

SITE-CONTROLLED SECURE MULTI-HOMING AND TRAFFIC ENGINEERING FOR IP

Randall Atkinson
Extreme Networks
RTP, NC, USA

Saleem Bhatti
University of St Andrews
St Andrews, UK

Steve Hailes
University College London
London, UK

ABSTRACT

Site multi-homing is an important capability in modern military networks. Resilience of a site is greatly enhanced when it has multiple upstream connections to the Global Information Grid, including the global Internet. Similarly, the ability to provide traffic engineering for a site can be important in reducing delays and packet loss over low-bandwidth and/or high-delay uplinks. Current approaches to site multi-homing and site traffic engineering (a) require assistance from a trusted network service provider; (b) inject significant additional routing information into the global Internet routing system. This approach reduces flexibility, does not scale and is a widespread concern today. The proposed Identifier-Locator Network Protocol (ILNP) offers backward compatible extensions for IPv6 to enable a site to (a) use multiple routing prefixes concurrently, without needing to advertise these more-specific site prefixes upstream to the site's service providers; (b) enables edge-site controlled traffic engineering and localised addressing, without breaking end-to-end connectivity. This feature combination provides both multi-homing and traffic engineering capabilities without any adverse impact on the routing system and does not require anything more than unicast routing capability in the provider network. ILNP enables concurrent multi-path transmission for a flow, without requiring multicast routing, to increase flow resilience to path interruptions. This technique has a secondary security benefit of reducing the risk of an adversary successfully blocking an ILNP flow via a Denial-of-Service attack on any single path or single link.

I. INTRODUCTION

Today, military networks need maximum flexibility and advanced capabilities in order to deliver different mission solutions. We have outlined previously a proposal for provision of a harmonised set of capabilities for site multi-homing, traffic engineering, end-to-end security, and support for mobile systems and networks [1]. In this paper, we explain in detail our new approach to multi-homing (MH) and to traffic engineering (TE). Our approach moves decisions about the choice of outbound links to the edge-network, and does not require additional routing

information to be introduced into the provider network. By decentralising TE and MH decisions to the end site network, these deployed networks have improved scalability and resilience, and could exercise control from the site itself, e.g. to counter the effects of jamming or link faults.

In our discussion, we chose the abstraction in Figure 1 for the site network because it maps to many real scenarios, e.g. a warship (mobile network) with multiple satellite uplinks; an infantry platoon (mobile network) with multiple radio links; or a military base with multiple, redundant external links. We show only two external links, for simplicity, but a larger number of external links can also be supported. (The 'coordination protocol' is not considered in this paper, but commercial systems exist today that provide such functionality and could be adapted for use.)




Fig. 1. General scenario: an example site network, with two site border routers (SBRs), each providing access to an independent external link.

Our approach enhances the IPv6 addressing format in a fully backwards-compatible manner. This ensures that existing deployed IPv6 routers can be used unchanged. By altering only the semantics of the IPv6 address bits at the edge of the network – at the site-border router (SBR) and within the site itself – we achieve a self-contained and incrementally-deployable improvement in multi-homing and traffic-engineering capability that can be used to support a variety of missions. Indeed, deployment of our proposed mechanism is invisible to well-behaved applications.

In this paper, Section II describes the motivation for our research, Section III provides an overview of the Identifier-Locator Network Protocol (ILNP), Section IV describes our solution for both Site Multi-Homing and Host Multi-

Homing, Section V describes our solution for Traffic Engineering, and Section VI discusses the security considerations for our approaches.

II. RATIONALE AND MOTIVATION

Mechanisms for multi-homing and traffic engineering are already available in commercial off-the-shelf (COTS) offerings, and it can be argued that such offerings are mature and already deployed. So, it is important to first state the motivation for revisiting these issues, both in general and also specifically in the context of military networks, as well as providing the rationale for our approach.

A. Multi-homing and Traffic Engineering Today

Today, site multi-homing for IP-based networks is provided, essentially, by adding additional information to routing tables all over the globe. [2] A common scenario is one that is described in Figure 2. In this example, the site's network is provisioned for external connectivity via two Internet Service Providers, ISP1 and ISP2. In this example, the site network might elect to use a provider-independent prefix or a provider-aggregatable prefix delegated from ISP1 (P_1) and/or from ISP2 (P_2).




Fig. 2. General multi-homing scenario: our example site network, with two site border routers (SBRs), each having provision through separate ISPs. We assume that, in this case, the site network has two routing prefixes, P_1 and P_2 .

Considering only P_1 , the value needs to be advertised separately to each upstream ISP. In turn, each upstream ISP needs to advertise that site-specific longer routing prefix into the *Default Free Zone (DFZ)* of backbone routers worldwide. While a given router typically will only have one entry for P_1 in its Forwarding Table, that router will need to have all advertisements for P_1 in its Routing Table.

Because IP routing relies upon the *Longest Prefix Match (LPM)* algorithm to select a path, a longer prefix delegated by ISP1 to the site must be carried separately by ISP1 and cannot be aggregated underneath a shorter IP routing prefix that belongs to ISP1.

