



Information Society
Technologies



ROUTING IN NEXT GENERATION

Title:	Document Version:
Deliverable D2.1 Alternative approaches to multihoming	1.2

Project Number:	Project Acronym:	Project Title:
035167	RiNG	Routing in Next Generation

Contractual Delivery Date:	Actual Delivery Date:	Deliverable Type* - Security**:
30/06/2007	23/09/2007	PU

* Type: P – Prototype, R – Report, D – Demonstrator, O – Other

** Security Class: PU- Public, PP – Restricted to other programme participants (including the Commission), RE – Restricted to a group defined by the consortium (including the Commission), CO – Confidential, only for members of the consortium (including the Commission)

Responsible and Editor/Author:	Organization:	Contributing WP:
Marcelo Bagnulo	UC3M	WP2

Authors (organizations):

Abstract:
<p>This document describes an alternative multihoming solution for IPv6 called SHIM6 that preserves the scalability of the routing system. This is achieved through the adoption of multiple provider aggregatable address blocks in the multihomed sites. In order to overcome the limitations resulting from this multi-addressed configuration, the SHIM6 multihoming architecture that defines a new SHIM6 Sub-Layer within the IP layer that maps the addresses presented to the higher layers and the addresses actually used for forwarding packets is proposed. In order to prevent the introduction of new vulnerabilities due to the adoption of this multihoming solution, a number of security measures are included. In particular, a new type of addresses, called HBA, is proposed. HBA are cryptographically in nature and they incorporate a hash of the prefix set available in the multihomed site in their interface identifier part. The result is that the addresses of each multihomed host will be inherently bound and they can be safely used interchangeably. The verification process associated with the usage of HBAs is extremely cost efficient since it does not require the usage of asymmetric cryptography. In addition, a signaling protocol for exchanging the address mapping information between communicating parties is defined and a failure detection and recovery procedure is specified.</p>

Keywords:
IPv6, Multihoming, Protocol Design, Security, Scalability, SHIM6

Revision History

The following table describes the main changes done in the document since its creation.

Revision	Date	Description	Author (Organization)
v1.0	27/08/2007	Document creation	Marcelo Bagnulo (UC3M)
v1.1	28/08/2007	Update document format	César Olvera (Consulintel)
v1.2	22/09/2007	Final review	Jordi Palet (Consulintel)

Executive Summary

In this deliverable we describe an alternative architecture for IPv6 multihoming called SHIM6 that preserves global routing system scalability thanks to the usage of provider aggregatable addressing. The proposed solution defines a new SHIM6 Multihoming Sub-Layer that performs mappings between identifiers and locators. The SHIM6 Multihoming Sub-Layer contains two planes: the data plane and the control plane. The data plane identifies the data packets that need to be translated and performs the conversion between identifiers and locators when it is required. In order to do that, the data plane uses information about mappings between identifiers and locators obtained through the signaling protocol of the control plane of the SHIM6 Multihoming Sub-Layer. The proposed signaling protocol allows the creation of session contexts in the parties involved in the communication. In these contexts, information about identifier to locator mappings is stored. The protocol defines the required exchanges to create, manage and dispose such mappings. In addition, the defined protocol has built-in security features that prevent the introduction of new vulnerabilities in the Internet because of the adoption of this multihoming architecture. In particular, the usage of a new type of addresses called HBA that are cryptographic in nature and that incorporate a one-way hash of the prefix set available in the multihomed site in their interface identifier part. The result is that all the addresses available in a multihomed host are inherently bound to each other, and the host can securely use them interchangeably. In addition, the proposed protocol integrates protection measures against DoS attacks, such as a 4-way handshake to establish the initial session context and flooding attack protection through the verification of the receiving parties' willingness to accept incoming packets using the REAP exchange.

The SHIM6 architecture preserves the scalability of the global routing system, it does not introduce new vulnerabilities in the Internet and it is easy to adopt since it does not require complex management in the end-site. In particular with respect to the last point, it should be noted that none of the presented mechanisms require manual configuration, allowing poorly managed sites to easily deploy the proposed solution. Moreover, as opposed to the multihoming technique currently deployed in IPv4, the fault tolerance capabilities of the solution do not require complex configuration of BGP or other protocols. In the presented approach, fault tolerance support is directly implemented in the end-hosts and it works without requiring user configuration. This enables the adoption of the presented solution in SOHO environments that lack of network administration.

