Network Working Group
Request for Comments: 1380
P. Gross
IESG Chair
P. Almquist
IESG Internet AD
November 1992

IESG Deliberations on Routing and Addressing

Status Of This Memo

Abstract

This memo summarizes issues surrounding the routing and addressing scaling problems in the IP architecture, and it provides a brief background of the ROAD group and related activities in the Internet Engineering Task Force (IETF).

It also provides a preliminary report of the Internet Engineering Steering Group (IESG) deliberations on how these routing and addressing issues should be pursued in the Internet Architecture Board (IAB)/IETF.

Acknowledgements

This note draws principally from two sources: the output from the ROAD group, as reported at the San Diego IETF meeting, and on numerous detailed discussions in the IESG following the San Diego IETF meeting. Zheng Wang, Bob Hinden, Kent England, and Bob Smart provided input for the "Criteria For Bigger Internet Addresses" section below. Greg Vaudreuil prepared the final version of the bibliography, based on previous bibliographies by Lyman Chapin and bibliographies distributed on the Big-Internet mailing list.

Table of Contents

1. INTRODUCTION

It seems unlikely that the designers of IP ever imagined at the time what phenomenal success the Internet would achieve. Internet connections were initially intended primarily for mainframe computers at sites performing DARPA-sponsored research. Now, of course, the Internet has extended its reach to the desktop and is beginning to extend into the home. No longer the exclusive purview of pure R&D establishments, the Internet has become well entrenched in parts of the corporate world and is beginning to make inroads into secondary and even primary schools. While once it was an almost exclusively U.S. phenomenon, the Internet now extends to every continent and within a few years may well include network connections in every country.

Over the past couple of years, we have seen increasingly strong indications that all of this success will stress the limits of IP unless appropriate corrective actions are taken. The supply of unallocated Class B network numbers is rapidly dwindling, and the amount of routing information now carried in the Internet is increasingly taxing the abilities of both the routers and the people who have to manage them. Somewhat longer-term, it is possible that we will run out of host addresses or network numbers altogether.

While these problems could be avoided by attempting to restrict the growth of the Internet, most people would prefer solutions that allow growth to continue. Fortunately, it appears that such solutions are possible, and that, in fact, our biggest problem is having too many possible solutions rather than too few.

This memo provides a preliminary report of IESG deliberations on how routing and addressing issues can be pursued in the IAB/IETF.

In following sections, we will discuss in more detail the problems confronting us and possible approaches. We will give a brief overview of the ROAD group and related activities in the IETF. We will then discuss possible courses of action in the IETF. Ultimately, the IESG will issue a recommendation from the IESG/IETF to the IAB.

2. ISSUES OF GROWTH AND EVOLUTION IN THE INTERNET

2.1 The Problems

The Internet now faces three growth-related problems:

- Class B network number exhaustion - Routing table explosion
- IP address space exhaustion

2.1.1 Class B Network Number Exhaustion

Over the last several years, the number of network numbers assigned and the number of network numbers configured into the Merit NSFnet routing database have roughly doubled every 12 months. This has led to estimates that, at the current allocation rate, and in the absence of corrective measures, it will take less than 2 years to allocate all of the currently unassigned Class B network numbers.

After that, new sites which wished to connect more than the number of hosts possible in a single Class C (253 hosts) would need to be assigned multiple Class C networks. This will exacerbate the routing table explosion problems described next.

2.1.2. Routing Table Explosion

As the number of networks connected to the Internet has grown, the amount of routing information that has to be passed around to keep track of them all is likewise growing. This is leading to two types of problems.

Hardware and Protocol Limits

Routing protocols must pass around this information, and routers must store and use it. This taxes memory limits in the routers, and can also consume significant bandwidth when older routing protocols are used, (such as EGP and RIP, which were designed for a much smaller number of networks).

The limits on the memory in the routers seem to be the most pressing. It is already the case that the routers used in the MILNET are incapable of handling all of the current routes, and most other service providers have found the need to periodically upgrade their routers to accommodate ever larger quantities of routing information. An informal survey of router vendors by the ROAD group estimated that most of the currently deployed generation of high-end routers will support O(16000) routes. This will be probably be adequate for the next 12 to 18 months at the current rate of growth. Most vendors have begun, or will soon begin, to ship routers capable of handling O(64000) routes, which should be adequate for an additional two years if the above Class B Network Number Exhaustion problem is solved.

