Network Working Group
Internet-Draft
Obsoletes: 3484 (if approved)
Intended status: Standards Track
Expires: September 6, 2012

D. Thaler, Ed.
Microsoft
R. Draves
Microsoft Research
A. Matsumoto
NTT
T. Chown
University of Southampton

March 5, 2012

Default Address Selection for Internet Protocol version 6 (IPv6)
draft-ietf-6man-rfc3484bis-01.txt

Abstract

This document describes two algorithms, for source address selection and for destination address selection. The algorithms specify default behavior for all Internet Protocol version 6 (IPv6) implementations. They do not override choices made by applications or upper-layer protocols, nor do they preclude the development of more advanced mechanisms for address selection. The two algorithms share a common context, including an optional mechanism for allowing administrators to provide policy that can override the default behavior. In dual stack implementations, the destination address selection algorithm can consider both IPv4 and IPv6 addresses - depending on the available source addresses, the algorithm might prefer IPv6 addresses over IPv4 addresses, or vice-versa.

All IPv6 nodes, including both hosts and routers, must implement default address selection as defined in this specification.

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 September 6, 2012.

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.

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

Table of Contents

1. Introduction 4
1.1. Conventions Used in This Document 5
2. Context in Which the Algorithms Operate 5
2.1. Policy Table 6
2.2. Common Prefix Length 8
3. Address Properties 8
3.1. Scope Comparisons 8
3.2. IPv4 Addresses and IPv4-Mapped Addresses 9
3.3. Other IPv6 Addresses with Embedded IPv4 Addresses 9
3.4. IPv6 Loopback Address and Other Format Prefixes 9
3.5. Mobility Addresses 10
4. Candidate Source Addresses 10
5. Source Address Selection 11
6. Destination Address Selection 14
7. Interactions with Routing 16
8. Implementation Considerations 16
9. Security Considerations 17
10. Examples 17
10.1. Default Source Address Selection 18
10.2. Default Destination Address Selection 18
10.3. Configuring Preference for IPv6 or IPv4 20
10.3.1. Handling Broken IPv6 20
10.4. Configuring Preference for Link-Local Addresses 21
10.5. Configuring a Multi-Homed Site 21
10.6. Configuring ULA Preference 23
10.7. Configuring 6to4 Preference 24
11. References 25
11.1. Normative References 25
11.2. Informative References 26
Appendix A. Acknowledgements 27
Appendix B. Changes Since RFC 3484 28
Authors' Addresses 29

1. Introduction

The IPv6 addressing architecture [RFC4291] allows multiple unicast addresses to be assigned to interfaces. These addresses may have different reachability scopes (link-local, site-local, or global). These addresses may also be "preferred" or "deprecated" [RFC4862]. Privacy considerations have introduced the concepts of "public addresses" and "temporary addresses" [RFC4941]. The mobility architecture introduces "home addresses" and "care-of addresses" [RFC6275]. In addition, multi-homing situations will result in more addresses per node. For example, a node may have multiple interfaces, some of them tunnels or virtual interfaces, or a site may have multiple ISP attachments with a global prefix per ISP.

The end result is that IPv6 implementations will very often be faced with multiple possible source and destination addresses when initiating communication. It is desirable to have default algorithms, common across all implementations, for selecting source and destination addresses so that developers and administrators can reason about and predict the behavior of their systems.

Furthermore, dual or hybrid stack implementations, which support both IPv6 and IPv4, will very often need to choose between IPv6 and IPv4 when initiating communication. For example, when DNS name resolution yields both IPv6 and IPv4 addresses and the network protocol stack has available both IPv6 and IPv4 source addresses. In such cases, a simple policy to always prefer IPv6 or always prefer IPv4 can produce poor behavior. As one example, suppose a DNS name resolves to a global IPv6 address and a global IPv4 address. If the node has assigned a global IPv6 address and a 169.254/16 auto-configured IPv4 address [RFC3927], then IPv6 is the best choice for communication. But if the node has assigned only a link-local IPv6 address and a global IPv4 address, then IPv4 is the best choice for communication. The destination address selection algorithm solves this with a unified procedure for choosing among both IPv6 and IPv4 addresses.

The algorithms in this document are specified as a set of rules that define a partial ordering on the set of addresses that are available for use. In the case of source address selection, a node typically has multiple addresses assigned to its interfaces, and the source address ordering rules in section 5 define which address is the "best" one to use. In the case of destination address selection, the DNS may return a set of addresses for a given name, and an application needs to decide which one to use first, and in what order to try others should the first one not be reachable. The destination address ordering rules in section 6, when applied to the set of addresses returned by the DNS, provide such a recommended ordering.

This document specifies source address selection and destination address selection separately, but using a common context so that together the two algorithms yield useful results. The algorithms attempt to choose source and destination addresses of appropriate scope and configuration status (preferred or deprecated in the RFC 4862 sense). Furthermore, this document suggests a preferred method, longest matching prefix, for choosing among otherwise equivalent addresses in the absence of better information.

This document also specifies policy hooks to allow administrative override of the default behavior. For example, using these hooks an administrator can specify a preferred source prefix for use with a destination prefix, or prefer destination addresses with one prefix over addresses with another prefix. These hooks give an administrator flexibility in dealing with some multi-homing and transition scenarios, but they are certainly not a panacea.

The selection rules specified in this document MUST NOT be construed to override an application or upper-layer's explicit choice of a legal destination or source address.

1.1. Conventions Used in This Document

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 BCP 14, RFC 2119 [RFC2119].

2. Context in Which the Algorithms Operate

Our context for address selection derives from the most common implementation architecture, which separates the choice of destination address from the choice of source address. Consequently, we have two separate algorithms for these tasks. The algorithms are designed to work well together and they share a mechanism for administrative policy override.

In this implementation architecture, applications use APIs [RFC3493] like getaddrinfo() that return a list of addresses to the application. This list might contain both IPv6 and IPv4 addresses (sometimes represented as IPv4-mapped addresses). The application then passes a destination address to the network stack with connect() or sendto(). The application would then typically try the first address in the list, looping over the list of addresses until it finds a working address. In any case, the network layer is never in a situation where it needs to choose a destination address from several alternatives. The application might also specify a source

address with bind(), but often the source address is left unspecified. Therefore the network layer does often choose a source address from several alternatives.

