

Routing in a Multi-provider Internet

Status of this Memo

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

Abstract

This document was prepared by the author on behalf of the Internet Architecture Board (IAB). It is offered by the IAB to stimulate discussion.

Over the past few years the Internet has undergone significant changes. Among them is the emergence of multiple Network Service Providers, where resources that provide Internet-wide IP connectivity (routers, links) are controlled by different organizations. This document presents some of the issues related to network layer routing in a multi-provider Internet, and specifically to the unicast routing.

1. Network Service Providers vs Network Service Subscribers

Within the current routing paradigm the service offered by a provider at the network layer (IP) is the set of destinations (hosts) that can be reached through the provider. Once a subscriber establishes direct connectivity to a provider, the subscriber can in principle reach all the destinations reachable through the provider. Since the value of the Internet-wide connectivity service offered by a provider increases with the number of destinations reachable through the provider, providers are motivated to interconnect with each other.

In principle a provider need not offer the same service (in terms of the set of destinations) to all of its subscribers -- for some of the subscribers the provider may restrict the services to a subset of the destinations reachable through the provider. In fact, for certain types of subscribers constrained connectivity could be seen as part of the service offered by a provider.

In a multi-provider environment individual providers may be driven by diverse and sometimes even conflicting goals and objectives. Some of the providers exist to provide connectivity to only a specific group

of Network Service Subscribers. Other providers place no constraints on the subscribers that can subscribe to them, as long as the subscribers pay the fee charged by the providers. Some of the providers place certain constraints on the reselling of the connectivity services by organizations (e.g., other providers) attached to the providers. Some of the providers may be operated by companies that are subject to specific regulations (e.g., regulated monopoly), while other providers are completely unregulated. The scope of geographical coverage among providers varies from a small region (e.g., county, town) to a country-wide, international, or even intercontinental.

There is no centralized control over all the providers in the Internet. The providers do not always coordinate their efforts with each other, and quite often are in competition with each other.

Despite all the diversity among the providers, the Internet-wide IP connectivity is realized via Internet-wide distributed routing, which involves multiple providers, and thus implies certain degree of cooperation and coordination. Therefore, there is a need to balance the providers' goals and objectives against the public interest of Internet-wide connectivity and subscribers' choices. Further work is needed to understand how to reach the balance.

2. Routing Requirements

Conceptually routing requirements can be classified into the following three categories: source preferences, destination preferences, and constraints on transit traffic. Source preferences allow an originator of a packet to exert control over the path to a destination. Destination preferences allow a destination to exert control over the path from a source to the destination. Constraints on transit traffic allow a provider to control the traffic that can traverse through the resources (routers, links) controlled by the provider.

From a conceptual point of view the requirements over the degree of control for source and destination preferences may vary from being able to just provide connectivity (regardless of the path), to being able to select immediate providers, to more complex scenarios, where at the other extreme a subscriber may want to have complete control over the path selection.

From a conceptual point of view the requirements over the degree of control for transit traffic may vary from control based only on the direct physical connectivity (controlling the set of organizations directly connected to the provider), to being able to restrict traffic to a particular set of sources or destinations, or a

combination of particular sources and destinations, or even take into account the paths to/from these sources and/or destinations.

In view of a potentially wide variety of routing requirements, we need to get a better understanding on the relative practical importance of various routing requirements. In practice organizations usually don't formulate their routing requirements in a vacuum. For example, since the primary role of a provider is to provide services to a set of subscribers, the provider usually formulates its routing requirements based on the set of the routing requirements of the subscribers the provider is expected to serve.

Support for various routing requirements should take into account the overhead and the scope of the overhead associated with those requirements. A situation where an organization can unilaterally impose routing information overhead on other organization (e.g., by requiring the other organization to maintain an additional routing information) should be viewed as undesirable. The cost of supporting a particular routing requirement should not be borne by organizations that do not benefit from supporting that requirement. Ideally the routing system should allow to shift the overhead associated with a particular routing requirement towards the entity that instigates the requirement (for example, there is a need to carefully balance the overhead associated with maintaining a state needed for multi-hop header compression vs carrying explicit forwarding information on a per packet basis). Organizations with simple routing requirements shouldn't bear the same routing information overhead as organizations with complex routing requirements.

