

Network Working Group
Internet-Draft
Intended status: Informational
Expires: April 18, 2013

T. Boot
Infinity Networks B.V.
October 15, 2012

BRDP for Homenet
draft-boot-homenet-brdp-00.txt

Abstract

This document describes the Border Router Discovery Protocol (BRDP) and all of its related components. BRDP enables multi-homing for small to medium sites, including Homenets, using Provider Aggregatable IPv6 addresses. It describes a mechanism for automated IP address configuration and renumbering, a mechanism for optimized source address selection and a new paradigm for packet forwarding, for support of multi-homed sites. BRDP prevents ingress filtering problems with multi-homed sites and supports load-balancing for multi-path transport protocols. This work also prevents routing scalability problems in the provider network and Internet Default Free Zone because small to medium multi-homed size sites would not need to request Provider Independent address blocks.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

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."

This Internet-Draft will expire on April 18, 2013.

Copyright Notice

Copyright (c) 2012 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 Simplified BSD License.

Table of Contents

1. Introduction	4
1.1. Detection of Homenet Perimeter interfaces on Border Routers	5
1.2. Propagation of Border Router information	5
1.3. Address configuration with Border Router information	6
1.4. Address selection with Border Router information	6
1.5. Routing based on Border Router information	7
1.6. Deployment of the Border Router Discovery Protocol	7
1.7. Requirements Language	8
2. Reference Scenarios	8
2.1. Single-homed site	8
2.2. Small multi-homed site or DMZ	10
2.3. Medium multi-homed site	12
2.4. Medium multi-homed site with ULAs and DHCP server	16
2.5. MANET site	18
3. Border Router Discovery Protocol (BRDP)	20
3.1. Border Router Information Option (BRIO)	21
3.2. BRDP processing	22
3.2.1. BRDP message generation and transmission	23
3.2.2. BRDP message reception	24
3.2.3. BRIO-Cache maintenance	25
3.2.4. BRDP loop prevention	26
3.3. Unified Path Metric (UPM)	27
4. BRDP based Address Configuration and Prefix Delegation	27
4.1. Border Router selection	28
4.1.1. Border Router Selection based on UPM	28
4.2. Address autoconfiguration	29
4.2.1. Address and prefix configuration with SLAAC or DHCP	29
4.2.2. Address generation and configuration for Routers	29
4.2.3. Support for Unique Local Addresses (ULA)	30
5. BRDP based Source Address Selection	30
5.1. Address Selection for dynamic DNS	30
6. BRDP based Routing	30
6.1. Problems with default gateway routing	31
6.2. Default gateway routing replaced with BRDP Based Routing	31
7. BRDP and IRTF RRG goals	32

7.1.	Scalability	33
7.2.	Traffic engineering	33
7.3.	Multi-homing	33
7.4.	Loc/id separation	33
7.5.	Mobility	34
7.6.	Simplified renumbering	34
7.7.	Modularity	34
7.8.	Routing quality	34
7.9.	Routing security	34
7.10.	Deployability	35
8.	Currently unaddressed issues	35
9.	Acknowledgements	35
10.	IANA Considerations	35
11.	Security Considerations	35
12.	Change log	36
13.	References	36
13.1.	Normative References	36
13.2.	Informative References	37
	Author's Address	38

1. Introduction

*** Note to the reader *** This Internet Draft is submitted as an early version for a proposal for the Homenet working group. This version is a merge from earlier documents. Now that is a single document, it is to be adjusted to comply to the Homenet scenario. This is work in progress.

IPv6 provides basic functionality for multi-homing, since nodes can have multiple addresses configured on their interfaces. However, it is difficult to utilize the advantages of this, as there is a strong tendency shielding the network topology from hosts and in general routing does not support multi-homing very well. As a result, it is difficult or impossible for a host to utilize available facilities of the network, such as multi-path. Also scalability of the Internet routing system is getting a problem due to a high demand of Provider Independent (PI) addresses.

The Border Router Discovery Protocol (BRDP) enhances the IPv6 model by enabling automated renumbering in dynamically changing multi-homed environments, such that routers and hosts cooperate on address configuration and path selection. BRDP utilizes Provider Aggregatable (PA) addresses and uses them as locator. Mapping identifiers to locators is out of scope of BRDP, also because other solutions exists or are being worked on. All these solutions work fine with BRDP, as long as they don't break IPv6.

BRDP applies to edge networks. These networks can be fixed, for example enterprise networks, small offices / home offices (SOHO) or home sites (Homenets). BRDP also can be used in wireless access networks, for example wireless access networks such as 3G or 4G, wireless LANs or mobile ad hoc networks (MANETs). A nice attribute of BRDP is that it supports multi-homing in heterogeneous networks, meaning that e.g. a Homenet network can have multiple wired broadband and 3G/4G connections to the Internet simultaneously.

In a multi-homed network, nodes are connected to the Internet via multiple exit points, possibly via multiple providers. [RFC5887] argues that if a site is multi-homed, using multiple PA routing prefixes, then the interior routers need a mechanism to learn which upstream providers and corresponding PA prefixes are currently reachable and valid. Next to that, these upstream providers or PA prefixes may change over time. This requires a dynamic renumbering mechanism that can handle planned or unplanned changes in the prefixes used. BRDP proposes a mechanism for automated renumbering in larger networks that goes beyond hosts in a single subnet.

BRDP uses the following key elements:

- o Propagation of available Border Routers and corresponding prefixes, described in Section 3;
- o Address autoconfiguration and prefix delegation, using BRDP provided hints, described in Section 4;
- o Source address selection, using BRDP provided hints, described in Section 5;
- o Packet forwarding to the Border Router that corresponds with the source address prefix, in case the destination address is not found in the routing domain, described in Section 6.

1.1. Detection of Homenet Perimeter interfaces on Border Routers

For fully automated deployment in Homenets, it is required that routers can discover automatically their uplink interfaces, that connect the Homenet to ISPs. Some mechanisms for automatic detection are described in [I-D.kline-default-perimeter].

After detection of an uplink interface to an ISP and reception of a prefix, the router starts acting as a border router. It starts acting as a DHCP server, with support of prefix delegation. It also configures at least an address out of the assigned prefix. This address is used as Border Router address and DHCP server address.

The BRDP protocol can also be used to assist perimeter detection. A router interface on which Border Router information is received should not be identified as an uplink interface to an ISP.

1.2. Propagation of Border Router information

The propagation of available Border Routers and corresponding prefixes is implemented as an extension on the Neighbor Discovery Protocol [RFC4861]. Border Router Information Options (BRIOS) are sent with Router Advertisements, and contain information about the Border Router, such as:

- o - the Border Router address;
- o - the prefix that corresponds with that Border Router;
- o - cost indication of the path via that Border Router to the core network, i.e. the Internet Default Free Zone (DFZ).

BRIOS are disseminated downstream through the network. All nodes store the information from BRIOS they receive in a BRIO cache.

Border routers with multiple prefixes send out a BRIO for each of these prefixes. In a multi-homed network, nodes will receive multiple prefix information, from multiple upstream Border Routers or from a Border Router with multiple prefixes.

1.3. Address configuration with Border Router information

Routers can generate IPv6 addresses, with regular SLAAC [RFC4862]. Generation is based on Prefix Information Option from upstream routers and optionally on information in the BRIO cache, e.g. using the prefix with the lowest cost to the Internet. In addition, routers may generate /128 IPv6 address-prefixes for a management interface, based on a Border Router prefix. Routers set up reachability to these addresses automatically, by adding the generated address or prefix in the routing protocol.

With BRDP, routers automatically learn Border Routers that act as DHCP server or relay agent [RFC3633]. When routers detect an alternate path to the DFZ, with no corresponding assigned address or prefix already, new prefixes are requested for using this alternate path.

Prefixes, of which the path to the DFZ is no longer available, are put 'out of service' by routers, meaning they are not used for address assignments anymore. Optionally, if the cost to the DFZ through a Border Router is far higher than via other available paths, a router can put the corresponding prefix out of service also. Prefixes that are out of service are released.

Prefixes that are in service are configured on interfaces with a 64-bit prefix length and advertised with a Prefix Information Option in Router Advertisements. The Prefix Information lifetime is copied from lifetime information in the BRIO cache.

Hosts can use the BRDP provided information together with the Prefix Information to autoconfigure addresses, based on IPv6 Stateless Address Autoconfiguration [RFC4862]. A host may also use DHCPv6 to get addresses or "Other configuration", using multicast or with unicast to the BRDP learned DHCP server address.

1.4. Address selection with Border Router information

Nodes with multiple configured addresses need to select a source address for outgoing connections. Default Address Selection for IPv6 [RFC6724] defines a mechanism, used as default behavior. It is open to more advanced mechanisms or site policies. BRDP provided information can be used for a more advanced mechanism, where the hosts select automatically a source address that corresponds with a path with the lowest cost to the DFZ. When multiple Border Routers are available, automatic load distribution and multi-path transport becomes available.