Human Limits

The number of routes does not merely tax the network's physical plant. Network operators have found that the inter-domain routing protocol mechanisms often need to be augmented by a considerable amount of configuration to make the paths followed by packets be consistent with the policies and desires of the network operators. As the number of networks increases, the configuration (and the traffic monitoring to determine whether the configuration has been done correctly) becomes increasingly difficult and time-consuming.

Although it is not possible to determine a number of networks (and therefore a time frame) where human limits will be exceeded, network operators view this as a significant short-term problem.

2.1.3. IP Address Exhaustion

If the current exponential growth rate continues unabated, the number of computers connected to the Internet will eventually exceed the number of possible IP addresses. Because IP addresses are divided into "network" and "host" portions, we may not ever fully run out of IP addresses because we will run out of IP network numbers first.

There is considerable uncertainty regarding the timeframe when we might exceed the limits of the IP address space. However, the issue is serious enough that it deserves our earliest attention. It is very important that we develop solutions to this potential problem well before we are in danger of actually running out of addresses.

2.1.4. Other Internetwork Layer Issues

Although the catalog of problems above is pretty complete as far as the scaling problems of the Internet are concerned, there are other Internet layer issues that will need to be addressed over the coming years. These are issues regarding advanced functionality and service guarantees in the Internet layer.

In any attempt to resolve the Internet scaling problems, it is important to consider how these other issues might affect the future evolution of Internet layer protocols. These issues include:

1) Policy-based routing
2) Flow control
3) Weak Quality-of-Service (QOS)
4) Service guarantees (strong QOS)
5) Charging

2.2 Possible Solutions

2.2.1. Class B Network Number Exhaustion

A number of approaches have been suggested for how we might slow the exhaustion of the Class B IP addresses. These include:

1) Reclaiming those Class B network numbers which have been assigned but are either unused or used by networks which are not connected to the Internet.

2) Modifying address assignment policies to slow the assignment of Class B network numbers by assigning multiple Class C network numbers to organizations which are only a little bit to big to be accommodated by a Class C network number.

> Note: It is already the case that a Class B number is assigned only if the applicant would need more than "several" Class C networks. The value of "several" has increased over time from 1 to (currently) 32.

3) Use the Class C address space to form aggregations of different size than the normal normal Class C addresses. Such schemes include Classless Inter-Domain Routing (CIDR) [Fuller92] and the C# scheme [Solen92].

2.2.2. Routing Table Explosion

As was described earlier, there are actually two parts to this problem. They each have slightly different possible approaches:

Hardware and Protocol Limits

1) More thrust. We could simply recognize the fact that routers which need full Internet routing information will require large amounts of memory and processing capacity. This is at best a very short-term approach, and we will always need to face this problem in the long-term.

2) Route servers (a variant of the "More Thrust" solution). Instead of requiring every router to store complete routing information, mechanisms could be developed to allow the tasks of computing and storing routes to be offloaded to a server. Routers would request routes from the server as needed (presumably caching to improve performance).

3) Topology engineering. Many network providers already try to design their networks in such a way that only a few of the routers need complete routing information (the others rely on default routes to reach destinations that they don't have explicit routes to). While this is inconvenient for network operators, it is a trend which is likely to continue.

It is also the case that network providers could further reduce the number of routers which need full routing information by accepting some amount of suboptimal routing or reducing alternate paths used for backup.

4) Charging-based solutions. By charging for network number advertisements, it might be possible to discourage sites from connecting more networks to the Internet than they get significant value from having connected.

5) Aggregation of routing information. It is fairly clear that in the long-term it will be necessary for addresses to be more hierarchical. This will allow routes to many networks to be collapsed into a single summary route. Therefore, an important question is whether aggregation can also be part of the short-term solution. Of the proposals to date, only CIDR could provide aggregation in the short-term. All longer-term proposals should aggregation.

Human Limits

1) Additional human resources. Network providers could devote additional manpower to routing management, or accept the consequences of a reduced level of routing management. This is obviously unattractive as anything other than a very short-term solution.

