

Internet Service Provider Address Coalitions (ISPACs)
<[draft-li-ispac-00.txt](#)>

Status of this Memo

This document is an Internet Draft. Internet Drafts are working documents of the Internet Engineering Task Force (IETF), its Areas, and its Working Groups. Note that other groups may also distribute working documents as Internet Drafts.

Internet Drafts are draft documents valid for a maximum of six months. Internet Drafts may be updated, replaced, or obsoleted by other documents at any time. It is not appropriate to use Internet Drafts as reference material or to cite them other than as a ``working draft'' or ``work in progress.''

Please check the `l1d-abstracts.txt` listing contained in the internet-drafts Shadow Directories on `nic.ddn.mil`, `nnsf.net`, `nic.nordu.net`, `ftp.nisc.sri.com`, or `munari.oz.au` to learn the current status of any Internet Draft.

1. Introduction

The introduction of CIDR [[1](#), [2](#)] substantially changed the way that address allocation was performed in the Internet. The result is that address allocation is provider-based: the address used by a host depends on which service provider is used by the host. One of the unfortunate implications is that if a host changes provider, it must renumber. In many cases this is awkward, expensive, and difficult.

This note describes an organization that may be formed to help reduce the need for user renumbering while still providing the benefits of CIDR to the remainder of the network. Such organizations do not currently exist, so this note is at best a discussion of the foreseeable future and is not a catalog of existing experience.

These organizations are formed by grouping Internet Service Providers (ISPs) into topologically connected sets, each of which shares one or more IP address prefixes. Such an organization is called an ISP Address Coalition (ISPAC, pronounced "ice pack"). To simplify the discussion, we assume that the ISPAC shares a single prefix. The discussion generalizes to multiple shared prefixes in the obvious manner.

2. Creation of an ISPAC

An ISPAC can be created from a set of ISPs which are willing to share an address prefix and are willing to meet certain technical and administrative requirements. There may also be additional legal requirements, such as anti-trust issues, that are beyond the scope of this document.

The primary technical requirement is that all members of the ISPAC be willing to cooperate and each treat the shared prefix as a common resource. Further, each member of the ISPAC should perform aggregation on the shared prefix address space when advertised outside of the ISPAC. Note that this aggregation may be performed in accordance with normal CIDR procedures, so this does not necessarily imply that only the aggregated prefix is advertised. This exceptional case is discussed in more detail below.

Because all members of an ISPAC are advertising the aggregated shared prefix, any member may attract traffic for portions of the shared prefix that need to be routed to another member. For this reason, it is essential that all members of the ISPAC be able to route to any other member of the ISPAC without leaving the ISPAC (i.e., the ISPAC is connected). One way that this requirement could be satisfied is for all members of the ISPAC to be present at a common interconnect.

The primary administrative requirement to form an ISPAC is the selection of an addressing authority for the shared prefix. A complete discussion of the issues in the selection of such an authority are beyond the scope of this document, but may include remuneration for providing for address space administration and impartiality of the authority. Once an addressing authority is selected, it is necessary to acquire the shared prefix. This would be done according to the normal address allocation procedures. The address space request from an ISPAC should be regarded as if it came from any ISP with the properties of the union of the members. No special privileges are accorded to requests from ISPACs, so normal justifications for address space would apply.

3. Operation of an ISPAC

The day to day operations of an ISPAC has certain special considerations which impact all of the members of the ISPAC. Because sites may change providers within the ISPAC, intra-ISPAC routing must be carefully managed. A user should be able to be assured that they are the only ones using a component in the shared prefix and that their routing reliability is comparable to that provided to users within the ISP assigned address prefixes. To accomplish this, the

members of the ISPAC must exchange routing information about components of the shared prefix. The exact mechanisms that are used to accomplish this are beyond the scope of this document, but there are many possibilities, including a common IGP, or via BGP. Members of the ISPAC should agree on a particular set of mechanisms for exchanging routing information about the shared prefix, including protocols and any administrative filtering for sanity checking.

Because components of the shared prefix are generally not advertised outside of the ISPAC, if the ISPAC itself is internally partitioned, it will cause an outage between users of the shared prefix. To prevent this, members of the ISPAC are strongly encouraged to insure that the ISPAC is biconnected or triconnected. One way that this can be accomplished is to have multiple interconnects at which all members of the ISPAC are present. These may be private or public interconnects. Alternately, increased connectivity can be provided by a number of pairwise interconnects.

