

Network Working Group
Request for Comments: 2902
Category: Informational

S. Deering
Cisco Systems
S. Hares
Merit Networks
C. Perkins
Nokia Research Center
R. Perlman
Sun Microsystems Laboratories
August 2000

Overview of the 1998 IAB Routing Workshop

Status of this Memo

This memo provides information for the Internet community. It does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (2000). All Rights Reserved.

Abstract

This document is an overview of a Routing workshop held by the Internet Architecture Board (IAB) during March 25-27, 1998. The major points of discussion are listed, along with some conclusions and action items for many of the points of discussion.

Table of Contents

1. Introduction	2
2. Conclusions and Action Items	3
2.1. Scaling of Unicast Routing and Addressing	3
2.1.1. Unicast Routing - Conclusions	3
2.1.2. Unicast Routing - Action Items	4
2.2. Levels of Addressing of Addressing and Routing	4
2.3. Network Address Translation (NAT) devices	5
2.3.1. NAT devices - Conclusions	5
2.3.2. NAT devices - Action Items	5
2.4. Multicast	5
2.4.1. Multicast - Conclusions	5
2.4.2. Multicast - Action Items	6
2.5. Routing Stability	6
2.5.1. Routing Stability - Conclusions	6
2.5.2. Routing Stability - Action Items	7
2.6. ToS/CoS/QoS	7

2.6.1. ToS/CoS/QoS - Action Items	8
2.7. Routing Protocol Security	8
2.7.1. Routing Security - Conclusions	8
2.7.2. Routing Security - Action Items	8
2.8. Routing Policy	8
2.8.1. Routing Policy - Conclusions	8
2.8.2. Routing Policy - Action Item	9
2.9. Network to Host Flow of Information	9
2.9.1. Host Information - Conclusions	9
2.9.2. Host Information - Action Items	9
2.10. Shorter Topics	9
2.10.1. Multi-strand Trunking	9
2.10.2. Routing Diagnostic and Development Tools	10
2.10.3. Anycast	10
2.10.4. Load Sensitive IGP routing for Best Effort Traffic	11
2.10.5. Geographical Addresses and Renumbering	11
3. Summary of Action items	11
3.1. Action Items for the IAB	11
3.2. Action Items for IETF Working Group Chairs	11
3.3. Action Items for the IRTF Routing Research Group	12
4. Security Considerations	12
A. Participants	12
References	13
Authors' Addresses	15
Full Copyright Statement	16

1. Introduction

March 25 to March 27, 1998 the Internet Architecture Board (IAB) held a workshop on Routing. The workshop focused on current problems within the Internet and the long term solutions that should be addressed. This document summarizes the discussions the group had on routing, and lists the conclusions reached by the workshop. Section 2 lists the conclusions reached by the participants of the workshop and the suggestions for additional work or redirection of current work. Sections 2.1-2.10 attempt to extract the major points of what was, in actuality, many multifaceted discussions, sometimes occurring all at the same time. Appendix A contains a list of the participants who attended the workshop. The full body of the report can be found at <http://www.iab.org>.

The topics covered at length during the IAB workshop were:

1. Scaling of Unicast Routing and Addressing (section 2.1)
2. Unicast Addressing Issues (Section 2.2)
3. The Effect of extending IP version 4 in the Internet by using Network Address Transformation boxes (Section 2.3)
4. Multicast Routing (Section 2.4)

5. Routing Instability (Section 2.5)
6. Quality of Service Routing (Section 2.6)
7. Routing Security (Section 2.7)
8. BGP Policy (Section 2.8)
9. Flows of information from network routing to hosts for improved services (Section 2.9)

In addition the following topics were briefly covered:

- a. Multi-strand trunking
- b. Better tools for monitoring and diagnosis of network problems
- c. Routing protocol bandwidth minimization
- d. Automatic renumbering and automatic organization
- e. Anycast
- f. Load-sensitive routing
- g. Geographical addressing

These shorter topics are contained in section 2.10.

It would be unrealistic to assume that the workshop had definitive answers to all the technical problems that were raised. The best that can be hoped is that we raised most of the relevant issues and gave opinions that were the best guess of the people at the meeting, keeping in mind that the attendees did not come armed with data to back up opinions. Much of the discussion amounted to an exploration of the intuition of the experts in attendance, intuition gained after years of experience in making the Internet work. More work is needed to validate the intuition and experience by way of scientific experimentation and analysis. Unfortunately, it's not so easy to find a spare collection of global Internets upon which one might perform controlled experiments.

