



IPv6 MTU Path Discovery

IPv6 MTU Path Discovery allows a host to dynamically discover and adjust to differences in the maximum transmission unit (MTU) size of every link along a given data path.

- [Information About IPv6 MTU Path Discovery, on page 1](#)
- [How to Configure IPv6 MTU Path Discovery, on page 2](#)
- [Configuration Examples for IPv6 MTU Path Discovery, on page 3](#)
- [Additional References, on page 3](#)

Information About IPv6 MTU Path Discovery

IPv6 MTU Path Discovery

As in IPv4, path MTU discovery in IPv6 allows a host to dynamically discover and adjust to differences in the MTU size of every link along a given data path. In IPv6, however, fragmentation is handled by the source of a packet when the path MTU of one link along a given data path is not large enough to accommodate the size of the packets. Having IPv6 hosts handle packet fragmentation saves IPv6 device processing resources and helps IPv6 networks run more efficiently.



Note In IPv6, the minimum link MTU is 1280 octets. We recommend using an MTU value of 1500 octets for IPv6 links.

With IPv6 path MTU discovery, a device originating IPv6 traffic has an MTU cache that contains MTU values received in ICMPv6 "toobig" messages. In order to prevent an attacker from filling the MTU cache, the device keeps track of the destinations to which it has originated (sent) traffic, and only accepts toobig ICMPv6 messages that have an inner destination matching one of these tracked destinations.

If a malicious device can learn to which destination the device is originating traffic, it could still send a toobig ICMPv6 message to the device for this destination, even if the attacker is not on the path to this destination, and succeeds in forcing his entry into the MTU cache. The device then starts fragmenting traffic to this destination, which significantly affects device performance.

Enabling flow-label marking for locally generated traffic can mitigate this attack. Originated packets are marked with a flow label (which is randomly generated and changed every minute), and toobig messages

received are checked against the values sent. Unless an attacker can snoop traffic, the attacker will not know which flow label to use, and its too big message will be dropped.

How to Configure IPv6 MTU Path Discovery

Enabling Flow-Label Marking in Packets that Originate from the Device

This feature allows the device to track destinations to which the device has sent packets that are 1280 bytes or larger.

SUMMARY STEPS

1. **enable**
2. **configure terminal**
3. **ipv6 flowset**
4. **exit**
5. **clear ipv6 mtu**

DETAILED STEPS

	Command or Action	Purpose
Step 1	enable Example: Device> enable	Enables privileged EXEC mode. <ul style="list-style-type: none"> • Enter your password if prompted.
Step 2	configure terminal Example: Device# configure terminal	Enters global configuration mode.
Step 3	ipv6 flowset Example: Device(config)# ipv6 flowset	Configures flow-label marking in 1280-byte or larger packets sent by the device.
Step 4	exit Example: Device(config)# exit	Exits global configuration mode, and places the device in privileged EXEC mode.
Step 5	clear ipv6 mtu Example: Device# clear ipv6 mtu	Clears the MTU cache of messages.

Configuration Examples for IPv6 MTU Path Discovery

Example: Displaying IPv6 Interface Statistics

In the following example, the **show ipv6 interface** command is used to verify that IPv6 addresses are configured correctly for FastEthernet interface 1/0. Information may also be displayed about the status of IPv6 neighbor redirect messages, IPv6 neighbor discovery messages, stateless autoconfiguration, and MTU size.

```
Device# show ipv6 interface fastethernet 1/0

Ethernet0 is up, line protocol is up
  IPv6 is stalled, link-local address is FE80::1
  Global unicast address(es):
    2001:DB8:2000::1, subnet is 2001:DB8:2000::/64
    2001:DB8:3000::1, subnet is 2001:DB8:3000::/64
  Joined group address(es):
    FF02::1
    FF02::2
    FF02::1:FF00:1
  MTU is 1500 bytes
  ICMP error messages limited to one every 100 milliseconds
  ICMP redirects are enabled
  ND DAD is enabled, number of DAD attempts: 1
  ND reachable time is 30000 milliseconds
  ND advertised reachable time is 0 milliseconds
  ND advertised retransmit interval is 0 milliseconds
  ND router advertisements are sent every 200 seconds
  ND router advertisements live for 1800 seconds
  Hosts use stateless autoconfig for addresses.
```

Additional References

Related Documents

Related Topic	Document Title
IPv6 addressing and connectivity	<i>IPv6 Configuration Guide</i>
Cisco IOS commands	Cisco IOS Master Commands List, All Releases
IPv6 commands	<i>Cisco IOS IPv6 Command Reference</i>
Cisco IOS IPv6 features	<i>Cisco IOS IPv6 Feature Mapping</i>

Standards and RFCs

Standard/RFC	Title
RFCs for IPv6	<i>IPv6 RFCs</i>

MIBs

MIB	MIBs Link
No new or modified MIBs are supported by this feature, and support for existing MIBs has not been modified by this feature.	To locate and download MIBs for selected platforms, Cisco IOS releases, and feature sets, use Cisco MIB Locator found at the following URL: http://www.cisco.com/go/mibs

Technical Assistance

Description	Link
The Cisco Support and Documentation website provides online resources to download documentation, software, and tools. Use these resources to install and configure the software and to troubleshoot and resolve technical issues with Cisco products and technologies. Access to most tools on the Cisco Support and Documentation website requires a Cisco.com user ID and password.	http://www.cisco.com/cisco/web/support/index.html