</a>	<a href="#">Design Goals for A New Routing Architecture . . . . .</a>	<a href="#">4</a>
<a href="#">3.1.</a>	<a href="#">Improved routing scalability . . . . .</a>	<a href="#">4</a>
<a href="#">3.2.</a>	<a href="#">Scalable support for traffic engineering . . . . .</a>	<a href="#">4</a>
<a href="#">3.3.</a>	<a href="#">Scalable support for multi-homing . . . . .</a>	<a href="#">4</a>
<a href="#">3.4.</a>	<a href="#">Scalable support for mobility . . . . .</a>	<a href="#">4</a>
<a href="#">3.5.</a>	<a href="#">Simplified renumbering . . . . .</a>	<a href="#">5</a>
<a href="#">3.6.</a>	<a href="#">Decoupling location and identification . . . . .</a>	<a href="#">5</a>
<a href="#">3.7.</a>	<a href="#">First-class elements . . . . .</a>	<a href="#">6</a>
<a href="#">3.8.</a>	<a href="#">Routing quality . . . . .</a>	<a href="#">6</a>
<a href="#">3.9.</a>	<a href="#">Routing security . . . . .</a>	<a href="#">6</a>
<a href="#">3.10.</a>	<a href="#">Deployability . . . . .</a>	<a href="#">6</a>
<a href="#">4.</a>	<a href="#">IANA Considerations . . . . .</a>	<a href="#">7</a>
<a href="#">5.</a>	<a href="#">Security Considerations . . . . .</a>	<a href="#">7</a>
<a href="#">6.</a>	<a href="#">References . . . . .</a>	<a href="#">7</a>
<a href="#">6.1.</a>	<a href="#">Normative References . . . . .</a>	<a href="#">7</a>
<a href="#">6.2.</a>	<a href="#">Informative References . . . . .</a>	<a href="#">7</a>
	<a href="#">Author's Address . . . . .</a>	<a href="#">7</a>
	<a href="#">Intellectual Property and Copyright Statements . . . . .</a>	<a href="#">9</a>



## **1. Introduction**

It is commonly recognized that the Internet routing and addressing architecture is facing challenges in scalability, mobility, multi-homing, and inter-domain traffic engineering. The Routing Research Group aims to design an alternate architecture to meet these challenges. This document presents a prioritized list of design goals for the architecture.

These goals should be taken as guidelines for evaluating possible architectural solutions. The expectation is that these goals will be applied with good judgment.

### **1.1. Requirements Language**

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC 2119](#) [[RFC2119](#)].

### **1.2. Priorities**

Each design goal in this document has been assigned a priority, which is one of 'required', 'strongly desired', 'desired', and 'optional'.

Required: The solution is REQUIRED to support this goal.

Strongly desired: The solution SHOULD support this goal unless there exist compelling reasons showing it is unachievable, extremely inefficient, or impractical.

Desired: The solution SHOULD support this goal.

Optional: The solution MAY support this goal.

It is possible that two design goals at the same priority level may be found to be in conflict with one another. If and when this happens, one of them may be subsequently re-prioritized to have the two in different priority levels.

## **2. General Design Goals Collected from Past**

[RFC 1958](#) [[RFC1958](#)] provides an excellent list of the original architectural principles of the Internet. We incorporate them here by reference, as part of our desired design goals.



### **3. Design Goals for A New Routing Architecture**

#### **3.1. Improved routing scalability**

Long experience with inter-domain routing has shown us that the global BGP routing table is continuing to grow rapidly [[BGPGrowth](#)]. Carrying this large amount of state in the control plane is expensive and places undue cost burdens on network participants that do not necessarily get value from the increases in the routing table size. Thus, the first required goal is to provide significant improvement to the scalability of the control plane. It is strongly desired to make the control plane scale independently from the growth of the Internet user population.

#### **3.2. Scalable support for traffic engineering**

Traffic engineering is the capability of directing traffic along paths other than those that would be computed by normal IGP/EGP routing. Inter-domain traffic engineering today is frequently accomplished by injecting more-specific prefixes into the global routing table, which results in a negative impact on routing scalability. A scalable solution for inter-domain traffic engineering is strongly desired.

#### **3.3. Scalable support for multi-homing**

Multi-homing is the capability of an organization to be connected to the Internet via more than one other organization. The current mechanism for supporting multi-homing is to let the organization advertise one or multiple prefixes into the global routing system, again resulting in a negative impact on routing scalability. More scalable solutions for multi-homing are strongly desired.

#### **3.4. Scalable support for mobility**

Mobility is the capability of a host, network, or organization to change its topological connectivity with respect to the remainder of the Internet, while continuing to receive packets from the Internet. Existing mechanisms to provide mobility support include (1) renumbering the mobile entity as it changes its topological attachment point(s) to the Internet; (2) renumbering and creating a tunnel from the entity's new topological location back to its original location; and (3) letting the mobile entity announce its prefixes from its new attachment point(s). The first approach alone is considered unsatisfactory as the change of IP address may break existing transport or higher level connections for those protocols using IP addresses as identifiers. The second requires the deployment of a 'home agent' to keep track of the mobile entity's

Li

Expires January 9, 2008

[Page 4]

current location and adds overhead to the routers involved, as well as adding stretch to the path of inbound packet. Neither of the first two approaches impacts the routing scalability. The third approach, however, injects dynamic updates into the global routing system as the mobile entity moves. Mechanisms that help to provide more efficient and scalable mobility support are desired, especially when they can be coupled with security and support topological changes at a high-rate.

