

Source and Destination Address Selection for IPv6

Joseph Davies

Introduction

Unlike typical IPv4 hosts, IPv6 hosts can have multiple addresses assigned to LAN or tunneling interfaces that correspond to the following:

- **Different scopes** For example, a LAN interface on an IPv6 host always has a link-local address (scoped for the local subnet) and typically has either a global address (scoped for the entire Internet) or a site-local address (scoped for the site of an organization). In some cases, a LAN interface might be configured with all three types of addresses (a link-local, a global, and a site-local address).
- **Different states** Autoconfigured stateless addresses can be in a non-deprecated (also known as preferred) or deprecated state.
- **Different uses** For global address prefixes, it is possible to have temporary addresses (with a randomly determined interface ID) and public addresses (with an interface ID based on the IEEE 802 address of the LAN adapter). For mobile IPv6 nodes, it is possible to have a home address and care-of addresses.

Additionally, an IPv6/IPv4 host has one or more tunnel interfaces that could have link-local, global, and site-local addresses assigned.

For a typical IPv4-only host that has a single interface assigned one IPv4 address and resolves names using the Domain Name System (DNS), the choice of which IPv4 addresses to use as the source and destination when initiating communication is straightforward. The source IPv4 address is the address assigned to the interface of the host. The destination addresses to which connections are attempted are the IPv4 addresses returned in the DNS Name Query Response message.

For a typical IPv6 host that has multiple IPv6 addresses assigned to multiple interfaces and multiple IPv6 addresses are returned in the DNS Name Query Response message, the choice of the source and destination IPv6 address is more complex. The source and destination IPv6 addresses should be matched in scope and purpose. For example, an IPv6 host should not choose a link-local source address when communicating with a global destination address. Additionally, the possible destination address should be sorted by preference.

To provide a standardized method to choose source and destination IPv6 addresses with which to attempt connections, RFC 3484 defines the following required algorithms:

- A source address selection algorithm to choose the best source address to use with a destination address.
- A destination address selection algorithm to sort the list of possible destination addresses in order of preference.

Because RFC 3484 defines a standard method of determining source and destination addresses, applications do not need to include their own address selection algorithms, reducing the development burden on IPv6-capable applications.

The source address selection algorithm is not used when the application specifies the source address. The destination address selection algorithm is not used when the application specifies the destination address, rather than a name. Additionally, the application can override the order of the destination addresses determined by the destination address selection algorithm.

Source and Destination Address Selection for IPv6

Joseph Davies

To allow administrative control over the preference of source or destination addresses based on address prefixes, RFC 3484 defines the use of a local policy table. The table consists of a series of entries containing the following:

- An address prefix Used to determine the best entry in the table for a given source or destination address. Just like the IPv6 route determination process, the best entry is based on the longest matching prefix.
- A precedence value Used to prefer one destination address to another. For example, when comparing two destination addresses (D1 and D2), if the precedence value for D1 is greater than the precedence value for D2, D1 will be preferred over D2.
- A label value Used to prefer a specific source address if the label of the source address matches the label of the destination address. For example, when comparing two source addresses (S1 and S2) against a destination address (D), if the label for S1 is the same as the label for D and the label for S2 is not same as the label for D, S1 is preferred for use over S2.

On computers running Windows XP with Service Pack 1 and later, Windows Server 2003, Windows Vista, or Windows Server 2008, you can view the prefix policy table with the netsh interface ipv6 show prefixpolicy command. The following is an example of the default prefix policy table.

Precedence	Label	Prefix
5	5	3ffe:831f::/32
10	4	::ffff:0:0/96
20	3	::/96
30	2	2002::/16
40	1	::/0
50	0	:::1/128

You can manage entries in the prefix policy table with the netsh interface ipv6 add prefixpolicy, netsh interface ipv6 set prefixpolicy, and netsh interface ipv6 delete prefixpolicy commands.

The following sections describe the source address selection and destination address selection algorithms in detail and provide an example of their use.

Source Address Selection Algorithm

The purpose of the source address selection algorithm is to identify a source IPv6 address of maximum preference for a given destination IPv6 address from a list of candidate source addresses. For hosts, the list of candidate source addresses consist of the unicast addresses assigned to the next-hop interface for the destination. For routers, the list of candidates can include addresses assigned to any forwarding interface of the router.