2) Better tools. Network operators and router vendors could work to develop more powerful paradigms and mechanisms for routing management.

The IETF has already undertaken some work in the areas of route filtering and route leaking.

2.2.3. IP Address Exhaustion

The following general approaches have been suggested for dealing with the possible exhaustion of the IP address space:

1) Protocol modifications to provide a larger address space. By enhancing IP or by transitioning to another protocol with a larger address space, we could substantially increase the number of available network numbers and addresses.

2) Addresses which are not globally unique. Several proposed schemes have emerged whereby a host's domain name is globally unique, but its IP address would be unique only within it's local routing domain. These schemes usually involve address translating

3) Partitioned Internet. The Internet could be partitioned into areas, such that a host's IP address would be unique only within its own area. Such schemes generally postulate application gateways to interconnect the areas. This is not unlike the approach often used to connect differing protocol families.

4) Reclaiming network numbers. Network numbers which are not used, or are used by networks which are not connected to the Internet, could conceivably be reclaimed for general Internet use. This isn't a long-term solution, but could possibly help in the interim if for some reason address exhaustion starts to occur unexpectedly soon.

3. PREPARING FOR ACTION

3.1 The IAB Architecture Retreats

In July 1991, the IAB held a special workshop to consider critical issues in the IP architecture (Clark91). Of particular concern were the problems related to Internet growth and scaling. The IAB felt the issues were of sufficient concern to begin organizing a special group to explore the issues and to explore possible solutions. Peter Ford (LANL) was asked to organize this effort. The IAB reconvened the architecture workshop in January 1992 to further examine these critical issues, and to meet jointly with the then-formed ROAD group (see below).

3.2 The Santa Fe IETF

At the November 1991 Santa Fe IETF meeting, the BGP Working Groups independently began a concerted exploration of the issues of routing table scaling. The principal approach was to perform route aggregation by using address masks in BGP to do "supernetting" (rather than "subnetting"). This approach would eventually evolve into CIDR. The BGP WG decided to form a separate subgroup, to be led by Phill Gross (ANS) to pursue this solution.

3.3 The ROAD Group and Beyond

At the Santa Fe IETF, the initially separate IAB and BGP WG activities were combined into a special effort, named the "ROuting and ADdressing (ROAD) Group", to be co-chaired by Ford and Gross.

The group was asked to explore possible near-term approaches for the scaling problems described in the last section, namely:

- Class B address exhaustion
- Routing table explosion
- IP address space exhaustion

The ROAD group was asked to report back to the IETF at the San Diego IETF (March 1992). Suggested guidelines included minimizing changes to hosts, must be incrementally deployable, and must provide support for a billion networks.

The ROAD group was not a traditional open IETF working group. It was always presumed that this was a one-time special group that would lead to the formation of other IETF WGs after its report in San Diego.

The ROAD group held several face-face meetings between the November 1991 (Santa Fe) and March 1992 (San Diego) IETF meetings. This included several times at the Santa Fe IETF meeting, December 1991 in Reston VA, January 1992 in Boston (in conjunction with the IAB architecture workshop), and January 1992 in Arizona). There was also much discussion by electronic mail.

The group produced numerous documents, which have variously been made available as Internet-Drafts or RFCs (see Bibliography in Appendix D).

As follow-up, the ROAD co-chairs reported to the IETF plenary in March 1992 in San Diego. Plus, several specific ROAD-related activities took place during the IETF meeting that week.

The Ford/Gross presentation served as a preliminary report from the ROAD group. The basic thrust was:

1. The Internet community needs a better way to deal with current addresses (e.g., hierarchical address assignments for routing aggregation to help slow Class B exhaustion and routing table

explosion). Classless Inter-Domain Routing (CIDR; also called "supernetting") was recommended. CIDR calls for:

- The development of a plan for hierarchical IP address assignment for aggregation in routing,
- Enhanced "classless" Inter-domain protocols (i.e., carry address masks along with IP addresses),
- Inter-Domain routing "Usage documents" for using addressing and routing plan with the enhanced inter-domain protocols, and for interacting with IGPs.

2. The Internet community needs bigger addresses for the Internet to stem IP address exhaustion. The ROAD group explored several approaches in some depth. Some of these approaches were discussed at the San Diego IETF. However, a final recommendation of a single approach did not emerge.

