Rough Cuts - CIDR Deployment

Bill Manning <bmanning@rice.edu>

Abstract

CIDR and its deployment changes the customer provider relationship, forcing closer association and more responsibility and accountability than has previously been thought. Some tools and techniques are summarized which should help people understand and deploy CIDR in a timely fashion.

I. What is CIDR

"CIDR is basically a hack which attempts to provide, with minimal impact on Internet hosts, better scaling of routing... If you read CIDR carefully, you will see that for it to work really well, fairly (tight) constraints must be placed on the topology of the Internet (i.e. most sites are singly-homed and provider mobility is relatively rare)."

--Vince Fuller

II. A brief History

In late 1991, the IAB met to discuss the state of the Internet. The net was showing all the signs of continued, explosive growth. The IAB selected a small group of people to examine issues and report on plans to deal with this growth. This group was constituted as the ROAD group... a team to look ahead and develop a road map for the near-term Internet. A few of the items that the group was attempting to deal with were address space exhaustion and route table explosion. In March of 1992, the concept of "supernetting" to the ROAD group. The ideas were fleshed out over the course of the next few months and became the CIDR rfcs.

III. Authorization

To make CIDR applicable, IANA had to make some policy changes. Among the first were the delegation of large blocks of "C" addresses into continental chunks. These blocks were assigned to regional NICs, like Internic, RIPE, and APNIC. Initially, there were no policies on assignment to NSPs from the NICs. After some pressuring, Network Service Providers (NSPs) did receive smaller blocks from these large, "super" blocks. They ranged in size from 128 "C"s to 2048 "C"s. It was clear that additional policy was needed. These policies were developed with input from the Network Service Provider community on allocation strategies. In the US, the Internic will assign up to 1024 "C" sized blocks to recognized NSPs. RIPE and APNIC have similar assignment policies.

NSPs, through the Intercontinental Engineering and Planning Group (IEGP) and participation in several mailing lists have adopted assignment strategies that made the first order cuts that reduced the number of addresses assigned. When we received our initial block, we used the traditional first-come, first-served (FCFS) methods that were in use from the inception of the Internet. This allowed us to gain firsthand knowledge of how -not- to manage a range of IP space. Even with the faults of FCFS, we gained by allowing us to meet the needs of our customers in a more timely fashion by having numbers immediately available for their use. On the other hand, we were aggravating the route table size.

After participation in these discussions with other NSPs, a rough set of guidelines emerged. Basic rules of thumb are that an NSP will:

- Allocate customer blocks on byte alignments.
- Try and keep blocks below 64 "C" nets.
- Allocate space in contiguous blocks.
- Escrow the top half of the requested space.

Use of these guidelines by NSPs and their customers allows for additional improvements in address conservation and strengthens the provider/consumer ties.

IV. Table Effects

None of these efforts did much to reduce the speed of growth in the routing tables; in fact they aggravated the problems since there were now larger numbers of natural networks being announced into the Internet. The External gateway protocols (EGP) nor the Internal gateway protocols (IGP) had support for network aggregation. There was a concerted effort to instantiate aggregation support in BGP version 4, which is currently the leading EGP. The primary intent was to reduce the speed with which the routing tables were growing.

BGP4 code was released to NSPs in third quarter 1993 from a number of vendors. A logical testbed was established over the Internet to test multivendor interoperability. By March 1994, most of the interoperability bugs had been found and fixed. Coincident with this effort, the major European and US NSPs, along with the main Internet exchange points were upgraded to be BGP4 & CIDR capable. The Internet "core" was declared CIDR capable on 15 April 1994 when it was determined by Havard Eidnes that all major transit networks were still reachable after more specific routes had been withdrawn. We had not kept pace however. The routing tables had grown large enough that a router, in a complex topology was unable to take a full table and still operate. Fortunately, at least two router manufacturers had early release equipment that was able to meet the need. This points out a serious problem in that the Internet is growing faster than hardware can be delivered to meet the growing demand. CIDR can help, but it is not enough to have assigned blocks and capable infrastructure.

V. NSP Response

Why were the tables growing? Although the exchange points and larger NSPs had the ability to send aggregated routes, the regional and national nets were not prepared to inject aggregated routes into the Internet. In part this was because the main focus on CIDR was to get aggregation support into BGP4 and get it deployed. IGPs exist that allow aggregation (OSPF, IISIS, RIPv2, EIGRP) and good old static routes work when all else fails. Unless everyone becomes CIDR aware, and designs their networks to be able to support aggregation, route table growth will force segmentation of the global Internet. To facilitate understanding of CIDR, tools and techniques need to be made available to everyone. These tools need to be able to facilitate hierarchical aggregation, promote densely packed address space, and allow for unexpected growth.

VI. Tools

The initial tool was to take the address blocks as allocated and start picking the closest match to a 256 boundary. This was an improvement over the automatic assignment of a 64 K block when customer needs exceeded a 256 block but was only good as a first pass approximation.

The second iteration was to take to total number of requested nets and inform the requester that although the total number of requested networks were allocated, that only half the requested number for the first year would be released. The others would be held in escrow until the requester was able to show that the first blocks had been assigned and were in use. Once this was done, the next series would be released. This model is a direct outgrowth of the RIPE-104 paper.