1.5. Routing based on Border Router information

Network Ingress Filtering [RFC2827] describes the need for ingress filtering, to limit the impact of distributed denial of service attacks, by denying traffic with spoofed source addresses access. It also helps ensure that traffic is traceable to its correct source network. Ingress Filtering for Multihomed Networks [RFC3704] provides solutions for multi-homed sites. However, the proposal applicable for PA addresses requires careful planning and configuration. It suggests routing based on source address, and a path on each Border Router to all ISPs in use, either with a direct connection or with tunnels between all Border Routers. It is hard to make such mechanisms work in an automated fashion, or mechanisms are not supported on Border Routers used today. As an evolutionary approach, BRDP provided information is to be used to forward packets to their destination without ingress filtering problems. The BRIO cache contains a mapping between Border Routers and the addresses that do pass ingress filtering. So the packet forwarding heuristic can be straightforward: send packets, where the destination is not in the routing domain itself, to the Border Router that owns the prefix of the source address.

Hosts use information in the Default Router List to select a default router. For selecting the best paths, hosts may use next hop selection based on source address and path costs to the corresponding Border Router, if such information is available to the host. Such next hop determination is useful for destinations outside the edge network, i.e. the destination address does not belong to a prefix in the BRIO cache.

1.6. Deployment of the Border Router Discovery Protocol

Enabling BRDP in an existing network is straightforward. First, all routers have to be updated for BRDP support. At this step, Border Router information is propagated in the network enabling BRDP assisted address autoconfiguration and prefix delegation and BRDP assisted source address selection. The second step is updating all routers with the BRDP based routing mechanism. To enable this mechanism the default gateway is removed from the routing table. This second step is a flag day operation. Rolling back is easy, by just re-inserting the default gateway.

After the update of the network, additional border routers can be added to the network and will be used automatically. Also a renumbering event will take place without any manual intervention.

BRDP does not provide session continuity when paths are broken. Mobility solutions are in place, or are work in progress. Recently,

interesting developments are work in progress, such as MPTCP [I-D.ietf-mptcp-multiaddressed] and ILNP [I-D.rja-ilnp-intro]. BRDP is very useful for both of these protocols.

1.7. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

2. Reference Scenarios

This section describes the use of BRDP in five different scenarios: a single homed Homenet, multi-homed site or DMZ, a medium multi-homed site, a medium multi-homed site with ULA with DHCP server and a MANET site.

2.1. Single-homed site

This scenario discusses BRDP operation for single-homed home networks. The scenario is taken from [I-D.ietf-homenet-arch].

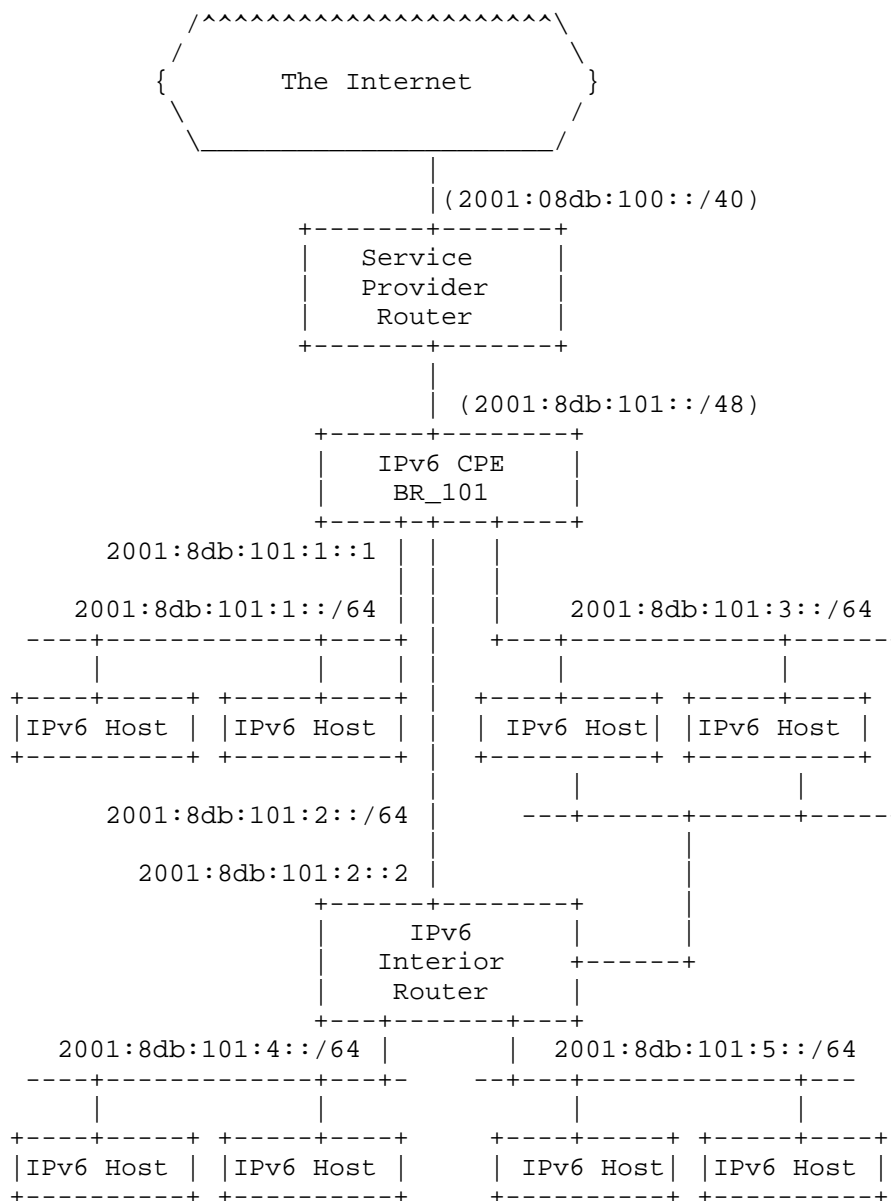


Figure 1: Scenario 1: Single-homed site

The CPE device has to discover the link to the ISP and has to get assigned an IPv6 prefix, in this scenario 2001:8db:101::/48. The CPE configures itself a global unique address prefix 2001:8db:101:1::1/64, assumed on the first interface in the homenet. It may

configure additional global unique address on other interfaces, but this is not required. This is existing functionality which is not updated by BRDP.

The CPE starts sending router advertisements. It also checks received router advertisements on already existing prefixes for the /48 prefix it has assigned by the ISP. In this scenario there is no other CPE, so no on-link prefixes exist. The CPE allocates and bind additional prefixes for all its interfaces, and send Router Advertisements with the Prefix Information Option. By then it has configured 2001:8db:101:1::/64, 2001:8db:101:2::/64 and 2001:8db:101:3::/64. The CPE router also acts as DHCP server, for the ISP provided prefix.

Now, the IPv6 hosts in the middle row learn these prefixes from Prefix Information Options sent by the CPE. They can configure IPv6 addresses, either with SLAAC or DHCP. Also, the IPv6 Interior Router can configure an IPv6 address, in this scenario on the link with prefix 2001:8db:101:2::/64.

The IPv6 Interior Router also receives the router advertisement with the onlink prefix 2001:8db:101:3::/64 on its interface on the right. It could configure an address in the prefix, but because it has already a globally unique address configured, there is no need for this. Question is if the router should echo the prefix as on-link. In this BRDP proposal, it doesn't. It is not the "delegated prefix holder".

Before the IPv6 hosts on the lower row can get their addresses, the IPv6 Interior Router has to be assigned two more prefixes. Here, BRDP starts playing its role. The CPE router advertises itself with Border Router Information Option, in its Router Advertisement. The IPv6 Interior Router learns this information, and gets the two needed prefixes from the CPE Router, using unicasted DHCP messages to the (CPE) Border Router. Two prefixes are assigned and configured, 2001:8db:101:4::/64 and 2001:8db:101:5::/64.

For full connectivity, the homenet uses an interior routing protocol. BRDP is agnostic on the routing protocol used.

2.2. Small multi-homed site or DMZ

This scenario discusses BRDP operation for multi-homed Small Office - Home Office (SOHO) networks and De-Militarized Zones (DMZ). The scenario is shown in Figure 2. Each provider assigns a PA /48 prefix to its customers. All addresses and prefixes are configured completely automatically. The feature of BRDP that adds value in this scenario is BRDP based Border Router selection for multi-homed

hosts. This is enabled by using BRDP based forwarding.