3. The Internet community needs to focus more effort on future directions for Internet routing and advanced Internet layer features.

Other ROAD-related activities at the San Diego IETF meeting included:

- Monday, 8:00 - 9:00 am, Report from the ROAD group on "Internet Routing and Addressing Considerations",

- Monday, 9:30-12:00pm, Geographical Addressing and Routing (during NOOP WG session),

- Monday, 1:30-3:30pm, Preliminary discussion of a CIDR routing and addressing plan (during ORAD session),

- Tuesday, 1:30-6:00pm, Internet Routing and Addressing BOF (to discuss ROAD results and to explore approaches for bigger Internet address space),

- Wednesday, 1:30-3:30pm, CIDR Supernetting BOF (joint with BGP WG),

- Thursday, 4:00-6:00pm, Summary of ROAD activities in San Diego followed by open plenary discussion.

The slides for the Monday presentation (Ford92), slides for the Thursday summary (and notes in the Chair's message) (Gross92), and notes for the other sessions are contained in the Proceedings of the Twenty-Third IETF (San Diego).

4. SETTING DIRECTIONS FOR THE IETF

4.1 The Need For Interim Solutions

Solutions to the problems of advanced Internet layer functionality are far from being well understood. While we should certainly encourage research in these areas, it is premature to start an engineering effort for an Internet layer which would solve not only the scaling problems but also those other issues.

Plus, most approaches to the problem of IP address space exhaustion involve changes to the Internet layer. Such approaches mean changes changes to host software that will require us to face the very difficult transition of a large installed base.

It is therefore not likely that we can (a) develop a single solution for the near-term scaling problems that will (b) also solve the longer-term problems of advanced Internet-layer functionality, that we can (c) choose, implement and deploy before the nearer-term problems of Class B exhaustion or routing table explosion confront us.

This line of reasoning leads to the inevitable conclusion that we will need to make major enhancements to IP in (at least) two stages.

Therefore, we will consider interim measures to deal with Class B address exhaustion and routing table explosion (together), and to deal with IP address exhaustion (separately).

We will also suggest that the possible relation between these nearer-term measures and work toward advanced Internet layer functionality should be made an important consideration.

4.2 The Proposed Phases

The IESG recommends that we divide the overall course of action into several phases. For lack of a better vocabulary, we will term these "immediate", "short-term", mid-term", and "long-term" phases. But, as the ROAD group pointed out, we should start all the phases immediately. We cannot afford to act on these phases consecutively!

In brief, the phases are:

- "Immediate". These are configuration and engineering actions that can take place immediately without protocol design, development, or deployment. There are a number of actions that can begin immediately. Although none of these will solve any of the problems, they can help slow the onset of the problems.

The IESG specifically endorses:

1) the need for more conservative address assignment policies,
2) alignment of new address assignment policies with any new aggregation schemes,
3) efforts to reclaim unused Class B addresses,
4) installation of more powerful routers by network operators at key points in the Internet, and
5) careful attention to topology engineering.

- "Short-term". Actions in this phase are aimed at dealing with Class B exhaustion and routing table explosion. These problems are deemed to be quite pressing and to need solutions well before the IP address exhaustion problem needs to be or could be solved. In this timeframe, changes to hosts can *not* be considered.

- "Mid-term". In the mid-term, the issue of IP address exhaustion must be solved. This is the most fundamental problem facing the IP architecture. Depending on the expected timeframe, changes to host software could be considered. Note: whatever approach is taken, it must also deal with the routing table explosion. If it does not, then we will simply be forced to deal with that problem again, but in a larger address space.

- "Long-term". Taking a broader view, the IESG feels that advanced Internet layer functionality, like QOS support and resource reservation, will be crucial to the long-term success of the Internet architecture.

Therefore, planning for advanced Internet layer functionality should play a key role in our choice of mid-term solutions.

In particular, we need to keep several things in mind:

1) The long-term solution will require replacement and/or extension of the Internet layer. This will be a significant trauma for vendors, operators, and for users. Therefore, it is particularly important that we either minimize the trauma involved in deploying the sort-and mid-term solutions, or we need to assure that the short- and mid-term solutions will provide a smooth transition path for the long-term solutions.

