Network Working Group
Request for Comments: 1467
Obsoletes: 1367

C. Topolcic
CNRI

August 1993

Status of CIDR Deployment in the Internet

Status of this Memo

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

Abstract

This document describes the current status of the development and deployment of CIDR technology into the Internet. This document replaces RFC 1367, which was a schedule for the deployment of IP address space management procedures to support route aggregation. Since all the milestones proposed in RFC 1367 except for the delivery and installation of CIDR software were met, it does not seem appropriate to issue an updated schedule. Rather, this document is intended to provide information about how this effort is proceeding, which may be of interest to the community.

1. Background

The Internet's exponential growth has led to a number of difficulties relating to the management of IP network numbers. The administrative overhead of allocating ever increasing volumes of IP network numbers for global users has stressed the organizations that perform this function. The volume of IP network numbers that are reachable through the Internet has taxed a number of routers' ability to manage their forwarding tables. The poor utilization of allocated IP network numbers has threatened to deplete the Class A and Class B address space.

During the past few years, a consensus has emerged among the Internet community in favor of a number of mechanisms to relieve these problems for the mid-term. These mechanisms are expected to be put into place in the short term and to provide relief for the mid-term. Fundamental changes to the Internet protocols to ensure the Internet's continued long term growth and well being are being explored and are expected to succeed the mid-term mechanisms.

The global Internet community have been cooperating closely in such forums as the IETF and its working groups, the IEPG, the NSF Regional Techs Meetings, INET, INTEROP, FNC, FEPG, and other assemblies in

order to ensure the continued stable operation of the Internet. Recognizing the need for the mid-term mechanisms and receiving support from the Internet community, the US Federal Agencies proposed procedures to assist the deployment of these mid-term mechanisms. These procedures were originally described in RFC 1366 [1], which was recently made obsolete by RFC 1466 [2]. In October 1992, a schedule was proposed for the implementation of the procedures, described in RFC 1367 [3].

2. Milestones that have been met

Most of the milestones of the proposed schedule were implemented on time. These milestones are shown below, essentially as they appear in [3], but with further comment where appropriate:

1) 31 October 92:

   The following address allocation procedures were continued:

   a) Initial set of criteria for selecting regional address registries were put into place, and requests from prospective regional registries were accepted by the IANA.

      The Reseaux IP Europeens Network Coordination Centre (RIPE NCC) requested to become a regional registry. As per the addressing plan of RFC 1366, the RIPE NCC was given the block 194.0.0.0 to 195.255.255.255 to administer for the European Internet community. The RIPE NCC had previously and independently obtained the block 193.0.0.0 to 193.255.255.255. Although this block had been allocated before RFC 1366, the RIPE NCC was able to manage it according to the guidelines in RFC 1366.

   b) Class A network numbers were put on reserve for possible future use. The unreserved Class A numbers became very difficult to obtain.

   c) Class B network numbers were issued only when reasonably justified. Whenever possible, a block of C's was issued rather than a B. The requirements for allocating a Class B became progressively more constrained until the date in step (3).

d) Class C network numbers were allocated according to the addressing plan of [1], now obsoleted by [2]. Allocation continued to be performed by the Internet Registry (IR) for regions of the world where an appropriate regional registry had not yet been designated by the IANA.

2) 14 February 93:

The schedule in [3] was re-evaluated, and there appeared to be no reason to readjust it, so it was continued as originally set out.

3) 15 April 93:

a) The IR began to allocate all networks according to the addressing plan of [1], now obsoleted by [2], in appropriately sized blocks of Class C numbers.

b) Class B network numbers became difficult to obtain, following the recommendation of the addressing plan and were only issued when justified.

Furthermore, throughout this time period, network service providers have requested blocks of network numbers from the Class C address space for the purpose of further allocating them to their clients. The network service providers were allocated such space by the RIPE NCC or the IR, acting for North America and the Pacific Rim. This process has started to distribute the function of address registration to a more regional level, closer to the end users. The process has operated as hoped for, with no major problems.

3. Milestone that has not been met

The proposed schedule of [3] stated that 6 June 1993 was the date when an address aggregation mechanism would be generally available in the Internet. Although this target date was based on the plans as stated by the router vendors and was reasonable at the time the schedule in [3] was formulated, it has slipped. Nevertheless, the continuation of that schedule has so far not added significantly to the problems of the Internet. The rest of this document looks at the current situation and what can be expected in the near future.

4. Current status of address aggregation mechanisms in commercial routers

Although RFCs 1366, 1466, and 1367 do not depend on any specific address aggregation technology, there is consensus in the Internet community to use Classless Inter-Domain Routing (CIDR) [4]. CIDR is

supported by BGP-4 and IDRP. Most router vendors are working on BGP-4, first, and there is a consensus to use BGP-4 to support the initial deployment of CIDR in the Internet.