As a consequence, we intend that implementations of getaddrinfo() will use the destination address selection algorithm specified here to sort the list of IPv6 and IPv4 addresses that they return. Separately, the IPv6 network layer will use the source address selection algorithm when an application or upper-layer has not specified a source address. Application of this specification to source address selection in an IPv4 network layer may be possible but this is not explored further here.

Well-behaved applications SHOULD iterate through the list of addresses returned from getaddrinfo() until they find a working address.

The algorithms use several criteria in making their decisions. The combined effect is to prefer destination/source address pairs for which the two addresses are of equal scope or type, prefer smaller scopes over larger scopes for the destination address, prefer non-deprecated source addresses, avoid the use of transitional addresses when native addresses are available, and all else being equal prefer address pairs having the longest possible common prefix. For source address selection, public addresses [RFC4941] are preferred over temporary addresses. In mobile situations [RFC6275], home addresses are preferred over care-of addresses. If an address is simultaneously a home address and a care-of address (indicating the mobile node is "at home" for that address), then the home/care-of address is preferred over addresses that are solely a home address or solely a care-of address.

This specification optionally allows for the possibility of administrative configuration of policy (e.g., via manual configuration or a DHCP option such as that proposed in [I-D.ietf-6man-addr-select-opt]) that can override the default behavior of the algorithms. The policy override takes the form of a configurable table that specifies precedence values and preferred source prefixes for destination prefixes. If an implementation is not configurable, or if an implementation has not been configured, then the default policy table specified in this document SHOULD be used.

2.1. Policy Table

The policy table is a longest-matching-prefix lookup table, much like a routing table. Given an address A, a lookup in the policy table produces two values: a precedence value Precedence(A) and a

classification or label Label(A).

The precedence value Precedence(A) is used for sorting destination addresses. If Precedence(A) > Precedence(B), we say that address A has higher precedence than address B, meaning that our algorithm will prefer to sort destination address A before destination address B.

The label value Label(A) allows for policies that prefer a particular source address prefix for use with a destination address prefix. The algorithms prefer to use a source address S with a destination address D if Label(S) = Label(D).

IPv6 implementations SHOULD support configurable address selection via a mechanism at least as powerful as the policy tables defined here. It is important that implementations provide a way to change the default policies as more experience is gained. Sections 10.3 and 10.4 provide examples of the kind of changes that might be needed.

If an implementation is not configurable or has not been configured, then it SHOULD operate according to the algorithms specified here in conjunction with the following default policy table:

Prefix	Precedence	Label
::1/128	50	0
::/0	40	1
::ffff:0:0/96	35	4
2002::/16	30	2
2001::/32	5	5
fc00::/7	3	13
::/96	1	3
fec0::/10	1	11
3ffe::/16	1	12

An implementation MAY automatically add additional site-specific rows to the default table based on its configured addresses, such as for ULAs and 6to4 addresses for instance (see Section 10.6 and Section 10.7 for examples). Any such rows automatically added by the implementation as a result of address acquisition MUST NOT override a row for the same prefix configured via other means. That is, rows can be added but never updated automatically. An implementation SHOULD provide a means for an administrator to disable automatic row additions.

One effect of the default policy table is to prefer using native source addresses with native destination addresses, 6to4 [RFC3056] source addresses with 6to4 destination addresses, etc. Another effect of the default policy table is to prefer communication using IPv6 addresses to communication using IPv4 addresses, if matching

source addresses are available.

Policy table entries for scoped address prefixes MAY be qualified with an optional zone index. If so, a prefix table entry only matches against an address during a lookup if the zone index also matches the address's zone index.

2.2. Common Prefix Length

We define the common prefix length CommonPrefixLen(S, D) of a source address S and a destination address D as the length of the longest prefix (looking at the most significant, or leftmost, bits) that the two addresses have in common, up to the length of S's prefix (i.e., the portion of the address not including the interface ID). For example, CommonPrefixLen(fe80::1, fe80::2) is 64.

3. Address Properties

In the rules given in later sections, addresses of different types (e.g., IPv4, IPv6, multicast and unicast) are compared against each other. Some of these address types have properties that aren't directly comparable to each other. For example, IPv6 unicast addresses can be "preferred" or "deprecated" [RFC4862], while IPv4 addresses have no such notion. To compare such addresses using the ordering rules (e.g., to use "preferred" addresses in preference to "deprecated" addresses), the following mappings are defined.

3.1. Scope Comparisons

Multicast destination addresses have a 4-bit scope field that controls the propagation of the multicast packet. The IPv6 addressing architecture defines scope field values for interface-local (0x1), link-local (0x2), subnet-local (0x3), admin-local (0x4), site-local (0x5), organization-local (0x8), and global (0xE) scopes [RFC4007].

Use of the source address selection algorithm in the presence of multicast destination addresses requires the comparison of a unicast address scope with a multicast address scope. We map unicast link-local to multicast link-local, unicast site-local to multicast site-local, and unicast global scope to multicast global scope. For example, unicast site-local is equal to multicast site-local, which is smaller than multicast organization-local, which is smaller than unicast global, which is equal to multicast global.

We write Scope(A) to mean the scope of address A. For example, if A is a link-local unicast address and B is a site-local multicast

address, then Scope(A) < Scope(B).

This mapping implicitly conflates unicast site boundaries and multicast site boundaries [RFC4007].

3.2. IPv4 Addresses and IPv4-Mapped Addresses

The destination address selection algorithm operates on both IPv6 and IPv4 addresses. For this purpose, IPv4 addresses should be represented as IPv4-mapped addresses [RFC4291]. For example, to lookup the precedence or other attributes of an IPv4 address in the policy table, lookup the corresponding IPv4-mapped IPv6 address.

IPv4 addresses are assigned scopes as follows. IPv4 auto-configuration addresses [RFC3927], which have the prefix 169.254/16, are assigned link-local scope. IPv4 private addresses [RFC1918], which have the prefixes 10/8, 172.16/12, and 192.168/16, are assigned global scope. IPv4 loopback addresses ([RFC1918], section 4.2.2.11), which have the prefix 127/8, are assigned link-local scope (analogously to the treatment of the IPv6 loopback address ([RFC4007], section 4)). Other IPv4 addresses are assigned global scope.

IPv4 addresses should be treated as having "preferred" (in the RFC 4862 sense) configuration status.

3.3. Other IPv6 Addresses with Embedded IPv4 Addresses

IPv4-compatible addresses [RFC4291], IPv4-mapped [RFC4291], IPv4-converted [RFC6145], IPv4-translatable [RFC6145], and 6to4 addresses [RFC3056] contain an embedded IPv4 address. For the purposes of this document, these addresses should be treated as having global scope.