2) The long-term solution will likely require globally unique endpoint identifiers with an hierarchical structure to aid routing. Any effort to define hierarchy and assignment mechanisms for short- or mid-term solutions would, if done well, probably have long-term usefulness, even if the long-term solution uses

radically different message formats.

3) To some extent, development and deployment of the interim measures will divert resources away from other important projects, including the development of the long-term solution. This diversion should be carefully considered when choosing which interim measures to pursue.

4.3 A Solution For Routing Table Explosion -- CIDR

The IESG accepted ROAD's endorsement of CIDR [Fuller92]. CIDR solves the routing table explosion problem (for the current IP addressing scheme), makes the Class B exhaustion problem less important, and buys time for the crucial address exhaustion problem.

The IESG felt that other alternatives (e.g., C#, see Solen92) did not provide both routing table aggregation and Class B conservation at comparable effort.

CIDR will require policy changes, protocol specification changes, implementation, and deployment of new router software, but it does not call for changes to host software.

The IESG recommends the following course of action to pursue CIDR in the IETF:

a. Adopt the CIDR model described in Fuller92.

b. Develop a plan for "IP Address Assignment Guidelines".

The IESG considered the creation of an addressing plan to be an operational issue. Therefore, the IESG asked Bernhard Stockman (IESG Operational Requirements Area Co-Director) to lead an effort to develop such a plan. Bernhard Stockman is in a position to bring important international input (Stockman92) into this effort because he is a key player in RIPE and EBONE and he is a co-chair of the Intercontinental Engineering Planning Group (IEPG).

A specific proposal [Gerich92] has now emerged. [Gerich92] incorporates the views of the IETF, RIPE, IEPG, and the Federal Engineering Planning group (FEPG).

c. Pursue CIDR extensions to BGP in the BGP WG.

This activity stated at the San Diego IETF meeting. A new BGP specification, BGP4, incorporating the CIDR extensions, is now available for public comment [Rekhter92a].

d. Form a new WG to consider CIDR-related extensions to IDRP (e.g., specify how run IDRP for IP inter-domain routing).

e. Give careful consideration to how CIDR will be deployed in the Internet.

This includes issues such as how to maintain address aggregation across non-CIDR domains and how CIDR and various IGPs will need to interact. Depending on the status of the combined CIDR activities, the IESG may recommend forming a "CIDR Deployment WG", along the same lines as the current BGP Deployment WG.

4.4 Regarding "Bigger Internet Addresses"

In April-May 1992, the IESG reviewed the various approaches emerging from the ROAD group activities -- e.g., "Simple CLNP" [Callon92a], "IP Encaps" [Hinden92], "CNAT" [Callon92b], and "Nimrod" [Chiappa91].

(Note: These were the only proposals under serious consideration in this time period. Other proposals, namely "The P Internet Protocol (PIP)" [Tsuchiya92b] and "The Simple Internet Protocol (SIP)" [Deering92] have since been proposed. Following the San Diego IETF deliberations in March, "Simple CLNP" evolved into "TCP and UDP with Bigger Addresses (TUBA)", and "IP Encaps" evolved into "IP Address Encapsulation (IPAE)" [Hinden92].)

The "Simple CLNP" approach perhaps initially enjoyed more support than other approaches.

However, the consensus view in the IESG was that the full impact of transition to "Simple CLNP" (or to any of the proposed approaches) had not yet been explored in sufficient detail to make a final recommendation possible at this time.

The feeling in the IESG was that such important issues as

- impact on operational infrastructure,
- impact on current protocols (e.g., checksum computation in TCP and UDP under any new IP-level protocol),
- deployment of new routing protocols,
- assignment of new addresses,
- impact on performance,
- ownership of change control
- effect of supporting new protocols, such as for address resolution,
- effect on network management and security, and
- the costs to network operators and network users who must

be trained in the architecture and specifics of any new protocols needed to be explored in more depth before a decision of this magnitude should be made.

At first the question seemed to be one of timing.

At the risk of oversimplifying some very wide ranging discussions, many in the IESG felt that if a decision had to be made *immediately*, then "Simple CLNP" might be their choice. However, they would feel much more comfortable if more detailed information was part of the decision.

