

End-site routing support for IPv6 multihoming¹☆

Marcelo Bagnulo*, Alberto García-Martínez, Juan Rodríguez, Arturo Azcorra

Universidad Carlos III de Madrid, Av. Universidad, 30. Leganés, Madrid, Spain

Available online 13 September 2005

Abstract

Multihoming is currently widely used to provide fault tolerance and traffic engineering capabilities. It is expected that, as telecommunication costs decrease, its adoption will become more and more prevalent. Current multihoming support is not designed to scale up to the expected number of multihomed sites, so alternative solutions are required, especially for IPv6. In order to preserve interdomain routing scalability, the new multihoming solution has to be compatible with Provider Aggregatable addressing. However, such addressing scheme imposes the configuration of multiple prefixes in multihomed sites, which in turn causes several operational difficulties within those sites that may even result in communication failures when all the ISPs are working properly. In this paper we propose the adoption of *Source Address Dependent* routing within the multihomed site to overcome the identified difficulties.

© 2005 Elsevier B.V. All rights reserved.

Keywords: Multihomed site; CIDR; Internet Service Providers

1. Introduction

Since the communications of a wide range of organizations rely on the Internet, access links are a critical resource to them. As a result, sites are improving the fault tolerance and QoS capabilities of their Internet access through *multihoming*, i.e. the achievement of global connectivity through several connections supplied by different Internet Service Providers (ISPs). However, the extended usage of the currently available IPv4 multi-homing solution is jeopardizing the future of the Internet since it has become a major contributor to the post-CIDR growth in the number of global routing table entries [1]. Therefore, in IPv6, the usage of *Provider Aggregatable (PA)* addressing is recommended for all sites, included multihomed ones, in order to preserve inter domain routing system scalability. While such addressing architecture reduces the amount of routing table entries in the *Default Free Zone* of the Internet, its adoption presents a fair number of challenges for the

end-sites, especially for those who multihome. Essentially, when a PA addressing is adopted, a multihomed site will have to configure multiple addresses, one per ISP, in every node of the site, in order to be reachable through all its providers. Such configuration pose quite a number of challenges for its adoption, since current hosts are not prepared to deal with multiple addresses per interfaces as it is required. In this paper, we will present how *Source Address Dependent (SAD)* routing can be adopted to deal with some of the difficulties imposed in this configuration.

The rest of this paper is structured as follows: First we will present the rationale for adopting SAD routing within multihomed sites. Then, we will detail the different configurations of SAD routing that may be required in different sites, including some trials performed and next we will present the capabilities of the resulting configuration. Finally we will present the conclusions of this work.

2. Rationale

2.1. Current IPv4 multihoming technique and capabilities

As mentioned above, a site is multi-homed when it obtains Internet connectivity through two or more service providers. Through multi-homing an end-site improves the fault tolerance of its connection to the global network and it can also perform *Traffic Engineering (TE)* techniques to

☆This work has been partly supported by the European Union under the E-Next Project FP6-506869 and by OPTINET6 project TIC-2003-09042-C03-01.

* Corresponding author.

E-mail addresses: marcelo@it.uc3m.es (M. Bagnulo), alberto@it.uc3m.es (A. García-Martínez), jrh@it.uc3m.es (J. Rodríguez), azcorra@it.uc3m.es (A. Azcorra).




Fig. 1. BGP route injection for the provision of multihoming in IPv4 networks.

select the path used to reach the different networks connected to the Internet.

In IPv4, the most widely deployed multi-homing solution is based in the announcement of the site prefix through all its providers. In this configuration, the site S obtains a *Provider Independent (PI)* prefix allocation directly from the Regional Internet Registry. Then, the site announces this prefix to its providers using BGP [2]. Then the multihomed site providers announce the prefix to its providers and so on, so that eventually the route is announced in the *Default Free Zone* (Fig. 1).

This mechanism provides fault tolerance capabilities, including preserving established connections throughout an outage. In addition, the following TE tools are available to the multihomed site: The multihomed site can define which one of the available exit paths will be used to carry outgoing traffic to a given destination, by proper configuration of the `LOCAL_PREFERENCE` attribute of BGP [3]. For incoming traffic, the multihomed site can influence the ISP through which it prefers to receive traffic by using the *AS prepending* technique, which consists in artificially making the route through one of the providers less attractive to external hosts by adding AS in the `AS_PATH` attribute of BGP [3]. (It should be noted that in this case, the ultimate decision of which ISP will be used to forward packets to the site belongs to the external site that is actually forwarding the traffic).

