A Survey of IPv6 Site Multihoming Proposals

Pekka Savola
CSC - Scientific Computing Ltd
PL405, 02100 Espoo
Finland
psavola@funet.fi

Tim Chown
School of Electronics and Computer Science
University of Southampton
Southampton SO17 1BJ, United Kingdom
tjc@ecs.soton.ac.uk

Abstract— **Site multihoming is a method by which an Internet end-site, for example an enterprise network, may connect to multiple service providers simultaneously. There are many reasons why multihoming is desirable, e.g. service resilience, network load balancing or provider independence. In the IPv4 Internet, multihoming has been achieved by use of relatively simple techniques, including networks advertising their network prefixes – whether such prefixes are independent of the Internet Service Providers (ISPs) or not – to the Internet global routing infrastructure. With the introduction of IPv6 the vast increase in the number of potential site prefixes means that for scalable site multihoming we cannot repeat such IPv4 multihoming practices. Thus new IPv6 multihoming solutions are required. In this paper we present an overview of currently proposed solutions and explore the challenges and motivations of site multihoming. Such a review is timely because multihoming remains a key perceived obstacle to widespread IPv6 deployment in mission-critical environments.**

I. Introduction

Multihoming is the process of obtaining simultaneous IP connectivity from multiple ISPs, which may be done for a number of reasons, such as protection against failures. Site multihoming is a subset of that: the case where an end-site, for example an enterprise, becomes multihomed. "Multiconnecting" or "multi-attaching", on the other hand, refers to obtaining simultaneous IP connectivity from the same ISP.

We describe the background, motivations, challenges and problems with IPv4 multihoming and then look at the different proposals for IPv6 multihoming. Solutions currently used for IPv4 do not scale for use with IPv6. For mission-critical networks, the time taken by the Internet Engineering Task Force (IETF) multi6 working group (WG) to reach a partial consensus has represented a barrier to widespread commercial IPv6 adoption.

This paper offers an overview and categorization of the issues involved and the current solution space. Detailed analysis is beyond the scope of this review.

Throughout this paper, familiarity with addressing, routing, Border Gateway Protocol (BGP), etc. is assumed.

II. Site Multihoming

This section gives background why site multihoming is such a difficult problem.

A. Motivations

There are a number of motivations why a site might multihome, some of them a lot more obvious than the others. These are [1], [2]:

1) Independence: being able to switch ISPs easily, without renumbering; being seen as independent also often has some "status value".
2) Redundancy: being able to protect yourself from a number of problems affecting the site's usability or availability, such as fiber cuts, hardware or software problems, specific configuration mistakes, etc. – this is a generic motivation for increasing resiliency against failures.
3) Load sharing: being able to distribute the incoming and outgoing traffic among different links or operators.
4) Performance: some traffic may have different requirements (e.g., low delay, packet loss, or jitter) and one may wish to obtain high-quality connectivity for that; on the other hand, some other traffic may not have these requirements, and could be satisfied using a bulk operator.
5) Policy: some organizations (e.g., universities) may have policies regarding which kind of traffic (e.g., commercial vs research) is allowed by the upstream provider.

In most cases, the most important motivation is redundancy. Independence is often also very desirable because it eliminates the need for renumbering. The last three motivations are not as common as the practical scenarios where these are absolutely required and cannot be accomplished any other way are rather rare.

B. The Relation of Addressing, Routing and Multihoming

To obtain Internet connectivity, sites typically get a single physical connection and a "loaned" share of their ISP's IP address space for the duration of the contract. These are so-called Provider Assigned (PA) addresses.

There are also ways to get official or de-facto Provider Independent (PI) addresses. Historic assignments were such, while enterprises may now enlist as Local Internet Registries (LIR) to obtain address space allocations which equate to PI. They can then avoid renumbering when changing ISP, because they do not have to change to use the PA address space of their new ISP.

Sometimes sites also negotiate with (ie. pay) their ISP to be able to keep their addresses after changing ISPs.

For the addresses to be useful, they must be routable in the Internet. In the first case above, an ISP advertises its own address aggregate routes. In the second, the site advertises its own PI prefix. In the third, the site

8th International Conference on Telecommunications - ConTEL 2005
ISBN: 953-184-081-4, June 15-17, 2005, Zagreb, Croatia

punches "holes" in the ISP's aggregate by advertising more specific routes. All of these route advertisements must reach the global Internet routing infrastructure, often referred to as the Default Free Zone (DFZ).