The IESG felt there needed to be an open and thorough evaluation of any proposed new routing and addressing architecture. The Internet community must have a thorough understanding of the impact of changing from the current IP architecture to a new one. The community needs to be confident that we all understand which approach has the most benefits for long-term internet growth and evolution, and the least impact on the current Internet.

The IESG considered what additional information and criteria were needed to choose between alternative approaches. We also considered the time needed for gathering this additional information and the amount of time remaining before it was absolutely imperative to make this decision (i.e., how much time do we have before we are in danger of running out of IP addresses *before* we could deploy a new architecture?).

This led the IESG to propose a specific set of selection criteria and information, and specific milestones and timetable for the decision.

The next section describes the milestones and timetable for choosing the approach for bigger Internet addresses.

The selection criteria referenced in the milestones are contained in Appendix B.

4.5 Milestones And Timetable For Making a Recommendation for "Bigger Internet Addresses"

In June, the IESG recommended that a call for proposals be made, with initial activities to begin at the July IETF in Boston, and with a strict timetable for reaching a recommendation coming out of the November IETF meeting [Gross92a].

Eventually, the call for proposals was made at the July meeting itself.

A working group will be formed for each proposed approach. The charter of each WG will be to explore the criteria described in Appendix B and to develop a detailed plan for IESG consideration.

The WGs will be asked to submit an Internet-Draft prior to the November 1992 IETF, and to make presentations at the November IETF. The IESG and the IETF will review all submitted proposals and then the IESG will make a recommendation to the IAB following the November 1992 IETF meeting.

Therefore, the milestones and timetable for the IESG to reach a recommendation on bigger Internet addresses are:

July 1992 -- Issue a call for proposals at the Boston IETF meeting to form working groups to explore separate approaches for bigger Internet addresses.

August-November 1992 -- Proposed WGs submit charters, create discussion lists, and begin their deliberations by email and/or face- to-face meetings. Redistribute the IESG recommendation (i.e., this memo). Public review, discussion, and modification as appropriate of the "selection criteria" in Appendix B.

October 1992 -- By the end of October, each WG will be required to submit a written description of the approach and how the criteria are satisfied. This is to insure that these documents are distributed as Internet-Drafts for public review well before the November IETF meeting.

November 1992 -- Each WG will be given an opportunity to present its findings in detail at the November 1992 IETF meeting. Based on the written documents, the presentations, and public discussions (by email and at the IETF), the IESG will forward a recommendation to the IAB after the November IETF meeting.

5. SUMMARY

The problems of Internet scaling and address exhaustion are fundamentally important to the continued health of the global Internet, and to the long-term success of such programs as the U.S. NREN and the European EBONE. Finding and embarking on a course of action is critical. However, the problem is so important that we need a deep understanding of the information and criteria described in Appendix B before a decision is made.

With this memo, the IESG re-affirms its earlier recommendation to the IAB that (a) we move CIDR forward in the IETF as described in section 4.3, and (b) that we encourage the exploration of other proposals for a bigger Internet address space according to the timetable in section 4.5.

Appendix A. FOR MORE INFORMATION

To become better acquainted with the issues and/or to follow the progress of these activities:

- Please see the documents in the Bibliography below.

- Join the Big-Internet mailing list where the general issues are discussed (big-internet-request@munnari.oz.au).

- Any new WG formed will have an open mailing list. Please feel free to join each as they are announced on the IETF mailing list. The current lists are:

PIP:	pip-request@thumper.bellcore.com
TUBA:	tuba-request@lanl.gov
IPAE:	ip-encaps-request@sunroof.eng.sun.com
SIP:	sip-request@caldera.usc.edu

- Attend the November IETF in Washington D.C. (where the WGs will report and the IESG recommendation will begin formulating its recommendation to the IAB).

Note: In order to receive announcements of:

- future IETF meetings and agenda,
- new IETF working groups, and
- the posting of Internet-Drafts and RFCs,

please send a request to join the IETF-Announcement mailing list (ietf-announce-request@nri.reston.va.us).

Appendix B. INFORMATION AND SELECTION CRITERIA FOR "BIGGER INTERNET ADDRESSES"

This section describes the information and criteria which the IESG felt that any new routing and addressing proposal should supply. As the community has a chance to comment on these criteria, and as the IESG gets a better understanding of the issues relating to selection of a new routing and addressing architecture, this section may be revised and published in a separate document.

