

Network Working Group
Internet-Draft
Intended status: Informational
Expires: August 19, 2010

J. Abley
Afiliias Canada
M. Bagnulo
A. Garcia-Martinez
UC3M
February 15, 2010

**Applicability Statement for the Level 3 Multihoming Shim Protocol
(Shim6)
draft-ietf-shim6-applicability-05**

Abstract

This document discusses the applicability of the Shim6 IPv6 protocol and associated support protocols and mechanisms to provide site multihoming capabilities in IPv6.

Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

This Internet-Draft will expire on August 19, 2010.

Copyright Notice

Copyright (c) 2010 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents

(<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the BSD License.

This document may contain material from IETF Documents or IETF Contributions published or made publicly available before November 10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow modifications of such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.

Table of Contents

1.	Introduction	3
2.	Deployment Scenarios	3
3.	Address Configuration	5
3.1.	Protocol Version (IPv4 vs. IPv6)	5
3.2.	Prefix Lengths	6
3.3.	Address Generation	6
3.4.	Use of CGA vs. HBA	6
4.	Shim6 and Ingress Filtering	7
5.	Shim6 Capabilities	9
5.1.	Fault Tolerance	9
5.1.1.	Establishing Communications After an Outage	9
5.1.2.	Short-Lived Communications	9
5.1.3.	Long-Lived Communications	10
5.2.	Load Balancing	10
5.3.	Traffic Engineering	10
6.	Application Considerations	11
7.	Interaction with Other Protocols	12
7.1.	Shim6 and Mobile IPv6	12
7.1.1.	Multihomed Home Network	12
7.1.2.	Shim6 Between the HA and the MN	15
7.2.	Shim6 and SeND	15
7.3.	Shim6 and SCTP	16
7.4.	Shim6 and NEMO	16
7.5.	Shim6 and HIP	17
8.	Security Considerations	18
8.1.	Privacy Considerations	19
9.	IANA Considerations	19
10.	Contributors	19
11.	Acknowledgements	19
12.	References	20
12.1.	Normative References	20
12.2.	Informative References	21
	Authors' Addresses	22

1. Introduction

Site multihoming is an arrangement by which a site may use multiple paths to the rest of the Internet to provide better reliability for traffic passing in and out of the site than would be possible with a single path. Some of the motivations for operators to multi-home their network are described in [\[RFC3582\]](#).

In IPv4, site multihoming is achieved by injecting into the global Internet routing system (sometimes referred to as the Default-Free Zone, or DFZ) the additional state required to allow session resilience over re-homing events [\[RFC4116\]](#). There is concern that this approach will not scale [\[RFC3221\]](#), [\[RFC4984\]](#).

In IPv6, site multihoming in the style of IPv4 is not generally available to end sites due to a strict policy of route aggregation in the DFZ. Site multihoming for sites without provider-independent (PI) addresses is achieved by assigning multiple addresses to each host, one or more from each provider. This multihoming approach provides no transport-layer stability across re-homing events.

Shim6 provides layer-3 support for making re-homing events transparent to the transport layer by means of a shim approach. State information relating to the multihoming of two endpoints exchanging unicast traffic is retained on the endpoints themselves, rather than in the network. Communications between Shim6-capable hosts and Shim6-incapable hosts proceed as normal, but without the benefit of transport-layer stability. The Shim6 approach is thought to have better scaling properties with respect to the state held in the DFZ than the IPv4 approach.

This note describes the applicability of the Level 3 multihoming (hereafter Shim6) protocol defined in [\[RFC5533\]](#) and the failure detection mechanisms defined in [\[RFC5534\]](#).

The terminology used in this document, including terms like locator, and ULID, is defined in [\[RFC5533\]](#).

2. Deployment Scenarios

The goal of the Shim6 protocol is to support locator agility in established communications: different layer-3 endpoint addresses may be used to exchange packets belonging to the same transport-layer session, all the time presenting a consistent identifier pair to upper-layer protocols.

In order to be useful, the Shim6 protocol requires that at least one

of the peers has more than one address which could be used on the wire (as locators). In the event of communications failure between an active pair of addresses, the Shim6 protocol will attempt to reestablish communication by trying different combinations of locators.

While other multi-addressing scenarios are not precluded, the scenario in which the Shim6 protocol is expected to operate is that of a multihomed site which is connected to multiple transit providers, and which receives an IPv6 prefix from each of them. This configuration is intended to provide protection for the end-site in the event of a failure in some subset of the available transit providers, without requiring the end-site to acquire PI address space or requiring any particular cooperation between the transit providers.