The source address selection algorithm compares two source addresses against the destination address and determines which of the two has a higher preference. By iteratively comparing the list of candidate source addresses, the algorithm determines the most preferred source address.

When comparing two possible source addresses S1 and S2 against a destination address D, the source address selection algorithm performs the following analysis:

1. Prefer the source address that equals the destination address.

Source and Destination Address Selection for IPv6

Joseph Davies

- If $S1 = D$, prefer $S1$.
 - If $S2 = D$, prefer $S2$.
 - If neither $S1$ nor $S2$ equal D , $S1$ and $S2$ are at the same level of preference.
2. Prefer the source address that has a scope appropriate for the destination address.
- If the scope of $S1$ is smaller than the scope of $S2$:
 - If the scope of $S1$ is smaller than the scope of the destination, prefer $S2$.
 - Otherwise, prefer $S1$.
 - If the scope of $S2$ is smaller than the scope of $S1$:
 - If the scope of $S2$ is smaller than the scope of the destination, prefer $S1$. Otherwise, prefer $S2$.
 - If both $S1$ and $S2$ have the same appropriate scope for D or a smaller scope than D , $S1$ and $S2$ are at the same level of preference.
3. Prefer an address that is not deprecated over one that is deprecated.
- If $S2$ is deprecated and $S1$ is not deprecated, prefer $S1$.
 - If $S1$ is deprecated and $S2$ is not deprecated, prefer $S2$.
 - If both $S1$ and $S2$ are not deprecated or both $S1$ and $S2$ are deprecated, $S1$ and $S2$ are at the same level of preference.
4. Prefer the use of a home address (for mobile IPv6 nodes).
- If $S1$ is both a home address and a care-of address and $S2$ is not, prefer $S1$.
 - If $S1$ is a home address and $S2$ is a care-of address, prefer $S1$.
 - If $S2$ is both a home address and a care-of address and $S1$ is not, prefer $S2$.
 - If $S2$ is a home address and $S1$ is a care-of address, prefer $S2$.
 - If neither $S1$ nor $S2$ are home addresses, $S1$ and $S2$ are at the same level of preference.
5. For routers, prefer the source address that is assigned to the next-hop interface for the destination.
- If $S1$ is assigned to the interface that will be used to send packets to D , prefer $S1$.
 - If $S2$ is assigned to the interface that will be used to send packets to D , prefer $S2$.
 - If neither $S1$ nor $S2$ are assigned to the outgoing interface for D or if both $S1$ and $S2$ are assigned to the outgoing interface for D , $S1$ and $S2$ are at the same level of preference.
6. Prefer the source address that has the same label in the prefix policy table as the destination address.

Source and Destination Address Selection for IPv6

Joseph Davies

- If the label of S1 matches the label of D and the label of S2 does not match the label of D, prefer S1.
 - If the label of S2 matches the label of D and the label of S1 does not match the label of D, prefer S2.
 - If neither S1 nor S2 have the same label as D or if both S1 and S2 have the same label as D, S1 and S2 are at the same level of preference.
7. Prefer the source address that is a public address over the source address that is a temporary address.
- If S1 is a public address and S2 is a temporary address, prefer S1.
 - If S2 is a public address and S1 is a temporary address, prefer S2.
 - If both S1 and S2 are public addresses or if both S1 and S2 are temporary addresses, S1 and S2 are at the same level of preference.
8. Prefer the source address that has the longest matching prefix with the destination.
- If the matching prefix length of S1 and D is greater than the matching prefix length of S2 and D, prefer S1.
 - If the matching prefix length of S2 and D is greater than the matching prefix length of S1 and D, prefer S2.
 - If both S1 and S2 have the same longest matching prefix length, S1 and S2 are at the same level of preference.

Destination Address Selection Algorithm

The purpose of the destination address selection algorithm is to sort the list of possible IPv4 and IPv6 destination addresses in order of highest to lowest preference. Destination IPv4 addresses are expressed as IPv4-mapped addresses (::ffff:w.x.y.z) and are scoped as global for public IPv4 addresses, site-local for private IPv4 addresses, and link-local for Automatic Private IP Addressing (APIPA) addresses.