So, for multiple site prefixes, N_p , and multiple upstream ISPs, N_I , the additional routing state due to site multi-homing is now $O(N_p \cdot N_I)$. This is consistent with prior BGP analysis done by others. [3]

Furthermore, because the prefix P_1 needs to be advertised to both ISPs, and that prefix might not be part of either ISP1's or ISP2's assigned address space, we need special co-operation, including manual administrative intervention, from ISP1 and ISP2 in order to make multi-homing work. This reduces flexibility in, for example, providing multi-homing for mobile networks (e.g. an aircraft carrier).

For traffic engineering, the situation is more complex and relies even more heavily on the co-operation of, and correct, timely configuration from, the ISP. Using approaches such as Multi-Protocol Label Switching (MPLS) [4], the site and the provider must agree *a priori* on policy mechanisms which include traffic descriptors that will allow the identification and correct handling of specific packets that will be offered differently provisioned paths. Again, there is heavy reliance on the ISP in order to provision this service within the provider's network.

It is likely that end-to-end (e.g. site-to-site, or host-to-host) security mechanisms, such as the military *High Assurance IP Encryptor (HAIPE)*, will be used to protect sensitive information, as a site usually cannot trust a third-party ISP. Within military-operated network segments, such as a tactical radio link, link encryption might also be used to provide additional protection (e.g. to reduce vulnerability of link control protocols to external attack).

Meanwhile, it remains important for military networks to have both MH and TE capabilities with high confidence in the provisioned IP service quality: at present, a site has little control over such functionality. Indeed, at present it is very difficult to provide both MH and site-selected TE. This complexity impedes a military site from enabling the flexible and adaptable connectivity that it needs. For example, providing dynamic network mobility for the site (e.g. a ship); or allowing easy migration of the site's external connectivity to a different upstream provider's network; or integrating additional external links for TE or MH capability.

B. Moving Control to the Site Network

Ideally, a military site could change its external connectivity dynamically, according to local site policy, configuring both TE mechanisms and site MH mechanisms. Today this is not possible. However, we believe this is a desirable and achievable goal.

Firstly, we note that the current approach, which adds routing state to the worldwide inter-domain routing tables (DFZ) does not scale and is not sustainable. As noted in [3] and [5], routing table growth is a growing concern. Indeed, quoting from RFC-4984, “... *the clear, highest-priority takeaway from the workshop is the need to devise a scalable routing and addressing system, one that is scalable in the face of multihoming, and that facilitates a wide spectrum of traffic engineering (TE) requirements.*”. [6]

Secondly, we note the previous observation about communications security: in order to preserve the integrity of the use of end-to-end cryptographic techniques, e.g. use of HAIPE, it makes sense to apply those security functions on an end-to-end basis (either site-to-site or host-to-host). Indeed, some have suggested that rather than securing routing protocols, which requires co-operation from (and complete trust of) the service provider, it is more important to simply secure the data delivery to a suitable level of protection [7], using end-to-end techniques. We believe the principle of site/edge control should also apply to MH and TE capabilities.

For example, in [8], a mechanism is proposed that allows multi-path routing, and so a limited form of traffic engineering, without the reliance on external service provider co-operation, but by use of pre-defined relay-points that together form an overlay network for allowing multiple traffic paths, whilst not adding any routing state overhead to the DFZ. These relays can be selected from the site networks. However, the relay points now become a performance bottleneck, an additional point of failure, and point of attack for anyone wishing to disrupt the service.

Also, if we consider the current organic growth of the Internet, there are some key technical features of the current topology and usage of the network that mean that use of multi-homing, and multi-path traffic for traffic engineering is indeed a viable option today. A good discussion of the issues is presented in [9] and we summarise here:

- *Existing path redundancy:* In [10], it is stated that empirical studies show that 90% of point-of-presence (PoP) pairs have four redundant (link-disjoint) paths between them. This path diversity is under-exploited in current IP deployments because of limitations in the current MH and TE mechanisms.
- *Inflexible routing dynamics:* In [11], again from empirical studies, it is estimated that 30%-80% of the time, a lower-delay, or lower-loss path exists for traffic, but it is not used. That is, better paths exist but are under-utilised due to current routing dynamics.
- *Multi-path provisioning in existing networks:* In [12],

measurement studies on a real provider backbone network show that many network operators already apply tuning mechanisms to link weights in order to provide equal cost paths. That is, many operators are already trying to load-balance traffic.

So, overall, an end-to-end approach that empowers site/edge networks to control the external network paths appears to be both feasible and strongly desirable.

III. OVERVIEW OF ILNP

In this section we present a brief overview of our proposed enhancements to the Internet Architecture, and also specifically to IPv6. We use the term *Identifier-Locator Network Protocol for IPv6 (ILNPv6)* to refer to our proposal, as it can be engineered as enhancements to IPv6. [1], [13]

A. Naming Problems in IP today

In our discussion below, we use the term *name* in a very general sense, to refer to any label that is attached to a network object. A summary is given in Table I.