Table of Contents

1. Introduction.....	6
2. The SHIM6 Multihoming Sub-Layer.....	8
2.1 The Data Plane.....	8
2.2 The Control Plane.....	9
2.2.1 Signaling protocol security.....	9
2.2.2 Session Context Creation.....	11
2.2.3 Re-Homing Procedure.....	13
2.2.4 Locator Set Management.....	14
3. Standardization of SHIM6.....	15
4. References.....	16

Table of Figures

Figure 1-1: PA addressing for multihoming 6

1. INTRODUCTION

As the costs associated with Internet access downtime continue to skyrocket [1] [2], sites are massively adopting configurations for high availability in order to insulate themselves from outages. Among such techniques, multihoming, i.e. the connection to the Internet through several providers, is becoming a preferred option. Multihoming is a redundant configuration that allows the site to preserve its Internet connectivity as long as one of its providers is available. In the multihoming solution currently deployed in the IPv4 Internet, the multihomed site announces a route to its address blocks through all the providers using BGP [3]. The result is that multiple routes towards the multihomed site are available in the inter-domain routing system. In particular, when one of the paths to the multi-homed site suffers an outage, the inter-domain routing system will recover and automatically will provide an alternative route to the multihomed site through another provider, if one is available. While this solution provides the fault tolerance and path selection features required to a multihoming solution, it presents limited scalability, since each multihomed site contributes with at least one routing table entry in the already oversized inter-domain routing tables. The impact of the current multihoming solution can be measured not only in additional router memory space and processing power but also in a delayed global BGP convergence, affecting the complete Internet community.

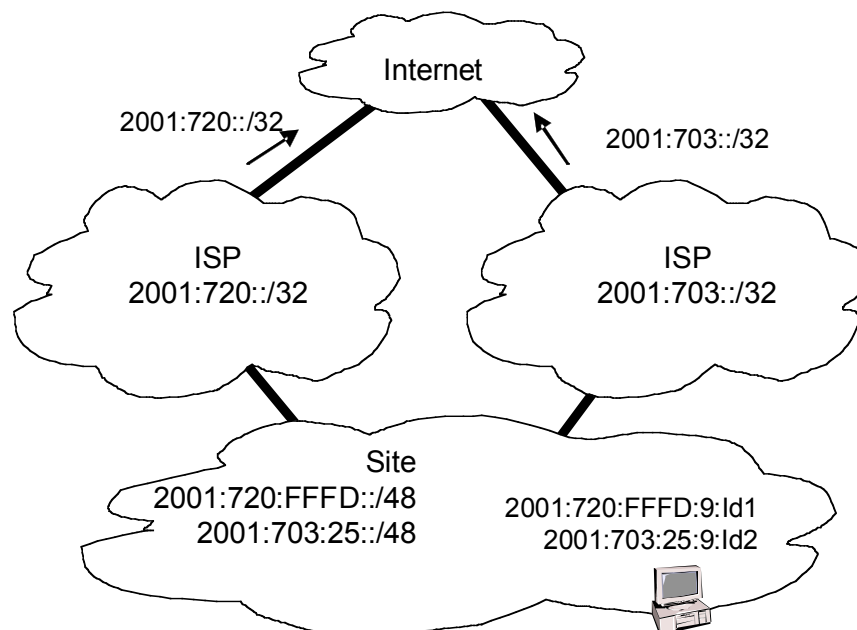


Figure 1-1: PA addressing for multihoming

Because the support of the multihoming solution currently deployed in the IPv4 Internet is becoming challenging even for the current number of multihomed sites, this approach is deemed unsuitable for the expected number of multihomed sites in the future IPv6 Internet, especially when considering that the wide adoption of low-budget broadband access technologies such as ADSL or CATV will enable multihoming in SOHO environments. As a consequence, an alternative multihoming solution for IPv6 is needed. The requirements imposed to the new solution essentially include all the benefits provided by the incumbent solution, i.e. fault tolerance and traffic engineering capabilities, and also an enhanced scalability with respect to the

number of multi-homed sites and other relevant Internet parameters [6]. In order to preserve routing system scalability, aggressive route aggregation can be achieved through provider-based aggregation [7], precluding the injection of routes associated with individual multi-homed end-sites. When Provider Aggregatable (hereafter PA) addressing is used, multi-homed sites obtain one prefix per each one of their providers. Consequently, as each provider will only announce its own prefix to the rest of the Internet, a given provider will be used to reach the multihomed site only when the destination addresses used belong to the prefix associated with the provider. So, in order to be reachable through all of the providers of the site, each host within the multihomed site will have to configure multiple addresses, one per provider. The resulting configuration is illustrated in the figure above.