IPv4-compatible, IPv4-mapped, and IPv4-converted addresses should be treated as having "preferred" (in the RFC 4862 sense) configuration status.

3.4. IPv6 Loopback Address and Other Format Prefixes

The loopback address should be treated as having link-local scope ([RFC4007], section 4) and "preferred" (in the RFC 4862 sense) configuration status.

NSAP addresses and other addresses with as-yet-undefined format prefixes should be treated as having global scope and "preferred" (in the RFC 4862) configuration status. Later standards may supersede this treatment.

3.5. Mobility Addresses

Some nodes may support mobility using the concepts of home address and care-of address (for example see [RFC6275]). Conceptually, a home address is an IP address assigned to a mobile node and used as the permanent address of the mobile node. A care-of address is an IP address associated with a mobile node while visiting a foreign link. When a mobile node is on its home link, it may have an address that is simultaneously a home address and a care-of address.

For the purposes of this document, it is sufficient to know whether or not one's own addresses are designated as home addresses or care-of addresses. Whether or not an address should be designated a home address or care-of address is outside the scope of this document.

4. Candidate Source Addresses

The source address selection algorithm uses the concept of a "candidate set" of potential source addresses for a given destination address. The candidate set is the set of all addresses that could be used as a source address; the source address selection algorithm will pick an address out of that set. We write CandidateSource(A) to denote the candidate set for the address A.

It is RECOMMENDED that the candidate source addresses be the set of unicast addresses assigned to the interface that will be used to send to the destination. (The "outgoing" interface.) On routers, the candidate set MAY include unicast addresses assigned to any interface that forwards packets, subject to the restrictions described below.

> Discussion: The Neighbor Discovery Redirect mechanism [RFC4861] requires that routers verify that the source address of a packet identifies a neighbor before generating a Redirect, so it is advantageous for hosts to choose source addresses assigned to the outgoing interface. Implementations that wish to support the use of global source addresses assigned to a loopback interface should behave as if the loopback interface originates and forwards the packet.

In some cases the destination address may be qualified with a zone index or other information that will constrain the candidate set.

For multicast and link-local destination addresses, the set of candidate source addresses MUST only include addresses assigned to interfaces belonging to the same link as the outgoing interface.

Discussion: The restriction for multicast destination addresses is necessary because currently-deployed multicast forwarding algorithms use Reverse Path Forwarding (RPF) checks.

For site-local destination addresses, the set of candidate source addresses MUST only include addresses assigned to interfaces belonging to the same site as the outgoing interface.

In any case, multicast addresses, and the unspecified address MUST NOT be included in a candidate set.

If an application or upper layer specifies a source address that is not in the candidate set for the destination, then the network layer MUST treat this as an error. The specified source address may influence the candidate set, by affecting the choice of outgoing interface. If the application or upper layer specifies a source address that is in the candidate set for the destination, then the network layer MUST respect that choice. If the application or upper layer does not specify a source address, then the network layer uses the source address selection algorithm specified in the next section.

On IPv6-only nodes that support SIIT [RFC6145], if the destination address is an IPv4-converted address then the candidate set MUST contain only IPv4-translatable addresses.

5. Source Address Selection

The source address selection algorithm produces as output a single source address for use with a given destination address. This algorithm only applies to IPv6 destination addresses, not IPv4 addresses.

The algorithm is specified here in terms of a list of pair-wise comparison rules that (for a given destination address D) imposes a "greater than" ordering on the addresses in the candidate set CandidateSource(D). The address at the front of the list after the algorithm completes is the one the algorithm selects.

Note that conceptually, a sort of the candidate set is being performed, where a set of rules define the ordering among addresses. But because the output of the algorithm is a single source address, an implementation need not actually sort the set; it need only identify the "maximum" value that ends up at the front of the sorted list.

The ordering of the addresses in the candidate set is defined by a list of eight pair-wise comparison rules, with each rule placing a

"greater than," "less than" or "equal to" ordering on two source addresses with respect to each other (and that rule). In the case that a given rule produces a tie, i.e., provides an "equal to" result for the two addresses, the remaining rules are applied (in order) to just those addresses that are tied to break the tie. Note that if a rule produces a single clear "winner" (or set of "winners" in the case of ties), those addresses not in the winning set can be discarded from further consideration, with subsequent rules applied only to the remaining addresses. If the eight rules fail to choose a single address, some unspecified tie-breaker should be used.

When comparing two addresses SA and SB from the candidate set, we say "prefer SA" to mean that SA is "greater than" SB, and similarly we say "prefer SB" to mean that SA is "less than" SB.

Rule 1: Prefer same address.
If SA = D, then prefer SA. Similarly, if SB = D, then prefer SB.

Rule 2: Prefer appropriate scope.
If Scope(SA) < Scope(SB): If Scope(SA) < Scope(D), then prefer SB and otherwise prefer SA. Similarly, if Scope(SB) < Scope(SA): If Scope(SB) < Scope(D), then prefer SA and otherwise prefer SB.

Rule 3: Avoid deprecated addresses.
The addresses SA and SB have the same scope. If one of the two source addresses is "preferred" and one of them is "deprecated" (in the RFC 4862 sense), then prefer the one that is "preferred."

Rule 4: Prefer home addresses.
If SA is simultaneously a home address and care-of address and SB is not, then prefer SA. Similarly, if SB is simultaneously a home address and care-of address and SA is not, then prefer SB. If SA is just a home address and SB is just a care-of address, then prefer SA. Similarly, if SB is just a home address and SA is just a care-of address, then prefer SB.

Implementations should provide a mechanism allowing an application to reverse the sense of this preference and prefer care-of addresses over home addresses (e.g., via appropriate API extensions such as [RFC5014]). Use of the mechanism should only affect the selection rules for the invoking application.

Rule 5: Prefer outgoing interface.
If SA is assigned to the interface that will be used to send to D and SB is assigned to a different interface, then prefer SA. Similarly, if SB is assigned to the interface that will be used to send to D and SA is assigned to a different interface, then prefer SB.

Rule 5.5: Prefer addresses in a prefix advertised by the next-hop
If SA or SA's prefix is assigned by the selected next-hop that will be used to send to D and SB or SB's prefix is assigned by a different next-hop, then prefer SA. Similarly, if SB or SB's prefix is assigned by the next-hop that will be used to send to D and SA or SA's prefix is assigned by a different next-hop, then prefer SB.