2. Conclusions and Action Items

The participants came to a number of conclusions after the discussions referred to in sections 2.1-2.10. These conclusions, presented in this document, provide summary statements and action items for the IETF community.

2.1. Scaling of Unicast Routing and Addressing

2.1.1. Unicast Routing - Conclusions

The participants of the workshop came to the following conclusions

1. Most of the current unicast routing stability problems can be fixed with improved implementation.

2. Some long term systemic issues that may eventually overwhelm the unicast routing are:
 - Flaps - which will only get worse unless work is undertaken
 - Multi-homing
3. We'd like more research into what's breaking; not just more data, but more analysis of the data

The group reviewed the following potential solutions:

- Architected NAT (improving the existing Network Address Translation schemes to provide better scaling)
- IPv6 (deploying an IP version 6 infrastructure)
- MAP/Encap (map to aggregatable addresses and encapsulate the original packet)
- Do nothing
- Aggressive renumbering (try to continue to encourage renumbering to improve utilization of the IP version 4 address space)
- Metro addressing (use a geographical or metropolitan based addressing scheme)

2.1.2. Unicast Routing - Action Items

We recommend that the IRTF Routing Research group should encourage more analysis of routing data, not just the collection of more data.

2.2. Levels of Addressing of Addressing and Routing

Levels of hierarchy do not matter to the customers. Address hierarchy must be distinguished from routing hierarchy. The group examined whether the current Internet has enough levels of hierarchy in Internet addresses or routing infrastructure. The group did not find that levels of hierarchy should be added to the Internet, at least for now. Flat routing at the AS level seems to be workable; if this changes in the future, hierarchy would need to be revisited, and studied with due consideration to convergence time for routing algorithms and trust management. There is no universal agreement that adding levels of hierarchy at this point in time provides a well-defined benefit. Furthermore, two levels is difficult for many people, and any more than that is difficult both to build and to use.

2.3. Network Address Translation (NAT) devices

2.3.1. NAT devices - Conclusions

Upon reviewing the NATs, the group

1. Noted that NAT devices are fairly widely deployed
2. Identified various problems with the use of NAT devices within the internet
3. Discussed the interaction between NAT devices and applications
4. Listed the following options regarding NAT devices:
 - Eliminate NATs
 - Fix NATs to interact better with the rest of the Internet
 - Fix applications to interact better with NAT boxes
 - Don't do certain things -- like IP Security (IPSec)

2.3.2. NAT devices - Action Items

1. Forward our concerns, problems and suggestions to the appropriate working groups
2. Note architectural work outside the NAT working group
3. Suggest to the IAB that it continue to be concerned about the issues involving NATs

2.4. Multicast

2.4.1. Multicast - Conclusions

Since the multicast model was created, many multicast applications have been tried over the Internet multicast routing fabric. The group began to discuss the multicast model in terms of enabling multicast applications to run efficiently, and scale favorably with future growth. Multicast applications place varying requirements on multicast routing.

Multicast applications may have a variable:

- number of sources,
- number of receivers,
- amount of data,
- amount of data in a burst, and length of quiet periods
- number of groups utilized per application or per set of cooperating applications, and
- amount of time during which the group exists
- topological distance between members of the group.
- volatility of membership

Multicast routing must provide the flexibility to support the varying requirements of different multicast applications. The current multicast model establishes multicast routing paths upon reception of a data packet. The discussion on the viability of the multicast model examined the viability of the model in terms of the uses of multicast routing by applications and the scalability to full Internet usage. For example, providing for many groups of small conferences (a small number of widely-dispersed people) with global topological scope scales badly given the current multicast model.

The group felt the existing multicast protocols and multicast should be evaluated in terms of the requirements listed above. The group suggested that the evaluation should include the multicast protocols DVMRP [12], MOSPF [8], PIM [4], CBT [2], and Express [5], as well as the following mechanisms used by multicast applications:

1. Registering with the core or the RP (Rendezvous Point),
2. Having the ID of the group include the core, and having joins specify the core
3. Having the ID of the group include the core, and having joins and data specify both
4. Sending data via unicast to all members, and
5. Sending data via unicast transport to the RP.

The group acknowledged that the current multicast model does not scale well for all scenarios that applications use.

The group noted that reliable multicast is surprisingly orthogonal to the issues about the scaling of the multicast model to all possible applications.

2.4.2. Multicast - Action Items

Encourage evaluation and written reports on these multicast protocols, and mechanisms for different types of protocols.