Even if this setup guarantees the scalability of the multihoming solution, such multi-addressed configuration is not without difficulties of its own when attempting to provide the additional features mentioned above. In particular, this configuration is not directly compatible with the ingress filtering techniques usually deployed in service providers [8]. The incompatibility is caused by the lack of coordination between the IPv6 source address selection mechanism, performed by the host, and the path selection mechanism, performed by the intra-site routing system. As long as outgoing packets are routed through the provider that has delegated the prefix contained in the source address, packets will flow freely; but when those packets are routed through a different ISP, they will be discarded by the ingress filtering mechanism due to source address incompatibility. It must be noted that because of this issue, packets may be discarded even in a scenario without failures.

Additional difficulties arise when providing fault tolerance capabilities. In order to preserve established communications through outages, the endpoints of the communication have to change the addresses used during the lifetime of the communication according to the available providers. Moreover, this address replacement has to be performed in a transparent fashion with respect to transport and application layers, in order to actually preserve the established communication. Current applications and transport layers, such as TCP and UDP, identify the endpoints of a communication through the IP addresses of the nodes involved, implying that the IP addresses selected at the communication establishment time must remain invariant through the lifetime of the communication. But as it has been presented earlier, once that an outage has occurred in one of the available ISPs, the associated address becomes unreachable, so an alternative address has to be used in order to convey packets to the multi-homed host. These two constraints impose that after an outage, packets must carry a different address, corresponding to an available ISP, but they have to be presented to transport and application layers as if they contained the original address, in order to be recognized as belonging to the established communication. Such approach requires additional mechanisms in both ends of the communication in order to preserve a coherent mapping between the IP addresses presented to the transport and application layers and those addresses actually contained in the packets.

We describe the Shim6 architecture for the provision of multihoming in IPv6 that deals with all the aforementioned concerns. The proposed IPv6 multihoming architecture introduces a new Shim6 Sub-Layer within the IP layer that will perform the required mapping between the addresses that are presented to the upper layer protocols and the addresses that are actually used for exchanging packets in the network. In addition the architecture proposes the adoption of an intra-site routing paradigm that takes into account the source address, so that source hosts can determine through the selection of the source address, the exit path of the packets. Such feature provides ingress filtering compatibility and provides the end host with the required tools to support end-to-end fault tolerance.

2. THE SHIM6 MULTIHOMING SUB-LAYER

As we describe in the introductory section, an additional SHIM6 multihoming sub-layer within the IP layer is needed in order to preserve established communications through outages in multihomed environments. This SHIM6 sub-layer performs the mapping between the address presented to the transport/application layer and the address actually used to forward packets. Architecturally, this new sub-layer will be placed between the IP routing sub-layer and the IP endpoint sub-layer [11] [12]. The IP routing sub-layer is where the forwarding functionalities like the mechanisms to determine the next hop of outgoing packets are located. The IP endpoint sub-layer contains end-to-end mechanisms such as IPSec [13]. The SHIM6 Sub-Layer will reside in between those two sub-layers, since it will provide the mapping between the addresses used for forwarding and the addresses used by the higher layers. The addresses used for forwarding will be referred hereafter as locators; the addresses presented to the higher layer will be referred as identifiers. It should be noted that a given address can be simultaneously a locator and an identifier, since it can be used for forwarding and also be presented to the higher layers. As a matter of fact, currently IP addresses are mostly used both as locators and identifiers in the communications.

The Multihoming Sub-Layer has two planes: the data plane and the control plane. The data plane performs the functions related with exchanging data packets. Those include the identification and translation of those packets that carry a locator that differs from the expected identifier, and the identifier and locator selection for exchanged packets.

The control plane performs the functions related to the establishment and management of the state required to properly perform the mapping between identifiers and locators. In the next sections we will describe in detail the protocols and mechanisms involved in each plane.

2.1 The Data Plane

The data plane of the SHIM6 multihoming layer will perform the translation between the identifiers and the locators used for exchanging packets. For that purpose, it has to properly identify the packets that need to be translated. This is so because not all packets will require the translation, since when no outage has occurred, a single address is used both as identifier and as locator. When a communication is first established, the initiating application will select the destination identifier and it may also select the source identifier. If this is not the case, the IP layer will perform the source address selection. As long as the selected addresses pair is operational, i.e. if packets containing the selected source and destination address reach the desired destination, those addresses will be used as locators in the packets exchanged. In this case, there is no need to perform any kind of mapping, since the identifiers and locators are the same.