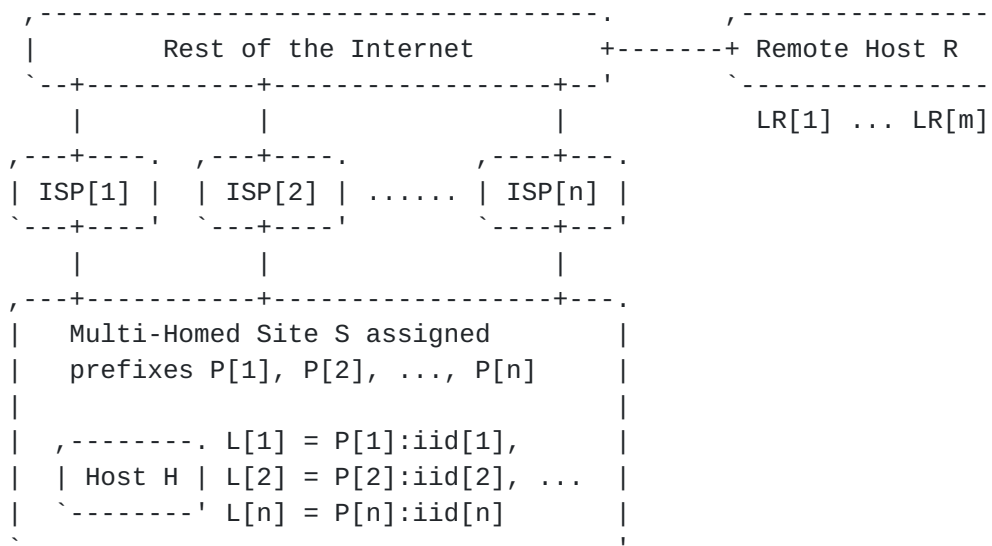


Figure 1

In the scenario illustrated in Figure 1 host H communicates with some remote host R. Each of the addresses $L[i]$ configured on host H in the multihomed site S can be reached through provider $ISP[i]$ only, since $ISP[i]$ is solely responsible for advertising a covering prefix for $P[i]$ to the rest of the Internet.

The use of locator $L[i]$ on H hence causes inbound traffic towards H to be routed through $ISP[i]$. Changing the locator from $L[i]$ to $L[j]$ will have the effect of re-routing inbound traffic to H from $ISP[i]$ to $ISP[j]$. This is the central mechanism by which the Shim6 protocol aims to provide multihoming functionality: by changing locators, host H can change the upstream ISP used to route inbound packets towards itself. Regarding to the outbound traffic to H, the path taken in

this case depends on both the actual locator LR[j] chosen by R, and the administrative exit selection policy of site S.

The Shim6 protocol has other potential applications beyond site multihoming. For example, since Shim6 is a host-based protocol, it can also be used to support host multihoming. In this case, a failure in communication between a multihomed host and some other remote host might be repaired by selecting a locator associated with a different interface.

3. Address Configuration

3.1. Protocol Version (IPv4 vs. IPv6)

The Shim6 protocol is defined only for IPv6. However, there is no fundamental reason why a Shim6-like approach could not support IPv4 addresses as locators, either to provide multihoming support to IPv4-numbered sites, or as part of an IPv4/IPv6 transition strategy. Some extensions to the Shim6 protocol for supporting IPv4 locators have been proposed in [[I-D.nordmark-shim6-esd](#)].

The Shim6 protocol, as specified for IPv6, incorporates cryptographic elements in the construction of locators (see [[RFC3972](#)], [[RFC5535](#)]). Since IPv4 addresses are insufficiently large to contain addresses constructed in this fashion, direct implementation of Shim6 as specified for IPv6 for use with IPv4 addresses might require protocol modifications.

In addition, there are other factors to take into account when considering the support of IPv4 addresses, in particular IPv4 locators. Using multiple IPv4 addresses in a single host in order to support Shim6 style of multihoming would result in an increased IPv4 address consumption, which with the current rate of IPv4 addresses would be problematic. Besides, Shim6 may suffer additional problems if locators become translated on the wire. Address translation is more likely to involve IPv4 addresses. IPv4 addressed can be translated to other IPv4 addresses (for example, private IPv4 address into public IPv4 address and vice versa) or to/from IPv6 addresses (for example, as defined by NAT64 [[I-D.ietf-behave-v6v4-xlate-stateful](#)]). When address translation occurs, a locator exchanged by Shim6 could be different to the address needed to reach the corresponding host, either because the translated version of the locator exchanged by Shim6 is not known or because the translation state does not exist any more in the translator device. Supporting these scenarios would require NAT traversal mechanisms which are not defined yet and which would imply additional complexity (as any other NAT traversal mechanism).

3.2. Prefix Lengths

The Shim6 protocol does not assume that all the prefixes assigned to the multihomed site have the same prefix length.

However, the use of CGA [[RFC3972](#)] and HBA [[RFC5535](#)] involve encoding information in the lower 64 bits of the locators. This imposes the requirement on address assignment to Shim6-capable hosts that all interface addresses should be able to accommodate 64-bit interface identifiers. It should be noted that this is imposed by [RFC4291](#) [[RFC4291](#)].

3.3. Address Generation

The security of the Shim6 protocol is based on the use of CGA and HBA addresses.

CGA and HBA generation process can use the information provided by the stateless auto-configuration mechanism defined in [[RFC4862](#)] with the additional considerations presented in [[RFC3972](#)] and [[RFC5535](#)].

