

Exchanging Routing Information Across Provider Boundaries
in the CIDR Environment

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.

1. Introduction

Classless Inter-Domain Routing (CIDR) has been adopted as a solution to the scaling problem in the Internet. The overall CIDR architecture is described in [1]. The architecture for IP address assignment with CIDR is covered in [2] and [3]. The inter-domain routing protocols that are capable of supporting CIDR are covered in [4], [5], and [6].

The purpose of this document is twofold. First, it describes various alternatives for exchanging inter-domain routing information across domain boundaries, where one of the peering domain is CIDR-capable and another is not. Second, it addresses the implications of running CIDR-capable inter-domain routing protocols (e.g., BGP-4, IDRP) on intra-domain routing.

This document is not intended to cover all the cases (either real or imaginable). Rather, it focuses on what are viewed to be the most common cases. We expect that individual service providers will use this document as guidelines in establishing their specific operational plans for the transition to CIDR.

The concepts of "network service provider" and "network service subscriber" were introduced in [3]. For the sake of brevity, we will use the term "provider" or "service provider" here to mean either "network service provider" or "network service subscriber", since for the most part, the distinction is not important to this discussion. Furthermore, we use the same terms to refer to the network and to the organization that operates the network. We feel that the context makes it amply clear whether we are talking about hardware or people, and defining different terms would only make this paper harder to read.

This document defines a CIDR-capable provider as the provider that can perform correct IP packet forwarding (both internally and to other adjacent providers) when the inter-domain routing information acquired by the provider is expressed solely in terms of IP address prefixes (with no distinction between A/B/C class of addresses).

This document defines CIDR-capable forwarding as the ability of a router to maintain its forwarding table and to perform correct forwarding of IP packets without making any assumptions about the class of IP addresses.

This document defines CIDR reachability information as reachability information that may violate any assumptions about the class of IP addresses. For instance, a contiguous block of class C networks expressed as a single IP address prefix constitutes CIDR reachability information.

2. Taxonomy of Service Providers

For the purpose of this document we partition all service providers into the following categories, based on the type and volume of inter-domain routing information a provider needs to acquire in order to meet its service requirements:

- Requirements imposed on a service provider preclude it from using Default inter-domain route(s) -- we'll refer to such a provider as a Type 1 provider.
- Requirements imposed on a service provider allow it to rely on using one or more Default routes for inter-domain routing, but this information must be supplemented by requiring the provider to acquire a large percentage of total Internet routing information -- we'll refer to such a provider as a Type 2 provider.
- Requirements imposed on a service provider allow it to rely on using one or more Default routes for inter-domain routing; however, to meet its service requirements the provider must supplement Default route(s) by acquiring a small percentage of total Internet routing information -- we'll refer to such a provider as a Type 3 provider.
- Requirements imposed on a service provider allow it to rely solely on using one or more Default routes for inter-domain routing; no other inter-domain routing information need to be acquired -- we'll refer to such a provider as a Type 4 provider.

3. Assumptions on Deployment of CIDR in the Internet

The document assumes that the CIDR deployment in the Internet will proceed as a three phase process.

In the first phase all the major service providers will become CIDR-capable. Specifically, all the providers that can't rely on using Default route(s) for inter-domain routing (Type 1 providers) are expected to deploy BGP-4 and transition to CIDR during this phase. It is expected that CIDR reachability information will first appear in the Internet upon transition of all Type 1 service providers to CIDR.

The second phase will commence upon completion of the first phase. During the second phase other service providers that are connected to the service providers that were transitioned to CIDR during the first phase will become CIDR-capable. Specifically, during the second phase it is expected that most of the providers that need to acquire a large percentage of the total Internet routing information (Type 2 provider) will become CIDR-capable. In addition, during the second phase some of the Type 3 providers may become CIDR-capable as well. This plan was agreed to by a number of major providers [8]. NSFNET's steps to implement this plan are described in [9].

Finally, during the third phase the rest of the Type 3 providers and most of the Type 4 providers will transition to CIDR.

It is expected that the duration of the first phase will be significantly shorter than duration of the second phase. Likewise, the duration of the second phase is expected to be shorter than the duration of the third phase.

This document addresses the need for service providers to exchange inter-domain routing information during the second and third phases of this deployment. During these phases, some providers will be CIDR-capable, and others will not. Hence this document considers routing exchanges where one of the peers is CIDR-capable and the other is CIDR-incapable.

4. Implications of CIDR on Interior Routing

A CIDR-capable service provider can use the following two techniques to distribute exterior routing information to all of its routers (both interior and border):

- utilize internal BGP/IDRP between all the routers
- use CIDR-capable IGPs (e.g., OSPF, IS-IS, RIP2)

The first technique doesn't impose any additional requirements on the IGP within the provider. Additional information on implementing the first option is presented in [5] (see Section A.2.4).