TABLE I
TERMINOLOGY USED IN THIS PAPER

Term	DNS Record	Definition
Address	AAAA, A	Name used both for locating and identifying a network entity
Locator	L	Name that locates, topologically, a sub-network
Identifier	I	Name that identifies a node, within the scope of a given locator

TABLE II
USE OF NAMES IN ILNP AND IP

Protocol layer	ILNP	IP
Application	FQDN	FQDN, IP address
Transport	Identifier, <i>I</i>	IP address
Network	Locator, <i>L</i>	IP address
Link	MAC address	MAC address

It is important to recognise that the IP address is currently used for two quite different functions – as a *locator* for naming a specific interface on a node (or a set of node interfaces on a common subnetwork), and as an *identifier* for naming the node itself (by virtue of the binding of the address bits to one of the node interfaces). The overloaded semantics of the IP address causes entanglement across these functions and across protocol layers. The current use of the IP address is within applications, within the transport protocols (e.g. within the TCP connection state and pseudo-header checksum), and also within the network

layer to route packets between the end nodes – see Table II. Again, quoting from RFC-4984, “... the so-called ‘locator/identifier’ overload of the IP address semantics is one of the causes of the routing scalability problem as we see today. Thus, a ‘split’ seems necessary to scale the routing system ...” [6].

B. Naming Enhancements

We replace the concept of the *address* with the concepts of an *Identifier* combined with a *Locator*. The *Locator* names an IP (sub)network: this is used only in routing, and not by the upper layers (e.g. TCP or UDP). The *Identifier* is only used for node identity (e.g. by TCP or UDP in their pseudo-header checksums).¹

The idea of an *Identifier/Locator* split is not a new idea, but our particular approach is new, and is specified in more detail than preceding proposals. [14]–[16] We believe that applications should use fully-qualified domain names (FQDNs) and not IP addresses directly, wherever possible, and this is consistent with the recommendations in RFC1958 [17]. A summary of the difference between the use of names in IP (v4 and v6) and the use in ILNP is given in Table II.

C. IPv6 Enhancements

While our approach above might seem abstract, we are implementing ILNP as an extension to IPv6, which we call *ILNPv6*. The syntactic similarities between the IPv6 packet header and the ILNPv6 packet header are deliberate. Essentially, the IPv6 address is already broken into two separate components: in ILNPv6, we simply term the top 64 bits as the *Locator* (L), and the bottom 64 bits as the *Identifier* (I), as shown in Figure 3. Significantly, the only difference in address usage occurs at the end nodes: in IPv6 the lower 64 bits act as an *Interface Identifier*, and the whole 128-bit address is bound directly to a given interface, while in ILNPv6, the lower 64 bits is used as a *Node Identifier* (I), and is *not* bound to an interface.

An end-system may use multiple I values and multiple L values simultaneously. For the duration of a given ILNP session, its I value should remain constant. For practical reasons, the Identifier is normally formed from one of the MAC addresses associated with the node. This is represented in the IEEE’s EUI-64 syntax, and so is very likely to be globally unique as well. Strictly, the I value needs to be

¹This will be implemented such that the BSD Sockets API hides these details from applications; hence, existing applications generally should work without modification. To simplify the creation of new applications, we propose the creation of an additional, more abstract, API.




Fig. 3. IPv6 address format (from RFC-3587 [18]) as used in ILNPv6. The top 64 bits retain the same semantics as in IPv6, i.e. that of a routing prefix, naming an IPv6 network. The lower 64 bits are now used as *node* identifier rather than an *interface* identifier. However, the lower 64 bits are not used for routing in the core network, so ILNPv6 has no impact on today’s IPv6 deployment.

unique only within the scope of the L value with which it is used. However, for practical purposes, having an I value that is likely to be globally unique is very useful, and allows us to dispense with IPv6 Duplicate Address Detection (DAD), which in turn greatly reduces the time required for a node to execute a location change.

Our definition for the ILNPv6 Locator is consistent with the IPv6 Addressing Architecture [19], specifically with section 2.5.4, which states that the sum of bits in the global routing prefix and the sub-net ID is 64 bits. Current IPv6 address allocation practices provide sites with IPv6 address blocks that are 48-bits long, which leaves 16 bits for intra-site sub-netting. As the ILNPv6 Locator is the same as an IPv6 routing prefix, ILNPv6 packets can travel across existing deployed IPv6 backbones. However, the host’s IPv6 stack has to be enhanced to enable ILNPv6 on that host (i.e. to deal with Node Identifier values). ILNPv6 Neighbour Discovery (ND) still uses the full 128-bits of the combined L:I value. So IPv6 ND also can be used without change. *In short, already deployed IPv6 routers will support ILNPv6 without any changes, and ILNPv6 can be deployed incrementally on the same network as IPv6.*

D. DNS Enhancements

To enable ILNPv6, several new DNS resource records are needed. We add the *I* record, which contains the unsigned 64-bit Identifier associated with a domain name. Similarly, the *L* record contains an unsigned 64-bit Locator associated with a domain name. As a node might have multiple Identifiers and multiple Locators, a given domain name also might have multiple *I* and multiple *L* records. The combination of a given *L* record and an associated *I* record is, effectively, equivalent to the current IPv6 address.

Reverse lookups can be done as today with IPv6. As a performance optimisation, we also have a pair of new DNS records that could be used for reverse lookups. The *PTRL* record names an authoritative DNS server for an

ILNPv6 sub-network, while the *PTRI* record is used to obtain the name of a node using a given Identifier on a given sub-network. This usage enables *PTRL* records to be cached, which is beneficial if performing reverse lookups for multiple nodes on the same sub-network.