Discussion: An IPv6 implementation is not required to remember which next-hops advertised which prefixes. The conceptual models of IPv6 hosts in Section 5 of [RFC4861] and Section 3 of [RFC4191] have no such requirement. Implementations that do not track this information shall omit rule 5.5.

Rule 6: Prefer matching label.
If Label(SA) = Label(D) and Label(SB) <> Label(D), then prefer SA. Similarly, if Label(SB) = Label(D) and Label(SA) <> Label(D), then prefer SB.

Rule 7: Prefer public addresses.
If SA is a public address and SB is a temporary address, then prefer SA. Similarly, if SB is a public address and SA is a temporary address, then prefer SB.

Implementations MUST provide a mechanism allowing an application to reverse the sense of this preference and prefer temporary addresses over public addresses (e.g., via appropriate API extensions such as [RFC5014]). Use of the mechanism should only affect the selection rules for the invoking application. This rule avoids applications potentially failing due to the relatively short lifetime of temporary addresses or due to the possibility of the reverse lookup of a temporary address either failing or returning a randomized name. Implementations for which privacy considerations outweigh these application compatibility concerns MAY reverse the sense of this rule and by default prefer temporary addresses over public addresses.

Rule 8: Use longest matching prefix.
If CommonPrefixLen(SA, D) > CommonPrefixLen(SB, D), then prefer SA. Similarly, if CommonPrefixLen(SB, D) > CommonPrefixLen(SA, D), then prefer SB.

Rule 8 may be superseded if the implementation has other means of choosing among source addresses. For example, if the implementation somehow knows which source address will result in the "best" communications performance.

Rule 2 (prefer appropriate scope) MUST be implemented and given high priority because it can affect interoperability.

6. Destination Address Selection

The destination address selection algorithm takes a list of destination addresses and sorts the addresses to produce a new list. It is specified here in terms of the pair-wise comparison of addresses DA and DB, where DA appears before DB in the original list.

The algorithm sorts together both IPv6 and IPv4 addresses. To find the attributes of an IPv4 address in the policy table, the IPv4 address should be represented as an IPv4-mapped address.

We write Source(D) to indicate the selected source address for a destination D. For IPv6 addresses, the previous section specifies the source address selection algorithm. Source address selection for IPv4 addresses is not specified in this document.

We say that Source(D) is undefined if there is no source address available for destination D. For IPv6 addresses, this is only the case if CandidateSource(D) is the empty set.

The pair-wise comparison of destination addresses consists of ten rules, which should be applied in order. If a rule determines a result, then the remaining rules are not relevant and should be ignored. Subsequent rules act as tie-breakers for earlier rules. See the previous section for a lengthier description of how pair-wise comparison tie-breaker rules can be used to sort a list.

Rule 1: Avoid unusable destinations.
If DB is known to be unreachable or if Source(DB) is undefined, then prefer DA. Similarly, if DA is known to be unreachable or if Source(DA) is undefined, then prefer DB.

> Discussion: An implementation may know that a particular destination is unreachable in several ways. For example, the destination may be reached through a network interface that is currently unplugged. For example, the implementation may retain for some period of time information from Neighbor Unreachability Detection [RFC4861]. In any case, the determination of unreachability for the purposes of this rule is implementation-dependent.

Rule 2: Prefer matching scope.
If Scope(DA) = Scope(Source(DA)) and Scope(DB) <> Scope(Source(DB)), then prefer DA. Similarly, if Scope(DA) <> Scope(Source(DA)) and Scope(DB) = Scope(Source(DB)), then prefer DB.

Rule 3: Avoid deprecated addresses.
If Source(DA) is deprecated and Source(DB) is not, then prefer DB.

Similarly, if Source(DA) is not deprecated and Source(DB) is deprecated, then prefer DA.

Rule 4: Prefer home addresses.
If Source(DA) is simultaneously a home address and care-of address and Source(DB) is not, then prefer DA. Similarly, if Source(DB) is simultaneously a home address and care-of address and Source(DA) is not, then prefer DB.

If Source(DA) is just a home address and Source(DB) is just a care-of address, then prefer DA. Similarly, if Source(DA) is just a care-of address and Source(DB) is just a home address, then prefer DB.

Rule 5: Prefer matching label.
If Label(Source(DA)) = Label(DA) and Label(Source(DB)) <> Label(DB), then prefer DA. Similarly, if Label(Source(DA)) <> Label(DA) and Label(Source(DB)) = Label(DB), then prefer DB.

Rule 6: Prefer higher precedence.
If Precedence(DA) > Precedence(DB), then prefer DA. Similarly, if Precedence(DA) < Precedence(DB), then prefer DB.

Rule 7: Prefer native transport.
If DA is reached via an encapsulating transition mechanism (e.g., IPv6 in IPv4) and DB is not, then prefer DB. Similarly, if DB is reached via encapsulation and DA is not, then prefer DA.

> Discussion: 6RD [RFC5969], ISATAP [RFC5214], and configured tunnels [RFC4213] are examples of encapsulating transition mechanisms for which the destination address does not have a specific prefix and hence can not be assigned a lower precedence in the policy table. An implementation MAY generalize this rule by using a concept of interface preference, and giving virtual interfaces (like the IPv6-in-IPv4 encapsulating interfaces) a lower preference than native interfaces (like ethernet interfaces).

Rule 8: Prefer smaller scope.
If Scope(DA) < Scope(DB), then prefer DA. Similarly, if Scope(DA) > Scope(DB), then prefer DB.

Rule 9: Use longest matching prefix.
When DA and DB belong to the same address family (both are IPv6 or both are IPv4): If CommonPrefixLen(Source(DA), DA) > CommonPrefixLen(Source(DB), DB), then prefer DA. Similarly, if CommonPrefixLen(Source(DA), DA) < CommonPrefixLen(Source(DB), DB), then prefer DB.

Rule 10: Otherwise, leave the order unchanged.
If DA preceded DB in the original list, prefer DA. Otherwise prefer DB.

Rules 9 and 10 may be superseded if the implementation has other means of sorting destination addresses. For example, if the implementation somehow knows which destination addresses will result in the "best" communications performance.

7. Interactions with Routing

This specification of source address selection assumes that routing (more precisely, selecting an outgoing interface on a node with multiple interfaces) is done before source address selection. However, implementations may use source address considerations as a tiebreaker when choosing among otherwise equivalent routes.