So, there is a clear link between PI, or similarly achieved de-facto PI, addressing and multihoming. Availability of PI address space makes multihoming much simpler for the user; a customer can be reachable through multiple ISP links using a single block of PI space. However, multihoming comes with a cost. Increasing the DFZ routing table size places the cost on the ISPs and Internet infrastructure. That cost might be better placed with the customer wanting multihoming.

C. Challenges

Designing a good site multihoming solution has a number of challenges; these make it difficult to find an approach without significant drawbacks.

1) Sites Want to Avoid Service Provider Lock-in: Especially larger sites hold independence from their service providers in high esteem – they want to be able to change their ISPs with relatively little ease, without renumbering, to be able to obtain competitive pricing and services. They also wish to be removed from concerns over their ISP going out of business.

2) Fine-grained Traffic Engineering is Complicated: Outbound traffic engineering is a relatively simple process, but inbound traffic engineering is very complicated. That is, to be able to affect decisions made by any node in the Internet, one has to distribute the traffic engineering information throughout the Internet. About the only way at the moment to do that is to use BGP to advertise a route (often a more specific route) with intended visibility to steer the traffic.

3) Connection Survivability is Important: When an outage happens and a site has to fall back to IP connectivity from another provider, existing TCP connections, UDP "sessions", etc. should continue to work without being reset – which would happen if the IP addresses changed and the protocol suite did not offer connection survivability. This is particularly important for long-lived and site-internal sessions.

4) Network Renumbering is Painful: It takes a lot of work to change IP addresses in all the nodes at the site – and also those hosts which are not at the site which have been configured to use the site's IP addresses! Therefore networks typically want to use either provider independent addresses, or Network Address Translators (NATs) (where applicable) to avoid the biggest renumbering pains if they would have to switch ISPs. It is vital to keep renumbering as simple as possible, as the other alternative is provider-independent addresses which have scalability issues. For more information about renumbering procedures, see [3], [4].

This can be mitigated slightly by trying to move to the direction of adding layers of indirection: instead of using IP addresses directly, use the Domain Name System (DNS) or other mappings. Then the worst problem is just managing the mappings during a renumbering event [4]. Fundamentally, this is a usability issue: the user should not have to care about renumbering, and should not need to deal with IP addresses.

5) The Internet Routing Infrastructure Must be Scalable: All of these challenges could be satisfied by assigning every site provider-independent addresses, and having those advertised to the whole Internet through multiple providers. However, this would not scale for multiple reasons. Such updates require 1) processing power, 2) memory to hold the number of prefixes, and 3) sufficient link bandwidth for updates. Especially the first can be a problem even with high-end equipment, particularly if failures could come in bursts as well.

The scalability is probably the most significant challenge because it is every router on the Internet that has to face the scalability burden, not the multihomed site itself. Therefore the site has no clear incentive to use a scalable mechanism, and the ISPs may have to accept such a mechanism for competitive reasons – because if they don't, the customer will likely just find someone else who will; we'll also discuss this in Section III-B.

Let's try to analyze the scalability concern: [2] makes some rough estimates. We look at two cases: the scenario where every enterprise of at least a) 50 employees, or b) 500 employees would have a multihoming solution affecting the global routing infrastructure. Calculating with a population mass of 1000 million (only), and enterprise density of a) 1000 and b) 50 per million people ([2] justifies why these numbers are reasonable), we would have a) 1,000,000 or b) 50,000 multihomed sites.

Depending on the estimated error rates in the different components and systems in the network (see e.g., [2] for an approximation), this might result in the order of O(100,000) updates per day, with bursts up to O(100,000) simultaneous changes when a failure occurs somewhere in the network; that's quite a bit of computation and message propagation for router CPUs.

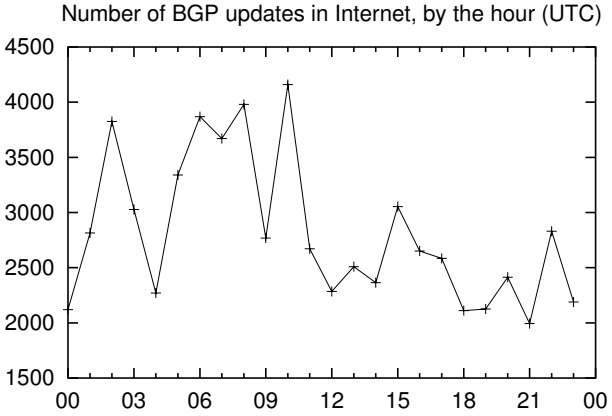

Fig. 1. Number of BGP Updates from Internet, by hour

Actually, this is probably an underestimation; figure 1 shows the measurements of the number of external BGP updates from the full Internet routes (around 140,000) entries as of April 2004. During 24 hours, there were

ConTEL 2005, ISBN: 953-184-081-4

about 64,000 updates (for about 10,000-20,000 sites), averaging 44 updates per minute[1]. As is expected, most instability occurs around the hours 02-10 (UTC), which seems to correspond to the maintenance windows after the office hours in North America. One may want to compare this to [5], [6]; in particular, in 2002, Sprint network reported higher churn for external BGP sessions, around 130 updates/minute [6]. Therefore, for 1,000,000 sites, even in the stable state, it would seem to be reasonable to estimate at least about a hundred-fold update rate (around 100-200 updates per second).