As a separate performance enhancement for managing site networks, we also introduce the *Locator Pointer (LP)* record. This record names an *L* record. Nodes that are attached to a site network (which could be a mobile network, for example) would typically have an *LP* record that pointed to the *L* record of that site network. So when the site network moves its point of Internet connection, only the network's own *L* record needs to be updated.

The existing Secure Dynamic DNS Update standard [20] would permit a mobile or multi-homed node (or whole network) to update its *L* record(s) when the node moves or its upstream connectivity changes (e.g. due to a link fault). Widely used systems software, such as Microsoft Windows and the BIND software used with UNIX, already include support for Secure Dynamic DNS Update. [21] Simulation results indicate that the existing IPv4 Internet does not cache *A* or *PTR* records nearly as well as commonly thought. [22] So wider use of existing dynamic DNS standards ought not have adverse operational consequences.

Separately, the DNS enhancements for ILNPv6 do not change the fundamental operation of the Domain Name System (DNS). So, the existing DNS Security (DNSsec) standards [23] can be used unchanged to authenticate these new DNS record types, and our proposed enhancements do not create any new security risks.

IV. MULTI-HOMING WITH ILNP

With ILNP, the new DNS *L* or *LP* records are used to advertise the current reachability for a node or site. New correspondents perform a DNS lookup, as at present, to determine how to send packets initially to the target node(s). Whenever a node's currently valid Locator(s) change, the node sends *ICMP Locator Update (LU)* control messages to its existing correspondents. (LU message are very similar to IPv6 Binding Update (BU) messages.) These messages can be authenticated either cryptographically using the IP Authentication Header, or non-cryptographically using a new ILNP Nonce, as appropriate for the node's threat environment. The correspondent receives this update, validates it, and then begins using the new Locator(s) to send packets to the original node.

A. IPv6 Site Multi-Homing

Using Figure 1, consider an IPv6 site network with two links to upstream Internet Service Providers (ISPs). Many sites are multi-homed today, so the situation shown in this diagram is quite common. A typical multi-homed site uses a single routing prefix delegated by one of its upstream ISPs. Because IP routing uses the Longest Prefix Match (LPM) algorithm to select the best path to a destination, a multi-homed site must advertise the same IP routing prefixes to each of its upstream providers. Therefore, each SBR has to advertise *all* of its IP routing prefix(es) on *each* upstream link. So the current approach to site multi-homing requires significant administrative effort (e.g. to coordinate the routing prefix advertisements among the site and its upstream providers and configure all of the site border routers appropriately), and also requires multiple site-specific routing prefix entries in the global routing table for each multi-homed site. More and more sites have become multi-homed this decade because of increased dependence on the Internet and related concerns about Internet availability. The current approach to site multi-homing is causing geometric growth rates in the global routing table, which is a major concern for the Internet community. [6]

B. ILNP Multi-Homing Concept

With ILNPv6, more-specific prefixes for a multi-homed site are not advertised globally and no special multi-homing coordination with the upstream ISPs is needed. As the Locator (routing prefix) for ILNPv6 is not part of the transport layer protocol state, any IPv6 routing prefix advertised by the provider from its existing aggregateable address space can be used. So, the routing table scalability issues of the current approach (whether IPv4 or IPv6) are eliminated.

With ILNPv6, each upstream ISP delegates a portion of that ISP's own IPv6 address space to the multi-homed site, i.e. one or more prefixes. Typically, each routing prefix delegated to a particular site will be of equal length. However, each upstream ISP need only advertise a single aggregated prefix globally covering all of its end sites.

In our example of Figure 4, ISP1 delegates a routing prefix L_1 to the site, while the ISP2 delegates a routing prefix L_2 to the site. The site uses Locator values L_1 and L_2 , respectively on external link 1 to ISP1 and external link 2 to ISP2. As the Locator values are not part of the transport protocol state, we can use both Locator values concurrently with a single transport session.

Each host picks both a Source Locator for itself and uses




Fig. 4. ILNP multi-homing scenario: our example site network, with two site border routers (SBRs), each having provision through separate ISPs. The site obtains Locators as required from each ISP: L_1 from ISP1’s address pool and L_2 from ISP2’s address pool. Both Locators can be used simultaneously and require *no* additional routing state whatsoever.

a Destination Locator for its correspondent, and includes both values, along with appropriate Identifier values, in the ILNPv6 packets that it originates. The destination Locator and Identifier values would, of course, be resolved from DNS. These packets are then routed normally within the site until they arrive at one of the site’s border routers.

With ILNP, the end hosts inside the multi-homed site may participate in the multi-homing, but are not required either to participate in or to be aware of the multi-homing. In the next two sections, first host multi-homing and then site multi-homing are explained in more detail.

C. ILNP Host Multi-Homing

With ILNP, the site might choose to use each externally delegated routing prefix within the site concurrently. So an ILNP node might have multiple valid Locator values on a single network interface, and these may be used concurrently for a single ILNP session. IPv6 Router Advertisements already permit multiple Locators to be advertised on a single subnetwork.

If one of the routing prefixes delegated to the site is no longer valid, perhaps because a fibre cut has eliminated the link to the corresponding upstream ISP or because of a change in ISP service agreements, then the site will cease to advertise the corresponding Locator (i.e. that Locator value will no longer be included in IPv6 Router Advertisements).