Stateful address auto-configuration using DHCP [[RFC3315](#)] is not currently supported, because there is no defined mechanism to convey the CGA Parameter Data Structure and other relevant information from the DHCP server to the host. The definition of such mechanism seems to be quite straightforward in the case of the HBA, since only the CGA Parameter Data Structure needs to be delivered from the DHCP server to the Shim6 host, and this data structure does not contain any secret information. In the case of CGAs, the difficulty is increased, since private key information should be exchanged as well as the CGA Parameter Data Structure. However, with appropriate extensions a DHCP server could inform to a host about the SEC value to use when generating an address, or DHCP could even be used by the host to delegate to the server the CPU-intensive task of computing a Modifier for a given <prefix, public key, SEC> combination [[I-D.ietf-csi-dhcpv6-cga-ps](#)].

3.4. Use of CGA vs. HBA

The choice between CGA and HBA is a trade-off between flexibility and performance.

The use of HBA is more efficient in the sense that addresses require less computation than CGA, involving only hash operations for both the generation and the verification of locator sets. However, the locators of an HBA set are determined during the generation process, and cannot be subsequently changed; the addition of new locators to that initial set is not supported, except by re-generation of the

entire set which will in turn cause all addresses to change.

The use of CGA is more computationally expensive, involving public key cryptography in the verification of locator sets. However, CGAs are more flexible in the sense that they support the dynamic modification of locator sets.

Therefore, CGAs are well suited to support dynamic environments such as mobile hosts, where the locator set must be changed frequently. HBAs are better suited for sites where the prefix set remains relatively stable.

It should be noted that, since HBAs are defined as a CGA extension, it is possible to generate hybrid HBA/CGA structures that incorporate the strengths of both: i.e. that a single address can be used as an HBA, enabling computationally-cheap validation amongst a fixed set of addresses, and also as a CGA, enabling dynamic manipulation of the locator set. For additional details, see [[RFC5535](#)].

4. Shim6 and Ingress Filtering

Ingress filtering [[RFC2827](#)] prevents address spoofing by dropping packets which come from customer networks with source addresses not belonging to the prefix assigned to them. The problem of deploying ingress filters with multihomed customers is discussed in [[RFC3704](#)], in particular considering the case in which non-PI addresses are used by customer networks. This is the case for IPv6 hosts in multihomed networks with PA, and also for a Shim6 host in a multihomed network. Note that this is also the case for other solutions supporting multihoming, such as SCTP [[RFC4960](#)], HIP [[RFC4423](#)], etc.

One solution to this problem is to make the providers aware of the alternative prefixes that can be used by a multihomed site, so that ingress filtering would not be applied to packets with source addresses belonging to these prefixes. This may be possible in some cases, but it cannot be assumed as the general case.

[[RFC3704](#)] proposes that non-PI addresses should ensure that each packet is delivered to the provider related with the prefix of its source address. To deliver packets to the appropriate outgoing ISP, some routers of the site must consider source addresses in their forwarding decisions, in addition to the usual destination-based forwarding. These routers maintain as many parallel routing tables as valid source prefixes are, and choose a route that is a function of both the source and the destination address. The way these routing tables are populated is out of the scope of this document.

As proposed in [[I-D.huitema-multi6-hosts](#)], it is required for site exit routers (at least) to be part of a single connected source based routing domain:




Figure 2

In this way, packets arriving to this connected source based routing domain would be delivered to the appropriate exit router.

Some particular cases of this generic deployment scenario are:

- a single exit router, in which the router chooses the exit provider according to the source address of the packet to be forwarded
- a site in which all routers perform source address based forwarding
- a site in which only site-exit routers perform source address based forwarding, and these site-exit routers are connected through point-to-point tunnels, so that packets can be tunneled to the appropriate exit router according to its source address

For hosts attached directly to networks of different providers, a host solution to ensure that packets are forwarded to the appropriate interface according to its source address must be provided. This problem is under discussion in the Multiple Interfaces (MIF) IETF Working Group.

Shim6 has no means to enforce neither host nor network forwarding for a given locator to be used as source address. If any notification is received from the router dropping the packets with legitimate source addresses as a result of ingress filtering, the affected locator could be associated to a low preference (or not being used at all). But even if such notification is not received, or not processed by the Shim6 layer, defective ingress filtering configuration will be treated as a communication failure, and Shim6 re-homing would finally

select a working path in which packets are not filtered, if this path exists. Note that this behavior results from the powerful end-to-end resilience properties exhibited by REAP.

5. Shim6 Capabilities

5.1. Fault Tolerance

5.1.1. Establishing Communications After an Outage

If a host within a multihomed site attempts to establish a communication with a remote host and selects a locator which corresponds to a failed transit path, bidirectional communication between the two hosts will not succeed. In order to establish a new communication, the initiating host must try different combinations of (source, destination) locator pairs until it finds a pair that works. The mechanism for this default address selection is described in [\[RFC3484\]](#). A commentary on this mechanism in the context of multihomed environments can be found in [\[I-D.bagnulo-ipv6-rfc3484-update\]](#).