While the IPv4 multihomed solution presented provides fairly good features regarding to fault tolerance and TE, it presents very limited scalability properties with respect to the interdomain routing system. Because of the usage of PI addressing by the multihomed sites, each multi-homed site using this solution contributes with routes to the Default Free Zone routing table, imposing additional stress to already oversized routing tables. Additionally, the increase of the number of ASs contributing to the BGP routing table resulting from PI addressing may affect to the time required for finding a valid path after a failure, since this time depends on the length of the longest path between the site originating the prefix and the site computing the path [4]. Finally, as the number of the ASs grows, the frequency of additions and withdraws of routes also increases, leading to a more unstable system. For this reason, more scalable




Fig. 2. Provider aggregation of end-site prefixes.

multi-homing solutions are being explored for IPv6 [5], in particular solutions that are compatible with the usage of PA addressing in multihomed sites, as it will be presented next.

2.2. Provider aggregation and multi-homing

In order to reduce the routing table size, the usage of PA addressing is required. This means that sites obtain prefixes which are part of their provider's allocation, so that its provider only announce the complete aggregate to their providers, and they do not announce prefixes belonging to other ISP aggregates, as presented in Fig. 2.

When provider aggregation of end-site prefixes is used, each end-site host interfaces obtains one IP address from each allocation, in order to be reachable through all the providers and benefit from multi-homing capabilities since ISPs will only forward traffic addressed to their own aggregates.

This configuration presents several concerns as it will be presented next.

- Ingress filtering [6] is a widely used technique for preventing the usage of spoofed addresses. However, in the described configuration, its usage presents additional difficulties for the source address selection mechanism and intra-site routing systems, since the exit path and source address of the packet must be coherent with the path, in order to bypass ingress filtering mechanisms.
- Difficulties in establishing communication in case of failure. When *Link1* or *Link3* becomes unavailable, addresses containing the *PXsite* prefix are unreachable from the Internet.
- Established connections will not be preserved in case of outage. If *Link1* or *Link3* fails, already established connections that use addresses containing *PXsite* prefix will fail, since packets addressed to the *PXsite* aggregate will be dropped because there is no route available for this destination. Note that an alternative path exists, but the routing system is not aware of it.

The presented difficulties show that additional mechanisms are needed in order allow the usage of PA addresses while still provide incumbent multi-homing solution

equivalent benefits. In this paper, we will explore the possibility of using Source Address Dependent routing as a tool to help to overcome the identified difficulties.

3. Source address dependent (SAD) routing

Source Address Dependent (SAD) routing essentially means that routers maintain as many routing tables as source address prefixes involved, and the packet is routed according to the routing table corresponding to the source address prefix that best matches the source address contained in the packet (Fig. 3). The idea of using SAD routing for site multihoming is presented in [7].

SAD routing can be used to provide ingress filtering compatibility when routing packets flowing from the multihomed site to the Internet. In this case, the source address of the exiting packets has been determined by the host that initiated the communication (the host in the multihomed site or the external host through the selection of the destination address of the initial packet) and then the routing system will forward the packet to the appropriate exit router in order to guarantee ingress filtering compatibility. The source address selection determines the ISP to be used for routing packets. Because of address filtering, the source address determines the forward path from the multihomed site to the rest of the Internet. Additionally, as the source address used in the initial packets will become the destination address of the reply packets, this selection determines the ISP to be used in the reverse path.

Since source address selection implies ISP selection, the adoption of SAD routing will also affect the mechanisms to be used in multihomed sites to define TE. In particular, it will shift TE capabilities from the routing system to the hosts themselves.

We will next evaluate the adoption of SAD routing in two typical multihomed configurations: sites running BGP but without redistributing the BGP information into IGP, and sites running an IGP to select the exit path. There is an additional possible configuration using static routes in




Fig. 4. SAD routing domain.

the multihomed site. However, this last configuration is fairly simple and several commercial routers already support it, so we will not provide a full description of it. Nevertheless, it should be noted that when SAD routing is used, it is possible to obtain fault tolerance and TE capabilities in sites with static routing since those features are now supported by the hosts themselves and not by the routing system.

In order to enable SAD routing in a multihomed site, SAD routing support is not required in all the routers within the site, but SAD routing has to be adopted creating a connected SAD routing domain that contains all the exit routers—see Fig. 4—as presented in [7].

Note that it is not necessary for the generic routing domain to be connected, i.e. it can be formed by a set of disconnected domains, all connected to the SAD routing domain.