A host, upon learning that the set of valid Locator values has changed, sends an *ICMP Locator Update* message to all of its correspondents, and also updates the valid set of Locators for itself in the DNS. These control messages only include the Locator values valid at that time. Of course, these control messages are authenticated for security reasons. Existing correspondents will update the set of valid Locator values for that host after receiving and successfully authenticating that control message. New correspondents

will discover the Locator values for that host from the (now updated) DNS, as usual.

This approach permits a host to control its own multi-homing.

D. ILNP Site Multi-Homing

Alternately, ILNP enables an intelligent Site Border Router (SBR) to provide site-wide multi-homing capabilities on behalf of all of the nodes inside the site. This approach is particularly useful if the site has deployed localised addressing, (e.g. an IPv6 unique locally assigned (ULA) routing prefix, $fc00::/7$ [24]) internally, rather than deploying global addressing internally. With this approach, the multi-homing is invisible to the hosts inside the site.

In this case, upon receiving a packet from inside the site for egress, the SBR will then apply any local policy about which upstream link should be used to forward the packet to its destination. The Source Locator field in the packet will be re-written by the SBR to be consistent with the upstream ISP’s delegated routing prefix. The Locator modification is critical, because prefix filters are widely deployed in the global Internet to reduce the threat of forged malicious IP packets. [25]

If we consider Localised Addressing, with the Locator value L_L used within the site network only, then all packets within the site will have TCP and IP state as shown in tuple (1), which represents a packet.

$$\langle TCP : I_S, P_S, I_D, P_D \rangle \langle ILNP : L_L, L_D \rangle \quad (1)$$

here, I and P denote Identifier and port number, respectively, and subscripts S and D denote, respectively, Source and Destination.

A packet that is to egress SBR1 will have the session state given in tuple (2) and packets to egress SBR2 will have session state as in tuple (3).

$$\langle TCP : I_S, P_S, I_D, P_D \rangle \langle ILNP : L_1, L_D \rangle \quad (2)$$

$$\langle TCP : I_S, P_S, I_D, P_D \rangle \langle ILNP : L_2, L_D \rangle \quad (3)$$

This is similar to a Network Address Translation (NAT) function, but unlike NAT for IPv4 or IPv6, rewriting Locator values has no impact on transport layer state, so the end-to-end integrity of the TCP connection is preserved. So, ILNPv6 can provide multi-homing to the site-network without requiring involvement by nodes inside that site.

Whilst ILNPv6 does not need to use the SBR locator rewriting to support multi-homing, it provides an engineering optimisation and a good point for network management.

All external traffic passes through the SBRs, thus making it easy for the SBRs to maintain per-session flow state, including the set of remote correspondents, and the Nonce values (see Section VI-C) for each current (or recent) ILNP session. If the set of valid Locators for the site changes, then these intelligent SBRs can send valid proxy *ICMP Locator Update* messages to the correspondents, as required.

V. TRAFFIC ENGINEERING WITH ILNP

The ILNP approach to site traffic engineering (TE) exploits the ability to use multiple Locator values and multiple uplinks. Today's policy-based mechanisms for site TE can be used to filter flows (e.g. based on network layer or transport layer headers) and associate a TE policy with each flow, as required, and the selection of the correct egress interface. The same approach can be used to provide obfuscation for the interior topology of a site. This last capability is commonly desired by Internet-connected end sites that have high threat profiles, such as military or homeland security sites. For ILNPv6, policies can use node identity regardless of location, making it easier to configure and maintain TE policy. However, Locator values could be used to give conditional policy, if required.

A. ILNP Site Traffic Engineering

With ILNP, Site Border Routers (SBRs) are permitted to rewrite either the Source Locator, the Destination Locator, or both. This is done after the SBR selects the egress interface for a packet, but before forwarding the packet out that interface. For traffic engineering, the SBRs apply locally configured policy as part of the outbound path selection (and hence egress interface selection) process.

If there are multiple SBRs in use at a single site, those SBRs will need to share session state among themselves. This is the same issue that arises for multi-homed sites with a firewall at each border. That distributed firewall synchronisation problem is already solved in commercially available products; the same approach can be applied here to synchronise session state among the set of SBRs.

For example, let us assume that, in order to exploit multi-path capability, our site network chooses to use two prefixes/Locators from ISP1 (L_1 and L_2), and two prefixes/Locators from ISP2 (L_3 and L_4), as shown in Figure 5. As IPv6 address allocation policies provide a /48 routing prefix for each end site or end organisation, 16 bits are




Fig. 5. ILNP traffic engineering scenario: our example site network, with two site border routers (SBRs), each having provision through separate ISPs. The site obtains Locators as required from the ISPs. L_1 and L_2 are both from ISP1's pool and L_3 and L_4 are both from ISP2's pool. All four Locators can be used simultaneously and require *no* additional routing state whatsoever.

available for local topology information. This also means that delegated routing prefixes normally will have equal length. As it is easily possible for the SBR to have access to some or all inter-domain routing data (e.g. via BGP), it can obtain knowledge of the network paths that might be followed by using these particular routing-prefix/Locator values.