Since a Shim6 context is normally only established between two hosts after initial communication has been set up, there is no opportunity for Shim6 to participate in the discovery of a suitable, initial (source, destination) locator pair. The same consideration holds for referrals, as it is described in [Section 6](#).

5.1.2. Short-Lived Communications

The Shim6 context establishment operation requires a 4-way packet exchange, and involves some overhead on the participating hosts in memory and CPU.

For short-lived communications between two hosts, the benefit of establishing a Shim6 context might not exceed the cost, perhaps because the protocols concerned are fault tolerant and can arrange their own recovery (e.g. DNS) or because the frequency of re-homing events is sufficiently low that the probability of such a failure occurring during a short-lived exchange is not considered significant.

It is anticipated that the exchange of Shim6 context will provide most benefit for exchanges between hosts which are long-lived. For this reason the default behaviour of Shim6-capable hosts is expected to employ deferred context-establishment. This default behaviour will be able to be overridden by applications which prefer immediate context establishment regardless of transaction longevity.

It must be noted that all the above considerations refer to the lifetime of the interaction between the peers and not about the lifetime of a particular connection (e.g. TCP connection). In other words, the Shim6 context is established between ULID pairs and it affects all the communication between these ULIDs. So, two nodes with multiple short-lived communications using the same ULID pair would benefit as much from the Shim6 features as two nodes having a single long-lived communication. One example of such scenario would be a web client software downloading web contents from a server over multiple TCP connections. Each TCP connection is short-lived, but the communication/contact between the two ULID could be long-lived.

5.1.3. Long-Lived Communications

As discussed in [Section 5.1.2](#), hosts engaged in long-lived communications will suffer lower proportional overhead, and greater probability of benefit than those performing brief transactions.

Deferred context setup ensures that session establishment time will not be increased by the use of Shim6.

5.2. Load Balancing

The Shim6 protocol does not support load balancing within a single context: all packets associated with a particular context are exchanged using a single locator pair per direction, with the exception of forked contexts, which are created upon explicit requests from the upper-layer protocol.

It may be possible to extend the Shim6 protocol to use multiple locator pairs in a single context, but the impact of such an extension on upper-layer protocols (e.g. on TCP congestion control) should be considered carefully.

When many contexts are considered together in aggregation, e.g. on a single host which participates in many simultaneous contexts or in a site full of hosts, some degree of load sharing should occur naturally due to the selection of different locator pairs in each context. However, there is no mechanism defined to ensure that this natural load sharing is arranged to provide a statistical balance between transit providers.

5.3. Traffic Engineering

The Shim6 protocol provides some lightweight traffic engineering capabilities in the form of the Locator Preferences option, which allows a host to inform a remote host of local preferences for locator selection.

This mechanism is only available after a Shim6 context has been established, and it is a host-based capability rather than a site-based capability. There is no defined mechanism which would allow use of the Locator Preferences option amongst a site full of hosts to be managed centrally.

6. Application Considerations

Shim6 provides multihoming support without forcing changes in the applications running on the host. The fact that an address has been generated according to the CGA or HBA specification does not require any specific action from the application, e.g. it can obtain remote CGA or HBA addresses as a result of a `getaddrinfo()` call to trigger a DNS Request. The storage of CGA or HBA addresses in DNS does not require also any modification of this protocol, since they are recorded using AAAA records. Moreover, neither the ULID/locator management [[RFC5533](#)] nor the failure detection and recovery [[RFC5534](#)] functions require application awareness.

However, a specific API [[I-D.ietf-shim6-multihome-shim-api](#)] is being developed for those applications which might require additional capabilities in ULID/locator management, such as the locator pair in use for a given context, or the set of local or remote locators available for it. This API can also be used to disable Shim6 operation when required.

It is worth to note that callbacks can benefit naturally from Shim6 support. In a callback, an application in B retrieves IP_A, the IP address of a peer A, and B uses IP_A to establish a new communication with A. As long as the address exchanged, IP_A is the ULID for the initial communication between A and B, and B uses the same address as in the initial communication, and this initial communication is alive (or the context has not been deleted), the new communication could use the locators exchanged by Shim6 for the first communication. In this case, communication could proceed even if the ULID of A is not reachable.

However, Shim6 does not provide specific protection to current applications when they use referrals. A referral is the exchange of the IP address IP_A of a party A by party B to party C, so that party C could use IP_A to communicate with party A [[I-D.ietf-multi6-app-refer](#)]. In a normal case, the ULID IP_A would be the only information sent by B to C as referral. But if IP_A is no longer valid as locator in A, C could have trouble in establishing a communication with A. Increased failure protection for referrals could be obtained if B exchanged the whole list of alternative locators of A, although in this case the application protocol should

be modified. Note that B could send to C the current locator of A, instead of the ULID of A, as a way of using the most recent reachability information about A. While in this case no modification of the application protocol is required, some concerns arise: host A may not accept one of its locator as ULID for initiating a communication, and if CGA are used, the locator may not be a CGA so a Shim6 context among A and C could not be created.

