ICMPv6

Like IPv4, IPv6 does not provide facilities for reporting errors. Instead, IPv6 uses an updated version of the Internet Control Message Protocol (ICMP) named ICMP version 6 (ICMPv6). ICMPv6 has the common IPv4 ICMP functions of reporting delivery or forwarding errors and providing a simple echo service for troubleshooting.

The ICMPv6 protocol also provides a framework for the following:

- Multicast Listener Discovery (MLD)
  MLD is a series of three ICMPV6 messages that replace version 2 of the Internet Group Management Protocol (IGMP) for IPv4 to manage subnet multicast membership. MLD is described in more detail in "Multicast Listener Discovery."
- Neighbor Discovery (ND)
  Neighbor Discovery is a series of five ICMPv6 messages that manage node-to-node communication on a link. Neighbor Discovery replaces Address Resolution Protocol (ARP), ICMPv4 Router Discovery, and the ICMPv4 Redirect message. Neighbor Discovery is described in more detail in "Neighbor Discovery."

ICMPv6 is required for an IPv6 implementation and is defined in RFC 4443.

Types of ICMPv6 Messages

There are two types of ICMPv6 messages:

1. Error messages
   Error messages are used to report errors in the forwarding or delivery of IPv6 packets by either the destination node or an intermediate router. The value of the 8-bit Type field in ICMPv6 error messages is in the range of 0 through 127 (the high order bit is set to 0). ICMPv6 error messages include Destination Unreachable, Packet Too Big, Time Exceeded, and Parameter Problem.
2. Informational messages
   Informational messages are used to provide diagnostic functions and additional host functionality such as MLD and Neighbor Discovery. The value of the Type field in ICMPv6 informational messages is in the range of 128 through 255 (the high order bit is set to 1). ICMPv6 informational messages are described in RFC 4443 and include Echo Request and Echo Reply.

ICMPv6 Header

An ICMPv6 header is indicated by setting the previous header's Next Header field to 58. Figure 30 shows the structure of all ICMPv6 messages.

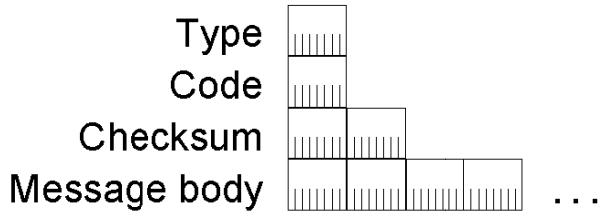

Figure 30 The ICMPv6 message structure

The fields in the ICMPv6 header are:

- **Type** – Indicates the type of ICMPv6 message. The size of this field is 8 bits. In ICMPv6 error messages, the high-order bit is set to 0. In ICMPv6 informational messages, the high-order bit is set to 1.

- **Code** – Differentiates among multiple messages within a given message type. The size of this field is 8 bits. If there is only one message for a given type, the Code field is set to 0.
- **Checksum** – Stores a checksum of the ICMPv6 message. The size of this field is 16 bits. The IPv6 pseudo-header is added to the ICMPv6 message when calculating the checksum.
- **Message body** – Contains ICMPv6 message-specific data.

ICMPv6 Error Messages

ICMPv6 error messages are used to report forwarding or delivery errors by either a router or the destination host. To conserve network bandwidth, ICMPv6 error messages are not sent for every error encountered. Instead, ICMPv6 error messages are rate limited. Rate limiting reduces the amount of bandwidth consumed by reporting errors.

Destination Unreachable

An ICMPv6 Destination Unreachable message is sent by either a router or a destination host when the packet cannot be forwarded to its destination. Figure 31 shows the ICMPv6 Destination Unreachable message.