The following paragraphs describe the implementation status and plans of software to support CIDR in various router vendors' products, listed in alphabetical order. Some speculation is necessarily involved in deriving these projections. See also the minutes of the July 1993 meeting of the BGP Deployment Working Group of the IETF [5].

3Com's BGP-4 code has been tested internally. They have code that accepts, forwards and manages aggregated routes properly, and they are ready to test it for interoperability with other vendors. They have yet to implement the code that forms the route aggregates. They expect to have Beta code done by September, and full release code shortly thereafter. The initial implementation will not support de-aggregation. Their plans here are not yet formulated. They will support de-aggregation if necessary.

ANS has a BGP-4 implementation that is being tested internally. It is stable enough to begin testing for interoperability with other vendors' implementations. Depending of the results of interoperability testing, this code could be deployed into the ANSNET by August. This delay is primarily because some routers are running older code, and they all need to be upgraded to GATED before they can all support BGP-4 internally. So the ability to support CIDR looks like it is about one to two months away. This code will not support controlled de-aggregation, but de-aggregation will be supported if necessary.

BBN plans to complete it's development of BGP-4 by early Summer 1994. Initial plans are to implement both aggregation and controlled de-aggregation with an early release of the software.

Cisco's BGP-4 implementation is under development at this time. There is pre-Beta code available for people to begin testing. It is expected that the code will be stable sometime during the summer of 1993 and will be made available for limited deployment at that time. This BGP-4 code will implement aggregation. It will not be part of the normal release cycle at this time. It will be available in a special software release based on the 9.21 release. This initial BGP-4 code will not implement controlled de-aggregation, but Cisco plans on implementing de-aggregation.

Proteon's BGP-4 code has been tested internally. They are ready to test it for interoperability with other vendors. If this works out reasonably well, then it is reasonable to expect that they can start

to deploy this as Beta code by August, with a target of full release in the fall. This initial implementation will not support aggregation or de-aggregation. Aggregation will be implemented soon thereafter, but their plans for de-aggregation are not yet formulated. They will implement de-aggregation if necessary.

Wellfleet is aiming at having beta code implementing BGP-4 roughly in early 1994. This code will include controlled de-aggregation.

5. Rate of growth

MERIT periodically publishes the number of networks in the NSFNET/ANSNET policy routing database. Analysis of this data suggests that the number of entries in this database is growing at approximately 8% per month, or doubling every nine or ten months [6].

Although there are currently over 13K networks in the NSFNET/ANSNET policy routing database, a number of them are not active. That is, they are not announced to the NSFNET/ANSNET Backbone. The 10K active network point was passed in late June. Assuming that the number of active networks continues to grow at the same rate as in the past, it can be projected that the 12K active network point will be reached sometime in approximately late September 1993 and that the 25K active network point will be reached sometime in mid-94 (two high water marks whose relevance will become apparent below).

The NSFNET/ANSNET routing database includes only those networks that meet the NSF Acceptable Use Policy (AUP) or the ANSNET CO+RE AUP. There are a number of networks connected to the Internet that do not meet these criteria. Although they are not in the NSFNET/ANSNET routing database, they are in the forwarding tables of a number of network providers. Currently, the number of networks that are connected to other known service providers but are not in the NSFNET/ANSNET routing database is significantly smaller than (less than 25% of) the number that are in the NSFNET/ANSNET database. There is no estimate available for the rate of growth of the number of such non-NSFNET/ANSNET networks. It is assumed here that the growth rate of these networks is approximately the same as that of AUP networks in the NSFNET/ANSNET routing database.

Analysis of the more than 13K networks in the NSFNET/ANSNET routing database, as well as the allocated but unconnected networks, suggests that CIDR deployment should have a significant impact on the number of forwarding table entries that any router needs to maintain, and its rate of growth. However, an in-depth study was begun at the July 1993 meeting of the BGP Deployment Working Group of the IETF [5] to (among other goals) evaluate the impact of CIDR on the growth rate of router forwarding tables.

6. Capacity of deployed networks

The following paragraphs describe the current occupancy of the forwarding tables of the routers of several transit network providers and their expected capacities and an estimate of the time when that capacity would be reached if the growth rate were to continue as today. This list is a subset of all relevant providers, but is considered approximately representative of the situation of other network providers. It is shown in alphabetical order.

ALTERNET nodes are Cisco routers, and currently carry approximately 11K to 12K routes, both AUP and non-AUP. With their current configuration, they have enough memory so that they are expected to support up to approximately 35K routes. If the rate at which the number of these routes is expected to grow is approximately the same as the rate that the NSFNET/ANSNET policy routing database is growing, then this number may be reached in late 1994. However, if the growth rate continues unchecked, it is expected that the processing capacity of the routers will be surpassed before their memory is exhausted. It is expected that CIDR will be in place long before this point is reached.