7. Interaction with Other Protocols

7.1. Shim6 and Mobile IPv6

We next consider some scenarios in which the Shim6 protocol and the MIPv6 protocol [[RFC3775](#)] might be used simultaneously.

7.1.1. Multihomed Home Network

In this case, the Home Network of the Mobile Node (MN) is multihomed. This implies the availability of multiple Home Network prefixes, resulting on multiple HoAs for each MN. Since the MN is a node within a multihomed site, it seems reasonable to expect that the MN should be able to benefit from the multihoming capabilities provided by the Shim6 protocol. Moreover, the MN needs to be able to obtain the multihoming benefits even when it is roaming away from the Home Network: if the MN is away from the Home Network while the Home Network suffers a failure in a transit path, the MN should be able to continue communicating using alternate paths to reach the Home Network.

The resulting scenario is the following:

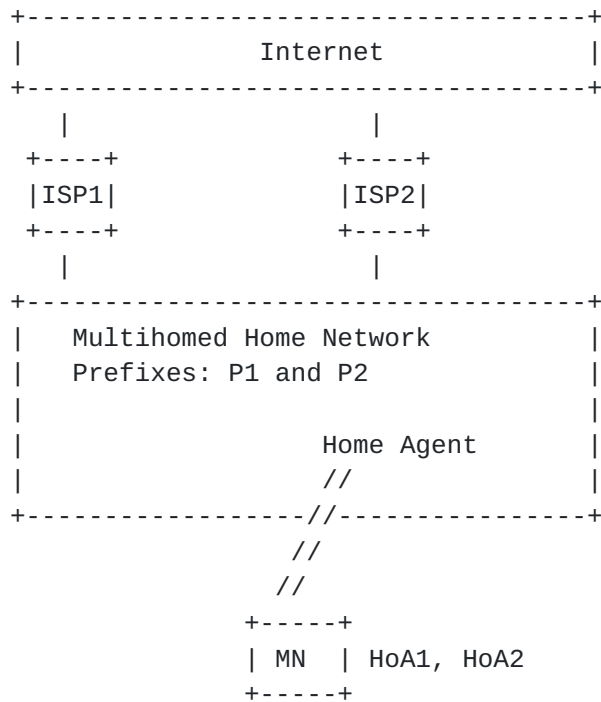


Figure 3

So, in this configuration, the Shim6 protocol is used to provide multiple communication paths to all the nodes within the multihomed sites (including the mobile nodes) and the MIPv6 protocol is used to support mobility of the mobile nodes of the multihomed site.

The proposed protocol architecture would be the following:

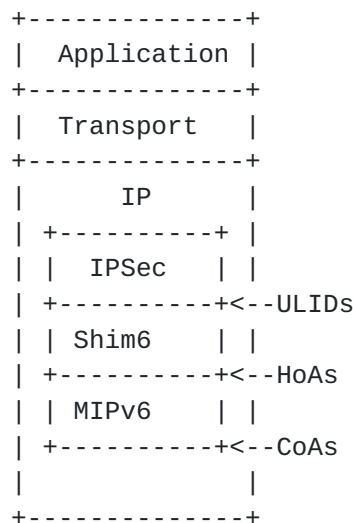


Figure 4

In this architecture, the upper layer protocols and IPSec would use ULIDs of the Shim6 protocol. Only the HoAs will be presented by the upper layers to the Shim6 layer as potential ULIDs. Two Shim6 entities will exchange their own available HoAs as locators. Therefore, Shim6 provides failover between different HoAs and allows preserving established communications when an outage affects the path through the ISP that has delegated the HoA used for initiating the communication (similarly to the case of a host within a multihomed site). The CoAs are not presented to the Shim6 layer and are not included in the local locator set in this case. The CoAs are managed by the MIPv6 layer, which binds each HoA to a CoA.

So, in this case, the upper layer protocols select a ULID pair for the communication. The Shim6 protocol translates the ULID pair to an alternative locator in case that is needed. Both the ULIDs and the alternative locators are HoAs. Next, the MIPv6 layer maps the selected HoA to the corresponding CoA, which is the actual address included in the wire.

The Shim6 context is established between the MN and the CN, and it would allow the communication to use all the available HoAs to provide fault tolerance. The MIPv6 protocol is used between the MN and the HA in the case of the bidirectional tunnel mode, and between the MN and the CN in case of the RO (Route Optimization) mode.

7.1.2. Shim6 Between the HA and the MN

Another scenario where a Shim6-MIPv6 interaction may be useful is the case where a Shim6 context is established between the MN and the HA in order to provide fault tolerance capabilities to the bidirectional tunnel between them.

Consider the case where the HA has multiple addresses (whether because the Home Network is multihomed or because the HA has multiple interfaces) and/or the MN has multiple addresses (whether because the visited network is multihomed or because the MN has multiple interfaces). In this case, if a failure affects the address pair that is being used to run the tunnel between the MN and HA, additional mechanisms need to be used to preserve the communication.