The destination address selection algorithm compares two destination addresses and determines which of the two has a higher preference. The algorithm sorts the list of destination addresses by iteratively comparing the list of candidate destination addresses.

When comparing two possible destination addresses D1 and D2, the destination address selection algorithm performs the following analysis:

1. Prefer the destination that is reachable to the one that is not.
 - If D2 is known to be unreachable or if the source address for D2 is not defined, prefer D1.
 - If D1 is known to be unreachable or if the source address for D1 is not defined, prefer D2.

Source and Destination Address Selection for IPv6

Joseph Davies

- If both D1 and D2 are both reachable or are both unreachable, D1 and D2 are at the same level of preference.
2. Prefer the destination that matches the scope of its source address.
 - If the scope of D1 is the same as the scope of its source address and the scope of D2 is not the same as the scope of its source address, prefer D1.
 - If the scope of D2 is the same as the scope of its source address and the scope of D1 is not the same as the scope of its source address, prefer D2.
 - If both D1 and D2 are at the same or at different scopes for their respective source addresses, D1 and D2 are at the same level of preference.
 3. Prefer destination addresses with source addresses that are not deprecated.
 - If the source address for D1 is not deprecated and the source address for D2 is deprecated, prefer D1.
 - If the source address for D2 is not deprecated and the source address for D1 is deprecated, prefer D2.
 - If D1 and D2 are either both deprecated or both not deprecated, D1 and D2 are at the same level of preference.
 4. Prefer destinations with source addresses that are home addresses.
 - If the source address for D1 is both a home address and a care-of address and the source address for D2 is not, prefer D1.
 - If the source address for D1 is a home address and the source address for D2 is a care-of address, prefer D1.
 - If the source address for D2 is both a home address and a care-of address and the source address for D1 is not, prefer D2.
 - If the source address for D2 is a home address and the source address for D1 is a care-of address, prefer D2.
 - If neither D1 nor D2 are home addresses or care-of addresses, D1 and D2 are at the same level of preference.
 5. Prefer the destination address that has the same label from the prefix policy table as its source address.
 - If the label of the source address for D1 matches the label of D1 and the label of the source address for D2 does not match the label of D2, prefer D1.
 - If the label of the source address for D2 matches the label of D2 and the label of the source address for D1 does not match the label of D1, prefer D2.

Source and Destination Address Selection for IPv6

Joseph Davies

- If both D1 and D2 match the labels of their respective source addresses or do not match the labels of their respective source addresses, D1 and D2 are at the same level of preference.
6. Prefer the destination address that has the highest precedence from the prefix policy table.
 - If the precedence of D1 is higher than the precedence of D2, prefer D1.
 - If the precedence of D2 is higher than the precedence of D1, prefer D2.
 - If D1 and D2 have the same precedence, D1 and D2 are at the same level of preference.
 7. Prefer a native IPv6 destination to an IPv6 transition technology destination.
 - If D2 is reached over an IPv6 transition technology and D1 is not, prefer D1.
 - If D1 is reached over an IPv6 transition technology and D2 is not, prefer D2.
 - If D1 and D2 are both native IPv6 addresses or are both IPv6 transition addresses, D1 and D2 are at the same level of preference.
 8. Prefer the destination address with the smallest scope.
 - If the scope for D1 is smaller than D2, prefer D1.
 - If the scope for D2 is smaller than D1, prefer D2.
 - If D1 and D2 have the same scope, D1 and D2 are at the same level of preference.
 9. Prefer the destination address that has the longest matching prefix length with its source address.
 - If the matching prefix length of D1 and its source address is larger than the matching prefix length of D2 and its source address, prefer D1.
 - If the matching prefix length of D2 and its source address is larger than the matching prefix length of D1 and its source address, prefer D2.
 - If both D1 and D2 have the same longest matching prefix length with their respective source addresses, D1 and D2 are at the same level of preference.

Leave the order of destination addresses unchanged.

Example of Using Address Selection

Host A is an IPv6/IPv4 host that has multiple interfaces and multiple addresses for each interface. Host A constructs and sends a DNS Name Query Request message for all resource records corresponding to the name `web.example.com`. The DNS Name Query Response message received by Host A contains a single IPv4 address record and multiple IPv6 address records.