So, it seems relatively obvious that this would not be scalable especially if mechanisms requiring more processing for updates would be adopted (such as Secure BGP or Secure Origin BGP). Instead, a different protocol, with more powerful data aggregation or computational facilities – for example, calculating the equivalence classes of prefixes based on the ISPs' Autonomous System (AS) numbers – would be necessary [7]; even better would be avoiding unscalable mechanisms in the first place.

III. Related Work

A. Multihoming Techniques

We describe the two most prominent ways for site multihoming with IPv4. We will explore the IPv6 site multihoming landscape at length in Section IV.

1) Multihoming with BGP: The most visible and complete form of multihoming is done with BGP, with the following steps:

- obtaining your own IP address space, or getting permission to advertise a more specific route of an ISP's aggregate,
- obtaining an Autonomous System (AS) number,
- obtaining physical connectivity to at least two ISPs,
- setting up at least two routers at the site as border routers,
- establishing BGP sessions between the ISPs and the site border router routers, advertising the address space, and
- selecting which links will be used for the incoming/outgoing traffic by configuring BGP.

There are a few shortcuts one can make (e.g., using a more specific prefix rather than getting your own addresses), but this is the most complete procedure for BGP multihoming.

2) Multihoming with NAT or Load Balancers: A partial solution to multihoming is using NAT and deploying a specific device at the border which picks the right ISP to use without having to run a routing protocol on the customer link [2]. This does not give the full benefits of multihoming, e.g., connection survability is missing, but nonetheless the NAT solutions have been deployed at some smaller sites.

There is also a similar class of solutions in load balancers which do not necessarily use NAT; these can be used to provide reasonably high availability for e.g., server clusters. These do not provide connection survivability, but for certain kind of services, that may not be a big problem.

These forms of multihoming are invisible to the rest of the Internet, so their popularity is impossible to measure.

B. Market Considerations

Site multihoming using an architecturally unscalable method, BGP, is too cheap: practically it costs nothing. Most costs are incurred from the equipment, the physical connectivity, and the expertise (e.g. consultants or your own staff). Compared to that, fees to Regional Internet Registries (RIRs) – which are not even required for end-sites – are not significant: in the order of a thousand euros per year. Compared to that, the expenses required for redundancy, e.g., two access links to the ISPs and two border routers seem much more significant. This leads to the "grazing the commons" effect: everyone serious about Internet use wishes to use the most complete site multihoming solution, BGP, and likely does not want to settle for less.

To fix that, there would have to be a fee for the use of the global routing infrastructure (e.g., 5,000 euros/year, plus 500 euros/year for every originated prefix – collected by RIRs and donated in full to (say) the Internet Society), but such a thing would be an administrative impossibility; one would have to answer questions such as: How would this be observed? By whom? What constitutes "global"? What would prevent someone from neglecting the fee but still advertising?

The only hope would be (1) developing alternative mechanisms so that they are usable (as is being done with IPv6 now in the IETF shim6 WG), to satisfy also e.g. traffic engineering and renumbering requirements, and (2) raising (artificially) the fees for the resources such as AS numbers so that they would only be used by those who really do need them. (This has a number of problems of its own, though.)

Living in a market society, one could however argue that as things work already (with some definition of "work"), the prices are probably right. The problem comes from the current payment model: the number registries take no stance on the routability of a prefix, and consequently the money flow from the number resources does not extend to the ISPs in whose networks the numbers are expected to be used. The ISPs are expected to route everything, or figure out their own policies, typically based on their business requirements. This would be different if the number registries would have to pay to the ISPs for taking a number block into use, or pay to the ISPs (in general) for routing a prefix or an AS, or if the sites would be generally expected to explicitly pay for each prefix or AS, and that cost would be transferred in part to the peers and upstreams.

C. The Status and Future Requirements

The IPv4 DFZ contained about 150,000 routes as of February 2005, up by almost 50,000 during the last 3 years. If there were no more specific routes, the total number would be 73,000.