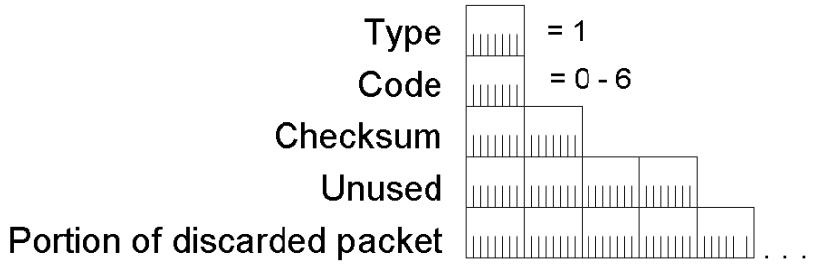

Figure 31 The ICMPv6 Destination Unreachable message

In the Destination Unreachable message, the Type field is set to 1 and the Code field is set to a value in the range of 0 through 4. After the Checksum field is the 32-bit Unused field and the portion of the discarded packet that makes the entire IPv6 packet containing the ICMPv6 message no larger than 1280 bytes (the minimum IPv6 MTU). The number of bytes of the discarded packet included in the message varies if there are IPv6 extension headers present. For an ICMPv6 message without extension headers, 1232 bytes of the discarded packet are included (1280 less a 40-byte IPv6 header and an 8-byte ICMPv6 Destination Unreachable header).
Table 5 shows the value of the Code field for the various Destination Unreachable messages.

Table 5 ICMPv6 Destination Unreachable Messages

Code value	Description
0	No route matching the destination was found in the routing table.
1	The communication with the destination is prohibited by administrative policy. This is typically sent when the packet is discarded by a firewall.
2	The address is beyond the scope of the source address.
3	The destination address is unreachable. This is typically sent because of an inability to resolve

	the destination's link layer address.
4	The destination port was unreachable. This is typically sent when an IPv6 packet containing a UDP message arrived at the destination but there were no applications listening on the destination UDP port.
5	The packet with this source address is not allowed due to inbound (ingress) or outbound (egress) packet filtering policies.
6	The packet matched a reject route and was discarded. A reject route is an address prefix configured on a router for traffic that the router must immediately discard.

Packet Too Big

An ICMPv6 Packet Too Big message is sent when the packet cannot be forwarded because the link MTU on the forwarding link is smaller than the size of the IPv6 packet. Figure 32 shows the ICMPv6 Packet Too Big message.

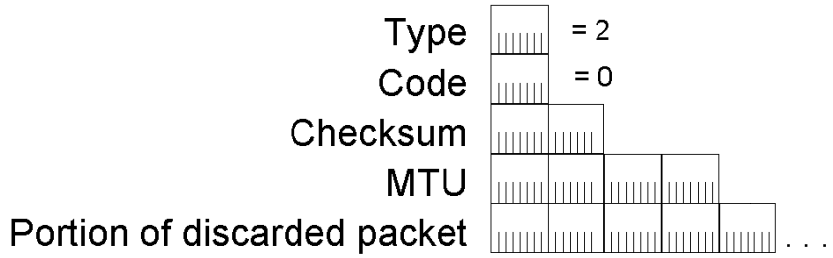

Figure 32 The ICMPv6 Packet Too Big message

In the Packet Too Big message, the Type field is set to 2 and the Code field is set to 0. After the Checksum field is the 32-bit MTU field that stores the link MTU for the link on which the packet was being forwarded. Next is the portion of the discarded packet that makes the entire IPv6 packet containing the ICMPv6 message no larger than the maximum length of 1280 bytes. The Packet Too Big message is used for the IPv6 Path MTU Discovery process described in "Path MTU Discovery."

Time Exceeded

An ICMPv6 Time Exceeded message is typically sent by a router when the Hop Limit field in the IPv6 header is zero, either upon receipt or after decrementing its value during the forwarding process. Figure 33 shows the ICMPv6 Time Exceeded message.

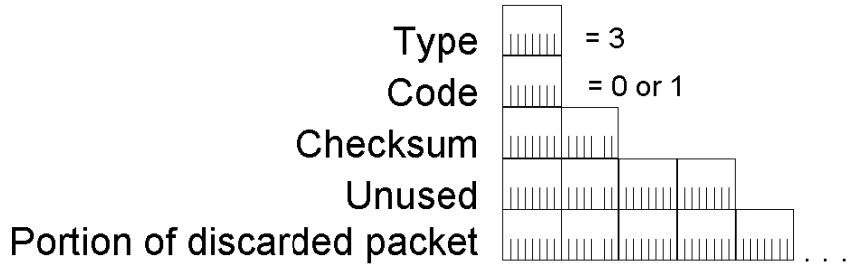

Figure 33 The ICMPv6 Time Exceeded message

In the Time Exceeded message, the Type field is set to 3 and the Code field is set to either 0 (when the Hop Limit field in the IPv6 header becomes 0) or 1 (when the fragmentation reassembly time of the destination host is exceeded). After the Checksum field is the 32-bit Unused field and the portion of the discarded packet that makes the entire IPv6 packet containing the ICMPv6 message no larger than 1280 bytes. The receipt of Time Exceeded messages for Code=0 indicates that either the Hop Limit of outgoing packets is not large enough to reach the destination or that a routing loop exists.