One possibility would be to use MIPv6 capabilities, by simply changing the CoA used as the tunnel endpoint. However, MIPv6 lacks of failure detection mechanisms that would allow the MN and/or the HA to detect the failure and trigger the usage of an alternative address. Shim6 provides such failure detection protocol, so one possibility would be re-using the failure detection function from the Shim6 failure detection protocol in MIPv6. In this case, the Shim6 protocol wouldn't be used to create Shim6 context and provide fault tolerance, but just its failure detection functionality would be re-used.

The other possibility would be to use the Shim6 protocol to create a Shim6 context between the HA and the MN so that the Shim6 detects any failure and re-homes the communication in a transparent fashion to MIPv6. In this case, the Shim6 protocol would be associated to the tunnel interface.

7.2. Shim6 and SeND

Secure Neighbor Discovery (SeND) [[RFC3971](#)] uses CGAs to prove address ownership for Neighbor Discovery [[RFC4861](#)]. The Shim6 protocol can use either CGAs or HBAs to protect locator sets included in Shim6 contexts. It is expected that some hosts will need to participate in both SeND and Shim6 simultaneously.

In the case that both the SeND and Shim6 protocols are using the CGA technique to generate addresses, then there is no conflict: the host will generate addresses for both purposes as CGAs, and since it will be in control of the associated private key, the same CGA can be used for the different protocols.

In the case that a Shim6-capable host is using HBAs to protect its locator sets, the host will need to generate hybrid HBA/CGA addresses

as defined in [[RFC5535](#)] and discussed briefly in [Section 3.4](#). In this case, the CGA Parameter Data Structure containing a valid public key and the Multi-Prefix extension are included as inputs to the hash function.

[7.3.](#) Shim6 and SCTP

The SCTP [[RFC4960](#)] protocol provides a reliable, stream-based communications channel between two hosts which provides a superset of the capabilities of TCP. One of the notable features of SCTP is that it allows the exchange of endpoint addresses between hosts, and is able to recover from the failure of a particular endpoint pair in a manner which is conceptually similar to locator selection in Shim6.

SCTP is a transport-layer protocol, higher in the protocol stack than Shim6, and hence there is no fundamental incompatibility which would prevent a Shim6-capable host from communicating using SCTP.

However, since SCTP and Shim6 both aim to exchange addressing information between hosts in order to meet the same generic goal, it is possible that their simultaneous use might result in unexpected behaviour, e.g. lead to race conditions.

The capabilities of SCTP with respect to path maintenance of a reliable, connection-oriented stream protocol are more extensive than the more general layer-3 locator agility provided by Shim6. Therefore, It is recommended that Shim6 is not used for SCTP sessions, and that path maintenance is provided solely by SCTP. There are at least two ways to enforce this behaviour. One option would be to make the stack, and in particular the Shim6 sublayer, aware of SCTP sockets and in this case refrain from creating a Shim6 context. The other option is that the upper layer, SCTP in this case, informs using a Shim6 capable API like the one proposed in [[I-D.ietf-shim6-multihome-shim-api](#)] that no Shim6 context must be created for this particular communication.

Note that the issues described here for SCTP may also arise for a multipath TCP solution.

[7.4.](#) Shim6 and NEMO

The NEMO [[RFC3963](#)] protocol extensions to MIPv6 allow a Mobile Network to communicate through a bidirectional tunnel via a Mobile Router (MR) to a NEMO-compliant Home Agent (HA) located in a Home Network.

If either or both of the MR or HA are multihomed, then a Shim6 context established preserves the integrity of the bidirectional

tunnel between them in the event that a transit failure occurs in the connecting path.

Once the tunnel between MR and HA is established, hosts within the Mobile Network which are Shim6-capable can establish contexts with remote hosts in order to receive the same multihoming benefits as any host located within the Home Network.

7.5. Shim6 and HIP

Shim6 and the Host Identity Protocol (HIP [[RFC4423](#)]) are architecturally similar in the sense that both solutions allow two hosts to use different locators to support communications between stable ULIDs. The signaling exchange to establish the demultiplexing context on the hosts is very similar for both protocols. However, there are a few key differences. First, Shim6 avoids defining a new namespace for ULIDs, preferring instead to use a routable locator as a ULID, while HIP uses public keys and hashes thereof as ULIDs. The use of a routable locator as ULID better supports deferred context establishment, application callbacks, and application referrals, and avoids management and resolution costs of a new namespace, but requires additional security mechanisms to securely bind the ULID with the locators. Second, Shim6 uses an explicit context header on data packets for which the ULIDs differ from the locators in use (this header is only needed after a failure/rehoming event occurs), while HIP compresses this context tag into the ESP SPI field of a BEET-mode security association BEET [[I-D.nikander-esp-beet-mode](#)]. Third, HIP as presently defined requires the use of public-key operations in its signaling exchange and ESP encryption in the data plane, while the use of Shim6 requires neither (if only HBA addresses are used). HIP by default provides data protection, while this is a non goal for Shim6.