However, in case of an outage, the locators used for the communication need to be changed while the initial identifiers have to be preserved. In this case, the SHIM6 Multihoming Sub-Layer will have to perform the translation between the identifiers and the locators. In order to properly perform this mapping, the Multihoming Sub-Layer needs to identify the packets that need to be translated. Considering that all the addresses available in a multihomed host can be used both as locators and as identifiers, and that it is possible that the same address is being simultaneously used as a locator in one communication and as an identifier in another one, the address contained in the received packets is not enough to identify the packets carrying addresses

that need to be translated. In order to overcome this difficulty, additional information needs to be carried in the packets themselves. There are several possible approaches to include the information needed to properly identify and de-multiplex the received packets [12], but in this approach we consider the usage of a new SHIM6 Multihoming Extension Header that carries the required information to perform the mapping. This approach is architecturally clean, because it does not overload the semantic of existent fields of the IPv6 header. The cost of this approach is the additional overhead that it imposes. However, it should be noted that, as we mentioned earlier, only some packets carry the new header, since it will only be used when an outage has occurred.

Summarizing, the data plane of the SHIM6 Multihoming Sub-Layer will perform the mapping between identifiers and locator. For incoming packets, it needs first to identify the packets that require to be translated. The identification is possible thanks to a new SHIM6 Multihoming Extension Header that will be carried in the packets that contain an address other than the identifier. For outgoing packets, the Multihoming Sub-Layer has to determine if the addresses used as identifiers are suitable to be used as locators or need to be translated into other locators. If they need to be translated, the Multihoming Sub-Layer will perform the translation and will include the SHIM6 Multihoming Extension Header.

For all this processing, the SHIM6 Multihoming Sub-Layer requires state information that stores the valid mappings between identifiers and locators, and the reachability status of each locator pair. All this information is acquired through the signaling protocol of the control plane of the Multihoming Sub-Layer that will be presented in the next section.

2.2 The Control Plane

Within the control plane of the SHIM6 multihoming sub-layer we can identify the following elements [14]:

- a mechanism to create the session context associated with the communication between two end-points,
- the tools required to manage such context,
- a mechanism to re-home the established communication, i.e. to change the locators actually used in the packet exchange, and
- a mechanism to close the session.

In this section we will describe each one of the aforementioned mechanisms. However, the separation of identifiers and locators can introduce new vulnerabilities in the resulting architecture. So in order to provide the required protection from such attacks, we will propose a security architecture for the multihoming protocol. Since such architecture will impact the whole protocol, we will start by performing the security analysis.

2.2.1 Signaling protocol security

We will first perform a threat analysis in order to identify the potential vulnerabilities that a multihoming signaling protocol can introduce in the architecture. Next we will describe the proposed security solution, based on a new type of cryptographic addresses called Hash Based Address, and we show that they provide proper protection to the defined protocols.

2.2.1.1 Threat analysis

The control plane will essentially manage the identifier-to-locator mapping functions within the SHIM6 Multihoming Sub-Layer. In particular though the associated signaling protocol, information about the different mappings between locators and identifiers will be exchanged. This signaling protocol could be used by an attacker, to launch different types of redirection attacks [15]. A redirection attack essentially consists in abusing of the multihoming mapping mechanisms to create false identifier-to-locator mappings. There are two possible types of redirection attacks: hijacking attacks and flooding attacks. In a hijacking attack, the attacker uses the multihoming protocol to induce a victim to associate one of the attacker's locators to the target identifier. This means that when the victim sends packets to the target identifier, he will be actually sending packets to the attacker. Through this attack, the attacker manages to steal the target's identity and hijack the communications between the target and the victim. In a flooding attack, the attacker starts a communication with a given party, for instance, downloading a heavy flow from a streaming server. Then the attacker uses the multihoming protocol to re-home the communication to a victim's locator, causing the server's flow to flood the victim.

Once that the threats have been identified, we should next determine the security level required for the solution. It seems wise to require that any new mechanisms adopted in the architecture must not introduce additional vulnerabilities to the network that are not currently present in the Internet. Redirection attacks are feasible in today's Internet when the attacker is along the path between the two communicating nodes. In other words, current communications are susceptible to man-in-the-middle attacks, and an attacker that is capable of intercepting the packets can hijack a communication if no additional security measures such as IPsec [13] or TLS [16] are adopted. This means that the prevention of man-in-the-middle attacks is not a goal of the multihoming protocol security mechanisms. On the other hand, current TCP/IP communications are not susceptible to the so-called time-shifted hijacking attacks [17]. In a time-shifted attack, the attacker launches the attack from an on-path position and then leaves, but the effects of the attack remains long after he has left. In the case of the multihoming protocol, the attacker would remain on-path the time required to create the false identifier-to-locator mapping in the victim's node and then he would leave. However, the victim will preserve the false mapping long after the attacker has left. Because this is a new vulnerability introduced by the multihoming protocol, the security mechanisms of the protocol must prevent such time-shifted hijacking attacks.