Near this time, both Charlie Kline (UIUC) and Dale Johnson (MERIT) released some whatif tools to assist in block size selection. The AGGIS tool from MERIT helps determine block alignment and runs on a *nix platform. The TREE tool provides a bit more information since it works on arbitrary sizes based on numbers of IP addressable elements. TREE works on both *nix and Apple Macintosh platforms. Bjorn Carlson (KTH) has written a couple of small programs that also help customers do whatif analysis and go one step further by creating output that can be used in at least one vendors router configuration process.

After the 29th IETF meeting, Tony Bates and Martin Terpstra of RIPE collected many of the pertinent CIDR documents and tools and have made them available. Try

FTP.RIPE.NET:/CIDR

for these tools, documentation, and current statistics. Although these tools assist in proper sizing of CIDR blocks, there are still several areas that need attention. Havard Eidnes has drafted a couple of excellent papers, one of which was published in the INET'93 proceedings that help the un-initiated NSP in understanding the benefits of CIDR and tools to take advantage of its features.

VII. What Next?

One item that needs attention by NICs and NSPs is proper documentation of assigned authority for delegated resources. The historical repository for this data has been a WHOIS server run by the INTERNIC. This was acceptable as long as the INTERNIC also did all the number assignments. Once others were given address administration authority, along with DNS registration authority, the INTERNIC was no longer able to mandate WHOIS registration in a centralized server. They have put some effort into correcting this problem by the development of a RemoteWhois server and client. This code allows anyone who has been assigned administrative authority over address space to keep their portion of the address space current in a local fashion. In this respect, this code follows the DNS model of delegation. It is my opinion that no CIDR allocation should be assigned unless the requester is willing to run a RWHOIS server for the space delegated.

Other items that need to be addressed are proxy aggregation, dynamic/intermittent connections, and DNS support.

The one item that is near term is proxy aggregation. In the event that a customer or down stream NSP is unable or unwilling to aggregate, it is possible for the upstream NSP to do the aggregation on behalf of the requester. It is required that both NSPs or the NSP and customer have a firm grasp of the actual topology in this event, since a third party NSP or multiple NSPs may have connections to a customer. If one NSP proxies for the customer, this will affect the efficacy of the other connections.

VIII. Captive Audience

What does this do the poor customer? At this time we do not have enough information to make accurate predictions. It does seem clear that a customer needs to add questions on CIDR allocation to the evaluation criteria in provider selection. In addition, customers can utilize these tools to make more efficient use of the address space that is allocated to them. It seems that providers may want to insist on retaining the integrity of their CIDR blocks. Such a strategy may leave a customer faced with the prospect of massive renumbering or being "locked" into a specific provider. The Internet

Engineering Task Force is working on techniques to facilitate automatic renumbering. This effort should mitigate many of these concerns.

The NSP, on the other hand is placed in a situation where some policies may be unenforceable. A case in point is the NSFnet AUP. This has been based on Network Prefix. With CIDR and BGP4, this seems to be no longer enforceable. The level of granularity becomes the CIDR prefix and the AS path.

IX. Summary

CIDR had the potential to lengthen the life expectancy of IPv4 well into the next century. To exploit this potential will require better informed and equipped Network Information Centers and Network Service Providers. It is no longer sufficient to simply provide data pipes. Customers will expect providers to expand the number and types of offerings. Among them may be network design assistance, with attendant understanding of techniques like CIDR. Use of CIDR demands accountability and responsibility be a component of Network Service Providers in this new era of the global internet.

X. Copyright

XI. References

[1] Robert Braden and Jon Postel. Requirements for internet gateways. RFC1009 - June1987 - Being revised

[2] Ross Callon. Use of OSI IS-IS for routing in TCP/IP and dual environments. RFC1195 - December 1990

[3] IAB, Lyman Chapin (chair). Applicability statement for OSPF. RFC1370 - October 1992

[4] Vince Fuller, Tony Li, Jessica Yu, Kannan Varadan. Classless Internet Domain Routing (CIDR): an address assignment and aggregation strategy. RFC1519 - September 1993

[5] Elise Gerich. Guidelines for management of IP address space. RFC1466 - May 1993

[6] Christian Huitema. IAB recommendation for an intermediate strategy to address the issue of scaling. RFC1481 - July 1993

[7] Havard Eidnes. Practical Considerations for Network Addressing using CIDR Block Allocation. Unpublished draft - April 1994

[8] Daniel Karrenberg, Marten Terpstra. European Internet Registry: IP Address Space Assignment Procedures. RIPE-104 - March 1994

[9] Mark Kosters, Scott Williamson. RWHOIS manual pages - March 1994

[10] Havard Eidnes. Practical Considerations for Network Addressing using a CIDR Block Allocation. - Proc. INET'93

[11] Braun, Ford, Rekhter. CIDR and Evolution of IP. - Proc.INET'93

Author Information

Bill Manning earned his degree in the School of Hard Knocks. He is currently employed by Rice University and spends most of his time in the care and feeding of SESQUINET, a regional network providing IP, DECnet, and OSI services in Texas.