The Shim6 working group was chartered to provide a solution to a specific problem, multihoming, which minimizes deployment disruption, while HIP is considered more of an experimental approach intended to solve several more general problems (mobility, multihoming and loss of end-to-end addressing transparency) through an explicit identifier/locator split. Communicating hosts that are willing and interested to run HIP (perhaps extended with Shim6's failure detection protocol) likely have no reason to also run Shim6. In this sense, HIP may be viewed as a possible long-term evolution or extension of the Shim6 architecture, or one possible implementation of the extended Shim6 design ESD [[I-D.nordmark-shim6-esd](#)].

8. Security Considerations

This section considers the applicability of the Shim6 protocol from a security perspective, i.e. which security features can expect applications and users of the Shim6 protocol.

First of all, it should be noted that the Shim6 protocol is not a security protocol, like for instance HIP. This means that as opposed to HIP, it is an explicit non goal of the Shim6 protocol to provide enhanced security for the communications that use the Shim6 protocol. The goal of the Shim6 protocol design in terms of security is not to introduce new vulnerabilities that were not present in the current non-Shim6 enabled communications. In particular, it is an explicit non goal of the Shim6 protocol security to provide protection from on-path attackers. On-path attackers are able to sniff and spoof packets in the current Internet, and they are able to do the same in Shim6 communications (as long as the communication flows through the path they are located on). So, summarizing, the Shim6 protocol does not provide data packet protection from on-path attackers.

However, the Shim6 protocol does use several security techniques. The goal of these security measures is to protect the Shim6 signaling protocol from new attacks resulting from the adoption of the Shim6 protocol. In particular, the use of HBA/CGA prevents on-path and off-path attackers to introduce new locators in the locator set of a Shim6 context, preventing redirection attacks [[RFC4218](#)]. Moreover, the usage of probes before re-homing to a different locator as a destination address prevents flooding attacks from off-path attackers.

In addition, the usage of a 4-way handshake for establishing the Shim6 context protects against DoS attacks, so hosts implementing the Shim6 protocol should not be more vulnerable to DoS attacks than regular IPv6 hosts.

Finally, many Shim6 signaling messages contain a Context Tag, meaning that only attackers that know the Context Tag can forge them. As a consequence, only on-path attackers can generate false Shim6 signaling packets for an established context. The impact of these attacks would be limited since they would not be able to add additional locators to the locator set (because of the HBA/CGA protection). In general the possible attacks have similar effects to the ones that an on-path attacker can launch on any regular IPv6 communication. The residual threats are described in the Security Considerations of the Shim6 protocol specification [[RFC5533](#)].

8.1. Privacy Considerations

The Shim6 protocol is designed to provide some basic privacy features. In particular, HBAs are generated in such a way, that the different addresses assigned to a host cannot be trivially linked together as belonging to the same host, since there is nothing in common in the addresses themselves. Similar features are provided when the CGA protection is used. This means that it is not trivial to determine that a set of addresses is assigned to a single Shim6 host.

However, the Shim6 protocol does exchange the locator set in clear text and it also uses a fixed Context Tag when using different locators in a given context. This implies that an attacker observing the Shim6 context establishment exchange or seeing different payload packets exchanged through different locators, but with the same Context Tag, can determine the set of addresses assigned to a host. However, this requires that the attacker is located along the path and that it can capture the Shim6 signaling packets. A more in depth analysis of the privacy of the Shim6 protocol can be found in [[I-D.bagnulo-shim6-privacy](#)].

9. IANA Considerations

This document has no actions for IANA.

10. Contributors

The analysis on the interaction between the Shim6 protocol and the other protocols presented in this note benefited from the advice of various people including Tom Henderson, Erik Nordmark, Hesham Soliman, Vijay Devarpalli, John Loughney and Dave Thaler.

11. Acknowledgements

Joe Abley's work was supported in part by the US National Science Foundation (research grant SCI-0427144) and DNS-OARC.

Marcelo Bagnulo worked on this document while visiting Ericsson Research Laboratory Nomadiclab.

Shinta Sugimoto reviewed this document and provided comments and text.

Iljitsch van Beijnum, Brian Carpenter, Sam Xia reviewed this document

and provided comments.

12. References

12.1. Normative References

- [RFC2827] Ferguson, P. and D. Senie, "Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing", [BCP 38](#), [RFC 2827](#), May 2000.
- [RFC3315] Droms, R., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", [RFC 3315](#), July 2003.
- [RFC3484] Draves, R., "Default Address Selection for Internet Protocol version 6 (IPv6)", [RFC 3484](#), February 2003.
- [RFC3704] Baker, F. and P. Savola, "Ingress Filtering for Multihomed Networks", [BCP 84](#), [RFC 3704](#), March 2004.
- [RFC3775] Johnson, D., Perkins, C., and J. Arkko, "Mobility Support in IPv6", [RFC 3775](#), June 2004.
- [RFC3963] Devarapalli, V., Wakikawa, R., Petrescu, A., and P. Thubert, "Network Mobility (NEMO) Basic Support Protocol", [RFC 3963](#), January 2005.
- [RFC3971] Arkko, J., Kempf, J., Zill, B., and P. Nikander, "SEcure Neighbor Discovery (SEND)", [RFC 3971](#), March 2005.
- [RFC3972] Aura, T., "Cryptographically Generated Addresses (CGA)", [RFC 3972](#), March 2005.
- [RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", [RFC 4291](#), February 2006.
- [RFC4423] Moskowitz, R. and P. Nikander, "Host Identity Protocol (HIP) Architecture", [RFC 4423](#), May 2006.
- [RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", [RFC 4861](#), September 2007.
- [RFC4862] Thomson, S., Narten, T., and T. Jinmei, "IPv6 Stateless Address Autoconfiguration", [RFC 4862](#), September 2007.
- [RFC4960] Stewart, R., "Stream Control Transmission Protocol",