[1]This is just a measurement of the global routing table on one day, as heard from AS2603, not a long-term average.

To multihome, a network may obtain an address prefix and advertise it with an associated allocated Autonomous System Number (ASN). Multiple prefixes may be associated with a unique ASN.

Statistics as of February 2005 suggest there are around 35,000 assigned ASNs out of a 16-bit number space. Of these assignments, only 18,900 (54%) are visible in the global routing table, and of those, about 13,300 (70%) seem to be origin-only (typically enterprise) networks.

An important question is how multihoming trends will change. Driven by the vastly increased address space of IPv6 – which has 128-bit addresses where IPv4 has 32-bit addresses – the number of IP-enabled devices online will likely grow from hundreds of thousands to billions by 2010.

One barometer of likely requirements is the size of commercial enterprises which may seek multihoming, as noted in Section II-C.5; we'd have millions of sites if for example 50-person (or even smaller) sites were to want to multihome.

Currently, the 16-bit ASN size is a limiter on multihoming ISPs, but the focus of the multihoming problem lies where it is more difficult to achieve, with the end customer. While ASNs may be expanded to 32 bits, the more difficult issue is how, or whether, to offer IPv6 PI address space as a multihoming method in the future – in that light, expanding the ASN space might even be a harmful thing to do.

The shift to always-on services in Small-to-medium enterprises (SME) and Small Office/Home Office (SOHO) networks is likely to increase the dependency on services that run into such networks where IPv6 is deployed (where IPv4 NAT currently prohibits such services). With expectation of connectivity, multihoming for such networks may become more commonplace, further increasing the multihoming pressure.

For IPv6 multihoming, a solution is needed in the near term to remove any barrier, whether perceived or real, to widespread commercial IPv6 deployment. Many initial IPv6 service deployments have been in academic networks, including Abilene (US) and GÉANT (Europe), where multihoming is not such a critical issue.

IV. IPv6 Multihoming Approaches

There are many classes of potential IPv6 multihoming solutions. There is unlikely to be a single "one size fits all" solution, thus we may expect to see some combination of techniques applied.

1) *Host-centric solutions*. Here, the multihoming support is enabled in the communicating end systems.
2) *Two-space identifier/locator solutions*. In this class, instead of using addresses as combined locators and identifiers, the locator ("where you are attached to the network") and identifier ("who you are") elements are split. These solutions can be either host-centric or routing-oriented, or have elements from both.
3) *Network and routing-oriented solutions*. Such solutions place the emphasis for multihoming in the routing infrastructure, e.g. through new IPv6 routing headers or options, or by specifying new protocols to pass, exchange, translate or map network prefixes and addresses.
4) *Geographical addressing*. Such schemes provide address aggregation based on geography.
5) *Temporary solutions*. This worst-case solution accepts a trade-off of some "temporary" measure which may be used until a scalable solution is developed and proven.

In the following subsections, we discuss examples of each class, highlighting advantages and disadvantages of their approaches.

Due to the number of proposals, we've omitted adding explicit references to each. More information and the references are available in [2], [8].

In Table I we have summarized the main differences relating to the main challenges (Section II-C) IPv4 multihoming techniques (Section III-A) and the main classes of IPv6 proposals.

TABLE I
Comparison of Multihoming Approaches

	Independ.	Conn. Surv.	TE	Scalab.
IPv4 with BGP	yes	yes	yes	no
IPv4 with NAT	mostly yes	no	no	yes
Host-c. + id/loc	no	yes	no	yes
Geo-PI or routing	yes	yes	no	no

A. Host-centric Solutions

Host-centric solutions refer to several classes of solutions, where the hosts become slightly more aware of the network. This builds on obtaining multiple addresses for each node.

This approach puts more "intelligence" in the hosts, making code for the hosts more complex, but with the benefit of not having to alter router devices, or introducing new special middleboxes for multihoming. In this way, the applications should not have to worry about the connectivity, this should be handled by the network or transport layer – an Application Programming Interface (API) extension may be defined to signal certain types of failure to the application.

1) Multiple Addresses: Any network connected to multiple providers, receiving PA address from those providers, distributes those prefixes to all hosts that the site wishes to be multihomed to. This can be achieved by multiple Router Advertisements. Hosts may then have multiple globally routable IPv6 addresses.

In such cases, hosts will need to select which source and destination addresses to use when communicating with (possibly multihomed) peers. The Default Address Selection (RFC3484) algorithm defines how this may be done, on a longest-match prefix basis. However, there are many associated issues, such as ingress filtering – when a host multihomed to providers A and B selects a source address from provider A, and sends the packet towards

ConTEL 2005, ISBN: 953-184-081-4

provider B, provider B may reject the packet because the source address is not within its own address space.