2.2.1.2 Hash Based Addresses

The security architecture proposed for the multihoming protocol is based in the use of Hash Based Addresses [18], [19]. Hash Based Addresses are a new type of global IPv6 addresses that incorporate into the interface identifier part a cryptographic one-way hash of the prefix-set available in the multihomed site. The result is that the binding between all the addresses of a multihomed host is encoded within the addresses themselves, providing hijacking protection. Through this tool, any node that is communicating with a multihomed node can efficiently verify that the alternative addresses proposed for continuing the communication are bound to the initial address through a simple hash calculation.

In order to benefit from the proposed security mechanism, the addresses of each multihomed host have to constitute an HBA set. In a general multihoming scenario, a multihomed host X attached to a link where N 64-bit prefixes [20] are available ($PX1::/64$, $PX2::/64$, ..., $PXN::/64$) generates the interface identifier part of each one of its addresses as a 64-bit hash of the prefix set available in the link and a random nonce. Including a random nonce enables the generation of multiple HBA sets associated to the same prefix set. After generating the interface identifier

parts, the addresses of the HBA set are generated by prepending the different prefixes of the prefix set with the interface identifier parts.

The output of the described procedure is a set of N HBAs that carry information about the prefixes available in the multihomed site within their interface identifier part. Each one of the generated addresses will have a different prefix from the input prefix set, while their interface identifier part will contain information about the complete prefix set in the form of a hash of the full prefix set. Because of their nature, each address contains information about all the other addresses of the set, and a receiver can easily verify if two addresses belong to the same set through a cost effective hash operation. After this verification, the receiver can securely use them interchangeably. In the next section we will describe how HBAs can be used to prevent time-shifted hijacking attacks in the Session Context Creation exchange.

2.2.2 Session Context Creation

Consider the case where one of the parties involved in a current or a future communication decides to create a multihoming session context in order to benefit from the enhanced fault tolerance capabilities of multihoming. In this context we will refer to the party that decides to initiate the session context creation process as the initiator and the other party involved in the communication as the receiver. (Note that the initiator may differ from the actual initiator of the communication itself).

We assume that at least one of the parties involved in the communication is multihomed and that the multihomed host(s) has generated its multiple addresses as an HBA set associated with the multiple prefixes available in its multihomed site.

The initiator must request the receiver the creation of a session context associated with a pair of identifiers. In order to do that, the initiator will issue an I1 packet that is a Session Context Initialization Request message. The goal of this first message is essentially to provide some form of Denial-of-Service (DoS) attack protection. Because session context creation implies the storage of session related information in the receiver, the Session Context Creation exchange can be used to launch DoS attacks against the receiver. In order to prevent such attacks, the receiver will refuse to create any session context related state until the initiator has proven its location through this preliminary packet exchange. This security measure does not fully preclude the possibility of DoS attacks, but at least it imposes an additional effort to the attacker, and provides some tracing capabilities. So, this message just informs the receiver about the initiator's will to establish a session context. Upon the reception of this message, the receiver can either ignore it and discard it (if there is no multihoming support or no interest in enabling multihoming for this particular communication) or continue with the Session Context Creation procedure. However, even in the latter case, the receiver will not create any state, but it will simply reply with a R1 message that is a Session Context Initialization Reply message containing the Session Specific Information that is generated as a hash of the pair of identifiers plus a secret of the receiver.

Once that the initiator has the Session Specific Information available, he can use it to prove to the receiver his location, and that he is ready to establish the session context. For that, the initiator sends an I2 message that is a Session Creation Request message containing the pair of identifiers and the locator set available at the initiator and the Session Specific Information. In addition, the context tag that will be carried in the SHIM6 Multihoming Extension Header to identify data packets that contain alternative locators will also be exchanged in the I2 message. In order to provide the required protection from time-shifted hijacking attacks, HBA features will be used. This means that in addition to the aforementioned information, relevant HBA parameters are also included in this exchange. Those include the full prefix set and the random nonce used in the HBA set generation. Note that the HBA parameter information and the locator

set information may be redundant, so it is not required to include it twice in the message. Just including the prefix set and not including the detailed locator set is enough, since the HBA set can be re-generated at the receiver side using the HBA set parameter information.

Upon the reception of the I2 message, the receiver will validate the received information, verifying that the initiator's identifier is included in the HBA set associated to the HBA set parameters received in the message. The verification process consists in regenerating the HBA set using the HBA generation procedure with the received parameters (prefix set and random nonce) and then verifying that the initiator's identifier is included in the generated set. If this verification is successful, the receiver creates the session context containing the received information, and it replies with a R2 message that is a Session Context Creation Acknowledgement message, in which the receiver will include its own locator set and the corresponding HBA validation information. The initiator verifies then that the receiver's identifier is included in the HBA set associated with the parameters received in the R2 message. If the verification is successful the initiator associates the received HBA set to the session context state.