It is expected that every proposal submitted for consideration should address each item below on an point-by-point basis.

B.1 Description of the Proposed Scheme

A complete description of the proposed routing and addressing architecture should be supplied. This should be at the level of detail where the functionality and complexity of the scheme can be clearly understood. It should describe how the proposal solves the basic problems of IP address exhaustion and router resource overload.

B.2 Changes Required

All changes to existing protocols should be documented and new protocols which need to be developed and/or deployed should be specified and described. This should enumerate all protocols which are not currently in widespread operational deployment in the Internet.

Changes should also be grouped by the devices and/or functions they affect. This should include at least the following:

- Protocol changes in hosts
- Protocol changes in exterior router
- Protocol changes in interior router
- Security and Authentication Changes
- Domain name system changes
- Network management changes
- Changes required to operations tools (e.g., ping, trace-route, etc.)
- Changes to operational and administration procedures

The changes should also include if hosts and routers have their current IP addresses changed.

The impact and changes to the existing set of TCP/IP protocols should be described. This should include at a minimum:

- IP
- ICMP
- DNS
- ARP/RARP
- TCP
- UDP
- FTP
- RPC
- SNMP

The impact on protocols which use IP addresses as data should be specifically addressed.

B.3 Implementation Experience

A description of implementation experience with the proposal should be supplied. This should include the how much of the proposal was implemented and hard it was to implement. If it was implemented by modifying existing code, the extent of the modifications should be described.

B.4 Large Internet Support

The proposal should describe how it will scale to support a large internet of a billion networks. It should describe how the proposed routing and addressing architecture will work to support an internet of this size. This should include, as appropriate, a description of the routing hierarchy, how the routing and addressing will be organized, how different layers of the routing interact (e.g., interior and exterior, or L1, L2, L3, etc.), and relationship between addressing and routing.

The addressing proposed should include a description of how addresses will be assigned, who owns the addresses (e.g., user or service provider), and whether there are restrictions in address assignment or topology.

B.5 Syntax and semantics of names, identifiers and addresses

Proposals should address the manner in which data sources and sinks are identified and addressed, describe how current domain names and IP addresses would be used/translated/mapped in their scheme, how proposed new identifier and address fields and semantics are used, and should describe the issues involved in administration of these id and address spaces according to their proposal. The deployment plan should address how these new semantics would be introduced and backward compatibility maintained.

Any overlays in the syntax of these protocol structures should be clearly identified and conflicts resulting from syntactic overlay of functionality should be clearly addressed in the discussion of the impact on administrative assignment.

B.6 Performance Impact

The performance impact of the new routing and addressing architecture should be evaluated. It should be compared against the current state of the art with the current IP. The performance evaluation for routers and hosts should include packets-per-second and memory usage projections, and bandwidth usage for networks. Performance should be evaluated for both high speed speed and low speed lines.

Performance for routers (table size and computational load) and network bandwidth consumption should be projected based on the following projected data points:

-Domains	10^3	10^4	10^5	10^6	10^7	10^8
-Networks	10^4	10^5	10^6	10^7	10^8	10^9
-Hosts	10^6	10^7	10^8	10^9	10^10	10^11

B.7 Support for TCP/IP hosts than do not support the new architecture

The proposal should describe how hosts which do not support the new architecture will be supported -- whether they receive full services and can communicate with the whole Internet, or if they will receive limited services. Also, describe if a translation service be provided between old and new hosts? If so, where will be this be done.

B.8 Effect on User Community

The large and growing installed base of IP systems comprises people, as well as software and machines. The proposal should describe changes in understanding and procedures that are used by the people involved in internetworking. This should include new and/or changes in concepts, terminology, and organization.

B.9 Deployment Plan Description

The proposal should include a deployment plan. It should include the steps required to deploy it. Each step should include the devices and protocols which are required to change and what benefits are derived at each step. This should also include at each step if hosts and routers are required to run the current and proposed internet protocol.

A schedule should be included, with justification showing that the schedule is realistic.

B.10 Security Impact

The impact on current and future security plans should be addressed. Specifically do current security mechanisms such as address and protocol port filtering work in the same manner as they do today, and what is the effect on security and authentication schemes currently under development.