To be able to influence or fine-tune the Default Address Selection algorithm in the hosts, a mechanism to distribute the policy information may have to be developed. This does not necessarily require a network protocol, though.

At a dual-stack site, where IPv4 and IPv6 connectivity may come from different ISPs, the choice of use of protocol is a multi-addressing issue. Here the address selection is typically determined by the A (IPv4) or AAAA (IPv6) DNS records returned, although the presence of a DNS record does not indicate connectivity per se.

2) Transport-layer Modifications: Connection survivability, that is, being able to recover and continue to use ongoing network connections when the IP address used changes, is often a desirable property of a multi-address solution.

A possibility is to use an alternative connection-oriented transport protocol, SCTP. SCTP passes initial endpoint addresses on startup, and then uses heartbeat messages to probe availability of each address combination. This allows SCTP to change the addresses being used. The drawbacks include the heartbeat overhead, issues for stateful firewalls where addresses change, handling of UDP, and the requirement to have SCTP enabled hosts and applications. SCTP may become popular for some specific purposes and applications, e.g. telephony, but probably will not become widely used outside such realms.

Similar but typically more simplified proposals to modify TCP have also been presented.

Requiring major TCP modifications and requiring that every other transport protocol is similarly modified has not generated much enthusiasm. Therefore, it seems to be better to cope with this particular issue on a more generic level by using identifier/locator separation solutions, described in section IV-B.

3) Use of Mobile IPv6: A host that is multihomed, and that comes online at one or more locations, could potentially use Mobile IPv6 protocols to handle the event if the locator addresses change.

There are a number of challenges with this approach. First, Home Agent deployment would be required for multihomed hosts. Second, such Home Agents would need to be multihomed for sufficient robustness, just shifting the problem around. Third, any host contacted would need to support some Correspondent Node functionality as full bidirectional tunneling is an unscalable approach. Fourth, any Binding Update message would need to be secured; with Mobile IPv6, this is done with return routability check – but when a network connectivity (through a network provider) has failed, this is no longer possible.

4) NAROS: Decoupling Traffic Engineering from Routing: De Launois et al have proposed an interesting idea [9] to address the traffic engineering part of the problem space. Hosts could implement a Name Address and Route system (NAROS) service, where the NAROS servers would tell the clients per destination which kind of addresses and routes they should use based on traffic engineering criteria.

The practical problem with this is that it requires implementation at every client and deployment of NAROS service support everywhere in the Internet, for guiding those clients which try to initiate communications toward the multihomed site. On the other hand, this kind of system could be very useful for the multihomed sites for helping the clients pick the right source and destination address pairs so that the return packets would travel through the desired paths.

B. Two-space Identifier/Locator Solutions

IP addresses are used for both locators of the nodes (in routing and forwarding), but also as identifiers of end-points (transport protocol bindings in hosts). These functions have traditionally been coupled as one for simplicity and security.

However, this model has problems when hosts have multiple IP addresses: a host may have multiple locators of different properties – one may work while the other may be suffering from network outage. For most purposes, it still has only one end-point identifier, and it would be very desirable that connections would not be tied to an interface or any particular address.

By splitting the problem into locator and identifier space, applications (and transport layer protocols) then handle identifiers, while locators become the responsibility of the network layer. Guaranteeing indentifier uniqueness is an important requirement, which may impact IPv6 stateless autoconfiguration. The separation of an address to a routing locator and a host identifier is by no means a trivial change, as that brings a large number of new security threats [10].

Almost 10 years ago people in the IETF felt that the locator-identifier solution approach would be worth serious exploration [11]. There has not been sufficient requirement going down that path, though – until now. As the current model has been rooted very deep in the Internet architecture, the separation is likely at the very least going to require some degree of application modifications and other changes.

The IETF multi6 WG [12] has seen a number of proposals for the separation, such as Weak Identifier Multihoming Protocol (WIMP), Multihoming without Identifiers (NOID), and Strong Identity Multihoming using 128 bit Identifiers (SIM). As far as we can tell, these have been abandoned when the work on the new solution, “shim6”, began in the new IETF shim6 WG. Therefore we do not list them for brevity.

1) Host Identity Protocol: As a two-space solution, the Host Identity Protocol (HIP) has the advantage of addressing implicitly the security of the locator-identifier mapping. It is not directly intended to be a multihoming solution, but can be used in scenarios employing multiple addresses.