3.1. Sites running BGP without redistribution of BGP information into IGP

Current IPv4 multihomed sites usually run BGP with their providers. Through BGP, they obtain reachability information from each of their ISPs. However, because operational issues, some sites do not redistribute the information obtained through BGP into the IGP [3]. So, in order to be able to properly select the intra site path towards an external destination, they include all the routers that are required to properly select the exit path in the IBGP mesh, including not only site exit routers, but also other internal routers that have access to multiple exit routers. This means that the IBGP cloud is wrapping the non-BGP aware routing domain, as presented in Fig. 5.

It should be noted that only the IBGP mesh must be connected, and that the non-BGP aware may be formed by multiple disconnected domains, only connected by the IBGP domain. It is clear that only the routers included in the IBGP mesh need to implement SAD routing in order to properly select the site exit path. So, since all these routers are running BGP, we can use BGP capabilities to provide SAD routing support.

In order to implement SAD routing, each exit router that is running EBGp has to attach a color tag to the routes received from the ISP, so that it is possible to




Fig. 3. SAD routing.




Fig. 5. IBGP full mesh.

identify the routes learned through each different ISP. Additionally, once the routing information is colored, it is necessary to map each of the colors to a source address prefix. Once that both color and its correspondent prefix information is available, it is possible to construct SAD routing tables that contain routing information per source prefix.

SAD routing can be implemented in this scenario using the BGP Communities [8] attribute to color the routing information. So, we assume a multihoming scenario where a multihomed site has n ISPs, each one of them has assigned $Pref_i$ to the multihomed site, with $i = 1, \dots, n$. So, in order to adopt SAD routing it is required that:

First, a private community value is assigned to each different ISP. Therefore, Com_i value is assigned to the routes obtained from ISP_i , being $i = 1, \dots, n$

Second, n routing tables are created in each of the routers involved, so that each router has one routing table per prefix in the site (i.e. per ISP). Additionally each router is configured to route packets containing a source address matching $Pref_i$ using the routing table i .

Third, BGP processing rules are configured in each router, so that routes containing a community attribute value equal to Com_i only affect routing table i .

Finally, each exit router that is peering with an external router in ISP_i is configured to attach the community value

Com_i to all the routes received from ISP_i , when announcing them through IBGP.

The resulting behavior is that each router within the IBGP mesh will have separate routing tables containing the information learned through each ISP. Packets containing a source address with the prefix of the ISP_i will be routed using the corresponding routing table. An example is shown in Fig. 6.

3.2. Sites using IGP

In this scenario, the multihomed site is using an IGP to inform about both internal and external destinations. The IGP learns about external destinations in one of the following three ways:

- Manually configured routes are imported into the IGP
- BGP redistribution into the IGP
- IGP exchange with the providers

As in the previous case, the whole multihomed site routing system is not required to support SAD routing but only a connected domain that has to contain all the exit routers. However, while BGP provides mechanisms to tag routing information so that the same protocol instance can be used to propagate information with different scopes, as presented in the previous section, current IGPs do not provide such capability.

So, in order to provide SAD routing support, different instances of the routing protocol run in parallel, each one of them associated with a source address prefix. In this way, different instances of the IGP will update different routing tables within the routers. The main difficulty with this approach is how to differentiate messages corresponding to the different instances of the IGP. Normally, different instances of the IGP run in different interfaces, so that each instance only receives its own messages. But in this case we want to run multiple instances of the IGP in the same




Fig. 6. SAD routing deployment using BGP communities.

interfaces, so we need a way to separate messages according to the instance of the IGP they belong to.

OSPF for IPv6 [9] explicitly supports running multiple instances in the same link and packets belonging to different instances are identified using the `Instance_ID` field in the OSPF header.

Another possibility is to send IGP messages using global addresses as source addresses. Usually, IGP messages are sent using link local addresses. But, since each router can be configured with multiple IP addresses, one per prefix, the router includes different source addresses in the messages corresponding to different instances of the IGP. This ship-in-the-night strategy would allow each IGP instance to believe that they are running alone in the link

3.3. Experimenting with SAD routing

We will next analyze the deployability of the approach by evaluating the available support for SAD routing in current implementations. In order to assess the deployment effort required to adopt the proposed solution we have built a testbed with widely available commercial routers and perform some trials in the framework of the Optinet6 research project. The testbed evaluated the capabilities to support SAD routing of Cisco 2500 routers, Cisco 7500 routers and Juniper M10 routers.