A situation where the overhead associated with supporting a particular routing requirement has to be carried by every entity (e.g., router, host) within an organization that would like to impose the requirement could be viewed as undesirable. An organization should be able to instantiate its routing requirements in a more or less central fashion, for example by utilizing just some of the routers.

Even if the scope of the routing information overhead is purely local, there is a need to perform a careful analysis of the tradeoff between the potential benefits and the cost associated with supporting various routing requirements.

3. Encapsulation

The technique of encapsulation allows for the creation of a "virtual" IP overlay over an existing IP infrastructure. This has certain implications for the Internet routing system.

In the presence of encapsulation, a provider may no longer be able to constrain its transit traffic to a particular set of ultimate sources and/or destinations, as a packet may be encapsulated by some router along the path, with the original source and/or destination addresses being "hidden" (via encapsulation) at the Network layer. Likewise, encapsulation may affect source and destination preferences, as a source (or a destination) may either (a) be unaware of the encapsulation, or (b) have little or no control over the encapsulated segment of a path.

Further work is needed to understand the implications of the overlay capabilities created via encapsulation on the semantics of routing requirements, as well as the interaction among the routing requirements by the entities that form the overlay and the entities that form the underlying infrastructure.

4. Price Structure and its Impact on Routing

Routing among providers, as well as between providers and subscribers may be influenced by the price structure employed by the providers, as well as the usage pattern of the subscribers. A provider can view routing as a mechanism that allows the provider to exert control over who can use the provider's services. A subscriber can view routing as a mechanism that allows the subscriber to exert control over the price it pays for the Internet connectivity.

The need to exert control has to be carefully balanced against the cost of the routing mechanisms needed to provide such control. In a competitive market one could question the viability of a mechanism whose incremental cost would be greater than the saving recovered by the mechanism -- competitive pressure or alternate mechanisms are likely to push providers and subscribers towards choosing the cheapest mechanism.

5. Scalability

One of the key requirements imposed on the Internet routing is its ability to scale. In addition to conventional metrics for scalability (e.g., memory, CPU, bandwidth), we need to take into account scalability with respect to the human resources required to operate the system. The need for deployment of CIDR already showed that a routing scheme that scales linearly with respect to the number of connected networks, or even to the number of connected organizations is unacceptable today, and is likely to be unacceptable in the long term. It is not clear whether routing that scales linearly with the number of providers is going to be acceptable in the long term.

Scaling implies that the Internet routing system needs to have powerful mechanisms to provide routing information aggregation/abstraction.

In the absence of Internet-wide coordination and in the presence of competition among the providers, the aggregation/abstraction mechanisms should minimize preconditions as well as limit the amount of required inter-provider coordination. Ideally the routing system should allow a provider to control the amount of its local resources needed to deal with the routing overhead based on considerations that are purely local to the provider.

One of the side effects of the routing information aggregation/abstraction is that some of the routing information is going to be lost. This may impact route optimality and even the ability to find an existing route. The need for routing information aggregation/abstraction also implies certain homogeneity of the information to be aggregated/abstracted. This needs to be counter-balanced against the potential diversity of routing requirements.

As a way to deal with the routing information loss due to aggregation/abstraction, we need to explore mechanisms that allow routing that is based on the on-demand acquisition of subsets of unaggregated information.

The overhead associated with supporting specific routing requirements has a direct impact on the overall scalability of the Internet routing system. We need to get a better understanding of how various routing requirements impact scalability. When the impact is significant, and the requirements have practical importance we need to develop mechanisms that allow the impact to be reduced.

6. Hierarchical Routing

Classless Inter-Domain Routing (CIDR) (RFC1518, RFC1519) that is used today for scalable Internet-wide routing is based on the technique of hierarchical routing. Essential to this technique is the assumption that Network layer addresses assigned to individual entities (e.g., hosts, routers) reflect the position of these entities within the network topology -- addresses are said to be "topologically significant". With CIDR addresses assigned to most of the individual sites are expected to reflect providers the sites are connected to -- CIDR uses "provider-based" addresses.

One of the fundamental consequences of using hierarchical routing is that in order to preserve topological significance of network addresses, changes in the network topology may need to be accompanied by the corresponding changes in the addresses. Presence of multiple

providers serving the same geographical area implies that a subscriber should be able to switch from one provider to another. Since such a switch implies changes in the Internet topology, it follows that to retain topological significance of the (provider-based) addresses within the subscriber, the subscriber has to change the addresses of all of its entities -- the process known as "renumbering". There are already tools to facilitate this process -- Dynamic Host Configuration Protocol (DHCP). However, DHCP is not yet widely deployed. Further work is needed to improve these tools, get them widely deployed, and to integrate them with Domain Name System (DNS).