Parameter Problem

An ICMPv6 Parameter Problem message is either sent by a router or by the destination. This occurs when an error is encountered in either the IPv6 header or an extension header, preventing IPv6 from performing further processing. Figure 34 shows the ICMPv6 Parameter Problem message.

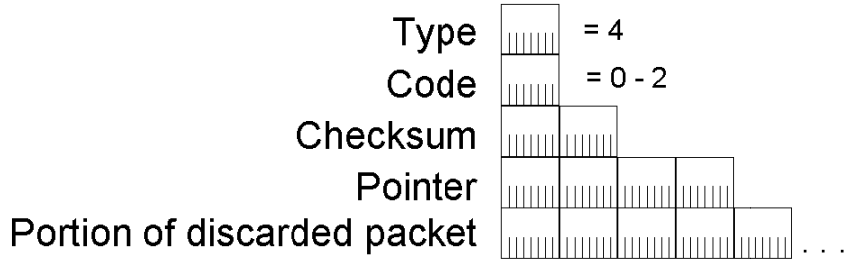

Figure 34 The ICMPv6 Parameter Problem message

In the Parameter Problem message, the Type field is set to 4 and the Code field is a value in the range of 0 through 2. After the Checksum field is the 32-bit Pointer field that indicates the byte offset in the offending IPv6 packet where the error was encountered. After the Pointer field is the portion of the discarded packet that makes the entire ICMPv6 message no larger than 1280 bytes. The Pointer field value is set to the correct offset even when the location of the error is not included in the portion of the discarded packet.
Table 6 shows the Code field values for Parameter Problem messages.

Table 6 ICMPv6 Parameter Problem Messages

Code value	Description
0	An error in a field within the IPv6 header or an extension header was encountered.
1	An unrecognized Next Header field value was encountered. This is equivalent to the IPv4 Destination Unreachable-Protocol Unreachable message.
2	An unrecognized IPv6 option was encountered.

ICMPv6 Informational Messages

ICMPv6 informational messages, defined in RFC 4443, provide diagnostic capabilities to aid in troubleshooting.

Echo Request

An ICMPv6 Echo Request message is sent to a destination to solicit an immediate Echo Reply message. The Echo Request/Echo Reply message facility provides a simple diagnostics function to aid in the troubleshooting of a variety of reachability and routing problems. Figure 35 shows the ICMPv6 Echo Request message.

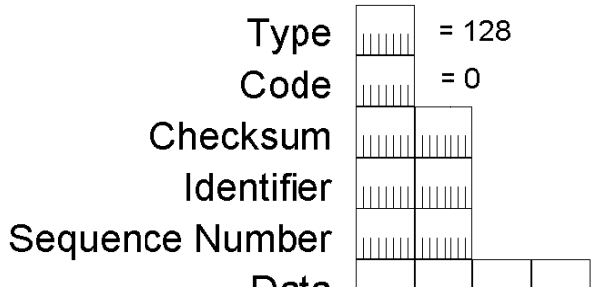

Figure 35 The ICMPv6 Echo Request message

In the Echo Request message, the Type field is set to 128 and the Code field is set to 0. After the Checksum field are the 16-bit Identifier and Sequence Number fields. The Identifier and Sequence Number fields are set by the sending host and used to match an incoming Echo Reply message with its corresponding Echo Request. The Data field is zero or more bytes of optional data that is also set by the sending host.

Echo Reply

An ICMPv6 Echo Reply message is sent in response to the receipt of an ICMPv6 Echo Request message. Figure 36 shows the ICMPv6 Echo Reply message.