All the tested routers support static SAD routing, i.e. routing based on the source address of the packets according to statically defined routes. However, the implementation of the SAD routing support differs considerably between them. Cisco IOS supports static SAD routing through manually defined rules that affect the processing of packets, called route-maps. In order to enable SAD routing, route-maps corresponding to each source address dependent route have to be defined. On the other hand, Juniper routers support multiple routing tables, so that it is possible to create as many routing tables as source address prefixes are involved, and then define the required rules so that the router will forward packets according to the routing table associated with the prefix contained in the source address. In the case of static SAD routing, the multiple routing tables are configured manually with the desired static routes.

Regarding dynamic SAD routing, the support provided by Cisco routers is very limited. Because SAD routing is supported as a manually defined route-map, and because route-map definition is mainly a manual process performed by the router operator, Cisco routers cannot update the routing information (i.e. route-maps) involved in the SAD routing. This means, that neither the BGP nor the IGP case are supported by this router vendor.

Because Juniper routers support multiple parallel routing tables, the support for dynamic SAD routing is provided more naturally. In the case of BGP, it is needed that different routing tables are updated depending on the values of the community attribute contained in the BGP route. While this seems pretty straightforward, it is not currently supported

by Juniper routers because of the existent constraint that imposes that a given instance of a routing protocol can only update a single routing table, making not viable that the BGP instance can update different routing tables based on the value of the community attribute. Such limitation does not apply for the IGP case, since the considered approach proposes the usage of multiple instances of the IGP running simultaneously, one per source prefix involved, and that each instance of the IGP updates its corresponding routing table. This configuration is currently supported in Juniper routers for OSPFv2 and also for BGP. It should be noted that this approach can be used as a temporary solution for the BGP case i.e. running multiple instances of BGP in parallel while the community based approach is not available.

4. Resulting capabilities

4.1. Fault tolerance capabilities

As we have presented earlier, the main motivation for multihoming is an improvement of the fault tolerance capabilities of the Internet access. The enhanced fault tolerance features must allow the site to communicate in the event of an outage affecting one of the available Internet accesses. In this section we will see that the availability of the SAD routing in the multihomed site enables the provision of the required fault tolerance capabilities.

Since the basic assumption behind adopting SAD routing for multihoming support is that the source address is determined by the initiating host, and that each source address prefix determines an exit ISP, fault tolerance capabilities will be provided by the hosts themselves. As described in the Host Centric Approach [7], such mechanisms are based on the trial and error procedure. Considering that each source address available in a host is bound to an exit path, the host can try different exit paths by changing the source address. The main difference between the fault tolerance provision of the different approaches for SAD routing presented in this paper is how fast the host can learn that a destination address is unreachable through the selected path.

When external routes are static, the intra site routing system has no external reachability information, so the packet will be forwarded outside the site and only when it reaches routers that have richer knowledge about the topology, it will be possible to determine whether the requested destination is reachable through the selected path. In the worst case, the initiating host will timeout and will retry with a different path.

When the multihomed site runs BGP or an IGP with its providers, reachability information is available closer to the host, i.e. in the site's routers, so in some cases the unreachability feedback will be provided faster than in the general case, where unreachability information is learned through timeouts. So, the host will attempt to use one of its source addresses to reach a certain destination.

The packet will be routed through the generic routing domain to the SAD routing domain. Once there, the routers will determine whether the selected destination is reachable with the selected source address. This means that a route to the selected destination exists in the routing table associated with the selected source address prefix. The possible resulting behaviors are:

If the selected destination is reachable through the selected source address, then the packet is forwarded towards the site exit router that leads to the ISP corresponding to the source address prefix selected.

If the selected destination is not reachable through the selected source address, but it is reachable through an alternative source address, then the packet is discarded and an *ICMP Destination Unreachable* with *Code 5* which means *Source Address Failed Ingress Policy* [10] is sent back to the host. The information about the proper source address prefix can be included in this message, for instance in the source address of the ICMP message. The host will then retry using the suggested source address.

If the selected destination is unreachable, the packet is discarded and an ICMP destination unreachable is sent back to the host. In this case, the host may retry if an alternative destination address is available.

4.2. Traffic engineering (TE) capabilities

As a consequence of using multiple prefixes in multihomed sites in conjunction with SAD routing, the party selecting the address of the multihomed host to be used during the communication is the party that determines the ISP to be used for the packets involved in this communication. So, TE mechanisms will have to influence such selection. It must be noted, that addresses used in a communication are determined by the party initiating the communication, so in this environment, policy mechanisms will not affect incoming and outgoing traffic separately as in the IPv4 case, but they will affect packets belonging to externally initiated communication and packets belonging to internally initiated communications differently. This is the first difference with the IPv4 case.