Local policy might require that traffic to some destination site be sent via ISP 1, rather than via ISP 2. The policy might be driven by desires to reduce cost (e.g. ISP 2 charges more than ISP 1), by quality-of-service concerns (e.g. link 1 to ISP 1 has more capacity or lower latency to the destination), by trust, or by some other factor.

Again, if we assume that the SBR has knowledge that the destination site is also using multiple Locator values (e.g. by snooping on DNS packets), in turn, that could be used to drive rewriting of the packet's Destination Locator value from the original value to a new equivalent value that uses a different path to the destination site than the original value would have used. So, in our example of Figure 5, by use of four Locators upstream, and knowledge of Locator values for remote sites, good path diversity could be achieved. Recalling the discussion of Section II-B, this effectively gives much greater ability to exploit different network paths from the end site network.

A primary driver for Source Locator rewriting is the widespread deployment of anti-forgery source filtering, both at end sites and within service provider edge routers. [25]. Such filters will cause a packet to be dropped if the Source Locator is inconsistent with the upstream ISP that is carrying the packet. So the SBR must ensure that the Source Locator of the packet sent out the egress interface is consistent with the Locator prefix that has been delegated by that upstream ISP to the end site that contains the SBR.

For example, consider two packet flows as in tuples (4) and (5). Each of these flows is from a separate host (subscripts a and b) in the site network of Figure 5, using a local Locator value L_L , and have separate destination endpoints, identified by subscripts J and U , respectively.

$$\langle TCP : I_a, P_a, I_J, P_J \rangle \langle ILNP : L_L, L_J \rangle \quad (4)$$

$$\langle TCP : I_b, P_b, I_U, P_U \rangle \langle ILNP : L_L, L_U \rangle \quad (5)$$

As these traverse the SBR, an internal policy decides that the first flow should traverse link 2, using Locator L_2 , as shown in tuple (6), and the second flow should traverse link 1, using Locator L_1 , as shown in tuple (7):

$$\langle TCP : I_a, P_a, I_J, P_J \rangle \langle ILNP : L_2, L_J \rangle \quad (6)$$

$$\langle TCP : I_b, P_b, I_U, P_U \rangle \langle ILNP : L_1, L_U \rangle \quad (7)$$

So, the SBR simply rewrites the local Locator value, L_L as required.

This approach to traffic engineering is complementary to existing mechanisms, such as MPLS. So, the site might make a TE decision about which upstream ISP to use for a particular outbound packet, but the upstream ISP can still perform its own TE within its portion of the network (e.g. using MPLS TE).

VI. SECURITY WITH ILNP

With deployed IP Security (*IPsec*) today, the IPsec Security Associations (SAs) are bound to full IP addresses at the local and remote sites, as a form of end-system identity. So, IPsec requires that the IP addresses at each end-point of the communication remain fixed – not varying over time and not varying between source and destination.

For supporting localised addressing (e.g. NAT), multi-homing, and mobility, IP addresses might change, so that requirement is not normally met. Since the High Assurance IP Encryptor (HAIPE) used to protect existing military IP networks is a US DoD profile of IETF standard IPsec, these issues directly apply to military IP networks.

A. Use of IPsec today

These issues forced the IETF to develop a workaround retrospectively to enable the IP Encapsulating Security Payload (ESP) to cope with these important capabilities in at least some deployments; the workaround encapsulates IPsec ESP within UDP, increasing the bandwidth lost and also increasing complexity within IPsec endpoints. [26]

Unfortunately, that workaround is not sufficient to resolve those same issues for the IP Authentication Header (AH). So AH remains unable to traverse NAT devices.

B. ILNP Security

However, ILNP Security Associations only include the Source Identifier and Destination Identifier, but **not** the Source Locator or Destination Locator. We have already described above that all ILNP multi-homing and traffic engineering functions rely only on the use and manipulation of the Locator values – Locator rewriting at the SBR. So, the use of IPsec with MH and TE functions is now orthogonal, whereas with IPv6 today, all these functions are entangled due to the use of the full 128 bits of the IPv6 address in these functions. This means that for ILNP sessions, the Locator values may be modified in transit from source to destination, and the Locator values may vary over time, without interfering with proper operation of IP Security for ILNP. Since the Locator fields are not included in the Authentication Header (AH) for ILNP, modifying the Locator values does not break the ability of AH to protect the critical information for the Internet session.

So, our proposed approach to site multi-homing or to site traffic engineering remains fully compatible with IPsec and so with the use of HAIPE.

C. ILNP Nonce

Not all threat environments require the computational expense and bandwidth overhead of cryptographic authentication (e.g. IP Authentication Header) for control traffic (e.g. ILNP Locator Update messages). So for those environments, we have devised the ILNP Nonce destination option. This option contains an unpredictable nonce value and protects ILNP control traffic from off-path attackers. This can be added to any ILNP packet, but is required for all ILNP Locator Updates. This mechanism ensures that ILNP without IPsec has security properties that are equivalent to IPv6 without IPsec.

Of course, for threat environments requiring stronger protection, the IP Authentication Header (AH) can also be used and will provide much stronger protections.