### **3.5. Simplified renumbering**

Today many of the end-sites receive their IP address assignments from their Internet Service Providers (ISP). When such a site changes providers, for routing to scale, the site must renumber into a new address block assigned by its new ISP. This can be costly, error-prone and painful. Automated tools, once developed, are expected to provide significant help in reducing the renumbering pain. It is not expected that renumbering will be wholly automated, as some manual reconfiguration is likely to be necessary for changing the last-mile link. However, the overall cost of this transition should be drastically lowered.

In addition to being configured into hosts and routers, where automated renumbering tools can help, IP addresses are also often used for other purposes such as access control lists. They are also sometimes hard-coded into applications used in environments where failure of the DNS could be catastrophic (e.g. certain remote monitoring applications). Although renumbering may be considered a mild inconvenience for some sites, and guidelines have been developed for renumbering a network without a flag day [[RFC4192](#)], for others, the necessary changes are sufficiently difficult so as to make renumbering effectively impossible. It is strongly desired that a new architecture allow end-sites to change providers with significantly less disruption.

### **3.6. Decoupling location and identification**

Numerous sources have noted that an IPv4 address embodies both host attachment point information and identification information. This overloading has caused numerous semantic collisions that have limited the flexibility of the Internet architecture. Therefore, it is desired that a solution separate the host location information namespace from the identification namespace.

Caution must be taken here to clearly distinguish the decoupling of host location and identification information, and the decoupling of end-site addresses from globally routable prefixes; the latter has been proposed as one of the approaches to a scalable routing



Li

Expires January 9, 2008

[Page 5]

architecture. Solutions to both problems, i.e. (1) the decoupling of host location and identification information and (2) a scalable global routing system (whose solution may, or may not, depend on the second decoupling) are required and it is required that their solutions are compatible with each other.

### **3.7. First-class elements**

If a solution makes use of a mechanism, it is strongly desired that the mechanism be a first-class element in the architecture [[FirstClass](#)]. More specifically, if a tunneling mechanism is used to provide network layering, connectivity virtualization, or a connection-oriented mode, then it is strongly desired that the tunneling mechanism be a first-class element in the architecture.

### **3.8. Routing quality**

The routing subsystem is responsible for computing a path from any point on the Internet to any other point in the Internet. The quality of the routes that are computed can be measured by a number of metrics, such as convergence, stability, and stretch.

The stretch of a routing scheme is the ratio of the maximum length of the routing path, on which a packet is delivered, to the length of the shortest path from the source to the destination node, over all source destination pairs.

[Editor's Note: A better definition of stretch, or a better reference would be much appreciated. This definition is derived from [[PODC06](#)].]

A solution is strongly desired to provide routing quality equivalent to what is available today or better.

### **3.9. Routing security**

Currently, the routing subsystem is secured through a number of protocol specific mechanisms of varying strength and applicability. Any new architecture is required to provide at least the same level of security as is deployed as of when the new architecture is deployed.

### **3.10. Deployability**

Since solutions that are not deployable are simply academic exercises, solutions are required to be deployable from a technical perspective. Furthermore, given the extensive deployed base of today's Internet, a solution is required to be incrementally

Li

Expires January 9, 2008

[Page 6]

deployable.

#### **4. IANA Considerations**

This memo includes no requests to IANA.

#### **5. Security Considerations**

All solutions are required to provide security that is at least as strong as the existing Internet routing and addressing architecture.

#### **6. References**

##### **6.1. Normative References**

- [RFC1958] Carpenter, B., "Architectural Principles of the Internet", [RFC 1958](#), June 1996.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC4192] Baker, F., Lear, E., and R. Droms, "Procedures for Renumbering an IPv6 Network without a Flag Day", [RFC 4192](#), September 2005.

##### **6.2. Informative References**

- [BGPGrowth] Huston, G., "BGP Routing Table Analysis Reports", <<http://bgp.potaroo.net/>>.
- [FirstClass] "First-class object", <[http://en.wikipedia.org/wiki/First-class\\_object](http://en.wikipedia.org/wiki/First-class_object)>.
- [PODC06] Konjevod, G., Andrea, R., and D. Xia, "Optimal-Stretch Name-Independent Compact Routing in Doubling Metrics", <<http://www.public.asu.edu/~dxia2/papers/PODC06.pdf>>.



Author's Address

Tony Li (editor)  
Cisco Systems, Inc.  
425 East Tasman Dr.  
San Jose, CA 95134  
USA

Phone: +1 408 853 1494  
Email: [tli@cisco.com](mailto:tli@cisco.com)

## Full Copyright Statement

Copyright (C) The IETF Trust (2007).

This document is subject to the rights, licenses and restrictions contained in [BCP 78](#), and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY, THE IETF TRUST AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

## Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in [BCP 78](#) and [BCP 79](#).

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at [ietf-ipr@ietf.org](mailto:ietf-ipr@ietf.org).

## Acknowledgment

Funding for the RFC Editor function is provided by the IETF Administrative Support Activity (IASA).