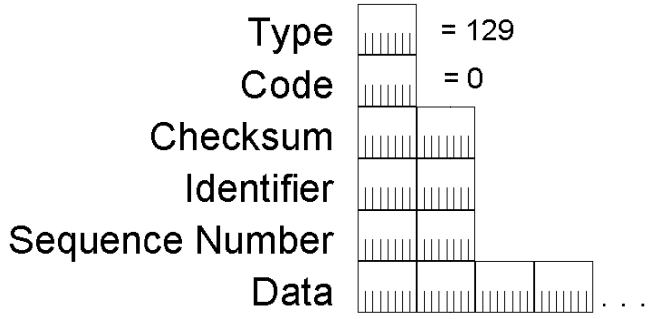

Figure 36 The ICMPv6 Echo Reply message

In the Echo Reply message, the Type field is set to 129 and the Code field is set to 0. After the Checksum field are the 16-bit Identifier and Sequence Number fields. The Identifier, Sequence Number, and Data fields are set with the same values as those in the Echo Request message that initially prompted the Echo Reply.

Comparing ICMPv4 and ICMPv6 Error Messages

Table 7 lists commonly used ICMPv4 error messages and their corresponding ICMPv6 equivalents.

Table 7 ICMPv4 Error Messages and Their Corresponding ICMPv6 Equivalents

ICMPv4 Message	ICMPv6 Equivalent
Destination Unreachable-Network unreachable (Type 3, Code 1)	Destination Unreachable-No route to destination (Type 1, Code 0)
Destination Unreachable-Host unreachable (Type 3, Code 1)	Destination Unreachable-Address unreachable (Type 1, Code 3)
Destination Unreachable-Protocol unreachable (Type 3, Code 2)	Parameter Problem-Unrecognized Next Header field (Type 4, Code 1)

Destination Unreachable-Port unreachable (Type 3, Code 3)	Destination Unreachable-Port unreachable (Type 1, Code 4)
Destination Unreachable-Fragmentation needed and DF set (Type 3, Code 4)	Packet Too Big (Type 2, Code 0)
Destination Unreachable-Communication with destination host administratively prohibited (Type 3, Code 10)	Destination Unreachable-Communication with destination administratively prohibited (Type 1, Code 1)
Time Exceeded-TTL expired in transit (Type 11, Code 0)	Time Exceeded-Hop Limit exceeded (Type 3, Code 0)
Time Exceeded-Fragmentation timer expired (Type 11, Code 1)	Time Exceeded-Fragmentation timer exceeded (Type 3, Code 1)
Parameter Problem (Type 12, Code 0)	Parameter Problem (Type 4, Code 0 or Code 2)
Source Quench (Type 4, Code 0)	This message is not present in IPv6.
Redirect (Type 5, Code 0)	Neighbor Discovery Redirect message (Type 137, Code 0). For more information, see "Neighbor Discovery."

Path MTU Discovery

The path MTU is the smallest link MTU of any link in the path between a source and a destination. IPv6 packets with a maximum size of the path MTU do not require fragmentation by the host and will be successfully forwarded by all routers on the path. To discover the path MTU, the sending node uses the receipt of ICMPV6 Packet Too Big messages.

The path MTU is discovered through the following process:

1. The sending node assumes that the path MTU is the link MTU of the interface on which the traffic is being forwarded.
2. The sending node sends IPv6 packets at the path MTU size.
3. If a router on the path is unable to forward the packet over a link with a link MTU that is smaller than the size of the packet, it discards the IPv6 packet and sends an ICMPV6 Packet Too Big message back to the sending node. The ICMPV6 Packet Too Big message contains the link MTU of the link on which the forwarding failed.
4. The sending node sets the path MTU for packets being sent to the destination to the value of the MTU field in the ICMPv6 Packet Too Big message.

The sending node starts again at step 2 and repeats steps 2 through 4 for as many times as are necessary to discover the path MTU. The path MTU is determined when either no additional ICMPv6 Packet Too Big messages are received or an acknowledgment is received from the destination.

In RFC 1981, it is recommended that IPv6 nodes support path MTU discovery. Those that do not must use the minimum link MTU of 1280 bytes as the path MTU.

Changes in Path MTU

Due to changes in routing topology, the path between source and destination might change over time. When a new path requires a lower path MTU, the earlier process begins at step 3 and repeats steps 2 through 4 until the new path MTU is discovered.

Decreases in path MTU are immediately discovered through the receipt of ICMPV6 Packet Too Big messages. Increases in path MTU must be detected by the sending node. As described in RFC 1981,

the sending node can attempt to send a larger IPv6 packet after a minimum of 5 minutes (10 minutes are recommended) upon receiving an ICMPv6 Packet Too Big message.