For example, suppose a node has interfaces on two different links, with both links having a working default router. Both of the interfaces have preferred (in the RFC 4862 sense) global addresses. When sending to a global destination address, if there's no routing reason to prefer one interface over the other, then an implementation may preferentially choose the outgoing interface that will allow it to use the source address that shares a longer common prefix with the destination.

Implementations that support Rule 5.5 also use the choice of router to influence the choice of source address. For example, suppose a host is on a link with two routers. One router is advertising a global prefix A and the other router is advertising global prefix B. Then when sending via the first router, the host may prefer source addresses with prefix A and when sending via the second router, prefer source addresses with prefix B.

8. Implementation Considerations

The destination address selection algorithm needs information about potential source addresses. One possible implementation strategy is for getaddrinfo() to call down to the network layer with a list of destination addresses, sort the list in the network layer with full current knowledge of available source addresses, and return the sorted list to getaddrinfo(). This is simple and gives the best results but it introduces the overhead of another system call. One way to reduce this overhead is to cache the sorted address list in the resolver, so that subsequent calls for the same name do not need to resort the list.

Another implementation strategy is to call down to the network layer to retrieve source address information and then sort the list of addresses directly in the context of getaddrinfo(). To reduce overhead in this approach, the source address information can be cached, amortizing the overhead of retrieving it across multiple calls to getaddrinfo(). In this approach, the implementation may not have knowledge of the outgoing interface for each destination, so it MAY use a looser definition of the candidate set during destination address ordering.

In any case, if the implementation uses cached and possibly stale information in its implementation of destination address selection, or if the ordering of a cached list of destination addresses is possibly stale, then it should ensure that the destination address ordering returned to the application is no more than one second out of date. For example, an implementation might make a system call to check if any routing table entries or source address assignments or prefix policy table entries that might affect these algorithms have changed. Another strategy is to use an invalidation counter that is incremented whenever any underlying state is changed. By caching the current invalidation counter value with derived state and then later comparing against the current value, the implementation could detect if the derived state is potentially stale.

9. Security Considerations

This document has no direct impact on Internet infrastructure security.

Note that most source address selection algorithms, including the one specified in this document, expose a potential privacy concern. An unfriendly node can infer correlations among a target node's addresses by probing the target node with request packets that force the target host to choose its source address for the reply packets. (Perhaps because the request packets are sent to an anycast or multicast address, or perhaps the upper-layer protocol chosen for the attack does not specify a particular source address for its reply packets.) By using different addresses for itself, the unfriendly node can cause the target node to expose the target's own addresses.

10. Examples

This section contains a number of examples, first of default behavior and then demonstrating the utility of policy table configuration. These examples are provided for illustrative purposes; they should not be construed as normative.

10.1. Default Source Address Selection

The source address selection rules, in conjunction with the default policy table, produce the following behavior:

Destination: 2001:db8:1::1
Candidate Source Addresses: 2001:db8:3::1 or fe80::1
Result: 2001:db8::1 (prefer appropriate scope)

Destination: ff05::1
Candidate Source Addresses: 2001:db8:3::1 or fe80::1
Result: 2001:db8:3::1 (prefer appropriate scope)

Destination: 2001:db8:1::1
Candidate Source Addresses: 2001:db8:1::1 (deprecated) or 2001:db8:2::1
Result: 2001:db8:1::1 (prefer same address)

Destination: fe80::1
Candidate Source Addresses: fe80::2 (deprecated) or 2001:db8:1::1
Result: fe80::2 (prefer appropriate scope)

Destination: 2001:db8:1::1
Candidate Source Addresses: 2001:db8:1::2 or 2001:db8:3::2
Result: 2001:db8:1:::2 (longest-matching-prefix)

Destination: 2001:db8:1::1
Candidate Source Addresses: 2001:db8:1::2 (care-of address) or 2001:db8:3::2 (home address)
Result: 2001:db8:3::2 (prefer home address)

Destination: 2002:c633:6401::1
Candidate Source Addresses: 2002:c633:6401::d5e3:7953:13eb:22e8 (temporary) or 2001:db8:1::2
Result: 2002:c633:6401::d5e3:7953:13eb:22e8 (prefer matching label)

Destination: 2001:db8:1::d5e3:0:0:1
Candidate Source Addresses: 2001:db8:1::2 or 2001:db8:1::d5e3:7953:13eb:22e8 (temporary)
Result: 2001:db8:1::2 (prefer public address)

10.2. Default Destination Address Selection

The destination address selection rules, in conjunction with the default policy table and the source address selection rules, produce the following behavior:

Candidate Source Addresses: 2001:db8:1::2 or fe80::1 or 169.254.13.78

Destination Address List: 2001:db8:1::1 or 198.51.100.121
Result: 2001:db8:1::1 (src 2001:db8:1::2) then 198.51.100.121 (src 169.254.13.78) (prefer matching scope)

Candidate Source Addresses: fe80::1 or 198.51.100.117
Destination Address List: 2001:db8:1::1 or 198.51.100.121
Result: 198.51.100.121 (src 198.51.100.117) then 2001:db8:1::1 (src fe80::1) (prefer matching scope)

Candidate Source Addresses: 2001:db8:1::2 or fe80::1 or 10.1.2.4
Destination Address List: 2001:db8:1::1 or 10.1.2.3
Result: 2001:db8:1::1 (src 2001:db8:1::2) then 10.1.2.3 (src 10.1.2.4) (prefer higher precedence)

Candidate Source Addresses: 2001:db8:1::2 or fe80::2
Destination Address List: 2001:db8:1::1 or fe80::1
Result: fe80::1 (src fe80::2) then 2001:db8:1::1 (src 2001:db8:1::2) (prefer smaller scope)

Candidate Source Addresses: 2001:db8:1::2 (care-of address) or 2001:db8:3::1 (home address) or fe80::2 (care-of address)
Destination Address List: 2001:db8:1::1 or fe80::1
Result: 2001:db8:1::1 (src 2001:db8:3::1) then fe80::1 (src fe80::2) (prefer home address)

Candidate Source Addresses: 2001:db8:1::2 or fe80::2 (deprecated)
Destination Address List: 2001:db8:1::1 or fe80::1
Result: 2001:db8:1::1 (src 2001:db8:1::2) then fe80::1 (src fe80::2) (avoid deprecated addresses)

Candidate Source Addresses: 2001:db8:1::2 or 2001:db8:3f44::2 or fe80::2
Destination Address List: 2001:db8:1::1 or 2001:db8:3ffe::1
Result: 2001:db8:1::1 (src 2001:db8:1::2) then 2001:db8:3ffe::1 (src 2001:db8:3f44::2) (longest matching prefix)