The second technique allows the provider to reduce the utilization of internal BGP/IDRP, but imposes specific requirements on the intra-domain routing. It also requires the ability to inject inter-domain routing information (acquired either via BGP or IDRP) into the intra-domain routing. Additional details on implementing the second option are provided in [7]. It is not expected that all the features enumerated in [7] will be widely needed. Therefore, it would be highly desirable to prioritize the features.

Note that both of these techniques imply that all the routers within a CIDR-capable service provider need to be capable of CIDR-based forwarding.

Discussion of which of the two techniques should be preferred is outside the scope of this document.

5. Exchanging Inter-Domain Routing Information

At each phase during the transition to CIDR one of the essential aspects of the Internet operations will be the exchange of inter-domain routing information between CIDR-capable providers and CIDR-incapable provider.

When exchanging inter-domain routing information between a CIDR-capable provider and a CIDR-incapable provider, it is of utmost importance to take into account the view each side wants the other to present. This view has two distinct aspects:

- the type of routing information exchanged (i.e., Default route, traditional (non-CIDR) reachability information, CIDR reachability information)
- routing information processing each side needs to do to maintain these views (e.g., ability to perform aggregation, ability to perform controlled de-aggregation)

The exchange of inter-domain routing information is expected to be controlled by bilateral agreements between the directly connected service providers. Consequently, the views each side wants of the other are expected to form an essential component of such agreements.

To facilitate troubleshooting and problem isolation, the bilateral agreements should be made accessible to other providers. One way to accomplish this is by placing them in a generally accessible

database. The details of how this can be implemented are outside the scope of this document. A possible way to accomplish this is described in [9].

Since the exchange of inter-domain routing information across provider boundaries occurs on a per peer basis, a border router is expected to provide necessary mechanisms (e.g., configuration) that will control exchange and processing of this information on a per peer basis.

In the following sections we describe possible scenarios for exchanging inter-domain routing information. It is always assumed that one side is CIDR-capable and the other is not.

5.1 Exchanging Inter-Domain Routing Information between CIDR-capable providers and CIDR-incapable Type 2 (default with large proportion of explicit routes) providers

We expect the border router(s) within a CIDR-capable provider to be capable of aggregating inter-domain routing information they receive from a CIDR-incapable Type 2 provider. The aggregation is expected to be governed and controlled via a bilateral agreement. Specifically, the CIDR capable provider is expected to aggregate only the information the other side (the CIDR-incapable provider) requests. In other words, the aggregation shall be done by the CIDR-capable provider (the receiver) and only when agreed to by the CIDR-incapable provider (the sender).

Passing inter-domain routing information from a CIDR-capable provider to a CIDR-incapable Type 2 provider will require an agreement between the two that would cover the following items:

- under what conditions the CIDR-capable provider can pass an inter-domain Default route to the CIDR-incapable provider
- exchange of specific non-CIDR reachability information
- controlled de-aggregation of CIDR reachability information

Agreements that cover the first two items are already implemented within the Internet. Thus, the only additional factor introduced by CIDR is controlled de-aggregation. A CIDR-capable provider may decide not to de-aggregate any CIDR reachability information, or to de-aggregate some or all of the CIDR reachability information.

If a CIDR-capable provider does not de-aggregate CIDR reachability information, then its non-CIDR Type 2 peer can obtain reachability information from it either as non-CIDR reachability information

(explicit Class A/B/C network advertisement) or as an inter-domain Default route. Since most of the current reachability information in the Internet is non-CIDR, a Type 2 provider would be able to acquire this information as explicit Class A/B/C network advertisements from the CIDR-capable provider, as it does now. Further, it is expected that at least on a temporary basis (until the completion of the second phase of the transition) in a majority of cases, Type 2 providers should be able to use an inter-domain Default route (acquired from the CIDR-capable provider) as a way of dealing with forwarding to destinations covered by CIDR reachability information.

Thus, it is expected that most of the cases involving a CIDR-capable Type 2 provider and a CIDR-capable provider that does not perform de-aggregation could be addressed by a combination of exchanging specific non-CIDR reachability information and an inter-domain Default route. Any inconvenience to a CIDR-incapable provider due to the use of an inter-domain Default route will be removed once the provider transitions to CIDR.

On the other hand, a CIDR-capable provider may decide to perform controlled de-aggregation of CIDR reachability information. Additional information on performing controlled de-aggregation can be found in [5] (Section 8). Special care must be taken when de-aggregating CIDR reachability information carried by a route with the ATOMIC_AGGREGATE path attribute. It is worth while pointing out that due to the nature of Type 2 provider (it needs to acquire a large percentage of total inter-domain routing information) it is expected that the controlled de-aggregation would result in substantial configuration at the border router that performs the de-aggregation.