The presented Session Context Establishment includes several security features, which were detailed during the description of the message exchange. In particular it is robust against DoS attacks thanks to the usage of the I1 and R1 message exchange. In addition, it is robust against time-shifted hijacking attacks due to the usage of the HBA technique. We will next illustrate this argument by presenting a possible attack and quantifying the effort required to perform it.

In the scenario presented in the previous section, a multihomed host HostX is communicating with another host HostY. HostX has generated its addresses through the HBA address set generation algorithm, resulting in $PX1:IX1, PX2:IX2, \dots, PXN:IXN$. HostY has a single address $PY:IY$. HostX and HostY are communicating using addresses $PXi:IXi$ and $PY:IY$ respectively. Consider now an attacker HostZ that has the intention of redirecting the communication to one of the addresses of HostX to an alternative address. We will assume that it is enough for the attacker to redirect the communication to any address of a given prefix, $PZ::/64$. The rationale behind this assumption is that HostZ has access to any address of the considered prefix.

So, in order to hijack the communication, HostZ must introduce a new prefix in the prefix set used for generating the HBA set of HostX. For that, HostZ is required to obtain a combination of prefix set and random nonce such as:

- 1- $PXi::/64$ and $PZ::/64$ are included in the prefix set
- 2- $PXi:IXi$ is included in the resulting HBA set

The other inputs apart from PXi and PZ prefixes may be changed at will by the attacker, for instance, the random nonce and the remaining prefixes of the prefix set can be altered. In any case, in order to obtain the desired HBA set, the attacker needs to try with different inputs, for instance with different random nonces, until the above two conditions are met. The expected number of times that the generation procedure will need to be repeated until the desired outcome is reached depends on the number of hash bits included in the interface identifier part of the HBAs. Since we are considering 64-bit long interface identifiers and that the "u" and the "g" bits are preserved as defined in the IPv6 specification [21], the expected number of iterations is $O(2^{62})$. We believe that the resulting security is enough for protecting regular traffic that flows unprotected through the network from potential redirection attacks introduced by the multihoming mechanisms, since the resulting protection is similar to the one offered by other network security protocols such as SeND [22]. However, as processing power increases, the protection provided by this mechanism decreases, since the amount of time required to try with

2^{62} different random nonces also decreases. Additional mechanisms can be used to improve the obtained protection, for instance artificially increasing the effort required for generating a valid HBA set, similar to the Sec mechanism used in CGA [23].

2.2.3 Re-Homing Procedure

When a failure occurs that affects the communication, the re-homing procedure is executed to divert the communication towards an alternative address pair. The re-homing procedure involves the following mechanisms: First a failure detection mechanism is required to identify when an outage has occurred. Once the failure has been identified, alternative locator pairs need to be explored. When a working locator pair is found, the communication is then diverted to it. In this section we will describe these mechanisms.

2.2.3.1 Failure Detection and Alternative Path Exploration Mechanism

The REAP protocol [24] defined in the Shim6 architecture [12] provides path failure detection and alternative path exploration capabilities between two multihomed hosts. It relies in two mechanisms, namely, the failure detection mechanism and the path exploration mechanism.

The failure detection mechanism is based on the Forced Bidirectional Detection (FBD) technique, which consists on making sure that whenever there is data traffic in one direction, there is also traffic in the other direction. This is accomplished by injecting additional control messages (called Keep Alive messages) when needed, which guarantee that the frequency of traffic in the reverse direction is above a predetermined threshold. The result is that when there is an ongoing data communication between two REAP peers, both peers can expect an incoming traffic frequency that is above the predetermined threshold defined by REAP. If the incoming traffic frequency is below this threshold, then this implies that a failure has occurred. In other words, after a given period of time no traffic has been received a failure on the path is assumed and the alternative path exploration mechanism is triggered.

The REAP protocol relies on two timers, the Keep Alive Timer and the Send Timer, and a control message, namely the Keep Alive message. The Keep Alive Timer TKA is started each time a node receives a data packet from its peer, and stopped and reset, each time the node sends a packet to the peer. When the Keep Alive Timer expires, a Keep Alive message is sent to the peer. The Send Timer TSend, defined roughly as three times the Keep Alive Timer plus a deviation to accommodate the Round Trip Time, is started each time the node sends a packet and stopped each time the node receives a packet from the peer. If no answer (either a Keep Alive or data packet) is received in the Send Timer period a failure is assumed and a locator path exploration is started. Consequently, the Send Timer reflects the requirement that when a node sends a payload packet there should be some return traffic within Send Timeout seconds. On the other hand, the Keepalive timer reflects the requirement that when a node receives a payload packet there should a similar response towards the peer within Keepalive seconds (if no traffic is interchanged, there is no Keep Alive signaling).