All ANSNET routers have recently been upgraded to AIX 3.2. This version supports up to 12K networks. These routers currently carry only the active networks in the NSFNET/ANSNET routing database. It is anticipated that the next version of router code will be deployed before September 1993, the projected date for when there will be 12K active networks. This version will support 25K active networks. Although there are no current plans for a version of router code that supports more than 25K networks, it is believed that CIDR will help this situation.

EBONE nodes are Cisco routers. They currently carry approximately 10K to 11K routes. With their current configuration, they may be able to support approximately 40K routes. However, the number of paths may be very relevant. The memory required for the BGP table (rather than the forwarding table) is a function of the number of paths. If a new transatlantic link were to be added, EBONE could receive all the North American routes through it. This would add a new set of paths. Each such transatlantic link would increase the memory required by approximately 20%. Due to the network topology between North America and Europe, new transatlantic links tend to result in new paths, and therefore significant memory requirements. It is very difficult to predict the addition of future transatlantic links because they result from business or political requirements, not bandwidth requirements.

ESNET uses Cisco routers. However, it is already in trouble, but not because of the size of the forwarding tables. The problem is its need to maintain considerable configuration information describing which networks it should or should not accept from its neighbors, and the fact that this information must be stored in a non-volatile memory of limited size. CIDR aggregation is expected to help this problem. Also, ESNET plans to deploy BGP-4 and CIDR only after it is in a full release, so does not plan to participate in the initial BGP-4 deployment. ESNET will upgrade their nodes to Cisco CSC-4's in the meantime.

All SPRINTLINK and ICM nodes have recently been upgraded to Cisco CSC-4 routers with 16MB of memory. They will carry full routing, including not only the routes that the NSFNET/ANSNET carries, but also routes to networks that do not comply with the NSF or CO+RE AUPs. The SPRINT routers currently carry approximately 11K to 12K routes, and it is expected that they will be able to support up to approximately 25K routes, as currently configured. The 25K announced network point may be reached in approximately mid-1994. Again, it is expected that CIDR deployment will have a significant impact on this growth rate, well before this time.

7. Acknowledgements

This report contains information from a number of sources, including vendors, operators, researchers, and organizations that foster cooperation in the Internet community. Specific organizations include the Intercontinental Engineering and Planning Group (IEPG), the BGP-4 Deployment Working Group of the IETF, the Federal Networking Council (FNC), and the FNC Engineering and Planning Group (FEPG). Specific individuals include, in alphabetical order, Arun Arunkumar, Tony Bates, Mary Byrne, Bob Collet, Mike Craren, Dennis Ferguson, Tony Hain, Elise Gerich, Mark Knopper, John Krawczyk, Tony Li, Peter Lothberg, Andrew Partan, Gary Rucinski, Frank Solensky, and Jessica Yu. This report would not have been possible without the willingness of these people to make their information public for the good of the community.

8. References

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

[2] Gerich, E., "Guidelines for Management of IP Address Space", RFC 1466, Merit, May 1993.

[3] Topolcic, C., "Schedule for IP Address Space Management Guidelines", RFC 1367, CNRI, October 1992.

[4] Fuller, V. et al, "Classless Inter-Domain Routing (CIDR): an Address Assignment and Aggregation Strategy", working draft obsoleting RFC 1338, BARRNet, February 1993.

[5] Yu, J., "Minutes of the BGP Deployment Working Group (BGPDEPL)", MERIT, July 1993.

[6] Solensky, F., Internet Growth Charts, "big-internet" mailing list, munnari.oz.au:big-internet/nsf-netnumbers-<yymm>.ps

9. Other relevant documents

Huitema, C., "IAB Recommendation for an Intermediate Strategy to Address the Issue of Scaling", RFC 1481, Internet Architecture Board, July 1993.

Knopper, M., "Minutes of the NSFNET Regional Techs Meeting", working draft, MERIT, June 1993.

Knopper, M., and Richardson, S., " Aggregation Support in the NSFNET Policy-Based Routing Database", RFC 1482, MERIT, June 1993.

Topolcic, C., "Notes of BGP-4/CIDR Coordination Meeting of 11 March 93", working draft, CNRI, March 1993.

Rekhter, Y., and Topolcic, C., "Exchanging Routing Information Across Provider/Subscriber Boundaries in the CIDR Environment", working draft, IBM Corp., CNRI, April 1993.

Rekhter, Y., and Li, T., "An Architecture for IP Address Allocation with CIDR", working draft, IBM Corp., cisco Systems, February 1993.

Gross, P., and P. Almquist, "IESG Deliberations on Routing and Addressing", RFC 1380, IESG, November 1992.

10. Security Considerations

Security issues are not discussed in this memo.

11. Author's Address

Claudio Topolcic
Corporation for National Research Initiatives
895 Preston White Drive, Suite 100
Reston, VA 22091

Phone: (703) 620-8990
EMail: topolcic@CNRI.Reston.VA.US