[RFC 4960](#), September 2007.

- [RFC4984] Meyer, D., Zhang, L., and K. Fall, "Report from the IAB Workshop on Routing and Addressing", [RFC 4984](#), September 2007.
- [RFC5533] Nordmark, E. and M. Bagnulo, "Shim6: Level 3 Multihoming Shim Protocol for IPv6", [RFC 5533](#), June 2009.
- [RFC5534] Arkko, J. and I. van Beijnum, "Failure Detection and Locator Pair Exploration Protocol for IPv6 Multihoming", [RFC 5534](#), June 2009.
- [RFC5535] Bagnulo, M., "Hash-Based Addresses (HBA)", [RFC 5535](#), June 2009.

12.2. Informative References

- [I-D.bagnulo-ipv6-rfc3484-update]
Bagnulo, M., "Updating [RFC 3484](#) for multihoming support", [draft-bagnulo-ipv6-rfc3484-update-00](#) (work in progress), December 2005.
- [I-D.bagnulo-shim6-privacy]
Bagnulo, M., "Privacy Analysis for the SHIM6 protocol", [draft-bagnulo-shim6-privacy-01](#) (work in progress), October 2006.
- [I-D.huitema-multi6-hosts]
Huitema, C. and R. Draves, "Host-Centric IPv6 Multihoming", [draft-huitema-multi6-hosts-03](#) (work in progress), February 2004.
- [I-D.ietf-behave-v6v4-xlate-stateful]
Bagnulo, M., Matthews, P., and I. Beijnum, "Stateful NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers", [draft-ietf-behave-v6v4-xlate-stateful-08](#) (work in progress), January 2010.
- [I-D.ietf-csi-dhcpv6-cga-ps]
Jiang, S., Shen, S., and T. Chown, "DHCPv6 and CGA Interaction: Problem Statement", [draft-ietf-csi-dhcpv6-cga-ps-01](#) (work in progress), December 2009.
- [I-D.ietf-multi6-app-refer]
Nordmark, E., "Multi6 Application Referral Issues",

[draft-ietf-multi6-app-refer-00](#) (work in progress),
January 2005.

[I-D.ietf-shim6-multihome-shim-api]

Komu, M., Bagnulo, M., Slavov, K., and S. Sugimoto,
"Socket Application Program Interface (API) for
Multihoming Shim", [draft-ietf-shim6-multihome-shim-api-12](#)
(work in progress), January 2010.

[I-D.nikander-esp-beet-mode]

Nikander, P. and J. Melen, "A Bound End-to-End Tunnel
(BEET) mode for ESP", [draft-nikander-esp-beet-mode-09](#)
(work in progress), August 2008.

[I-D.nordmark-shim6-esd]

Nordmark, E., "Extended Shim6 Design for ID/loc split and
Traffic Engineering", [draft-nordmark-shim6-esd-01](#) (work in
progress), February 2008.

[RFC3221] Huston, G., "Commentary on Inter-Domain Routing in the
Internet", [RFC 3221](#), December 2001.

[RFC3582] Abley, J., Black, B., and V. Gill, "Goals for IPv6 Site-
Multihoming Architectures", [RFC 3582](#), August 2003.

[RFC4116] Abley, J., Lindqvist, K., Davies, E., Black, B., and V.
Gill, "IPv4 Multihoming Practices and Limitations",
[RFC 4116](#), July 2005.

[RFC4218] Nordmark, E. and T. Li, "Threats Relating to IPv6
Multihoming Solutions", [RFC 4218](#), October 2005.

Authors' Addresses

Joe Abley
Afilias Canada, Inc.
Suite 204
4141 Yonge Street
Toronto, Ontario M2P 2A8
Canada

Phone: +1 416 673 4176
Email: jabley@ca.afilias.info
URI: <http://afilias.info/>

Marcelo Bagnulo
U. Carlos III de Madrid
Av. Universidad 30
Leganes, Madrid 28911
Spain

Phone: +34 91 6248814
Email: marcelo@it.uc3m.es
URI: <http://www.it.uc3m.es/>

Alberto Garcia Martinez
U. Carlos III de Madrid
Av. Universidad 30
Leganes, Madrid 28911
Spain

Phone: +34 91 6248782
Email: alberto@it.uc3m.es
URI: <http://www.it.uc3m.es/>