Multi-level hierarchical routing allows for recapturing additional routing information (routing entropy) due to the mismatch between addresses and topology at a particular level in the routing hierarchy at some higher level in the hierarchy (e.g., at an exchange point among providers). This enables the routing system to contain the scope of entities impacted by the mismatch. Containing the scope of entities could be an important factor to facilitate graceful renumbering. Further work is needed to develop appropriate deployment strategies to put these capabilities in place.

It is important to emphasize that the requirement to maintain topologically significant addresses doesn't need to be applied indiscriminately to all the organizations connected to the Internet -- hierarchical routing requires that most, but not all addresses be topologically significant. For a large organization it could be sufficient if the set of destinations within the organization can be represented within the Internet routing system as a small number of address prefixes, even if these address prefixes are independent of the providers that the organization uses to connect to the Internet ("provider-independent" addresses). The volume of routing information that a large organization would inject into the Internet routing system would be comparable to the (aggregated) routing information associated with a large number of small organizations.

Existence of multiple providers allows a subscriber to be simultaneously connected to more than one provider (multi-homed subscribers). CIDR offers several alternatives for handling such cases. We need to gain more operational experience as well as better understand tradeoffs associated with the proposed alternatives.

An alternative to CIDR address assignment is to assign addresses based purely on the geographical location. However, address assignment that reflects geographical location of an entity implies that either (a) the Internet topology needs to be made sufficiently congruent to the geography, or (b) addresses aren't going to be topologically significant. In the former case we need to understand

the driving forces that would make the topology congruent to the geography. In the latter case techniques other than hierarchical routing need to be developed.

7. Routing Information Sharing

While ensuring Internet-wide coordination may be more and more difficult, as the Internet continues to grow, stability and consistency of the Internet-wide routing could significantly benefit if the information about routing requirements of various organizations could be shared across organizational boundaries. Such information could be used in a wide variety of situations ranging from troubleshooting to detecting and eliminating conflicting routing requirements. The scale of the Internet implies that the information should be distributed. Work is currently underway to establish depositories of this information (Routing Registries), as well as to develop tools that analyze, as well as utilize this information.

8. Summary

In this section we enumerate some of the issues that the IAB thinks should be brought to the attention of the Internet community.

The following two tasks require the most immediate attention:

- further work is needed to develop technologies that facilitate renumbering
- further work is needed to investigate feasibility of routing information aggregation above the direct (immediate) provider level

The following tasks are viewed as medium term:

- further work is needed to get a better understanding on the relative practical importance of various routing requirements
- further work is needed to understand of how various routing requirements impact scalability of the routing system
- further work is needed to investigate alternatives to hierarchical routing

Finally, the following tasks are viewed as long term:

- further work is needed to understand and utilize the benefits of routing information sharing

- further work is needed to understand the implications of virtual overlays created via encapsulation
- further work is needed to understand how different price structures influence routing requirements
- further work is needed to understand how to balance the providers' goals and objectives against the public interest of Internet-wide connectivity and subscribers' choices.

9. Conclusions

This document presents some of the issues related to routing in a multi-provider Internet. There are no doubt routing-related areas that are not covered in this document. For instance, such areas as multicast routing, or routing in the presence of mobile hosts, or routing in the presence of a large shared media (e.g., ATM) aren't discussed here. Further work is needed to understand the implications of a multi-provider Internet on these areas.

The impact of multi-provider Internet goes well beyond just routing, and percolates into such areas as network management, troubleshooting, and others. Further work is needed to assess the implications of multi-provider environment on these areas, as well as to understand the interaction among all these areas from a system-wide perspective.

10. Acknowledgments

Many thanks to all the IAB members, and especially to Brian Carpenter, Robert Elz, Christian Huitema, Paul Mockapetris, and Lixia Zhang for their contributions to this document.

Security Considerations

Security issues are not discussed in this memo.

Editor's Address

Yakov Rekhter
T.J. Watson Research Center IBM Corporation
P.O. Box 704, Office H3-D40
Yorktown Heights, NY 10598

Phone: +1 914 784 7361
EMail: yakov@watson.ibm.com