D. Topology Obfuscation

With Locator rewriting at the SBRs enabled, it is also possible to obfuscate details of the network topology within the site. Since node identifiers are highly likely to be globally unique, the site border routers can cache the internal location of internal nodes, rewrite the Source Locator field

on egress packets, and rewrite the Destination Locator field on ingress packets. This makes it harder for an external adversary to learn about the internal details of the site network. This scheme can work independently of (and in conjunction with) the IPv6 privacy extension [27] for generating identifiers for internal nodes and providing node identifier agility.

E. DoS Resistance

Denial of Service (DoS) attacks are an increasing concern in the deployed Internet. ILNP's Locator agility is central to the multi-homing and traffic engineering capabilities described above. This same mechanism also enables a new approach to resisting DoS attacks.

If a multi-homed ILNP host or site is being targeted with a DoS attack, the affected host or site can change the active Locator set to exclude link(s) and Locator value(s) that are impacted by that attack. Locator(s), and associated link(s), being targeted by the adversary are removed from the working set for the targeted host or site, which enables existing sessions to remain up despite the DoS attack. Further, new sessions will use the new Locator(s), and associated link(s), and so will also be able to evade the DoS attacks. If the adversary decides to target all Locators for a given node or site, this technique forces the adversary to either spend more resources mounting the attack or to split existing resources across the several different Locator values. So at a bare minimum, this approach increases the work function required for the adversary to have an attack with equivalent impact as attacks on sites that are multi-homed today with the existing IPv4 or IPv6 multi-homing techniques.

VII. RELATED WORK

ILNP has its roots in the earlier GSE/8+8 proposal to the IETF. [16] Atkinson's activity within the IRTF Name Space Research Group (NSRG) on an architecture derived from GSE/8+8 later turned into ILNP. However, the concept of identifier/locator separation goes back decades. [14] The Internet Architecture Board (IAB) has encouraged research into Identifier/Locator architectures for more than a decade, including creating the IRTF NSRG. [28] [6]

Most recently, the Routing Research Group of the Internet Research Task Force has been studying possible evolution of the Internet's routing architecture. The current Routing RG charter was driven by the most recent IAB Workshop on Routing and Addressing. [6] Since the Routing RG re-charter, multiple proposals been presented within the Routing RG. Most proposals have focused narrowly on the

site multi-homing issues, rather than also trying to address a broader set of routing challenges (e.g. host mobility, site mobility, host multi-homing). ILNP is trying to address the broader set of issues, including (but not limited to) site multi-homing.

Aside from ILNP, Six/One [29] and the Host Identity Protocol [30] also crisply distinguish between identifiers and locators. Both ILNP and HIP have their roots in the IRTF NSRG discussions, so in a sense they could be considered siblings. Six/One is focused on site multi-homing through address translation by site border routers, but does not provide host multi-homing or host mobility features. Like ILNP, HIP requires end system updates. The current set of HIP specifications requires that all HIP data packets are cryptographically-protected, although this might change in future. HIP addresses both host multi-homing and mobility. [31] The LISP proposal involves changes to routers rather than hosts, does not always use non-routable identifiers, and currently does not support host multi-homing or host mobility. [32] LISP's *Endpoint Identifier (EID)* is a scoped address, used both for transport-layer identity, and for interior routing & packet forwarding within the end site. At present, none of these are being standardised by the IETF. Several proposals either have developed or are developing Experimental RFCs, within either the IETF or the IRTF.

VIII. CONCLUSION

Our ongoing work on the Identifier Locator Network Protocol (ILNP) provides crisp separation of Locators and Identifiers in IPv6. Since the Locator uses the same syntax and semantics as the top 64 bits of the IPv6 address format of RFC-3587 [18], it can work with existing IPv6 routers and across current IPv6 deployed networks. This enables both incremental deployment and backwards compatibility.

ILNP allows, but does not require, Locator values to be re-written at site border routers (SBRs). This means that multi-homing (MH) can be managed at the SBRs. Also, multiple Locator values, i.e. routing prefixes, can be used to manage traffic engineering (TE) at the SBRs also. ILNP is designed to integrate MH and TE as native capabilities.

Meanwhile, the Locator/Identifier separation allows the Identifier to be used in IPsec/HAIPE in harmony with localised addressing (e.g. NAT), multi-homing, and Traffic Engineering. Furthermore, the Locator agility of ILNP, coupled with already defined mechanisms for Identifier agility in RFC-4941 [27] gives the potential for excellent DoS resistance, as well as permitting obfuscation of site-specific addressing details.