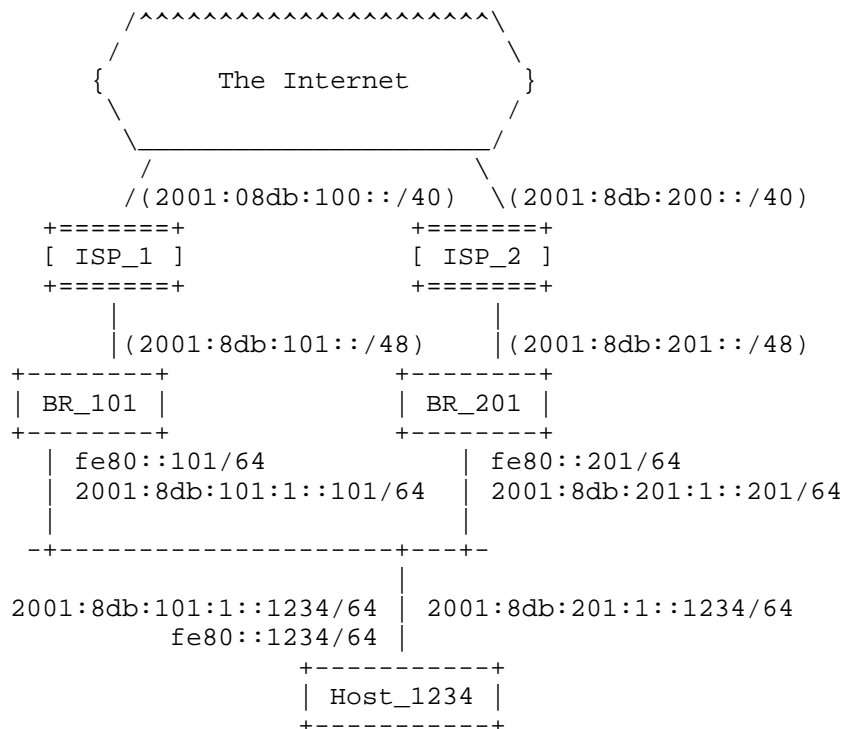


Figure 2: Scenario 2: multi-homed Small Office - Home Office (SOHO) network or DMZ

In this scenario, Host_1234 has configured two addresses using SLAAC [RFC4862], one with prefix 2001:8db:101:1::/64 from Border Router BR_101 and one with prefix 2001:8db:201:1::/64 from Border Router BR_201. Host_1234 has learned these prefixes from Prefix Information Options sent by both Border Routers according to [RFC4861]. The host has learned via BRIOs that these prefixes belong to Border Routers. The host can use optimal paths by selecting BR_101 as default router for all packets with a source address with prefix 2001:8db:101:1::/64 and default gateway BR_201 for all packets with a source address with prefix 2001:8db:201:1::/64. Non-optimal default router selection on hosts is handled by the routers, "misdirected" packets are forwarded to the correct Border Router.

BRDP enables routers to deliver non-optimal directed packets from attached hosts towards a Border Router that owns the prefix of the source address, if such a Border Router exists. In the above scenario, a packet sent from Host_1234 with source address 2001:8db:

201:1::1234 to default router BR_101 would be dropped due to on an ingress filter, when no mechanism is in place to redirect the packet. BRDP based forwarding provides such a mechanism automatically. Instead of dropping the packet, BR_101 forwards it to BR_201.

2.3. Medium multi-homed site

This scenario discusses BRDP operation for medium sized multi-homed networks. The difference with the previous scenario is that the network paths between hosts and the Border Routers have intermediate routers. The scenario is shown in Figure 3. The added value of BRDP in this scenario is the discovery of Border Routers for hosts and routers beyond the first hop as well as Border Router Selection for hosts and routers.

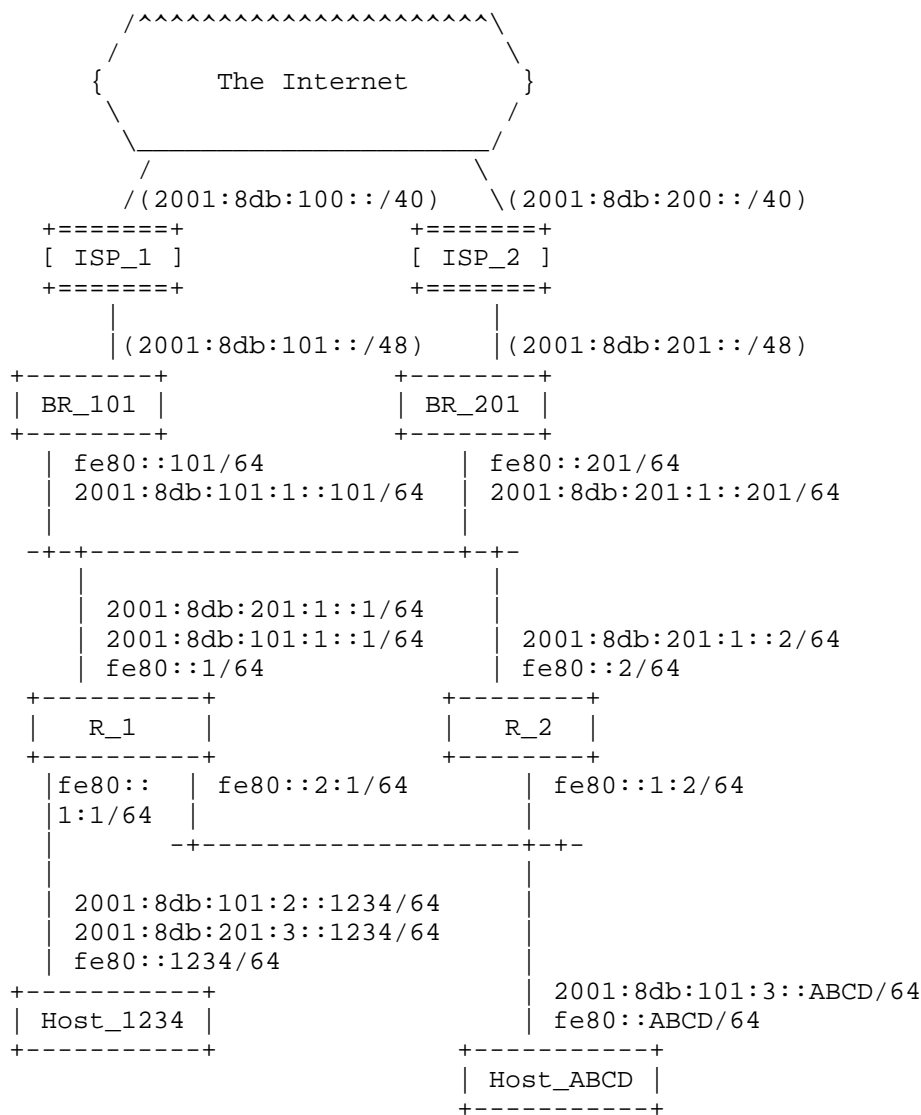


Figure 3: Scenario 3: medium sized multi-homed network

Routers can learn advertised on-link prefixes automatically via the Prefix Information Option in IPv6 ND RAs. In this scenario, routers R_1 and R_2 learn prefix 2001:8db:101:1::/64 from BR_101 and prefix 2001:8db:201:1::/64 from BR_201. Routers may autoconfigure addresses on their interfaces. In this example, R_1 has configured addresses from both providers on its upstream interface, R_2 only configured an address based on the prefix of BR_201. If the routers run a routing

protocol, the learned prefixes are made reachable in the network. In the next steps of the autoconfiguration process, the prefixes and addresses on the other links are automatically configured, but first we discuss the BRDP messages that are disseminated through the network.

Routers automatically learn Border Routers and mapping between prefixes and Border Routers using BRDP. The diagram in Figure 4 depicts BRIO message dissemination in scenario 2. The two Border Routers advertise their own address and corresponding prefix with an address prefix. Nothing prevents them from forwarding each other's BRIO message, although resending BRIO information on non-MANET interfaces is not useful. Both routers R_1 and R_2 forward both Border Router address prefixes, using separate BRIOs in RAs, on downstream interfaces. In this way all routers and hosts in the network are aware of all reachable Border Routers and corresponding prefixes.

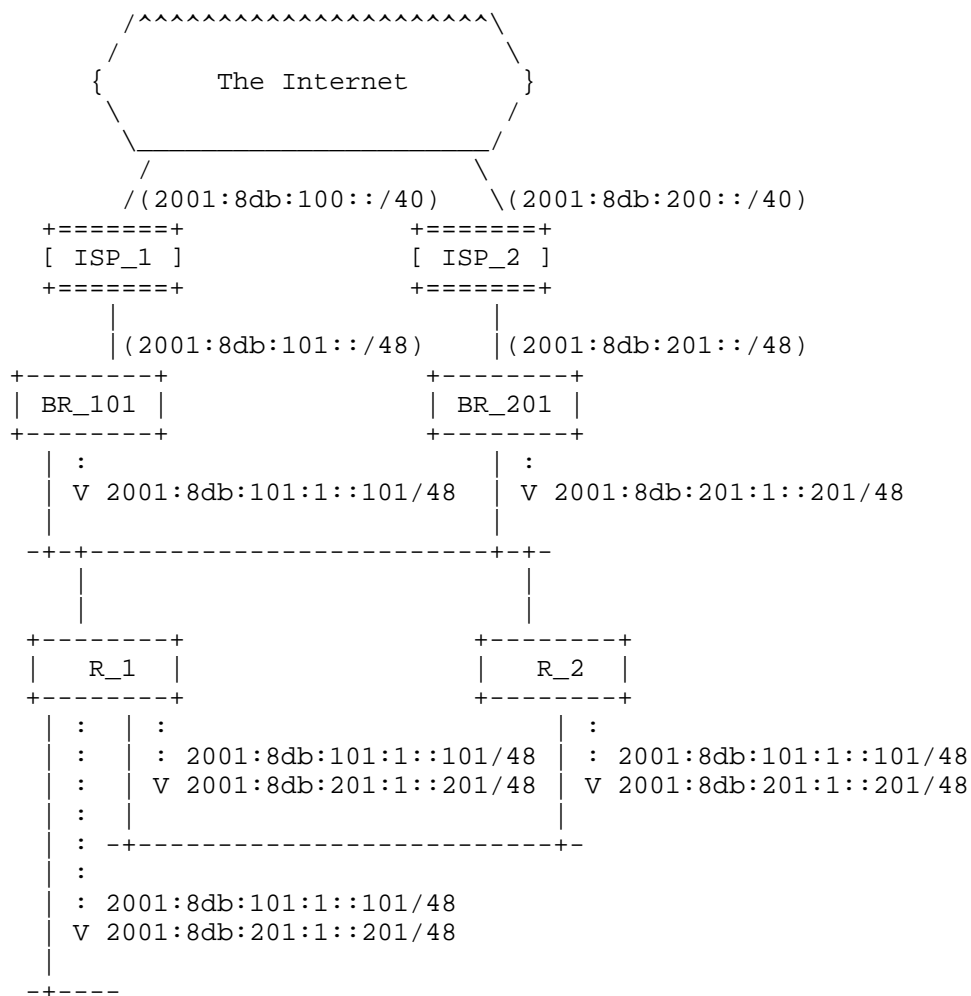


Figure 4: BRIO dissemination in Scenario 3

Routers are not required to configure global addresses on each interface. In the example, only the interface pointing to the Internet has configured global addresses. Routers may also use a (logical) management interface for global reachability.