Based on the contents of the DNS Name Query Response message, IPv6 on Host A does the following:

1. Use the source address selection algorithm to determine the most preferred source address to use with each destination IPv6 address.
2. Use the destination address selection algorithm to determine the preference order of the destination IPv4 and IPv6 addresses.

Source and Destination Address Selection for IPv6

Joseph Davies

The ordered set of destination addresses with their corresponding source addresses are provided to the application. The application can then try communicating using each destination and source address combination until communications are established.

For example, Host A is configured with the following addresses.

LAN interface:

- 2001:db8:21a5:a454:2aa:ff:fe21:5c2f (global address, non-deprecated state, public use)
- 2001:db8:21a5:a454:20da:3198:2c50:1a57 (global address, non-deprecated state, temporary use)
- 2001:db8:21a5:a454:1d15:9c:8e4c:902b (global address, deprecated state, temporary use)
- fec0:3a4f:78ea:a454:2aa:ff:fe21:5c2f (site-local address, non-deprecated state)
- fe80::2aa:ff:fe21:5c2f (link-local address, non-deprecated state)
- 157.60.17.211 (public IPv4 address [global scope], non-deprecated state)

ISATAP tunneling interface:

- 2001:db8:21a5:a499:200:5efe:157.60.17.211 (global ISATAP address, non-deprecated state)
- fe80::200:5efe:157.60.17.211 (link-local address, non-deprecated state)

The DNS Name Query Response message for the name web.example.com returns the following addresses:

- 207.73.118.98 (public IPv4 address [global scope])
- 2001:db8:21a5:a4ca:2aa:ff:fe35:2c1a (global address)
- 2001:db8:21a5:a499:200:5efe:207.73.118.98 (ISATAP global address)
- fec0:3a4f:2a34:1aa7:2aa:ff:fe35:2c1a (site-local address)

The result of the source address selection algorithm is the following:

- For destination address 2001:db8:21a5:a4ca:2aa:ff:fe35:2c1a, the preferred source address to use is 2001:db8:21a5:a454:2aa:ff:fe21:5c2f (global address, non-deprecated state, public use) because it has a matching scope, it is a public address, and has the longest prefix match with the destination.
- For destination address 2001:db8:21a5:a499:200:5efe:207.73.118.98, the preferred source address to use is 2001:db8:21a5:a499:200:5efe:157.60.17.211 (global ISATAP address, non-deprecated state) because it has a matching scope and the longest prefix match with the destination.

Source and Destination Address Selection for IPv6

Joseph Davies

- For destination address fec0:3a4f:2a34:1aa7:2aa:ff:fe35:2c1a, the preferred source address to use is fec0:3a4f:78ea:a454:2aa:ff:fe21:5c2f (site-local address, non-deprecated state) because it has a matching scope and the longest prefix match with the destination.

The result of the destination address selection algorithm is the following destination address in preferred order:

1. fec0:3a4f:2a34:1aa7:2aa:ff:fe35:2c1a (smallest scope)
2. 2001:db8:21a5:a4ca:2aa:ff:fe35:2c1a (native addresses preferred over ISATAP addresses)
3. 2001:db8:21a5:a499:200:5efe:207.73.118.98 (ISATAP addresses preferred over IPv4-mapped addresses)
4. 207.73.118.98

Therefore, Host A attempts to establish communications with the web.example.com server using the following sets of source and destination addresses:

1. Source address fec0:3a4f:78ea:a454:2aa:ff:fe21:5c2f, destination address fec0:3a4f:2a34:1aa7:2aa:ff:fe35:2c1a
2. Source address 2001:db8:21a5:a454:2aa:ff:fe21:5c2f with destination address 2001:db8:21a5:a4ca:2aa:ff:fe35:2c1a
3. Source address 2001:db8:21a5:a499:200:5efe:157.60.17.211, destination address 2001:db8:21a5:a499:200:5efe:207.73.118.98
4. Source address 157.60.17.211, destination address 207.73.118.98

IPv6 addresses are preferred over IPv4 addresses because the entry in the default prefix policy table for IPv4 addresses (the address prefix ::ffff:0:0/96) has a low precedence value.

Note An application, such as the Ping tool included with Windows, can skip address selection and begin using the addresses as they are returned in the DNS Name Query Response message.