The path exploration mechanism starts whenever the node has not received any packet during a fixed period of time (Send Timer). A path may become invalid either because one of the locators used may become invalid or inoperational, or the pair itself has been declared as inoperational. A full exploration mechanism should check all possible pairs of source/destination locators until at least one working locator pair is found. Instead of using a request/response approach the first of both sides which detects the failure tries each of the peer's locators sending probes through each of its interfaces. Each probe carries information about the current state of the communication and the probes which have been received so far through the rest of the interfaces. The state of the

connection can be one of three possible states: a) Operational, when both peers can see each other, b) Exploring, when one of the peers have detected a problem and has currently not seen any traffic from the peer or c) Inbound_OK, when the node sees traffic from the peer but the peer does not see any traffic from the node. The information related with the rest of the probes received, which is carried on every probe allows the end hosts to be able to know which are the locator pairs working in the outgoing direction, on the case there are multiple probes. The path exploration mechanism ends when both peers have received a probe confirming that the peer can see them. It should be noted that the defined exploration mechanism is capable of discovering locators pairs that are working in only one direction (i.e. unidirectional reachability) thanks to the information about all received probes contained in all the reply probes.

In the current implementation once a node detects a failure, it starts the path exploration mechanism. A Probe message is sent to test the current locator pair, and if no responses are obtained during a period of time called Retransmission Timer (TRTx), the nodes start sending Probes testing the rest of the available address pairs, using all possible source/destination address pairs. Once a probe is received by the node, it sends another probe to the peer indicating that it is seeing traffic from it (Inbound_OK). After receiving this last probe the peer will answer confirming the reception of this last probe and indicating that the new locator pair is in the Operational state.

These Probe messages are used to confirm reachability and can be used as an address verification mechanism to modify the state of the locator being probe to ACTIVE. Note that each end point of the communication explores unidirectional reachability and based on its observations decides the pair of locators to use in a not coordinated way. Therefore the pair of locators selected by each end host may be different.

At the end of the path exploration mechanism, each host will have a pair of ACTIVE locators which can be used to continue the communication.

2.2.4 Locator Set Management

From a general perspective, it may be required to change the available locator set during the lifetime of the session. For that reason an Update message and an Update Acknowledgment messages are defined. The Update message contains the complete set of available locators for the session. It should be noted that the semantics of the operation is to substitute the available locator set for the session with the one included in the message.

Because of security issues, the new locator set must be contained in the HBA set associated with the identifier. The result is that this message will be used to withdraw locators when they fail and reconstitute them back once they are available again. The reasons for withdrawing a locator are essentially local. For instance if a failure occurs and one of the prefixes available in the multihomed site is deprecated, then it makes sense to withdraw the correspondent locator from the pool of available locators for the established sessions. It should be noted that such address is removed as a locator, but it can still be used as an identifier. Once that the peer receives an Update message, it will confirm it sending an Update Acknowledgment message.

3. STANDARDIZATION OF SHIM6

The solution described in this deliverable is being standardized by the IETF in the following documents:

- E. Nordmark, M. Bagnulo, Shim6: Level 3 Multihoming Shim Protocol for IPv6 Internet-Draft (work in progress), draft-ietf-shim6-proto-08.
- J. Arkko, I. Beijnum, Failure Detection and Locator Pair Exploration Protocol for IPv6 Multihoming, Internet-Draft (work in progress), draft-ietf-shim6-failure-detection-06.
- M. Bagnulo, Hash Based Addresses (HBA), Internet-Draft (work in progress), draft-ietf-shim6-hba-02.