So, the one-hop neighbours of BR_101 and BR_201, being R_1 and R_2, have learned the prefixes and configured addresses on their upstream interfaces. And all nodes in the network have learned the Border Router prefixes. The next step is to get configured addresses on the hosts in Figure 2. This is done by using DHCP Prefix Delegation. R_1 and R_2 request a prefix from either or both BR_101 and BR_201

for binding as on-link prefix on the links, and advertise those using Prefix Information Options to the hosts. This will result in a maximum of four prefixes that are advertised on the downlink of R_1 and R_2. Having multiple prefixes from the same ISP bound on a link is not useful. So a router requests a prefix from a Border Router only if no other prefix of that Border Router is advertised already by another router on this network segment.

In this example, R_1 has been delegated two prefixes by DHCP PD for the link with host Host_1234; 2001:8db:101:2::/64 and 2001:8db:201:3::/64. No other router is on this link. R_1 or R_2 has also been delegated a prefix on the link to host Host_ABCD; 2001:8db:101:3::/64. It cannot be seen in Figure 2 which router has been delegated the prefix, nor if another prefix for this link has been delegated. No redundant prefix is delegated, as the routers learned with RA PIO already delegated prefixes for known Border Routers.

Now, Host_1234 and Host_ABCD can autoconfigure addresses for their interfaces. Host_1234 configures two addresses, one for each Border Router. Host_ABCD chooses not to use ISP_2.

Nodes R_1 and Host_1234 can use both providers, by using two configured global addresses. Any multi-path facility can be used, either on an application layer or with a multi-path transport protocol.

Host_ABCD may forward packets to the Internet via router R_1 or R_2. If R_2 is selected as default router, R_2 forwards the packets to BR_101 as this Border Router corresponds to the prefix of the source address 2001:8db:101:3::ABCD. This works well, even in this case where R_2 hasn't configured an address with a BR_101 prefix for itself, and selected a global address from the BR_201 prefix only.

2.4. Medium multi-homed site with ULAs and DHCP server

In this example, the scenario 2 is extended by adding Unique Local Addresses (ULA) for communication within the site itself. For simplicity there is only one ISP present. The ULA IP configuration, with prefix fd00:8db::/48, is managed by DHCPv6 server DHCP_201. The scenario is shown in Figure 5.

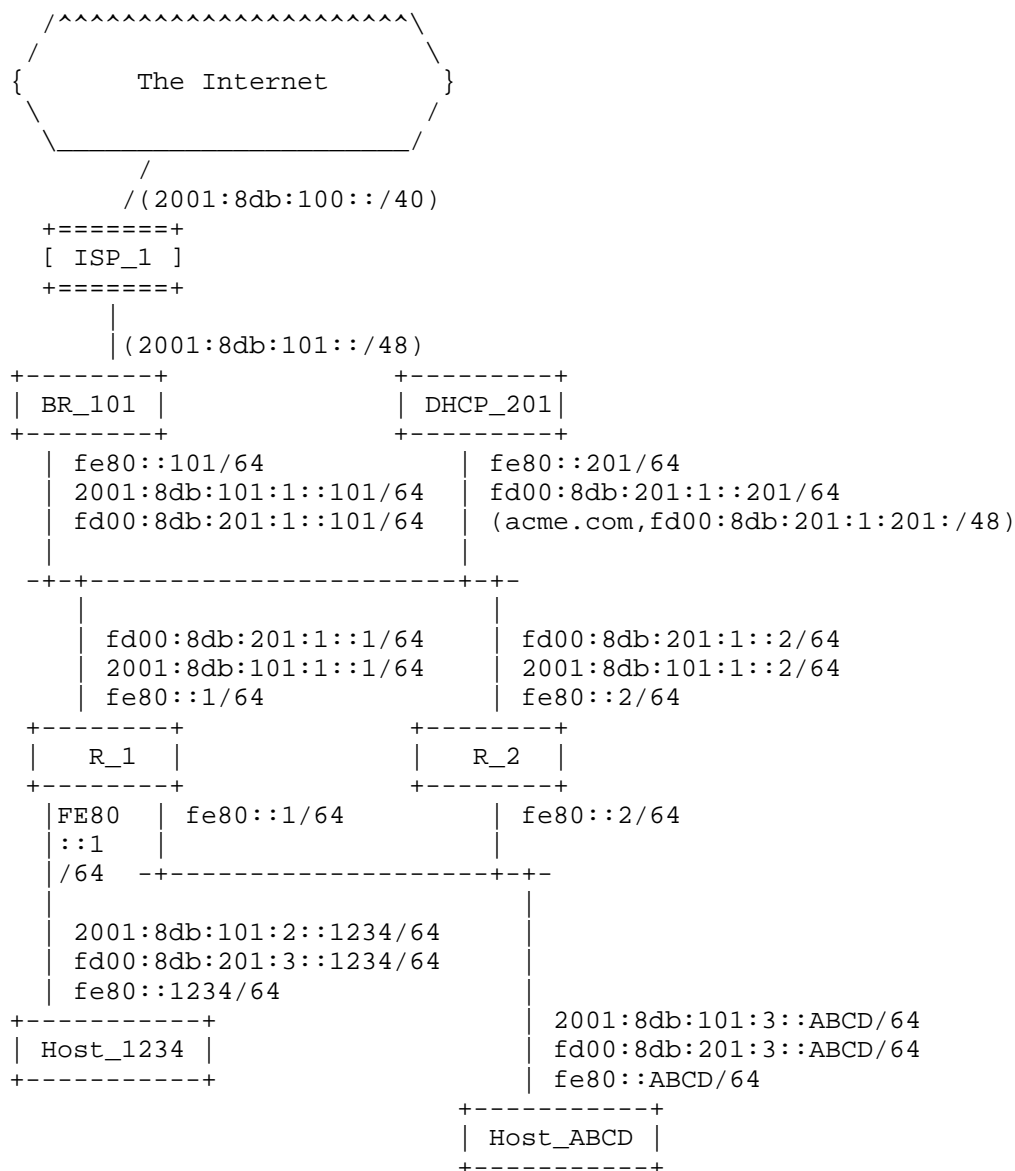


Figure 5: Scenario 4: a medium sized multi-homed site with ULAs and DHCP server.

In this scenario, all nodes have configured a ULA and a Global Unicast Address using prefix delegation in the way that was described in Section 2.2. ULA prefix delegation is automated just like PA addresses. The DHCP server is therefore implemented on a router, in

this case DHCP_201. This router advertises the ULA prefix with BRDP, here fd00:8db:201::/48.

Although BRDP provides automatic prefix and address configuration for ULA, a network administrator is free to configure it manually, along using BRDP for global addresses.

BRDP based ULA configuration with BRDP based routing would result in routing packets with ULA destinations outside the site to the originator of the ULA prefix, in this case router DHCP_201. DHCP_201 is not connected to the Internet or another site owning the ULA, so packets to non-existing destinations are dropped. DHCP_201 indicates such with the BRIO F-bit set, meaning the Border Router is floating.

This scenario, it is demonstrated that BRDP and DHCPv6 cooperate in address configuration. BRDP provides announcements of Border Routers and DHCP servers. Routers request prefixes with DHCP, and can request other parameters also. Such parameters are disseminated to other nodes, either with router advertisements or acting as DHCP server itself. Routers may also act as DHCP relay, redirecting address requests to the Border Router(s). The Router Advertisement M-bit and O-bit indicates availability of DHCPv6 services to attached nodes.

Difficulties may arise when both ULA and global addresses are used for Internet connectivity, e.g. when address translation is used. To distinguish, the Border Router not providing Internet connectivity informs nodes in the network using Service Selection suboption, similar to "Service Selection for Mobile IPv6" [RFC5149]. This procedure helps also for extranet connectivity. In this scenario, the ULA is used within the ACME Corporation, nodes are made aware by adding "acme.com" in the BRIO Service Selection Option.

It is for the reader to work out extensions for this scenario, where the ULA prefix originator is a Border Router to another site, e.g. a link from a branch office to a head quarter, or a ULA-only side connected to the Internet with NAT66.