Candidate Source Addresses: 2002:c633:6401::2 or fe80::2
Destination Address List: 2002:c633:6401::1 or 2001:db8:1::1
Result: 2002:c633:6401::1 (src 2002:c633:6401::2) then 2001:db8:1::1 (src 2002:c633:6401::2) (prefer matching label)

Candidate Source Addresses: 2002:c633:6401::2 or 2001:db8:1::2 or fe80::2
Destination Address List: 2002:c633:6401::1 or 2001:db8:1::1
Result: 2001:db8:1::1 (src 2001:db8:1::2) then 2002:c633:6401::1 (src 2002:c633:6401::2) (prefer higher precedence)

10.3. Configuring Preference for IPv6 or IPv4

The default policy table gives IPv6 addresses higher precedence than IPv4 addresses. This means that applications will use IPv6 in preference to IPv4 when the two are equally suitable. An administrator can change the policy table to prefer IPv4 addresses by giving the ::ffff:0.0.0.0/96 prefix a higher precedence:

Prefix	Precedence	Label
::1/128	50	0
::/0	40	1
::ffff:0:0/96	100	4
2002::/16	30	2
2001::/32	5	5
fc00::/7	3	13
::/96	1	3
fec0::/10	1	11
3ffe::/16	1	12

This change to the default policy table produces the following behavior:

Candidate Source Addresses: 2001::2 or fe80::1 or 169.254.13.78
Destination Address List: 2001::1 or 198.51.100.121
Unchanged Result: 2001::1 (src 2001::2) then 198.51.100.121 (src 169.254.13.78) (prefer matching scope)

Candidate Source Addresses: fe80::1 or 198.51.100.117
Destination Address List: 2001::1 or 198.51.100.121
Unchanged Result: 198.51.100.121 (src 198.51.100.117) then 2001::1 (src fe80::1) (prefer matching scope)

Candidate Source Addresses: 2001::2 or fe80::1 or 10.1.2.4
Destination Address List: 2001::1 or 10.1.2.3
New Result: 10.1.2.3 (src 10.1.2.4) then 2001::1 (src 2001::2) (prefer higher precedence)

10.3.1. Handling Broken IPv6

One problem in practice that has been recently observed occurs when a host has IPv4 connectivity to the Internet, but has "broken" IPv6 connectivity to the Internet in that it has a global IPv6 address, but is discconnected from the IPv6 Internet. Since the default policy table prefers IPv6, this can result in unwanted timeouts.

This can be solved by configuring the table to prefer IPv4 as shown above. An implementation that has some means to detect that it is not connected to the IPv6 Internet MAY do this automatically. An

implementation could instead treat it as part of its implementation of Rule 1 (avoid unusable destinations).

10.4. Configuring Preference for Link-Local Addresses

The destination address selection rules give preference to destinations of smaller scope. For example, a link-local destination will be sorted before a global scope destination when the two are otherwise equally suitable. An administrator can change the policy table to reverse this preference and sort global destinations before link-local destinations:

Prefix	Precedence	Label
::1/128	50	0
::/0	40	1
::ffff:0:0/96	35	4
fe80::/10	33	1
2002::/16	30	2
2001::/32	5	5
fc00::/7	3	13
::/96	1	3
fec0::/10	1	11
3ffe::/16	1	12

This change to the default policy table produces the following behavior:

Candidate Source Addresses: 2001::2 or fe80::2
Destination Address List: 2001::1 or fe80::1
New Result: 2001::1 (src 2001::2) then fe80::1 (src fe80::2) (prefer higher precedence)

Candidate Source Addresses: 2001::2 (deprecated) or fe80::2
Destination Address List: 2001::1 or fe80::1
Unchanged Result: fe80::1 (src fe80::2) then 2001::1 (src 2001::2) (avoid deprecated addresses)

10.5. Configuring a Multi-Homed Site

Consider a site A that has a business-critical relationship with another site B. To support their business needs, the two sites have contracted for service with a special high-performance ISP. This is in addition to the normal Internet connection that both sites have with different ISPs. The high-performance ISP is expensive and the two sites wish to use it only for their business-critical traffic with each other.

Each site has two global prefixes, one from the high-performance ISP

and one from their normal ISP. Site A has prefix 2001:aaaa:aaaa::/48 from the high-performance ISP and prefix 2007:0:aaaa::/48 from its normal ISP. Site B has prefix 2001:bbbb:bbbb::/48 from the high-performance ISP and prefix 2007:0:bbbb::/48 from its normal ISP. All hosts in both sites register two addresses in the DNS.

The routing within both sites directs most traffic to the egress to the normal ISP, but the routing directs traffic sent to the other site's 2001 prefix to the egress to the high-performance ISP. To prevent unintended use of their high-performance ISP connection, the two sites implement ingress filtering to discard traffic entering from the high-performance ISP that is not from the other site.

The default policy table and address selection rules produce the following behavior:

Candidate Source Addresses: 2001:aaaa:aaaa::a or 2007:0:aaaa::a or fe80::a
Destination Address List: 2001:bbbb:bbbb::b or 2007:0:bbbb::b
Result: 2007:0:bbbb::b (src 2007:0:aaaa::a) then 2001:bbbb:bbbb::b (src 2001:aaaa:aaaa::a) (longest matching prefix)

In other words, when a host in site A initiates a connection to a host in site B, the traffic does not take advantage of their connections to the high-performance ISP. This is not their desired behavior.

Candidate Source Addresses: 2001:aaaa:aaaa::a or 2007:0:aaaa::a or fe80::a
Destination Address List: 2001:cccc:cccc::c or 2006:cccc:cccc::c
Result: 2001:cccc:cccc::c (src 2001:aaaa:aaaa::a) then 2006:cccc:cccc::c (src 2007:0:aaaa::a) (longest matching prefix)

In other words, when a host in site A initiates a connection to a host in some other site C, the reverse traffic may come back through the high-performance ISP. Again, this is not their desired behavior.

This predicament demonstrates the limitations of the longest-matching-prefix heuristic in multi-homed situations.