B.11 Future Evolution

The proposal should describe how it lays a foundation for solving emerging internet problems such as security/authentication, mobility, resource allocation, accounting, high packet rates, etc.

Appendix C. BIBLIOGRAPHY

-Documents and Information from IETF/IESG:

[Ford92] Ford, P., and P. Gross, "Routing And Addressing Considerations", Proceedings of the Twenty-Third IETF, March 1992.

[Gross92] Gross, P., "Chair's Message and Minutes of the Open IETF Plenary", Proceedings of the Twenty-Third IETF, March 1992.

[Gross92a] Gross, P., "IESG Deliberations on Routing and Addressing", Electronic mail message to the Big-Internet mailing list, June 1992.

-Documents directly resulting from the ROAD group:

[Fuller92] Fuller, V., Li, T., Yu, J., and K. Varadhan, "Supernetting: an Address Assignment and Aggregation Strategy", RFC 1338, BARRNet, cisco, Merit, OARnet, June 1992.

[Hinden92] Hinden, B., "New Scheme for Internet Routing and Addressing (ENCAPS)", Email message to Big-Internet mailing list, March 16, 1992.

[Callon92a] Callon, R., "TCP and UDP with Bigger Addresses (TUBA), A Simple Proposal for Internet Addressing and Routing", RFC 1347, DEC, June 1992

[Deering92] Deering, S., "City Codes: An Alternative Scheme for OSI NSAP Allocation in the Internet", Email message to Big-Internet mailing list, January 7, 1992.

[Callon92b] CNAT

-Related Documents:

[Hinden92b] Hinden, R., and D. Crocker, "A Proposal for IP Address Encapsulation (IPAE): A Compatible version of IP with Large Addresses", Work in Progress, June 1992.

[Deering92b] Deering, S., "The Simple Internet Protocol", Big-Internet mailing list.

[Stockman92] Karrenberg, D., and B. Stockman, "A Proposal for a Global Internet Addressing Scheme", Work in Progress, May 1992.

[Rekhter92] Rekhter, Y., and T. Li, "Guidelines for IP Address Allocation", Work in Progress, May 1992.

[Rekhter92b] Rekhter, Y., and T. Li, "The Border Gateway Protocol (Version 4)", Work in Progress, September 1992.

[Rekhter92c] Rekhter, Y., and P. Gross, "Application of the Border Gateway Protocol", Work in Progress, September 1992.

[Gerich92] Gerich, E., "Guidelines for Management of IP Address Space", RFC 1366, Merit, October 1992.

[Solen92] Solensky, F., and F. Kastenholz, "A Revision to IP Address Classifications", Work in Progress, March 1992.

[Wang92] Wany, Z., and J. Crowcroft, "A Two-Tier Address Structure for the Internet: A Solution to the Problem of Address Space Exhaustion", RFC 1335, University College London, May 1992.

[Callon91] Callon, R., Gardner, E., and R. Colella, "Guidelines for OSI NSAP Allocation in the Internet", RFC 1237, NIST, Mitre, DEC, July 1991.

[Tsuchiya92a] Tsuchiya, P., "The IP Network Address Translator (NAT): Preliminary Design", Work in Progress, April 1991.

[Tsuchiya92b] Tsuchiya, P., "The 'P' Internet Protocol", Work in Progress, May 1992.

[Chiappa91] Chiappa, J., "A New IP Routing and Addressing Architecture", Work in Progress, July 1991.

[Clark91] Clark, D., Chapin, L., Cerf, V., Braden, R., and R. Hobby, "Towards the Future Internet Architecture", RFC 1287, MIT, BBN, CNRI, ISI, UCDavis, December 1991.

Security Considerations

Security issues are discussed in sections 4.4, B.2, B.10, and B.11.

Authors' Addresses

Phillip Gross, IESG Chair
Advanced Network & Services
100 Clearbrook Road
Elmsford, NY

Phone: 914-789-5300
EMail: pgross@ans.net

Philip Almquist
Stanford University
Networking Systems
Pine Hall 147
Stanford, CA 94305

Phone: (415) 723-2229
EMail: Almquist@JESSICA.STANFORD.EDU