2.5. MANET site

BRDP was developed for address autoconfiguration in MANETs. This scenario, see Figure 6 demonstrates the powerful multi-homing facilities provided to the MANET nodes.

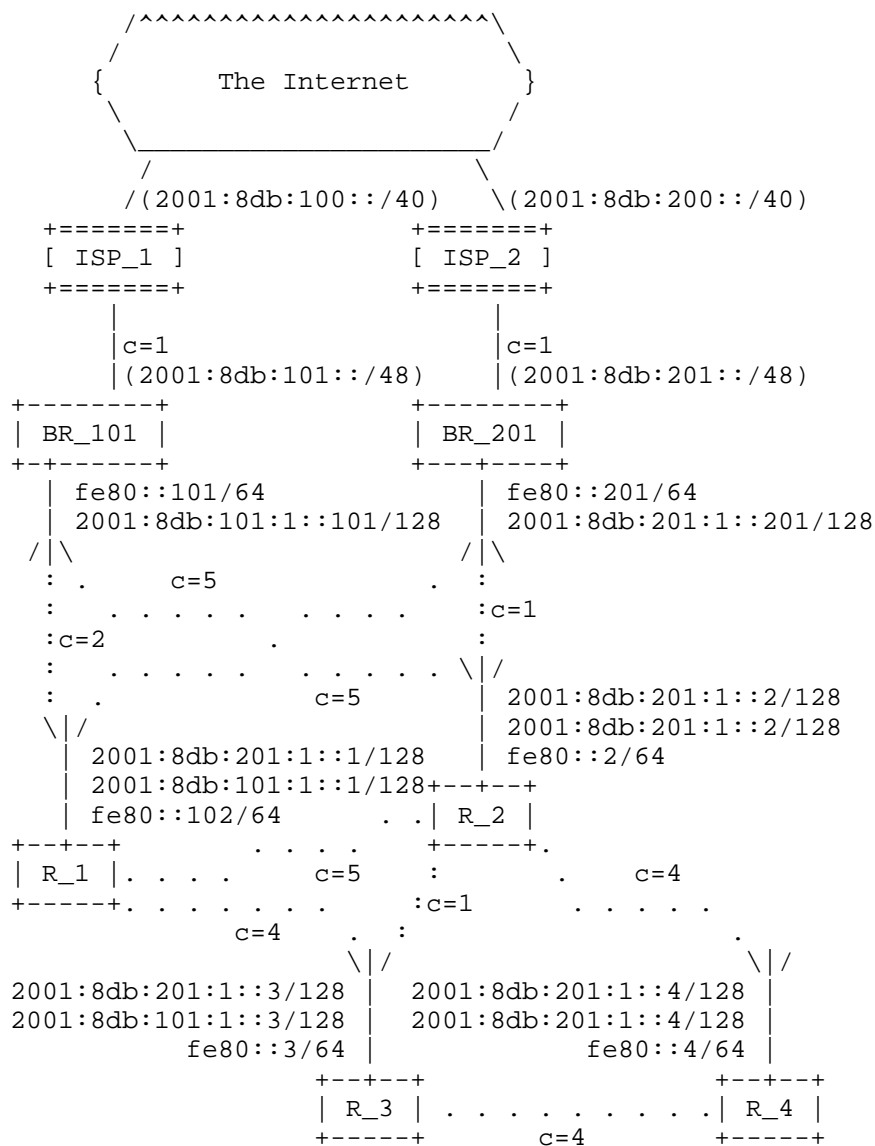


Figure 6: Scenario 5: a MANET site

On the MANET interfaces, addresses are configured using a 64-bit prefix provided by BRDP, appending it with a 64-bit Interface Identifier according to BRDP based address autoconfiguration. This creates a 128-bit prefix length as recommended in IP Addressing Model in Ad Hoc Networks [RFC5889]. Each MANET node has configured two global addresses, one for each ISP. With BRDP, the nodes are aware

of the cost of the path to the DFZ, defined as dimensionless metric for both directions of the patch. This enables optimized source address selection, and as an implicit result a Border Router and ISP selection. In the scenario, R_1 is near to BR_101 and the cost via this Border Router is lower than via BR_201. The table below shows costs to the DFZ for all nodes, via both ISPs. Paths with lowest costs are marked with *.

Costs to DFZ	Via ISP_1	Via ISP_2
BR_101	1*	7
BR_201	7	1*
R_1	3*	6
R_2	6	2*
R_3	7	3*
R_4	10	6*

The optimized source address selection facility is also of utility in the other scenarios. For example, the cost of the link to the ISP could be set depending of bandwidth and optionally on utilization. Nodes would use a near uplink to an ISP, and as a result some form of load distribution is enabled. Note that nodes still can use the alternative addresses, in fact it is recommended to use multi-path transport protocols for better load balancing and improved robustness.

For isolated MANETs, a DHCP server election mechanism can be used. Nodes may initiate to advertise a self-generated ULA. In such cases, it is recommended that a prefix is used with a 56-bit random ULA identifier (including random 16-bit Subnet ID) and 64-bit prefix length. Other nodes join this prefix, although some may wish to start or continue using their own prefix. The latter would occur in cases of a merge of previous isolated MANETs.

3. Border Router Discovery Protocol (BRDP)

BRDP is an extension to the IPv6 ND mechanism [RFC4861] that provides information about the reachability, availability, prefix information, quality and cost of upstream providers, and enables automated (re)numbering of possibly multi-homed routers and hosts.

BRDP adds the Border Router Information Option (BRIO) to the Router Advertisement (RA) of IPv6 ND. A BRIO contains all relevant information of an upstream Border Router and the corresponding

provider.

Border Routers initiate sending BRIO messages, other routers in the network disseminate the messages downstream through the network. Nodes store the information from received BRIOS in a BRIO cache, to be used for address generation, DHCP server discovery, address selection or packet forwarding.

A BRIO cache entry records reception of a BRIO for a single advertised prefix, received via a neighbor router. Border Routers that need to advertise multiple prefixes simply use multiple BRIOS, each with its own address prefix. For further processing of BRIO entries, only the entry with the lowest cost to a Border Router is used, for each Border Router.

When a node is multi-homed, it will receive BRIOS from multiple upstream Border Routers.

3.1. Border Router Information Option (BRIO)

The Border Router Information Option carries information that allows a nodes in the edge network to select and utilize a Border Router.

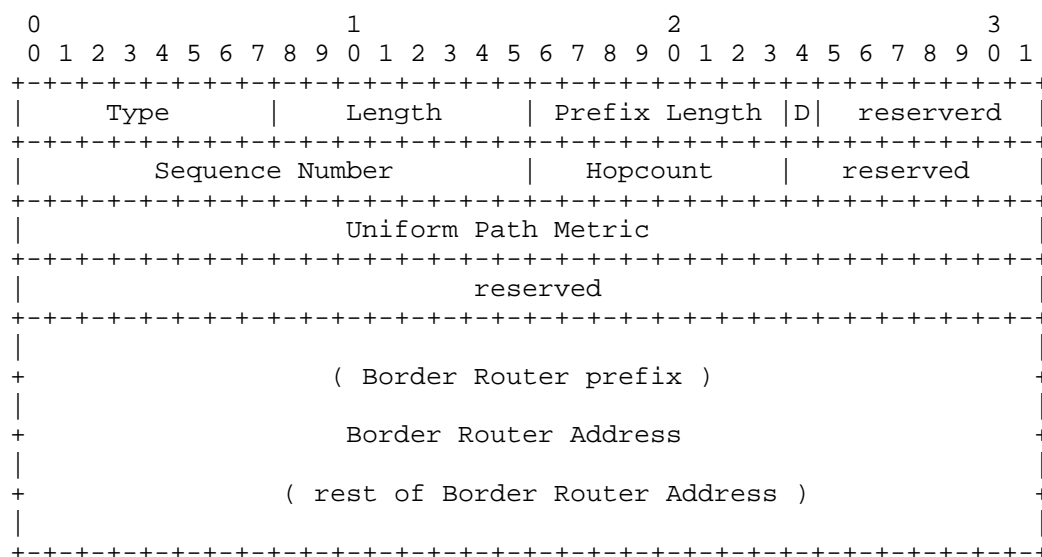


Figure 7: BRIO base option

Fields:

Type:

8-bit identifier of the Border Router Information Option type.
The value of this option identifier is to be determined.

Length:

8-bit unsigned integer. The length of the option (including the type and length fields) in units of 8 octets. A BRIO has a length value of 4.

Prefix Length:

8-bit unsigned integer. The number of leading bits in the Border Router Address, that indicates the assigned prefix for that Border Router. The Prefix Length is used for BRDP Based Routing, as described in Section 6.

DHCP (D):

When the D-flag is set, the Border Router is acting as a DHCP server or DHCP relay agent [RFC3315].

reserved:

Reserved bits. Currently unused, set to 0.

Sequence Number:

16-bit unsigned integer. It is set by the Border Router and incremented with each new BRIO it sends on a link. When forwarding downstream, the sequence number is not changed.

Hopcount:

8-bit field registering the number of hops from the advertizing Router to the Border Router. Border Routers send the initial BRIO with its Hopcount set to zero. Routers increment the Hopcount by one when forwarding a BRIO.