For inter-ISPAC routing, each member of the ISPAC must advertise the shared prefix. Normally, this will be done with BGP, so the shared prefix may appear to have multiple source autonomous systems. For this reason, advertising the existence and membership of the ISPAC and the particular prefix to the administrators of other domains is useful to help eliminate false alarms in interdomain sanity checking systems.

By default, simple advertisement of the shared prefix will cause all traffic to enter the ISPAC at the nearest member. This will only be optimal for sites which have their optimal connection to that neighbor. In many cases, it would be preferable to determine optimal entry points for components of the shared prefix. At the same time, it is essential to not advertise all components to the remainder of the network, as this would defeat the aggregation of the prefix. To preserve aggregation and provide more optimal routing, there is a useful technique which can be borrowed from normal interdomain routing. In this application, this technique a member of the ISPAC may advertise some of the components of the shared prefix, but the member must insure that the scope of these advertisements is carefully restricted. The restriction of scope is usually accomplished by administrative filtering and careful coordination the receiver of the advertisements.

Because any advertisement of the shared prefix by a member of the ISPAC may attract traffic, it may be the case that a member may end up providing transit service to other members of the ISPAC. It is up to the members of the ISPAC to determine their own policy regarding this transit traffic. It may be that such traffic is equally distributed, in which case responsibilities are shared equally. Alternately, due to customer dissimilarities, the traffic loads may be highly skewed. In this case, the shared prefix may only be advertised to external providers at interconnects where all members

and external providers are present. In this case, it is likely that members would not advertise the shared prefix across their private interconnects. This will affect the quality of the routing for sites within the shared prefix.

Note that nothing in this discussion precludes a provider from joining multiple ISPACs. Further, these ISPACs may share members in any fashion. ISPACs may be disjoint, may overlap, may be tangential or may intersect freely.

4. Benefits of an ISPAC

The primary benefit of an ISPAC is to reduce the requirement on renumbering when a site changes provider. This benefit only occurs when the site changes providers within the ISPAC. As a result, a site gains the most benefit from the ISPAC when there are many providers within the ISPAC. Further, a site may have substantial costs if changing to a provider which is distant. Thus, a site derives the most benefit if there are multiple providers within the same geographic area. Note that this is most effective if the different members of the ISPAC are closely comparable competitors. The site will consider reliability, bandwidth and service when selecting an alternate provider, in addition to considering membership in the ISPAC.

Note that the end site itself is not a member of the ISPAC and has no obligation whatsoever to change to another provider within the ISPAC.

Another benefit of an ISPAC is for sites that wish to be multi-homed through multiple providers. If such sites are numbered from the shared prefix and are connected through multiple providers that are all members of the ISPAC, then no extra routing information needs to be injected into global routing. This is a significant benefit because even with CIDR, the scalability of global routing can fall victim due to the number of prefixes injected due to multi-homed sites.

Another benefit that aids in the aggregation of routing information is the ability of the ISPAC to take multiple smaller providers and advertise only the shared prefixes. Without the ISPAC, each ISP would be forced to inject their own routes. The additional aggregation resulting from membership in the ISPAC should greatly improve the level of aggregation that can happen across small providers, both by allowing the ISPAC a larger block of addresses and by automatically aggregating the ISPAC shared prefixes at the border of the ISPAC. An ISPAC also allows a smaller provider to receive significant amounts of address space which is not allocated through an upstream provider. This is advantageous as the smaller provider may wish to be associated with one or more larger providers. An association with any single larger provider might result in inferior

routing, and an inability to leave that large provider without renumbering. With an ISPAC, such a member would instead be dependent on the continued existence of the ISPAC.

5. Drawbacks of an ISPAC

Along with these many benefits, there are also a number of drawbacks to membership in an ISPAC. The primary difficulty stems from the inability to perform aggregation on the shared prefix within the ISPAC. As noted above, this implies that members of the ISPAC must be willing to devote a certain amount of resources to supporting ISPAC routing. Also as mentioned above, the quality of routing for a shared prefix may suffer as compared to a private prefix due to the aggregation at the ISPAC boundary. Due to this aggregation, traffic will tend to enter the nearest border of the ISPAC. Without aggregation, traffic would seek the closest entry to the ISP. In cases in which these two entries differ, the routing may be suboptimal.