Notify the IRTF Routing Research Group of the need to charter activity in this area.

2.5. Routing Stability

2.5.1. Routing Stability - Conclusions

Damping the effects of route updates enhances stability, but possibly at the cost of reachability for some prefixes. A prefix can be damped and reachable via another path, so that for such prefixes the effects of damping are less serious than for other prefixes. The performance of various algorithms for enhancing stability should be

measured by recording whether the affected route prefixes are reachable or not reachable. Using current damping approaches, approximately 1% of the prefixes are affected at any one point in time. We should try to find out how many prefixes are unreachable because of damping.

2.5.2. Routing Stability - Action Items

The conclusion is that this effort merits continued investigation.

The IRTF Routing Research Group should measure how stable things are, and if stability is an issue, to study methods of making them more stable.

2.6. ToS/CoS/QoS

The group noted that the terms Type of Service (ToS), Class of Service (CoS), and Quality of Service (QoS) are imprecise as currently used. The discussion started by defining the terminology as follows:

- ToS: hop by hop routing based on destination plus ToS bits [9]
- CoS: classes of service based on service contracts. These classes of service are enabled by a variety of mechanisms which include queueing, and multiple physical or link level paths.
- QoS: managing routes that meet certain quality of service constraints, and involving the following steps:
 - * routing the resource requests
 - * setting up a path that satisfies the constraints
 - * routing the data

There is no smooth dividing line between between ToS and QoS. ToS is relative. QoS is absolute. The group discussed whether there is a demand for ToS, CoS and QoS. Differentiated-services [3] as discussed in the IETF is ToS++.

The group also discussed a more general concept of "Constraint Based Routing" which was defined as traffic engineering on large aggregated flows. Constraint based routing allows the providers to better utilize the bandwidth in their network to handle traffic requests from users. Besides enabling policy management techniques, constraint based routing allows providers to route traffic based on the characteristics of the traffic flows.

2.6.1. ToS/CoS/QoS - Action Items

We recommend that IETF should look into the issue of Constraint Based Routing.

2.7. Routing Protocol Security

2.7.1. Routing Security - Conclusions

After a lengthy discussion of the various problems of network security, the group notes that:

1. Routers need intrinsic system security as good as or better than any host computer.
2. Improving router security will not solve all problems.
3. Console access to the router can do everything.
4. One compromised router can create disaster.
5. ISPs and vendors should consider taking some control traffic out of band, due to lack of wire speed authentication.
6. We discussed other issues that will be passed on to the appropriate people involved with network security.
7. Identified areas of work to improve things (e.g., wire speed authentication).

2.7.2. Routing Security - Action Items

The IETF should encourage work on "wire speed" authentication, pairwise authentication of routers in routing protocols, and Byzantine robustness [6] in routing protocols.

2.8. Routing Policy

2.8.1. Routing Policy - Conclusions

During our discussion on routing policy the group reviewed what could be done with BGP. The group noted that:

1. Some routing policies requested by ISPs or NSPs are not solvable with BGP. Some of these "unsolvable" routing policies can be put into effect using tunnels and static configuration.
2. BGP is only a mechanism for announcing reachability
3. BGP routing control traffic direction without regard to traffic volume.
4. BGP policy management is too delicate, too easy to mess up, and fragile.

5. Router Configuration Language is very complex and error-prone
6. We can't count on symmetric routing, so ISPs/NSPs/Enterprise nets should deal with it.

The group concluded the Internet needed a better routing policy specification language.

2.8.2. Routing Policy - Action Item

Pass the concerns about the Routing Policy Syntax Language (RPSL) [1] to chairs of the Routing Policy Syntax (RPS) working group [11].

2.9. Network to Host Flow of Information

2.9.1. Host Information - Conclusions

Publishing information about traffic statistics along backbone routes could improve the way Internet services replicate data for retrieval from various sites. This replication could be especially important for the retrieval of information off the web. Currently, web pages refer people to caches local to their sites; for instance, a European site might be used for United Kingdom customers and a North American site for North American customers. Proponents of web caches want to auto-configure the locations of web caches so a user's web browser can automatically discover the local cache. Other applications share this need for finding the best cache for a particular service.

2.9.2. Host Information - Action Items

The group recommends a BOF be held on Measuring Path Characteristics. Measurement of path characteristics should include:

- format for exchange of measurement data
- mechanisms for distribution of measurement data

IPPM working group [7] is dealing with issues within the measurement problem space.