Uniform Path Metric (UPM):

A measure for the cost of the bi-directional path between the upstream Router and the Default Free Zone of the Internet. Uniform Path Metric is set to some initial value by the Border Router and is incremented by each Router forwarding the BRIO.

Border Router Address:

128-bit address of the Border Router. For reachability, the Border Router is expected to add this own address (prefix) in the routing system.

3.2. BRDP processing

The main BRDP processing functions of a Router are BRDP message generation, transmission and reception and the maintenance of a BRIO-

Cache. Routers forward BRDP messages using ICMP ND Router Advertisements.

3.2.1. BRDP message generation and transmission

A BRDP message is part of a Router Advertisement and includes a set of BRIOs. It provides the current state of (paths to) the Border Routers listed in the set of BRIOs. BRIOs originate from a Border Router, and contain initially metric information on connectivity to the Internet. BRIOs are forwarded downstream in the edge network.

When a Router sends a ICMP ND Router Advertisement, it SHOULD include a set of BRIOs by appending them to the message. The maximum number of BRIOs in a single BRDP message is a Router configuration parameter. BRIO selection for advertisement is done based on the information stored in the BRIO-Cache. BRIOs that do not pass the loop prevention check described in Section 3.2.4 SHOULD NOT be selected.

The UPM and Hopcount fields of the advertised BRIOs are updated. An UPM-increment, based on uniformed bi-directional link metrics, is added to the UPM and the Hopcount is incremented by 1. UPM-increment MAY be governed by a hysteresis and dampening mechanism. Also forecasted information MAY be used.

Each BRIO originating from a Border Router has an increased Sequence Number. This BRIO is forwarded in the edge network and refreshes entries in BRIO-Caches of downstream Routers.

Router Advertisements are sent in response to Router Solicitation messages or unsolicited with a uniformly-distributed random interval between MinRtrAdvInterval and MaxRtrAdvInterval [RFC4861]. The MaxRtrAdvInterval falls between a minimum of 30 milliseconds, specified in [RFC6275] and a maximum of 1800 seconds, specified in [RFC4861]. In addition, the Router MAY send a Router Advertisement when an important change in a to be sent BRIO would occur.

When a Router sends Router Advertisements more frequently than an upstream Router, this Router MAY repeatedly send BRIOs with a constant Sequence Number but with an updated UPM or Hopcount.

The ICMP ND Router Advertisement MAY include the Advertisement Interval Option [RFC6275]. This option contains the interval at which the sending router sends unsolicited multicast Router Advertisements.

A Router SHOULD inform downstream Routers in case the path to a previous advertised Border Router is lost, by at least 3 times

retransmitting the previously sent BRIO with a UPM value of 4294967295.

In case a Border Router loses its connection to the infrastructure it will lose its Border Router functionality and become a normal Router. In that case it performs the same procedure as a Router that has lost the path to a previous advertised Border Router.

For each Border Router listed in the BRIO-Cache, the UPM-loop-prevention-threshold and the Hopcount-loop-prevention-threshold variables are maintained. These variables are used by the loop prevention mechanism described in Section 3.2.4. The thresholds are set or updated when sending BRDP messages. When sending a BRIO with a higher Sequence Number than the previously sent BRIO for that Border Router, the threshold variables are set to the UPM and Hopcount values in BRIO to be sent. When sending a BRIO with the same Sequence Number as the previously sent BRIO, the loop-prevention-thresholds are independently updated if either the UPM or Hopcount of the outgoing BRIO is lower than their thresholds.

A Router that detects an attractive candidate BRIO but is prohibited from using it because of the loop prevention check, MAY send a (unicast) Router Solicitation message to the Border Router. The Border Router responds to such a Router Solicitation message with a new BRIO. Sending Router Solicitations MUST be rate limited. A next version of this document would include a specification for sending the unicast Router Solicitation message.

3.2.2. BRDP message reception

When a BRDP message is received, the Sequence Number fields of the contained BRIOS are checked; the Sequence Number of a received BRIO MUST be equal to or higher than the Sequence Number in the cache for an existing entry in the cache, with wrap-around checking. Otherwise, the BRIO will be discarded.

BRIO messages do not need to be forwarded at fixed time intervals, because the RA intervals on different Routers are not synchronized. Therefore, large gaps in Sequence Numbers may occur. Increment values between 0 and 65000 are accepted. Increment values between 65001 and 65535 are rejected.

Information in received BRIOS is stored in a BRIO-Cache table. Other information is stored as well, such as the BRIO upstream node, a timestamp indicating when the most recent message was received and the measured or signaled RA interval.

3.2.3. BRIO-Cache maintenance

Each Router maintains a BRIO-Cache that stores all information on Border Routers. Unique cache entries are maintained on (Border Router Address, address of the upstream router that forwarded the BRIO) tuples. This information is obtained by receiving BRIOs, or, in case of a Border Router, by getting information from the interface that connects to the Internet. The BRIO-Cache also maintains context information for the BRIO such as the BRIO sender, link metrics and UPM-increment for this sender, history, statistics and status information. History information includes a timestamp indicating when the most recent message was received and a measured or signaled RA interval. Status information includes the BRIO selection outcome for BRIO forwarding as explained in Section 3.2.1 and the Border Router selected for address generation as explained in Section 4.

BRIO entries in the BRIO-Cache stay valid for a certain period of time. During this period, they can be used for Border Router selection by the Router, for forwarding BRIOs and for address generation. BRIO-Cache information could also be useful for source address selection [RFC6724]. The lifetime of a BRIO is determined by using the timing information sent along with the RA ([RFC6275], section 7.3) or statistics of received BRIOs.

Some values in the BRIO-Cache can be updated independent of incoming BRDP messages. A Router MAY update the UPM-increment based on link quality measurements performed in an environment with changing link metrics. A Router SHOULD indicate in its BRIO-Cache which BRIO entries are currently selected for forwarding and for address generation. Border Router Selection MAY take place after the UPM of a BRIO entry has been updated.

In case the link to the Router from which a BRIO has been received is broken, the UPM and the Hopcount of the BRIO entry in the cache are set to the maximum value, i.e. 4294967295 and 255.

A cache cleanup routine SHOULD run at regular intervals to get rid of stale entries. Stale entries are removed when the entry is not updated for 5400 seconds or all of the following conditions are met:

- o The stale entry is not used by the Router itself for address generation.
- o The stale entry was not selected for forwarding in the last three Router Advertisement.
- o The stale entry was not recently updated by a received BRIO. In this context, recently is defined as the maximum of a) three times its own unsolicited multicast Router Advertisements interval and b) three times the senders unsolicited multicast Router Advertisements interval.

Cache entries MAY also be removed, under the condition that the BRIO-Cache has reached a configured maximum number of entries and a new, to be stored BRIO is received. A removal candidate is selected based on:

- o The candidate entry is not used by the Router itself.
- o The candidate entry was not selected for forwarding in the last Router Advertisement.
- o The candidate entry is redundant; other information for the same Border Router is stored in the cache with a better UPM and / or was received more recently.
- o The candidate entry is redundant; other information for the same Service Selection Identifier is stored in the cache with a better UPM and / or was received more recently.
- o The candidate entry is less attractive; other Border Routers are stored in the cache with better UPM and / or were received more recently.

3.2.4. BRDP loop prevention

A BRDP loop check mechanism prevents that a Router forwards an earlier advertized BRIO.

BRDP loop-free operation is guaranteed as long as at least one of the following conditions is true:

- o The to be sent BRIO has a higher Sequence Number than a BRIO for this Border Router that was sent before. The loop check mechanism uses wrap-around logic. Increments up to 32768 are acceptable (wrap-around logic needs checking).
- o The to be sent BRIO is generated from the same BRIO-Cache entry as the BRIO that was sent most recently.
- o The to be sent BRIO has the same Sequence Number as the BRIO for this Border Router that was sent before but the BRIO-Cache entry UPM is equal to or lower than the UPM-loop-prevention-threshold for this Border Router.
- o The to be sent BRIO has the same Sequence Number as the BRIO for this Border Router that was sent before but the BRIO-Cache entry Hopcount is equal to or lower than the Hopcount-loop-prevention-threshold for this Border Router.

In some circumstances, a Router would select a BRIO for forwarding that fails the loop prevention check. For example, the link to the upstream neighbor is lost and an alternative path is available, with a higher UPM and a higher Hopcount or with a lower Sequence Number. The Router cannot assure this candidate BRIO is not reflecting its own advertized message, therefore it should not send this BRIO. Instead, it sends a unicast Router Solicitation message to that Border Router.

3.3. Unified Path Metric (UPM)

Unified Path Metric (UPM) is a measure for the cost of the path between the Router and the Internet Default Free Zone. It is a united metric for both inbound and outbound paths. On each hop, the UPM is incremented with an UPM-increment, which is derived from the routing protocol and / or is obtained from lower layers.

It is on forehand not known what is more important; Border Router selection based on path metric to the Border Router or the path metric for the reverse path. In BRDP, UPM is used for optimizing Border Router selection for both the inbound and the outbound traffic. Note that actual traffic will use the path provided by the routing protocols, not by BRDP.

Since the UPM uses 32 bits, its maximum value is 4294967295. On each hop, an UPM-increment is calculated for each Router from which a BRIO has been received. UPM-increments have a value between 1 and 16777215, to support a 255 hop path, with maximum UPM increments.