HIP includes a new Host Identity namespace, and a new Host Identity layer, between the network and transport layers. The Host Identifier (HI) is cryptographic, being

the public key of an asymmetric key pair. The HI would usually be published in the HIP Rendezvous Service or the DNS. The packets in transit use the locator, but the endpoints use the hashes of the HIs at the transport layer. Security is established via a four-way handshake.

HIP does not require transport layer modifications, but does require changes in the IP layer in end systems. The encryption used may, as per any end-to-end encryption, cause problems for firewalls or middleboxes seeking to inspect the packet contents. A key issue is that identifier-locator mappings need to be securely managed; at some point in the future DNS Security may be the solution. At the time of writing, there are a number of experimental implementations and the base specification is expected to complete soon.

2) Site Multihoming by IPv6 Intermediation (shim6): The combined identifier/locator split proposal ("shim6") [13] was formed very recently as a result of the IETF multi6 design team work. It combines a number of interesting elements.

The proposal adds a shim layer just above the basic IP processing at the destination node. The applications use an upper-layer identifier (ULID) which is translated, if necessary, by the shim layer. The ULIDs may also be routable, which eases the use in case of more complex applications.

Use of Hash-based Addresses (HBA) allow one to embed information about the equivalent prefixes in the interface identifier of an address. This allows the redirection of a session from one address to another to be secured when connectivity fails.

The session survivability can be negotiated at the desired time even later on; there is no need for additional roundtrips before connections can be established.

There is no additional overhead for packets in general; only in some cases (e.g. negotiation) might there be an impact on the packet size.

The proposal, although still very much in progress, seems to be a good combination of the strengths of different identifier/locator split solutions. Obviously, it still does not solve certain other requirements such as traffic engineering, but looks like a very useful tool especially for smaller sites and the IPv6 protocol suite in general.

3) Location Independent Addressing: Location Independent Addressing for IPv6 (LIN6) aims to achieve the same as HIP by splitting the 128-bit IPv6 address into locator and identifier. LIN6 is primarily a mobility mechanism, but it includes multihoming as a byproduct, with mapping agents handling the locator-identifier associations. LIN6 is under commercial development, being subject to patent applications.

4) 8+8, "GSE" and revisions: Where HIP offers separation of namespace in the host, it is also possible to provide that separation in the network, by modifying locator information in network devices (routers).

The original 8+8 definition split the IPv6 address into a 64-bit globally unique identifier and a 64-bit locator. In an alternate addressing architecture for IPv6 ("GSE"), Site Border Routers (SBRs) rewrite the locator part of the IPv6 address as packets enter and leave the site network. The 8+8 scheme became 6+2+8, with 6 bytes for the locator, 2 bytes for site subnetting, and 8 bytes for identity. Site-local locators are used inside the network. As a packet exits a site, it has its site-local source address locator bytes rewritten to match the SBR locator, while packets entering the site have their destination address locator bytes rewritten to the site-local locator.

This mapping means that external routing policy can be applied at the edge. It also prevents forging of source addresses by hosts, as routers assert the locator point of injection. However, GSE does not consider failover scenarios, which is a fairly significant omission. GSE also requires a hierarchy for the locator addresses (to avoid potentially 2^{48} routes being injected into the IPv6 DFZ), which could include major ISPs. Finally, the system requires a secure method to look up locators and identifiers.

There have also been proposals for so-called "16+16", where the identifier is stored separately, e.g., in an extension header. This causes additional overhead and makes firewalling more complicated as the filters need to skip over such headers to check, e.g., TCP port numbers.

In summary, the network-driven approaches to identifier/locator separation have issues and seem to have been abandoned.

C. Network and Routing-related Solutions

Other network-oriented approaches have been proposed. These are designed in such a way that hosts are unaware of the multihoming support in the network, and need no code changes.

Some people have also proposed the temporary solutions listed under Section IV-E also as permanent solutions, with the "justification" that router processing, memory and storage capacities might grow faster than the use.

1) Multi Homing Aliasing Protocol: The Multi Homing Aliasing Protocol (MHAP) is a routing-oriented approach to multihoming. It places the "intelligence" in intermediate devices in the network. Multihomed traffic is transformed into single-homed traffic at a device (router) called the MHAP Client which is located close to the source, and then transformed back into multihomed traffic at a device (router) called the MHAP Endpoint at the endpoint site.

MHAP splits the IPv6 routing table into the (current) aggregated table, plus a new MHAP table which is based on PI and geographic blocks. The MHAP Clients communicate in turn with Rendezvous Points (RPs) that act as aggregators.