4.2.1. TE for externally initiated communications

When a host outside the multihomed hosts attempts to initiate a communication with a host within the multihomed site, it first obtains the set of destination addresses, and then it selects one according to the Default Address Selection procedure [11]. It seems then that the only point where the multihomed site can express TE considerations is through the DNS server replies. The DNS server can be configured in order to modify the order of the addresses returned to express some form of TE constraint.

This mechanism can work fine to provide some form of load balancing and load sharing. The DNS server can be configured so that $x\%$ of the queries are replied with an address with prefix of ISPA first and the rest of the times

$(100-x)\%$ are replied with an address with prefix of ISPB first. In addition SRV [12] records can be used to provide enhanced capabilities by those applications that support them. When the host receives the list of addresses, it will process them according to RFC3484. If none of the rules described works, the list is unchanged and the first address received is tried first. Note that the list may be changed by the address selection algorithm because of the host policies.

4.2.2. TE for internally initiated communications

For internally initiated communications, the exit ISP is determined by the source address included in the initiating packet. This means that the source address selection mechanism [11] will determine the exit ISP. RFC 3484 defines a policy table that can be configured in order to express TE considerations. The policy table allows a fine-grained policy definition where a source address can be matched with a destination address/prefix, allowing most of the required policy configurations.

5. Conclusions

In this paper we have presented the case for the adoption of SAD routing in multihomed environments. The scalability limitations of the current multihoming solution based on the usage of Provider Independent addressing have been largely acknowledged by the Internet community, and only a new multihoming solution compatible with PA addressing will preserve IPv6 inter domain routing system scalability. However, the adoption of PA addressing in multihomed environments implies that multihomed sites need to internally configure as many prefixes as providers they multihome to, causing several difficulties, such as incompatibilities with ingress filtering, incapability to preserve established connections through outages and so on and so forth. This is basically due to the fact that when multiple PA prefixes are present in the multihomed site, the source address selection process determines the ISP to be used in the communication. This is so because in order to preserve ingress filtering compatibility, the packet has to be forwarded through the ISP that is compatible with the selected source address. Current destination address based routing does not take into account the source address of the packet, making it unsuitable to provide ingress filtering compatibility that is source address related. SAD routing is then the natural option to overcome the difficulties caused by ingress filtering. Moreover, once that SAD routing is available on the multihomed site, it is possible to obtain additional benefits such as fault tolerance and traffic engineering capabilities with a reduced complexity. SAD routing is not a new technology and it is available in some limited form in current router implementation, which facilitates its adoption and deployment. However, SAD routing currently is a special feature whose applicability was limited to very specific scenarios.

References

- [1] G. Huston, Commentary on inter-domain routing in the internet, RFC 3221, 2001.
- [2] Y. Rekhter, T. Li, A border gateway protocol 4 (BGP-4), RFC 1771, 1995.
- [3] I. Van Beijnum, BGP, Oreilly, 2002.
- [4] C. Labovitz, R. Wattenhofer, S. Venkatachary, A. Ahuja, The impact of internet policy and topology on delayed routing convergence Proceedings of INFOCOM 2001, Anchorage, Alaska, 2000.
- [5] J. Abley, B. Black, V. Gill, Goals for IPv6 site-multihoming architectures, RFC 3582, 2003.
- [6] P. Ferguson, D. Senie, Network ingress filtering: defeating denial of service attacks which employ IP source address spoofing, RFC 2827, 2000.
- [7] C. Huitema, R. Draves, M. Bagnulo, Host-Centric IPv6 Multihoming, Internet-Draft (Work in progress).
- [8] R. Chandra, P. Traina, T. Li, BGP communities attribute, RFC 1997, 1996.
- [9] R. Coltun, D. Ferguson, J. Moy, OSPF for IPv6, RFC 2740, 1999.
- [10] A. Conta, S. Deering, Internet control message protocol (ICMPv6) for the internet protocol version 6 (IPv6) specification, RFC 2463, 1998.
- [11] R. Draves, Default address selection for internet protocol version 6 (IPv6), RFC 3484, 2003.
- [12] A. Gulbrandsen, P. Vixie, L. Esibov, A DNS RR for specifying the location of services (DNS SRV), RFC 2782, 2003.