Further discussion on metrics and how the UPM-increment value is determined is outside the scope of this document.

4. BRDP based Address Configuration and Prefix Delegation

BRDP supports stateless address autoconfiguration [RFC4862], DHCP managed IP configuration [RFC3315] and DHCP Prefix Delegation [RFC3633]. Routers can also use a variant of stateless address autoconfiguration, where BRDP provided information is used to configure Router management interfaces or used to configure off-link addresses, used in ad hoc networks [RFC5889].

BRDP adds topology awareness in address configuration. Nodes can configure multiple addresses, each to support a different facility. ULAs can be used for site internal traffic. Global addresses are mandatory for access to the Internet, assuming address translation is not used.

A node that is offered multiple prefixes for stateless address autoconfiguration or multiple addresses by DHCP chooses to configure one or more addresses. BRDP provides information for the candidate addresses. An important criterium is the costs of the path to the Internet DFZ. A node would prefer addresses with lower costs.

BRDP does not modify stateless address autoconfiguration and DHCP protocols, except that in a edge network, Routers may perform stateless address autoconfiguration from the Border Router

Information Option (BRIO), for their management or MANET interfaces. This enables edge network-wide address configuration, because BRIOs are disseminated over multiple hops in the edge network, while PIOs are link local messages only.

When a BRIO is stored in the BRIO cache table, the node checks if a corresponding address already exists for the Border Router from which this BRIO originates. If not, and a corresponding address for that Border Router is beneficial, address generation for that Border Router is triggered.

4.1. Border Router selection

When a node needs to communicate to nodes on the Internet, it MUST select a (set of) Border Router(s) for address generation. A node MAY generate multiple addresses for smooth handover implementing make-before-break or distributing traffic over multiple Border Routers. A description how Border Routers can be used concurrently is out-of-scope for this document.

Information concerning available Border Routers is kept in the BRIO-Cache.

The Border Router selection mechanism MAY be triggered by received BRDP messages, changes in metrics on links to neighbors advertising BRDP messages, changes in costs to Border Routers used or on a time-driven basis.

The Border Router selection algorithm SHOULD be based on UPM. UPM is used for selecting the Border Router with the best connectivity to the Internet. The Border Router selection algorithm MAY be extended with any other information. Future defined BRIO suboptions could provide additional information, such authorization and service selection. Border Router selection MAY be based on the type of the Border Router Address, e.g. a globally unique address or a unique local address.

Border Router selection does not provide nor select a routing path to the Border Router.

4.1.1. Border Router Selection based on UPM

The node uses the UPM for Border Router selection preferring the best bi-directional paths between the node and the Internet. Note that the BRIO UPM includes the initial metric set by the Border Router and is not solely a metric between the node and the Border Router. The initial metric set by Border Routers can be used for Border Router preference and for load balancing.

In order to use an up-to-date UPM in the selection procedure the UPM-increment is calculated by the node before selecting a Border Router. UPM is discussed in Section 3.3.

4.2. Address autoconfiguration

Nodes should use a topologically correct address when communicating with corresponding nodes on the Internet. Topologically correct addresses should be configured for each Border Router used.

4.2.1. Address and prefix configuration with SLAAC or DHCP

Nodes can use existing IPv6 address configuration protocols, such as SLAAC [RFC4862] and DHCP [RFC3315]. Nodes can use SLAAC based on prefix information, provided by the upstream router. Nodes may use DHCP multicast and neighbor routers will relay those packets to selected Border Routers with D-flag set or reply with DHCP parameters it has received from a Border Router before for itself.

Nodes using SLAAC may also query a DHCP server on a Border Router themselves for additional parameters, using the BRDP learned address of the DHCP server.

A Router should request a prefix for attached subnetworks, with DHCP-PD [RFC3633], where there is at that moment no on-link prefix for a selected Border Router.

4.2.2. Address generation and configuration for Routers

A generated address for a Router management interface or a MANET has a /128 prefix. It is constructed from a 64-bit Interface Identifier and a 64-bit prefix from the Border Router Address. The generated 128-bit address SHOULD be advertised in the routing system. The generated address may be used for user traffic, either inside the edge network or traffic towards the Internet.

For the Interface Identifier used, the BRDP-based Address Generation MUST implement a mechanism for generating a highly unique Interface Identifier. Known mechanisms are:

- o Modified EUI-64 format-based Interface Identifier, [RFC4291], based on IEEE 802 48-bit MAC address or IEEE EUI-64 identifier. However, this method does not guarantee identifiers are unique as duplicate MAC addresses can occur.
- o Generation of randomized Interface Identifiers, [RFC4941].
- o Well-distributed hash function, [RFC3972].

After Address Generation, RFC4429 Optimistic Duplicate Address Detection [RFC4429] should be used. A passive Duplicate Address

Detection, based on information in the routing protocol information bases could be used as an alternative. Still, uniqueness is not fully guaranteed. Main reasons for non-uniqueness are merging of edge network segments, node movement, node misbehavior or address spoofing attacks. Details on handling a duplicate address condition are out-of-scope for this document.

A generated addresses clean-up routine SHOULD run at regular intervals to get rid of stale addresses.

4.2.3. Support for Unique Local Addresses (ULA)

Address generation for globally unique addresses and unique local addresses (ULA) [RFC4193] is equivalent. If no BRIO for a unique local addresses is available, a router may start as a Border Router and DHCP server for a self generated 48-bit ULA prefix.

5. BRDP based Source Address Selection

As a next step, multi-homed nodes perform source address selection for new, self-initiated connections. The algorithm described in Default Address Selection for IPv6 [RFC6724] uses the concept of a "candidate set" of potential source addresses. Rule 8 of source address selection is "Uses longest match prefix". The goal of this rule is to select the address with good communications performance. If other means of choosing among source addresses for better performance is available, that should be used.

BRDP provides attributes for prefix, such as a cost metric to the Internet. This information can be used to select the "best" source address. For multi-path transport protocols, it is also important to have a mechanism to select alternative addresses. For example, rule 4 gives preference to a Home Address. Alternate addresses can be used for route optimization and to avoid overhead of the Mobile IP tunnel.

5.1. Address Selection for dynamic DNS

BRDP provided information can also be utilized by address lookup protocols such as DNS. A node can register its addresses dynamically, with support of preference and load balancing if the mechanism used support such.

6. BRDP based Routing

BRDP introduces a new paradigm for packet forwarding for multi-homed

sites, where forwarding to a default gateway is replaced by source address based forwarding towards a corresponding Border Router. This enforces that packets will be sent via the selected upstream provider, without the need of tunneling. As such, it prevents problems with ingress filters in multi-homed edge networks [RFC3704].

The BRDP Based Routing mechanism provides basic support for load distribution over multiple Border Routers. BRDP Based Routing forwards the packets to the Border Router that corresponds with the source address. As a result, nodes can utilize multiple paths, if available. Standardization of this load balancing functionality is work in progress in the IETF MPTCP working group.

When a router forwards a packet to a next-hop node, via the interface where this packet was received, and the next-hop address was selected using BRDP based routing, then the router should not send an ICMP redirect message to that host. This is because the upstream node would cache the redirect for the destination address, while the forwarding decision was based on the source address.

6.1. Problems with default gateway routing

Usually, the nexthop selection is based on the destination address. In case of default gateway routing and multiple exit routers to multiple providers, the source has no influence on what exit router is used. In case of ingress filtering and lack of a mechanism to redirect packets to exit routers that correspond to the source address, packets may be dropped.

This default gateway routing behavior blocks incremental enhancement of the Internet, e.g. through the addition of support for more dynamic networks and / or host based load distribution mechanisms. In a MANET, it also prevents the use of make-before-break [RFC3753] mechanisms.

6.2. Default gateway routing replaced with BRDP Based Routing

Default gateway based routing for IPv4 is defined in [RFC1812], section 5.2.4.3:

- (5) Default Route: This is a route to all networks for which there are no explicit routes. It is by definition the route whose prefix length is zero.

With BRDP Based Routing, another type of route is introduced:

- (6) BRDP Route: This is a route to all networks for which there are no explicit routes, and a default route is not used. The nexthop IP address is found by means of a Border Router Information Cache (BRIO-Cache) lookup based on the source address and, if a matching BRIO-Cache entry is found, a subsequent FIB lookup based on the selected Border Router address.

Note that route types (3) and (4) are not defined in RFC1812.

BRDP Based Routing can be turned on and off with the existence of a default route in the IGP. This switch function might be useful in migration scenarios towards BRDP Based Routing.

The Border Router should run the IGP on the interface with the BRDP advertized Border Router address, to make sure this address is reachable in the edge network.

In the edge network, all interior routers should run BRDP and BRDP Based Routing. All interior routers will have a BRIO-Cache with information for selecting Border Routers as exit points to the Internet. A BRIO-Cache entry contains a Border Router address and a summary prefix assigned to that Border Router. BRIO-Cache lookup follows the longest-match rule.

Forwarding is solely based on FIB lookups, the nexthop IP address is found either by a FIB lookup with the destination address or by a FIB lookup with the address of the Border Router that corresponds with the source address. If the nexthop IP address lookup fails, the packet is discarded.