One of the drawbacks of MHAP is that all (current) single-homed hosts wishing to contact multi-homed targets need to communicate via an MHAP Client, somewhere upstream. This is a significant deployment "bootstrap" issue. Security considerations also need to be made, e.g. for DoS attacks on the RPs. The proposal seems to have been abandoned, but is listed here for completeness.

2) Multi-Connecting: Where a site multi-connects to a single ISP via different paths, the multihoming support can be facilitated by the ISP. Such a measure only protects against a certain subset of local failures, but retains the advantage that a single ISP address block can be used at the site. This is not a full form of multihoming, but deserves to be mentioned here.

D. Geographical Addressing

Due to the size of the IPv6 address space, it has been proposed that (part of) the address space be mapped (literally) to geographic locations. There are a couple of variants of this proposal.

Geographical addressing schemes allow some degree of aggregation within a city or similar area. Less specific geographic prefixes would be seen on regional, national or international interconnects.

The main drawback of these schemes is that an ISPs' topology and interconnections (and interconnection policies) do not follow geography. Therefore there would have to be lots of more specific routes, and is it not clear who would advertise the aggregate for a country because that Internet Exchange (IX) or ISP would attract traffic that they may receive no income for.

1) Addressing by Population: The first is the Geographically Aggregatable Provider Independent Address Space proposal (GAPI), which features 13 subcontinental allocations with /32 size prefixes allocated to cities and metro areas in countries. The scheme weights allocations by population densities.

2) Addressing by Coordinates: The second is the IPv6 Provider-Independent Global Unicast Address Format, which uses the WGS-84 standard latitude and longitude to derive address blocks that cover squares on the Earth's surface that are 6.4m wide. The scheme includes a 44-bit reference ID and a 16 bit SLA, to which the 64-bit host part of the address may be concatenated.

3) Addressing by Internet Exchanges: One can argue that geography meets routing at an exchange point. It may thus be appropriate to consider aggregation at an exchange (as per the obsoleted RFC 2374). In such a model customers would take address space from the exchange, rather than the ISP connecting to the exchange, making a local change of ISP a more simple task.

This does not solve the actual problem, because an Internet Exchange Point (IXP) assigning addresses to the sites would effectively be similar to sites multiconnecting to an ISP; the IXP would just be an ISP or an association of ISPs with a different name.

E. Temporary solutions

In the absence of an immediately deployable multihoming solution, a temporary measure could be offered until "proper" solution(s) emerge.

One possibility is to offer /48 or /32 size PI prefixes to any site, based on some qualifying measure (e.g., a certain number of employees or customers), and to time limit the allocations. At the time of writing, such a proposal has been made to ARIN, for sites that could qualify for an ASN.

Another is to use a simpler ASN-based approach, by just giving a /32 prefix to anyone with an ASN. This could create an ASN "landrush" however.

Finally, an IPv6 PI space could be given to anyone with a demonstrable existing IPv4 PI address block.

While none of these would be an immediate problem, as the number of IPv6 prefixes is only just over 1,000, the drawback with temporary solutions is that they become a future "swamp" from which it can be hard to step back. "Temporary" becomes permanent, and "must be justified" becomes "must be available to everyone without discrimination".

It may be difficult, as we have found with IPv4, to undo early "generous" or "temporary" allocations of address blocks.

V. Discussion

In Table II, we try to capture a very short summary of each proposed solution.

The multi-addressing approach can be deployed in a small site, with multiple exit routers. For this to be effective, there are some minimum requirements that should be handled. First, where outbound packets fail ingress filters, packets should be redirected to the right exit for the source address, e.g., through a tunnel. Where a SBR is down, internal prefixes associated with that router should be deprecated, such that hosts do not attempt to establish new connections through that link. At the same time, new connections inbound via that router could be "redirected" by removing the DNS entry based on that prefix for the internal hosts(s) (assuming the DNS operates on a very short time-to-live (TTL)).

Research in this area is still lacking. The most simple cases could start from multi-address solutions, applied to a single SBR, and then extended to multiple SBRs (which would require ingress filtering to be catered for).

Exploration of such techniques may give useful results in the short term for a specific set of smallish sites.

The shim6 effort also looks like a promising approach for connection survivability. It is in particular useful for those smaller and medium-sized sites which do not have complex traffic engineering requirements.

Meanwhile, very large enterprises could be given /32 prefixes, subject to conditions (e.g. number of employees), such that host-based solutions would not need to be deployed in networks with many dozens of thousands of hosts. The number of such allocations would not be large, but the drawback is that likely such policies would soon be extended for smaller and smaller enterprises at the price of scalability. So, this is a very "slippery slope" which may be better avoided if possible.