Multicast Listener Discovery

Multicast Listener Discovery (MLD) is the IPv6 equivalent of Internet Group Management Protocol version 2 (IGMPv2) for IPv4. MLD is a set of messages exchanged by routers and nodes, enabling routers to discover the set of multicast addresses for which there are listening nodes for each attached interface. Like IGMPv2, MLD only discovers the list of multicast addresses for which there is at least one listener, not the list of individual multicast listeners for each multicast address. Multicast Listener Discovery (MLD) is documented in RFC 2710.

MLD Messages

Unlike IGMPv2, MLD uses ICMPv6 messages instead of defining its own message structure. All MLD messages are ICMPv6 messages types 130, 131, and 132. The three types of MLD messages are:

1. Multicast Listener Query

Multicast Listener Query is used by a router to query a link for multicast listeners. There are two types of Multicast Listener Query messages: The General Query and the Multicast-Address-Specific Query. The General Query is used to query for multicast listeners of all multicast addresses. The Multicast-Address-Specific Query is used to query for multicast listeners of a specific multicast address. The two message types are distinguished by the multicast destination address in the IPv6 header and a multicast address within the Multicast Listener Query message.

2. Multicast Listener Report

Multicast Listener Report is used by a multicast listener to either report interest in receiving multicast traffic for a specific multicast address or to respond to a Multicast Listener Query.

3. Multicast Listener Done

Multicast Listener Done is used by a multicast listener to report that it is no longer interested in receiving multicast traffic for a specific multicast address.

An MLD message packet consists of an IPv6 header, a Hop-by-Hop Options extension header, and the MLD message. The Hop-by-Hop Options extension header contains the IPv6 Router Alert Option documented in RFC 2711. It is used to ensure that routers process MLD messages sent to multicast addresses on which the router is not listening. Figure 37 shows the format of an MLD message packet.

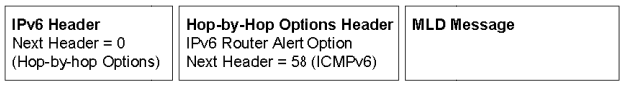

Figure 37 The Format of an MLD message packet

Multicast Listener Query

An MLD Multicast Listener Query message is equivalent to the IGMPv2 Host Membership Query message. It is used by a router to query an attached link for listening hosts.

In the IPv6 header, the source address is the link-local address of the interface on which the query is being sent. The Hop Limit field is set to 1. For the General Query, the destination address is the link-local scope all-nodes multicast address (FF02::1). For the Multicast-Address-Specific Query, the destination address is the specific multicast address being queried.

Figure 38 shows the MLD Multicast Listener Query message.

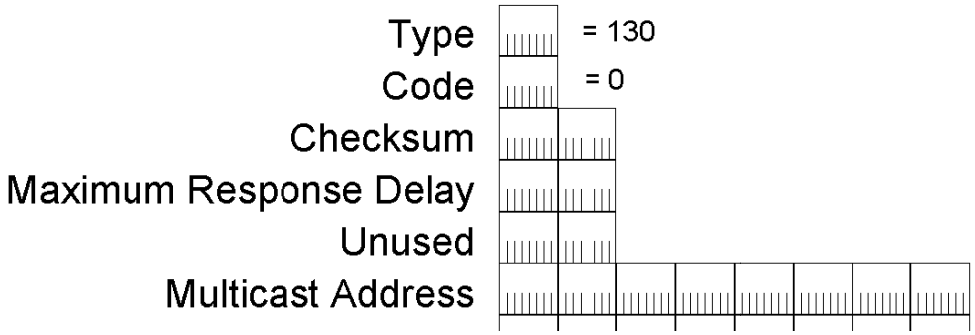

Figure 38 The MLD Multicast Listener Query message

In the MLD Multicast Listener Query message, the Type field is set to 130 and the Code field is set to 0. After the Checksum field are the 16-bit Maximum Response Delay and Reserved fields. The Maximum Response Delay is the maximum amount of time in milliseconds within which a multicast group member must report its membership using an MLD Multicast Listener Report message. In the General Query, the Multicast Address field is set to the unspecified address (::). In the Multicast-Address-Specific Query, the Multicast Address field is set to the specific multicast address that is being queried.