However, the administrators of sites A and B can achieve their desired behavior via policy table configuration. For example, they can use the following policy table:

Prefix	Precedence	Label
::1/128	50	0
2001:aaaa:aaaa::/48	43	6
2001:bbbb:bbbb::/48	43	6
::/0	40	1
::ffff:0:0/96	35	4
2002::/16	30	2
2001::/32	5	5
fc00::/7	3	13
::/96	1	3
fec0::/10	1	11
3ffe::/16	1	12

This policy table produces the following behavior:

Candidate Source Addresses: 2001:aaaa:aaaa::a or 2007:0:aaaa::a or fe80::a
Destination Address List: 2001:bbbb:bbbb::b or 2007:0:bbbb::b
New Result: 2001:bbbb:bbbb::b (src 2001:aaaa:aaaa::a) then 2007:0:bbbb::b (src 2007:0:aaaa::a) (prefer higher precedence)

In other words, when a host in site A initiates a connection to a host in site B, the traffic uses the high-performance ISP as desired.

Candidate Source Addresses: 2001:aaaa:aaaa::a or 2007:0:aaaa::a or fe80::a
Destination Address List: 2001:cccc:cccc::c or 2006:cccc:cccc::c
New Result: 2006:cccc:cccc::c (src 2007:0:aaaa::a) then 2001:cccc:cccc::c (src 2007:0:aaaa::a) (longest matching prefix)

In other words, when a host in site A initiates a connection to a host in some other site C, the traffic uses the normal ISP as desired.

10.6. Configuring ULA Preference

RFC 5220 [RFC5220] sections 2.1.4, 2.2.2, and 2.2.3 describe address selection problems related to ULAs [RFC4193]. By default, global IPv6 destinations are preferred over ULA destinations, since an arbitrary ULA is not necessarily reachable:

Candidate Source Addresses: 2001:db8:1::1 or fd11:1111:1111:1::1
Destination Address List: 2001:db8:2::2 or fd22:2222:2222:2::2
Result: 2001:db8:2::2 (src 2001:db8:1::1) then fd22:2222:2222:2::2 (src fd11:1111:1111:1::1) (prefer higher precedence)

However, a site-specific policy entry can be used to cause ULAs within a site to be preferred over global addresses as follows.

Prefix	Precedence	Label
::1/128	50	0
fd11:1111:1111::/48	45	14
::/0	40	1
::ffff:0:0/96	35	4
2002::/16	30	2
2001::/32	5	5
fc00::/7	3	13
::/96	1	3
fec0::/10	1	11
3ffe::/16	1	12

Such a configuration would have the following effect:

Candidate Source Addresses: 2001:db8:1::1 or fd11:1111:1111:1::1
Destination Address List: 2001:db8:2::2 or fd22:2222:2222:2::2
Unchanged Result: 2001:db8:2::2 (src 2001:db8:1::1) then fd22:2222:
2222:2::2 (src fd11:1111:1111:1::1) (prefer higher precedence)

Candidate Source Addresses: 2001:db8:1::1 or fd11:1111:1111:1::1
Destination Address List: 2001:db8:2::2 or fd11:1111:1111:2::2
New Result: fd11:1111:1111:2::2 (src fd11:1111:1111:1::1) then 2001:
db8:2::2 (src 2001:db8:1::1) (prefer higher precedence)

Since ULAs are defined to have a /48 site prefix, an implementation might choose to add such a row automatically on a machine with a ULA.

It is also worth noting that ULAs are assigned global scope. As such, the existence of one or more rows in the prefix policy table is important so that source address selection does not choose a ULA purely based on longest match:

Candidate Source Addresses: 2001:db8:1::1 or fd11:1111:1111:1::1
Destination Address List: ff00:1
Result: 2001:db8:1::1 (prefer matching label)

10.7. Configuring 6to4 Preference

By default, NAT'ed IPv4 is preferred over 6to4-relayed connectivity:

Candidate Source Addresses: 2002:836b:4179::2 or 10.1.2.3
Destination Address List: 2001:db8:1::1 or 203.0.113.1
Result: 203.0.113.1 (src 10.1.2.3) then 2001:db8:1::1 (src 2002:836b:
4179::2) (prefer matching label)

However, NAT'ed IPv4 is now also preferred over 6to4-to-6to4 connectivity by default. Since a 6to4 prefix might be used natively within an organization, a site-specific policy entry can be used to

cause native IPv6 communication (using a 6to4 prefix) to be preferred over NAT'ed IPv4 as follows.

Prefix	Precedence	Label
::1/128	50	0
2002:836b:4179::/48	45	14
::/0	40	1
::ffff:0:0/96	35	4
2002::/16	30	2
2001::/32	5	5
fc00::/7	3	13
::/96	1	3
fec0::/10	1	11
3ffe::/16	1	12

Such a configuration would have the following effect:

Candidate Source Addresses: 2002:836b:4179:1::1 or 10.1.2.3
Destination Address List: 2002:836b:4179:2::2 or 203.0.113.1
New Result: 2002:836b:4179:2::2 (src 2002:836b:4179:1::1) then 203.0.113.1 (sec 10.1.2.3) (prefer higher precedence)

Since 6to4 addresses are defined to have a /48 site prefix, an implementation might choose to add such a row automatically on a machine with a native IPv6 address with a 6to4 prefix.

11. References

11.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC3056] Carpenter, B. and K. Moore, "Connection of IPv6 Domains via IPv4 Clouds", RFC 3056, February 2001.

[RFC3701] Fink, R. and R. Hinden, "6bone (IPv6 Testing Address Allocation) Phaseout", RFC 3701, March 2004.

[RFC3879] Huitema, C. and B. Carpenter, "Deprecating Site Local Addresses", RFC 3879, September 2004.

[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.

[RFC4380] Huitema, C., "Teredo: Tunneling IPv6 over UDP through Network Address Translations (NATs)", RFC 4380, February 2006.

[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.

[RFC6145] Li, X., Bao, C., and F. Baker, "IP/ICMP Translation Algorithm", RFC 6145, April 2011.

11.2. Informative References

[I-D.ietf-6man-addr-select-opt] Matsumoto, A., Fujisaki, T., Kato, J., and T. Chown, "Distributing Address Selection Policy using DHCPv6", draft-ietf-6man-addr-select-opt-03 (work in progress), February 2012.

[RFC1794] Brisco, T., "DNS Support for Load Balancing", RFC 1794, April 1995.

[RFC1918] Rekhter, Y., Moskowitz, R., Karrenberg, D., Groot, G., and E. Lear, "Address Allocation for Private Internets", BCP 5, RFC 1918, February 1996.

[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.

[RFC3493] Gilligan, R., Thomson, S., Bound, J., McCann, J., and W. Stevens, "Basic Socket Interface Extensions for IPv6", RFC 3493, February 2003.