7. BRDP and IRTF RRG goals

The IRTF Routing Research Group (RRG) was chartered to explore solutions for problems on routing and addressing, when the Internet continues to evolve. It has explored a number of proposed solutions, but did not reach consensus on a single, best approach [I-D.irtf-rrg-recommendation]. In fulfillment of the routing research group's charter, the co-chairs recommend that the IETF pursue work in three areas, "Evolution" [I-D.zhang-evolution], "Identifier/Locator Network Protocol (ILNP)" [I-D.rja-ilnp-intro] and "Renumbering" [RFC5887]. BRDP fits in all three approaches.

BRDP is an evolution in IPv6 address configuration and address selection, as well as forwarding to destinations outside the routing domain. As a result, it removes a demand for Provider Independent

addresses for (small) multi-homed edge networks. BRDP enables sites to use multiple Provider Aggregatable address blocks, while being able to utilize multi-homing for improved redundancy of communications and enlarged capacity. Each site that continues to use Provider Aggregatable addresses when getting multi-homed, instead of using its own Provider Independent address space, reduces the growth of the routing tables in the Default Free Zone.

BRDP can cooperate or live next many other solutions. ILNP is a good example for cooperation, BRDP provides multi-path transport capabilities to ILNP nodes. This multi-path transport capability applies to many other approaches also, such as map&enap and nat66.

Because BRDP provides automatic address and prefix configuration, Renumbering is far less problematic. That said, legacy (IPv4) hosts, applications and network equipment is not BRDP enabled and manual address configuration will be used for many years to come.

In Design Goals for Scalable Internet Routing [I-D.irtf-rrg-design-goals], a number of design goals are defined. The role BRDP can play for these goals are briefly described in the next sections.

7.1. Scalability

Because BRDP is implemented in edge networks, and not in the core, scalability of BRDP is less an issue. BRDP solves the Internet routing problem at the source, by reducing the demand for PI addresses.

7.2. Traffic engineering

BRDP provides traffic engineering options to end-nodes. End-nodes can configure multiple addresses and use these for utilizing multi-path capabilities of the network. Using multi-path is being worked on by the IETF MPTCP working group.

7.3. Multi-homing

The core function of BRDP is providing support for IPv6 multi-homing, without any problems caused by ingress filtering [RFC3704].

7.4. Loc/id separation

BRDP does not mandate any approach for location / identification. For packet forwarding, addresses are used as locator. If addresses are used as identifiers also, for example in Mobile IP, BRDP supports route optimization where traffic uses the Home Address as identifier

and care-of addresses as locator. MPTCP provides the route optimization capability.

7.5. Mobility

BRDP was defined as a solution for address autoconfiguration for ad hoc networks. With BRDP, nodes can easily configure topology correct addresses in a multi-homes ad hoc network. BRDP does not provide session continuity functions. Mobility solutions are already in place or new approaches are proposed. All approaches should work well with BRDP, as BRDP does not modify the IPv6 protocol.

7.6. Simplified renumbering

BRDP makes site renumbering fully automatic. This applies to node address configuration on the IPv6 stack and prefix delegation and configuration on routers. IP addresses could be configured on many other places, either manually or using specific protocols for such purpose. Complete automatic numbering is possible if all mechanisms in use support dynamic addresses. There is definitely more work to do [RFC5887].

7.7. Modularity

BRDP is a small, but important piece of the puzzle. It applies to edge networks only. It helps other mechanisms to work well in a multi-homed network using PA addresses, but also provides multi-path capabilities in multi-homed networks with PI addresses or multi-homing with connections to Extranets.

7.8. Routing quality

BRDP is not a routing protocol, so it has no influence on routing quality. But the functionality of routing to a default gateway is changed. BRDP based routing supports paths to multiple Border Routers, where hosts can select which Border Router to use. In such scheme, nodes can select the route to use, based on quality of available routes. MPTCP provides this route selection functionality.

7.9. Routing security

BRDP doesn't update any routing protocols. BRDBP based routing modifies the default gateway heuristic, the route to prefix `::/0` is replaced by a route to a Border Router, which corresponds with the source address of a packet. As a result, ingress filtering is distributed over all routers in the edge network and invalid packets are dropped as near to the source as possible.

The BRDP protocol runs on IPv6 NDP and inherits all security aspects. BRDP messages are disseminated in the edge network, which may enlarge the needs for protection. Implementing SeND [RFC3971] is recommended.

7.10. Deployability

BRDP deployment takes place edge network by edge network. Each network that migrates to BRDP, instead of getting a PI address block, reduces the load on the Internet routing infrastructure.

For implementing BRDP on an edge network, all routers in the network must support BRDP. BRDP support for hosts is optional. Enterprise networks can migrate site by site.

8. Currently unaddressed issues

BRDP based routing may have impact on multicast routing, e.g. selecting the route to a RP.

It is not fully understood how BRDP may influence host behavior on RA M and O bits, and may bypass a 1-hop router DHCP relay server for getting information for a BRDP-learned DHCP server.

Currently unaddressed issues are to be addressed in a next version of this document.

9. Acknowledgements

BRDP is inspired by MANEMO technology; thanks to all who contributed to it. Thanks to Arjen Holtzer (TNO), co-author of earlier Internet drafts on BRDP. Thanks to Ran Atkinson, who guided me towards a BRDP Based Routing mechanism that does not rely on routing headers or encapsulation.

10. IANA Considerations

TBD

11. Security Considerations

TBD

12. Change log

This -00 version is gathering the material of BRDP, produced for Autoconf and RRG. It is a bit cleaned up, with removal of some details for MANET and with removal of options for emergency services, service selection and authorization.

13. References

13.1. Normative References

- [RFC1812] Baker, F., "Requirements for IP Version 4 Routers", RFC 1812, June 1995.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC3753] Manner, J. and M. Kojo, "Mobility Related Terminology", RFC 3753, June 2004.
- [RFC3972] Aura, T., "Cryptographically Generated Addresses (CGA)", RFC 3972, March 2005.
- [RFC4193] Hinden, R. and B. Haberman, "Unique Local IPv6 Unicast Addresses", RFC 4193, October 2005.
- [RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 4291, February 2006.
- [RFC4429] Moore, N., "Optimistic Duplicate Address Detection (DAD) for IPv6", RFC 4429, April 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.
- [RFC4941] Narten, T., Draves, R., and S. Krishnan, "Privacy Extensions for Stateless Address Autoconfiguration in IPv6", RFC 4941, September 2007.
- [RFC6724] Thaler, D., Draves, R., Matsumoto, A., and T. Chown, "Default Address Selection for Internet Protocol Version 6 (IPv6)", RFC 6724, September 2012.

13.2. Informative References

- [I-D.ietf-homenet-arch]
Chown, T., Arkko, J., Brandt, A., Troan, O., and J. Weil,
"Home Networking Architecture for IPv6",
draft-ietf-homenet-arch-04 (work in progress), July 2012.
- [I-D.ietf-mptcp-multiaddressed]
Ford, A., Raiciu, C., Handley, M., and O. Bonaventure,
"TCP Extensions for Multipath Operation with Multiple
Addresses", draft-ietf-mptcp-multiaddressed-10 (work in
progress), October 2012.
- [I-D.irtf-rrg-design-goals]
Li, T., "Design Goals for Scalable Internet Routing",
draft-irtf-rrg-design-goals-06 (work in progress),
January 2011.
- [I-D.irtf-rrg-recommendation]
Li, T., "Recommendation for a Routing Architecture",
draft-irtf-rrg-recommendation-16 (work in progress),
November 2010.
- [I-D.kline-default-perimeter]
Kline, E., "Default Perimeter Identification",
draft-kline-default-perimeter-00 (work in progress),
July 2012.
- [I-D.rja-ilnp-intro]
Atkinson, R., "ILNP Concept of Operations",
draft-rja-ilnp-intro-11 (work in progress), July 2011.
- [I-D.zhang-evolution]
Zhang, B. and L. Zhang, "Evolution Towards Global Routing
Scalability", draft-zhang-evolution-02 (work in progress),
October 2009.
- [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.
- [RFC3633] Troan, O. and R. Droms, "IPv6 Prefix Options for Dynamic
Host Configuration Protocol (DHCP) version 6", RFC 3633,
December 2003.

- [RFC3704] Baker, F. and P. Savola, "Ingress Filtering for Multihomed Networks", BCP 84, RFC 3704, March 2004.
- [RFC3971] Arkko, J., Kempf, J., Zill, B., and P. Nikander, "SEcure Neighbor Discovery (SEND)", RFC 3971, March 2005.
- [RFC5149] Korhonen, J., Nilsson, U., and V. Devarapalli, "Service Selection for Mobile IPv6", RFC 5149, February 2008.
- [RFC5887] Carpenter, B., Atkinson, R., and H. Flinck, "Renumbering Still Needs Work", RFC 5887, May 2010.
- [RFC5889] Baccelli, E. and M. Townsley, "IP Addressing Model in Ad Hoc Networks", RFC 5889, September 2010.
- [RFC6275] Perkins, C., Johnson, D., and J. Arkko, "Mobility Support in IPv6", RFC 6275, July 2011.

Author's Address

Teco Boot
Infinity Networks B.V.

Email: teco@inf-net.nl