4. REFERENCES

- [1] T. Wilson, "The Cost Of Downtime", 1999, Internet Week.
- [2] S. Yin, K. Twist, "The Coming Era of Absolute Availability", 2003, RHK.
- [3] I. Van Beijnum, "BGP", 2002, Oreilly.
- [4] G. Huston, "Commentary on Inter-Domain Routing in the Internet", 2001, Internet Engineering Task Force (IETF), RFC 3221.
- [5] C. Labovitz, A. Ahuja, A. Bose, "Delayed Internet Routing Convergence", 2000, SIGCOMM 2000.
- [6] J. Abley, B. Black, V. Gill, "Goals for IPv6 Site-Multihoming Architectures", 2003, Internet Engineering Task Force (IETF), RFC 3582.
- [7] V. Fuller, T. Li, J. Yu, K. Varadhan, "Classless Inter-Domain Routing (CIDR): an AddressAssignment and Aggregation Strategy", 1993, Internet Engineering Task Force (IETF), RFC 1519.
- [8] P. Ferguson, D. Senie. "Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing.", 1998, Internet Engineering Task Force (IETF), RFC 2267.
- [9] M. Bagnulo, A. García Martínez, J. Rodríguez, A. Azcorra. "The Case for Source Address Dependent Routing in Multihoming", 2003, Proceeding of the International Workshop on Multimedia Interactive Protocols and Systems, Lecture Notes of Computer Science. Springer-Verlag.
- [10] C. Huitema, R. Draves, M. Bagnulo, "Ingress filtering compatibility for IPv6 multihomed sites", 2004, Internet Engineering Task Force (IETF), Internet Draft (work in progress).
- [11] G. Huston, "Architectural Approaches to Multi-Homing for IPv6", 2004, Internet Engineering Task Force (IETF), Internet Draft (work in progress).
- [12] E. Nordmark, M. Bagnulo, Shim6: Level 3 Multihoming Shim Protocol for IPv6 Internet-Draft (work in progress), draft-ietf-shim6-proto-08.
- [13] S. Kent, R. Atkinson, "Security Architecture for the Internet Protocol", 1998, Internet Engineering Task Force (IETF), RFC 2401.
- [14] M. Bagnulo, J. Arkko, "Functional decomposition of the M6 protocol", 2004, Internet Engineering Task Force (IETF), Internet Draft (work in progress).
- [15] E. Nordmark, T. Li, "Threats relating to IPv6 multihoming solutions", 2004, Internet Engineering Task Force (IETF) RFC4218.
- [16] S. Blake-Wilson, M. Nystrom, D. Hopwood, J. Mikkelsen, T. Wright, "Transport Layer Security (TLS) Extensions", 2003, Internet Engineering Task Force (IETF), RFC 3546.

- [17] P. Nikander, J. Arkko, T. Aura, G. Montenegro, E. Nordmark, “Mobile IP version 6 Route Optimization Security Design Background”, 2004, Internet Engineering Task Force (IETF), RFC4225.
- [18] M. Bagnulo, A. García-Martínez, A. Azcorra. “Efficient Security for IPv6 Multihoming” ACM Computer Communications Review, pp 61-68. Vol. 35, nº 2. April 2005.
- [19] M. Bagnulo, “Hash Based Addresses (HBA) ”, 2007, Internet Engineering Task Force (IETF), Internet Draft (work in progress).
- [20] R. Hinden, S. Deering, “Internet Protocol Version 6 (IPv6) Addressing Architecture”, 2003, Internet Engineering Task Force (IETF), RFC 3513.
- [21] S. Deering, R. Hinden, “Internet Protocol, Version 6 (IPv6) Specification”, 1998, Internet Engineering Task Force (IETF), RFC 2460.
- [22] J. Arkko, J. Kempf, B. Sommerfeld, B. Zill, P. Nikander, “SEcure Neighbor Discovery (SEND)” , Internet Engineering Task Force (IETF), RFC3971, 2005.
- [23] T. Aura, “Cryptographically Generated Addresses (CGA)”, Internet Engineering Task Force (IETF), RFC3972, 2005.
- [24] J. Arkko, I. Beijnum, Failure Detection and Locator Pair Exploration Protocol for IPv6 Multihoming, Internet-Draft (work in progress), draft-ietf-shim6-failure-detection-06, 2007.
- [25] R. Draves, “Default Address Selection for Internet Protocol version 6 (IPv6)”, 2004, Internet Engineering Task Force (IETF), RFC 3484.
- [26] R. Moskowitz, P. Nikander, P. Jokela, T. Henderson, “Host Identity Protocol”, 2004, Internet Engineering Task Force (IETF), Internet Draft (work in progress).
- [27] E. Nordmark, “Multi6 Application Referral Issues”, 2004, Internet Engineering Task Force (IETF), Internet Draft (work in progress).
- [28] D. Johnson, C. Perkins, J. Arkko, “Mobility Support in IPv6”, 2004 Internet Engineering Task Force (IETF), RFC 3775.
- [29] M. Bagnulo, A. Garcia-Martinez, I. Soto, “Application of the MIPv6 protocol to the multihoming problem”, 2003, Internet Engineering Task Force (IETF), Internet Draft (work in progress).
- [30] E. Nordmark, “Multihoming without IP Identifiers”, 2004, Internet Engineering Task Force (IETF), Internet Draft (work in progress).
- [31] M. Bagnulo, A. García Martínez, A. Azcorra, C. de Launois. “An Incremental Approach to IPv6 Multihoming”. Computer Communications Journal, in press.