Supporting multiple provider-based addresses must be made simpler; avoiding renumbering is one of the most difficult-to-handle issues, thus documenting renumbering procedures [4] and making renumbering gradually easier is very worthwhile. Fortunately, there is already work in progress relating to this.

As a missing piece, research should be conducted on the traffic engineering requirements and methods which

would not require the use of BGP and subsequent pollution of the DFZ. For larger sites, this is possibly going to be the most important work of all.

TABLE II
SUMMARY OF SPECIFIC SOLUTION PROPOSALS

Approach	Short summary
Multiple Addresses	Used in identifier/locator solutions
Transport-layer Mods	Too difficult to change all the transports
Use of Mobile IPv6	Just shifts the problem around; abandoned
NAROS	Helps with scalable traffic engineering
WIMP, NOID, SIM	Abandoned for shim6
HIP	Requires too many changes for site mh
Shim6	Active work; no independence and TE
LIN6	Some deployment but no real adoption
8+8, GSE, etc.	Abandoned
MHAP	Abandoned
Multi-connecting	Gives only a degree of redundancy
Pop. Addressing	Abandoned
Geo. Addressing	Still being proposed but has issues
IX Addressing	Renaming "ISP" to "IX" doesn't help
"PI to everyone"	Unscalable
"PI to ASN holders"	Unscalable, ASN landrush problems

VI. CONCLUSIONS

Site multihoming with IPv4 is an issue which has generally been performed with a method that shows little care for Internet routing architecture. When considering how to deal with multihoming in IPv6, we must ensure that adopted solutions are scalable and architecturally sound.

This review illustrates the wide variety of proposed solutions – solutions used for IPv4 are not applicable or acceptable to IPv6. This variety, and the tardiness of forming an IETF consensus has meant that progress towards agreed solution(s) has been slow. Moreover, the solutions do not address the full spectrum of operational requirements, so more work is to be expected.

Even though the IPv6 deployment is still at a relatively early stage, and no immediate scalability problems are seen at the moment, it seems impossible or very difficult to withdraw "temporary" solutions when they start becoming unscalable – IPv4 is proof enough for that.

VII. ACKNOWLEDGEMENTS

The authors acknowledge the contributions and discussions of those in the IETF multi6 WG.

REFERENCES

[1] J. Abley, B. Black, and V. Gill, "Goals for IPv6 Site-Multihoming Architectures," RFC 3582, Aug. 2003.
[2] P. Savola, "Examining Site Multihoming in Finnish Networks," Master's thesis, Helsinki University of Technology, Finland, 2003. [Online]. Available: http://staff.csc.fi/psavola/di.ps
[3] H. Berkowitz, "Router Renumbering Guide," RFC 2072, Jan. 1997.
[4] F. Baker, E. Lear, and R. Droms, "Procedures for Renumbering an IPv6 Network without a Flag Day," draft-ietf-v6ops-renumbering-procedure-04.txt, Feb. 2005, work in progress.
[5] C. Labovitz, G. R. Malan, and F. Jahanian, "Internet Routing Instability," *IEEE/ACM Transactions on Networking*, vol. 6, no. 5, pp. 515–528, 1998.
[6] C. Chuah, S. Bhattacharyya, and C. Diot, "Measuring I-BGP Updates and Their Impact on Traffic," Sprint ATL Technical Report TR02-ATL-051099, Oct 1999.
[7] A. Broido and k claffy, "Analysis of Route Views BGP data: policy atoms," *Proceedings of Network-related data management (NRDM) workshop*, 2001.
[8] M. Dunmore, "Evaluation of Multihoming Solutions," Feb. 2005. [Online]. Available: http://www.6net.org/publications/deliverables/D4.5.3.pdf
[9] C. de Launois, O. Bonaventure, and M. Lobelle, "The NAROS Approach for IPv6 Multihoming with Traffic Engineering." [Online]. Available: citeseer.ist.psu.edu/delaunois03naros.html
[10] E. Nordmark and T. Li, "Threats Relating to IPv6 Multihoming Solutions," draft-ietf-multi6-multihoming-threats-03.txt, Jan. 2005, work in progress.
[11] B. Carpenter, J. Crowcroft, and Y. Rekhter, "IPv4 Address Behaviour Today," RFC 2101, Feb. 1997.
[12] IETF. Site Multihoming in IPv6 (multi6) charter. [Online]. Available: http://www.ietf.org/html.charters/multi6-charter.html
[13] E. Nordmark and M. Bagnulo, "Multihoming L3 Shim Approach," draft-ietf-multi6-l3shim-00.txt, Jan. 2005, work in progress.