2.10. Shorter Topics

2.10.1. Multi-strand Trunking

PPP did multi-link in a way that required too much computation and could not be used for faster links. Internet technology should treat multiple parallel trunks as 1 link at the IP layer, but with multi-dimensional metrics.

Multi-strand Trunking - Action Items

There is design and development work at layer two which should be done to support the multiple parallel trunks. This layer two work is outside the scope of the IETF. Layer three routing should support richer metrics in OSPF.

2.10.2. Routing Diagnostic and Development Tools

2.10.2.1. Routing Diagnostics - Conclusions

1. It would be nice to have an Authoritative Database listing those prefixes permitted from each AS. The authoritative data base was attempted before without success, but the group felt it might be useful to try again.
2. SNMP version 3 should be deployed in order to make use of its improved authentication, scope and rate limiting
3. Remotely-controlled traffic monitors should be used to measure traffic
4. Better tools are needed for preventative problem detection

2.10.2.2. Routing Diagnostics - Action Items

1. Encouraged an authoritative database within the Internet
2. Notify SNMP version 3 working groups regarding needs for authentication, scope, and rate limiting.
3. Encourage funding of better tools for remotely controlled traffic sources and pro-active problem detection.

2.10.3. Anycast

2.10.3.1. Anycast - Conclusions

1. We need to describe the advantages and disadvantages of anycast.
2. Local-scoped well-known anycast addresses will be useful to applications.

2.10.3.2. Anycast - Action Items

A BOF should be held to plan work on anycast.

If a working group forms, a paper on the advantages and disadvantages of anycast should be included as part of the charter.

2.10.4. Load Sensitive IGP routing for Best Effort Traffic

2.10.4.1. Load Sensitive IGP - Conclusions

While load sensitive routing is interesting in some ways, it cannot be considered until certain problems are worked out. Currently, constraint based routing is assigning administrative metrics to allow routing to adapt to different traffic patterns. Load sensitive routing may increase oscillation and instability of routes. This instability of routes, sometimes called churn, may affect the ability of the routing infrastructure to scale.

Load sensitive routing would allow IGPs to better utilize links. Past and current efforts in load sensitive routing include: QoS OSPF [10], Q-OSPF [10], and load sensitive routers developed by BBN.

2.10.4.2. Load Sensitive IGP - Action items

The IRTF Routing Research group chair and Routing Area Director should discuss this subject and determine what techniques from Load Sensitive IGP routing are ready for IETF, and what requires additional research.

2.10.5. Geographical Addresses and Renumbering

This topic was discussed, but without any conclusions or action items.

3. Summary of Action items

3.1. Action Items for the IAB

1. The IAB should be concerned about the issues involving NATs
2. Authoritative Database (for addresses within domains) should be encouraged within the Internet
3. Encourage funding of better tools for remotely controlled traffic sources and pro-active problem detection.

3.2. Action Items for IETF Working Group Chairs

1. NAT: Forward our concerns, problems and suggestions to the appropriate working groups
2. We recommend that IETF should work the issue of Constraint Based Routing.
3. The IETF should encourage work on "wire speed" authentication, pair-wise authentication of routers in routing protocols, and Byzantine robustness in routing protocols.

4. Concerns about the Routing Policy Specification Language (RPSL) should go to the Routing Policy Systems (RPS) working group chair.
5. The group recommends a BOF be held on Measuring Path Characteristics. The BOF should consider the data exchange format of measurement and mechanisms to distribution of data mechanism. It is noted that the IPPM working group is dealing with issues within the measurement problem space.
6. There is layer two work which should be done to support the multiple parallel trunks which is outside the scope of the IETF. Layer three routing should support richer metrics in OSPF.
7. SNMP version 3 working groups should be notified about the issues about authentication, scope, and rate limiting.
8. A BOF should be held to plan work on anycast. A document on anycast should be part of the proposed working group charter.

3.3. Action Items for the IRTF Routing Research Group

1. We recommend that the IRTF Routing Research working group try to encourage more analysis of routing data, not just the collection of more data.
2. Encourage evaluation and written reports on the evaluation of multicast protocols and mechanisms for different types of protocols
3. The IRTF Routing Research group chair and the Routing Area Director should discuss Load Sensitive IGP routing and determine whether it is ready for the IETF.

4. Security Considerations

Security considerations were an important part of the discussions at the workshop, but the workshop decided not to publish a summary of these discussions. Other documents that address the issues of routing infrastructure security have recently been published.