Another extremely important point about ISPACs is that they do not eliminate the need to renumber. If a site chooses to leave the ISPAC, they will have to renumber, just as if they left a non-ISPAC provider. As with the normal renumbering case, certain accommodations may be made to ease the transition to the site's new prefix. These circumstances and operations are part of normal CIDR prefix management and are not detailed here.

Similarly, if an ISP leaves an ISPAC, then any of the ISP's customers which are numbered out of the shared prefix will have to make a difficult decision. If they choose to remain with the ISP or change to a ISP not in the ISPAC, then they will have to renumber. Alternately, they can change to another provider within the ISPAC. In both this and the above case, the end user ends up renumbering. It is important to note these cases as the primary benefit of an ISPAC is to decrease, but not eliminate, the need for site renumbering.

If an ISPAC is composed of ISPs of differing sizes, there may be some difficulty with inequities of responsibilities to the ISPAC. To see this clearly, consider a case where a very large provider and a very small provider form an ISPAC. In this case, the small provider may have to provide transit service for the very large provider's traffic. Yet in return, the small provider may receive little transit traffic through the very large provider. Similarly, the very large provider may consume a much larger percentage of the shared prefix.

6. Generalization

So far, we have described the simplest and smallest forms of an ISPAC. These consist of a small number of providers aligned around a small number of interconnects. In this section we consider some of the situations which can occur with larger ISPACs. Note that these examples are used solely to illustrate the scalability of the concept and are not a recommendation for particular ways of using ISPACs.

An obvious opportunity to use ISPACs is in metropolitan areas, grouping providers around population centers. If the metropolitan area has a sufficiently large population, it probably also has a significant amount of traffic, one or more interconnects, and a number of ISPs. An ISPAC in this situation bears a superficial resemblance to geographic addressing. However, there are significant differences. Unlike geographic addressing, participation in the ISPAC is wholly voluntary. This maximizes the choices that the site has, and so this should be considered a significant advantage. Further, this allows competition between ISPACs, which is vital in ensuring that the ISPAC provides quality service to its sites. Finally, the ISPAC would easily allow private interconnects from members to non-members, which may or may not advertise the shared prefix. Scaling the number of interconnects is key to insuring high connectivity in the topology, for distributing traffic, and for optimizing routing.

ISPACs may be further generalized along topological or geographic boundaries, although neither of these directions is required. For example, if there are several providers which are at a number of major interconnects, they may wish to form an ISPAC, even though the interconnects may be widely separated. Such an ISPAC should be more robust since it would not partition without the failure of several interconnects. With this size of ISPAC, it is also possible for a company to place both its headquarters and its remote offices within the ISPAC shared prefix. This is beneficial in that the company does not inject multiple prefixes into global routing.

If we continue this generalization to its logical extreme, we end up with a global ISPAC which all providers belong to. The shared prefix is similar to the provider independent address space that exists today. However, there is one subtle difference in that the address space is independently managed under the direction of the ISPAC.

7. Conclusions

To summarize, Internet Service Provider Address Coalitions (ISPACs) are ad hoc organizations which can be formed by groups of ISPs. There are significant technical and administrative difficulties to be overcome in forming and maintaining these organizations. However, the benefits that result from these organizations can be significant both to end users and providers alike. At the same time, if executed

correctly, ISPACs are not detrimental to the scalability of Internet routing, and may, in some cases, aid that scalability.

The primary benefit of an ISPAC is a reduction in the need for end sites to renumber. This need is not eliminated with the application of ISPACs, but can be significantly reduced, depending on the exact deployment of ISPACs. The benefits to the ISP include the ability to offer a desired service to interested sites, the ability to utilize address space that is independent of larger providers, and the ability to decrease the number of prefixes that it advertises.

7. Security Considerations

Security issues are not discussed in this document.

8. Acknowledgments

Much of this memo is the result of ongoing fruitful discussions with Yakov Rekhter and the work done on stratum based aggregation. Joel Halpern (NewBridge) and Bill Manning (ISI) were kind enough to review and comment on a previous version of this paper.

9. Author's addresses

Tony Li
Juniper Networks, Inc.
3260 Jay St.
Santa Clara, CA 95051
Phone: +1 408 327 1906
Fax: +1 408 327 1910
email: tli@juniper.net

10. References

- [1] Fuller, V., Li, T., Yu, J., and Varadhan, K., 'Supernetting: an Address Assignment and Aggregation Strategy', [RFC1519](#), September 1993
- [2] Rekhter, Y., Li, T., "An Architecture for IP Address Allocation with CIDR", [RFC1518](#), September 1993