5.2 Exchanging Inter-Domain Routing Information between CIDR-capable providers and CIDR-incapable Type 3 (Default with few explicit routes) providers

In this case, as in the case described in Section 5.1, it is expected that a border router in a CIDR-capable provider would be able to aggregate routing information it receives from a CIDR-incapable Type 3 provider. The aggregation is expected to be governed and controlled via a bilateral agreement. Specifically, the CIDR capable provider is expected to aggregate only the information the CIDR-incapable provider requests.

The only difference between this case and the case described in Section 5.1 is the fact that a CIDR-incapable provider requires just a small percentage of total inter-domain routing information. If this information falls into a non-CIDR category, then a Type 3 provider would be able to acquire it from a CIDR-capable provider. If this is CIDR reachability information, then in a majority of cases it is

expected that forwarding to destinations covered by this information could be handled via an inter-domain Default route.

It is still expected that a border router in a CIDR-capable provider would be able to aggregate routing information it receives from a CIDR-incapable Type 3 provider. The aggregation is expected to be governed and controlled via a bilateral agreement. Specifically, the CIDR capable provider is expected to aggregate only the information the other side (the CIDR-incapable provider) requests.

5.3 Exchanging Inter-Domain Routing Information between CIDR-capable providers and CIDR-incapable Type 4 (Default only) providers

Again, it is still expected that a border router in a CIDR-capable provider would be able to aggregate routing information it receives from a CIDR-incapable Type 4 provider. The aggregation is expected to be governed and controlled via a bilateral agreement. Specifically, the CIDR capable provider is expected to aggregate only the information the CIDR-incapable provider requests.

The only difference between this case and the case described in Section 5.1 is the fact that CIDR-incapable provider would not require any inter-domain routing information, other than the Default inter-domain route. Therefore, controlled de-aggregation of CIDR reachability information is not an issue.

6. Conclusions

It is expected that the reduction in the global volume of routing information will begin immediately upon completion of the first phase of the transition to CIDR. The second phase will allow simpler bilateral arrangements between connected service providers by shifting the responsibility for routing information aggregation towards the parties that are better suitable for it, and by significantly reducing the need for routing information de-aggregation. Thus, most of the gain achieved during the second phase will come from simplifying bilateral agreements. The third phase it intended to complete the goals and objectives of the second phase.

7. Acknowledgments

This document was largely stimulated by the discussion that took place during the Merit/NSFNET Regional Tech Meeting in Boulder, January 21-22, 1993. We would like specifically acknowledge contributions by Peter Ford (Los Alamos National Laboratory), Elise Gerich (NSFNET/Merit), Susan Hares (NSFNET/Merit), Mark Knopper (NSFNET/Merit), Bill Manning (Sesquinet/Rice University), and John Scudder (NSFNET/Merit).

8. References

- [1] Fuller, V., Li, T., Yu, J., and K. Varadhan, "Classless Inter-Domain Routing (CIDR): An Address Assignment and Aggregation Strategy", RFC 1519, BARRNet, cisco, Merit, and OARnet, September 1993.
- [2] Gerich, E., "Guidelines for Management of IP Address Space", RFC 1466, Merit, May 1993.
- [3] Rekhter, Y., and T. Li, "An Architecture for IP Address Allocation with CIDR", RFC 1518, T.J. Watson Research Center, IBM Corp., cisco Systems, September 1993.
- [4] Rekhter, Y., and T. Li, "A Border Gateway Protocol 4 (BGP-4)", Work in Progress, June 1993.
- [5] Rekhter, Y., and P. Gross, "Application of the Border Gateway Protocol in the Internet", Work in Progress, September 1992.
- [6] Hares, S., "IDRP for IP", Work in Progress, March 1993.
- [7] Varadhan, K., "BGP4 OSPF Interaction", Work in Progress, March 1993.
- [8] Topolcic, C., "Notes on BGP-4/CIDR Coordination Meeting of 11 March 93", Informal Notes, March 1993.
- [9] Knopper, M., "Aggregation Support in the NSFNET Policy Routing Database", RFC 1482, Merit, June 1993.

9. Security Considerations

Security issues are not discussed in this memo.

10. Authors' Addresses

Yakov Rekhter
T.J. Watson Research Center, IBM Corporation
P.O. Box 218
Yorktown Heights, NY 10598

Phone: (914) 945-3896
EMail: yakov@watson.ibm.com

Claudio Topolcic
Corporation for National Research Initiatives
1895 Preston White Drive, Suite 100
Suite 100
Reston, VA 22091

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