A. Participants

(Email addresses as of the meeting date.)

Harald Alvestrand	Harald.Alvestrand@maxware.no
Fred Baker	fred@cisco.com
Jeff Burgan	burgan@corp.home.net
Brian Carpenter	brian@hursley.ibm.com
Noel Chiappa	jnc@ginger.lcs.mit.edu
Rob Coltun	rcoltun@fore.com
Steve Deering	deering@cisco.com
Deborah Estrin	estrin@usc.edu

Dino Farinacci	dino@cisco.com
Paul Francis	francis@slab.ntt.co.jp
Elise Gerich	epg@home.net
Joel Halpern	jhalpern@newbridge.com
Sue Hares	skh@merit.edu
Cyndi Jung	cmj@3Com.com
Dave Katz	dkatz@jnx.com
Tony Li	tli@juniper.net
Peter Lothberg	roll@stupi.se
Louis Mamakos	louie@uu.net
Dave Meyer	dmm@cisco.com
Keith Moore	moore@cs.utk.edu
Bob Moskowitz	rgm@htt-consult.com
Thomas Narten	narten@raleigh.ibm.com
Vern Paxson	vern@ee.lbl.gov
Charles E. Perkins	cperkins@eng.sun.com
Radia Perlman	Radia.Perlman@East.Sun.COM
Yakov Rekhter	yakov@cisco.com
Allyn Romanow	allyn@MCI.NET
Martha Steenstrup	msteenst@bbn.com
George Swallow	swallow@cisco.com

References

- [1] Alaettinoglu, C., Bates, T., Gerich, E., Karrenberg, D., Meyer, D., Terpstra, M. and C. Villamizar, "Routing Policy Specification Language (RPSL)", RFC 2280, January 1998.
- [2] Ballardie, A., "Core Based Trees (CBT) Multicast Routing Architecture", RFC 2201, September 1997.
- [3] Blake, S., Black, D., Carlson, M., Davies, E., Wang, Z. and W. Weiss, "An Architecture for Differentiated Service", RFC 2475, December 1998.
- [4] Estrin, D., Farinacci, D., Helmy, A., Thaler, D., Deering, S., Handley, M., Jacobson, V., Liu, C., Sharma, P. and L. Wei, "Protocol Independent Multicast-Sparse Mode (PIM-SM): Protocol Specification", RFC 2362, June 1998.
- [5] Holbrook, H., Cheriton, D, "EXPRESS Multicast", SIGCOMM 99, September 1999.
- [6] Charlie Kaufman, Radia Perlman, and Mike Speciner. Network Security: Private Communication in a Public World, pages 462--465. Prentice-Hall, Inc., 1995.

- [7] W. Leland and M. Zekauskas (chairs). IP Performance Metrics (IPPM), October 1997. <http://www.ietf.org/html.charters/ippm-charter.html>.
- [8] Moy, J., "Multicast Extensions to OSPF", RFC 1584, March 1994.
- [9] Nichols, K., Blake, S., Baker, F. and D. Black, "Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers", RFC 2474, December 1998.
- [10] H. Sandick and E. Crawley (chairs). QoS Routing (qosr), April 1997. <http://www.ietf.org/html.charters/qosr-charter.html>.
- [11] C. Villamizar and C. Alaettinoglu (chairs). Routing Policy Syntax (RPS), July 1995. <http://www.ietf.org/html.charters/rps-charter.html>.
- [12] Waitzman, D., Partridge, C. and S. Deering, "Distance Vector Multicast Routing Protocol", RFC 1075, November 1988.

Authors' Addresses

Questions about this memo can be directed to:

Stephen E. Deering
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA

Phone: +1 408 527-8213
EMail: deering@cisco.com

Susan Hares
Merit, Inc.
1071 Beal Avenue,
Ann Arbor, MI 48109
USA

Phone: +1 313 936-2095
EMail: skh@nexthop.com

Radia Perlman
Sun Microsystems Laboratories
2 Elizabeth Drive
Chelmsford, MA 01824
USA

Phone: +1 978 442-3252
EMail: Radia.Pperlman@sun.com

Charles E. Perkins
Nokia Research Center
313 Fairchild Drive
Mountain View, CA 94043
USA

Phone: +1 650 625-2986
EMail: Charles.Perkins@nokia.com
Fax: +1 650-625-2502

Full Copyright Statement

Copyright (C) The Internet Society (2000). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS 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.

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.