REFERENCES

- [1] R. Atkinson, S. Bhatti, and S. Hailes, "Harmonised Resilience, Security and Mobility Capability for IP," in *27th IEEE Military Communications Conference*. San Diego, CA: IEEE, Nov. 2008.
- [2] J. Abley, K. Lindqvist, *et al.*, "IPv4 Multihoming Practices and Limitations," IETF, RFC 4116, July 2005.
- [3] X. Meng, Z. Xu, B. Zhang, G. Huston, S. Lu, and L. Zhang, "IPv4 Address Allocation and the BGP Routing Table Evolution," *ACM Computer Communications Review*, vol. 35, no. 1, pp. 71–80, 2005.
- [4] E. Rosen, A. Viswanathan, and R. Callon, "Multiprotocol Label Switching Architecture," IETF, RFC 3031, Jan. 2001.
- [5] T. Bu, L. Gao, and D. Townsley, "On Characterizing BGP Routing Table Growth," *Computer Networks*, vol. 45, no. 1, pp. 45–54, 2004.
- [6] D. Meyer, L. Zhang, and K. Fall, "Report from the IAB Workshop on Routing and Addressing," IAB, RFC 4984, Sept. 2007.
- [7] D. Wendlandt, I. Avramopoulos, D. G. Anderson, and J. Rexford, "Don't Secure Routing Protocols, Secure Data Delivery," in *5th ACM Workshop on Hot Topics in Networks (Hotnets-V)*, Nov. 2006.
- [8] W. Xu and J. Rexford, "MIRO: Multi-path Interdomain Routing," in *ACM SIGCOMM 2006*. ACM, 2006, pp. 171–182.
- [9] J. He and J. Rexford, "Toward Internet-Wide Multipath Routing," *IEEE Network*, vol. 22, no. 2, pp. 16–21, Mar/Apr 2008.
- [10] R. Teixeira, K. Marzullo, S. Savage, and G. M. Voelker, "Characterizing and Measuring Path Diversity of Internet Topologies," in *SIGMETRICS 2003*. New York, NY, USA: ACM, 2003, pp. 304–305.
- [11] S. Savage, A. Collins, E. Hoffman, J. Snell, and T. Anderson, "The End-To-End Effects of Internet Path Selection," in *SIGCOMM 1999*. New York, NY, USA: ACM, 1999, pp. 289–299.
- [12] G. Iannaccone, C. Chen-Nee, S. Bhattacharyya, and C. Diot, "Feasibility of IP Restoration in a Tier 1 Backbone," *IEEE Network*, vol. 18, no. 2, pp. 13–19, Mar/Apr 2004.
- [13] R. Atkinson, S. Bhatti, and S. Hailes, "A Proposal for Unifying Mobility with Multi-Homing, NAT, and Security," in *5th ACM International Workshop on Mobility Management and Wireless Access - MOBIWAC2007*. Chania, Crete, Greece: ACM, Oct. 2007.
- [14] C. Bennett, S. Edge, and A. Hinchley, "Issues in the Interconnection of Datagram Networks," ARPA Network Working Group, Internet Experiment Note (IEN) 1, July 1977.
- [15] I. Castineyra, N. Chiappa, and M. Steenstrup, "The Nimrod Routing Architecture," IETF, RFC 1992, Aug. 1996.
- [16] M. O'Dell, "GSE - An Alternate Addressing Architecture for IPv6," IETF, Internet-Draft draft-ipng-gseaddr-00.txt, Feb. 1997.
- [17] B. Carpenter, "Architectural Principles of the Internet," IETF, RFC 1958, June 1996.
- [18] R. Hinden, S. Deering, and E. Nordmark, "IPv6 Global Unicast Address Format," IETF, RFC 3587, Aug. 2003.
- [19] R. Hinden and S. Deering, "IP Version 6 Addressing Architecture," IETF, RFC 4291, Feb. 2006.
- [20] B. Wellington, "Secure Domain Name System (DNS) Dynamic Update," IETF, RFC 3007, Nov. 2000.
- [21] C. Liu and P. Albitz, *DNS and BIND, 5th Edition*. Sebastopol, CA, USA: O'Reilly and Associates, May 2006.
- [22] J. Jung, E. Sit, H. Balakrishnan, and R. Morris, "DNS Performance and the Effectiveness of Caching," in *Proceedings of 1st ACM SIGCOMM Internet Measurement Workshop*. San Francisco, CA, USA: ACM, Nov. 2001.
- [23] R. Arends, R. Austein, M. Larson, D. Massey, and S. Rose, "DNS Security Introduction and Requirements," RFC 4033, IETF, RFC 4033, Mar. 2005.
- [24] R. Hinden and B. Haberman, "Unique Local IPv6 Unicast Addresses," IETF, RFC 4193, Oct. 2005.
- [25] P. Ferguson and D. Senie, "Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing," IETF, RFC 2827, May 2000.
- [26] A. Huttunen, B. Swander, V. Volpe, L. DiBurro, and M. Stenberg, "UDP Encapsulation of IPsec ESP Packets," IETF, RFC 3948, Jan. 2005.
- [27] T. Narten, R. Draves, and S. Krishnan, "Privacy Extensions for Stateless Address Autoconfiguration in IPv6," IETF, RFC 4941, Sept. 2007.
- [28] M. Kaat, "Overview of 1999 IAB Network Layer Workshop," IAB, RFC 2956, Oct. 2000.
- [29] C. Vogt, "Six/One Router: A Scalable and Backwards-Compatible Solution for Provider-Independent Addressing," in *Proceedings of 3rd ACM International Workshop on Mobility in the Evolving Internet Architecture*. Seattle, WA, USA: ACM, Aug. 2008.
- [30] R. Moscovitz *et al.*, "Host Identity Protocol," IRTF, RFC 5201, Apr. 2008.
- [31] P. Nikander, T. Henderson, C. Vogt, and J. Arkko, "End-Host Mobility and Multihoming with the Host Identity Protocol," IRTF, RFC 5206, Apr. 2008.
- [32] D. Farinacci *et al.*, "Locator-Identifier Separation Protocol (LISP)," IRTF, Internet-Draft draft-farinacci-lisp-12.txt, Mar. 2009.