Multicast Listener Report

An MLD Multicast Listener Report message is equivalent to the IGMPv2 Host Membership Report message. It is used by a listening node to either report its interest in receiving multicast traffic at a specific multicast address or respond to an MLD General or Multicast-Address-Specific Query message.

In the IPv6 header, the source address is the link-local address of the interface on which the report is being sent. The Hop Limit field is set to 1 and the destination address is the specific multicast address being reported.

Figure 39 shows the MLD Multicast Listener Report message.

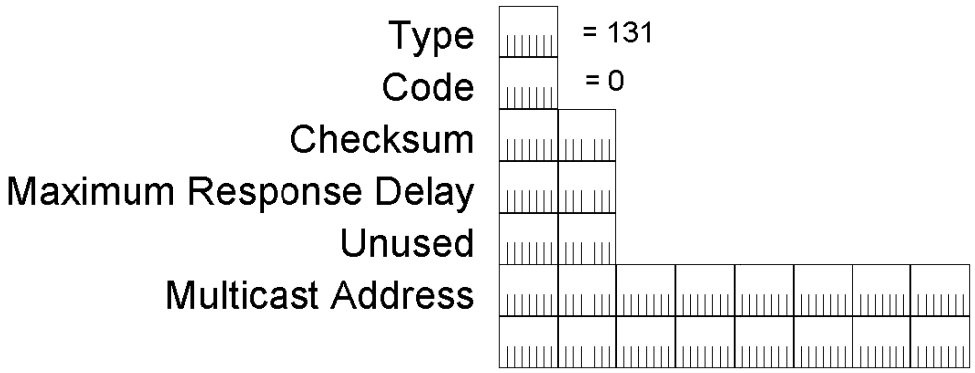

Figure 39 The MLD Multicast Listener Report message

In the MLD Multicast Listener Report message, the Type field is set to 131 and the Code field is set to 0. The Maximum Response Delay field is not used in a Multicast Listener Report message and is set to 0. The Multicast Address field is set to the specific multicast address that is being reported.

Multicast Listener Done

An MLD Multicast Listener Done message is equivalent to the IGMPv2 Leave Group message. It is used by a multicast group member to inform local routers that it might be the last group member on the subnet.

In the IPv6 header, the source address is the link-local address of the interface on which the report is being sent. The Hop Limit field is set to 1 and the destination address is the link-local scope all-routers multicast address (FF02::2).

Figure 40 shows the MLD Multicast Listener Done message.

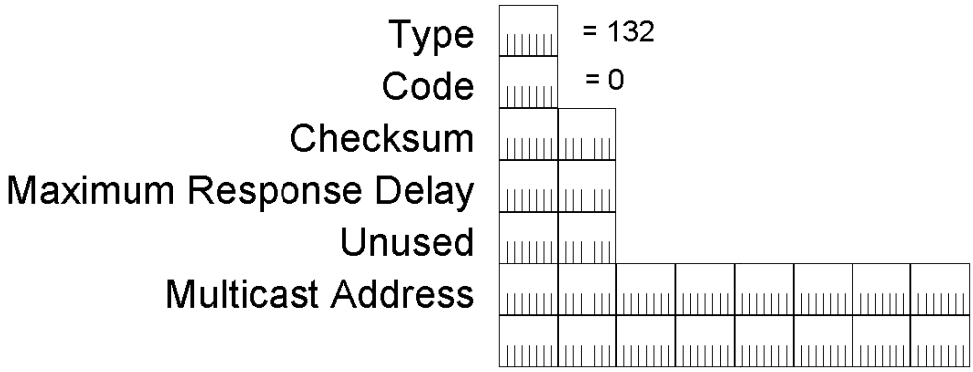

Figure 40 The MLD Multicast Listener Done message

In the MLD Multicast Listener Done message, the Type field is set to 132 and the Code field is set to 0. The Maximum Response Delay field is not used in a Multicast Listener Done message and is set to 0. The Multicast Address field is set to the specific multicast address for which the sending node is informing local routers that it is no longer as listener.

MLDv2

Windows Vista and Windows Server “Longhorn” also support Multicast Listener Discovery version 2 (MLDv2), specified in RFC 3810, which allows IPv6 hosts to register interest in source-specific multicast traffic with their neighboring routers. A host running Windows Vista or Windows Server “Longhorn” can register interest in receiving IPv6 multicast traffic from only specific source addresses (an include list) or from any source except specific source addresses (an exclude list).