[RFC3927] Cheshire, S., Aboba, B., and E. Guttman, "Dynamic Configuration of IPv4 Link-Local Addresses", RFC 3927, May 2005.

[RFC4007] Deering, S., Haberman, B., Jinmei, T., Nordmark, E., and B. Zill, "IPv6 Scoped Address Architecture", RFC 4007, March 2005.

[RFC4191] Draves, R. and D. Thaler, "Default Router Preferences and More-Specific Routes", RFC 4191, November 2005.

[RFC4213] Nordmark, E. and R. Gilligan, "Basic Transition Mechanisms for IPv6 Hosts and Routers", RFC 4213, October 2005.

[RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", RFC 4861, September 2007.

[RFC5014] Nordmark, E., Chakrabarti, S., and J. Laganier, "IPv6 Socket API for Source Address Selection", RFC 5014, September 2007.

[RFC5214] Templin, F., Gleeson, T., and D. Thaler, "Intra-Site Automatic Tunnel Addressing Protocol (ISATAP)", RFC 5214, March 2008.

[RFC5220] Matsumoto, A., Fujisaki, T., Hiromi, R., and K. Kanayama, "Problem Statement for Default Address Selection in Multi-Prefix Environments: Operational Issues of RFC 3484 Default Rules", RFC 5220, July 2008.

[RFC5969] Townsley, W. and O. Troan, "IPv6 Rapid Deployment on IPv4 Infrastructures (6rd) -- Protocol Specification", RFC 5969, August 2010.

[RFC6052] Bao, C., Huitema, C., Bagnulo, M., Boucadair, M., and X. Li, "IPv6 Addressing of IPv4/IPv6 Translators", RFC 6052, October 2010.

[RFC6275] Perkins, C., Johnson, D., and J. Arkko, "Mobility Support in IPv6", RFC 6275, July 2011.

Appendix A. Acknowledgements

RFC 3484 acknowledged the contributions of the IPng Working Group, particularly Marc Blanchet, Brian Carpenter, Matt Crawford, Alain Durand, Steve Deering, Robert Elz, Jun-ichiro itojun Hagino, Tony Hain, M.T. Hollinger, JINMEI Tatuya, Thomas Narten, Erik Nordmark, Ken Powell, Markku Savela, Pekka Savola, Hesham Soliman, Dave Thaler, Mauro Tortonesi, Ole Troan, and Stig Venaas. In addition, the anonymous IESG reviewers had many great comments and suggestions for clarification.

This revision was heavily influenced by the work by Arifumi Matsumoto, Jun-ya Kato, and Tomohiro Fujisaki in a working draft that made proposals for this revision to adopt, with input from Pekka Savola, Remi Denis-Courmont, Francois-Xavier Le Bail, and the 6man Working Group. Dmitry Anipko, Mark Andrews, and Ray Hunter also

provided valuable feedback on this revision.

Appendix B. Changes Since RFC 3484

Some changes were made to the default policy table that were deemed to be universally useful and cause no harm in every reasonable network environment. In doing so, care was taken to use the same preference and label values as in RFC 3484 whenever possible, and for new rows to use label values less likely to collide with values that might already be in use in additional rows on some hosts. These changes are:

1. Added the Teredo [RFC4380] prefix (2001::/32), with the preference and label values already widely used in popular implementations.
2. Added a row for ULAs (fc00::/7) below native IPv6 since they are not globally reachable, as discussed in Section 10.6.
3. Added a row for site-local addresses (fec0::/10) in order to depreference them, for consistency with the example in Section 10.3, since they are deprecated [RFC3879].
4. Depreferenced 6to4 (2002::/32) below native IPv4 since 6to4 connectivity is less reliable today (and is expected to be phased out over time, rather than becoming more reliable). It remains above Teredo since 6to4 is more efficient in terms of connection establishment time, bandwidth, and server load.
5. Depreferenced IPv4-Compatible addresses (::/96) since they are now deprecated [RFC4291] and not in common use.
6. Added a row for 6bone testing addresses (3ffe::/16) in order to depreference them as they have also been phased out [RFC3701].

Similarly, some changes were made to the rules, as follows:

1. Changed the definition of CommonPrefixLen() to only compare bits up to the source address's prefix length. The previous definition used the entire source address, rather than only its prefix. As a result, when a source and destination addresses had the same prefix, common bits in the interface ID would previously result in overriding DNS load balancing [RFC1794] by forcing the destination address with the most bits in common to be always chosen. The updated definition allows DNS load balancing to continue to be used as a tie breaker.
2. Added Rule 5.5 to allow choosing a source address from a prefix advertised by the chosen next-hop for a given destination. This allows better connectivity in the presence of BCP 38 [RFC2827] ingress filtering and egress filtering. Previously, RFC 3484 had issues with multiple egress networks reached via the same interface, as discussed in [RFC5220].

3. Removed restriction against anycast addresses in the candidate set of source addresses, since the restriction against using IPv6 anycast addresses as source addresses was removed in Section 2.6 of RFC 4291 [RFC4291].
4. Changed mapping of RFC 1918 [RFC1918] addresses to global scope in Section Section 3.2. Previously they were mapped to site-local scope. However, experience has resulted in current implementations already using global scope instead. When they were mapped to site-local, Destination Address Selection Rule 2 (Prefer matching scope) would cause IPv6 to be preferred in scenarios such as that described in Section 10.7. The change to global scope allows configurability via the prefix policy table.

Finally, some editorial changes were made, including:

1. Changed global IP addresses in examples to use ranges reserved for documentation.
2. Added additional examples in Section 10.6 and Section 10.7.
3. Added Section 10.3.1 on "broken" IPv6.
4. Updated references.

Authors' Addresses

Dave Thaler (editor)
Microsoft
One Microsoft Way
Redmond, WA 98052

Phone: +1 425 703 8835
Email: dthaler@microsoft.com

Richard Draves
Microsoft Research
One Microsoft Way
Redmond, WA 98052

Phone: +1 425 706 2268
Email: richdr@microsoft.com

Arifumi Matsumoto
NTT SI Lab
Midori-Cho 3-9-11
Musashino-shi, Tokyo 180-8585
Japan

Phone: +81 422 59 3334
Email: arifumi@nttv6.net

Tim Chown
University of Southampt on
Southampton, Hampshire SO17 1BJ
United Kingdom

Email: tjc@ecs.soton.ac.uk