

Understanding Baby Giant/Jumbo Frames Support on Catalyst 4000/4500 with Supervisor III/IV

Document ID: 29805

Contents

Introduction

Before You Begin

- Conventions
- Prerequisites
- Components Used

Baby Giant and Jumbo Support in Supervisor I and II

Baby Giant and Jumbo Frame Support in Supervisor III/IV

Baby Giants

- Software Support
- Hardware Support
- Configuration
- Verification
- Caveats

Jumbo Frames

- Software Support
- Hardware Support
- Configuration
- Verification
- Consideration when configuring Jumbo Frames on SVI
- Consideration of Configuring Jumbo Frames with Port-Channel

Related Information

Introduction

This document discusses the support of varying Maximum Transmit Units (MTUs) on Catalyst 4000/4500 series switches with Supervisor III and IV.

Standard Ethernet frame MTU is 1500 bytes. This does not include the Ethernet header and Cyclic Redundancy Check (CRC) trailer, which is 18 bytes in length, to make the total Ethernet frame size of 1518. In this document, MTU size or packet size refers only to Ethernet payload. Ethernet frame size refers to the whole Ethernet frame, including the header and the trailer. Baby giant frames refer to Ethernet frame size up to 1600 bytes, and jumbo frame refers to Ethernet frame size up to 9216 bytes.

Baby and Jumbo Frame Support on Catalyst 4000 Switches

Catalyst 4000 Switch Software	Baby Giant	Jumbo Frame	Max Frame size
CatOS ⁽¹⁾	Not Supported	Not Supported	1522
IOS ⁽²⁾	Yes (12.1(12c)EW)	Yes (12.1(13)EW)	9216

⁽¹⁾ Catalyst 4000s running CatOS refers to Supervisor I and II based Catalyst 4000 switches, Catalyst 2948G, Catalyst 2980G, and Catalyst 4912G. Refer to the Baby Giant and Jumbo Support in Supervisor I and II

section of this document for further details.

⁽²⁾ Catalyst 4000s running IOS refers to Supervisor III or IV based Catalyst 4000/4500 switches. Please refer to following sections for understanding the feature support and caveats.

Before You Begin

Conventions

For more information on document conventions, see the Cisco Technical Tips Conventions.

Prerequisites

There are no specific prerequisites for this document.

Components Used

The information in this document is based on the software and hardware versions below.

- Catalyst 4500 with Supervisor IV Engine
- Cisco IOS @12.1(13)EW

Baby Giant and Jumbo Support in Supervisor I and II

Supervisor I and II based Catalyst 4000/4500 switches, which includes the WS-C2948G, WS-C2980G, and the WS-C4912G fixed-configuration switches, do not support baby giant or jumbo frames due to an Application-Specific Integrated Circuit (ASIC) limitation.

A possible workaround is to force the switchport to accept an extra four bytes of data by configuring it as a trunk.

When a port is enabled for 802.1q trunking (Inter-Switch Link (ISL) encapsulation is not supported on Supervisor I and II based switches), the switch will automatically assume that there is an extra four bytes of data appended on, incrementing the frame size of the Layer 2 (L2) packet. Therefore, for implementations that require exactly only one tag to be carried (either 802.1q or Multiprotocol Label Switching (MPLS), but not both), it is possible to force the switchport to accept an extra four bytes of data by configuring it as a trunk port.

For example, if a port needs to carry an MPLS label, configure the port as an 802.1q trunk by changing the native VLAN to be the one desired to carry the traffic.

Baby Giant and Jumbo Frame Support in Supervisor III/IV

Baby giant refers to Ethernet frames up to 1600 bytes on the Catalyst 4000/4500 platform, or packet sizes of (MTU size) of 1552 bytes (without any header or trailer bytes). The table below lists sample protocols that can use the baby giant feature, and the configuration that is required.

Protocol/Application	Number of Header Bytes	Total Frame size	Command
----------------------	------------------------	------------------	---------

802.1q trunking	4	1500 + 4 + 18 = 1522	No MTU commands required.
QinQ pass-through (802.1q inside 802.1q, useful for ISPs to segregate customer traffic)	4 + 4	1500 + 8 + 18 = 1526	system mtu 1504
MPLS VPN pass-through (two 4-byte labels)	4 + 4	1500 + 8 + 18 = 1526	system mtu 1508
UTI/L2TPV3 pass-through (encapsulate one Ethernet packet in another Ethernet packet with a tunneling header. Useful to transport any payload such as IP/IPX and so on over an IP backbone).	18+ 20+12	1500 + 50 + 18 = 1568	system mtu 1550

Jumbo frames refer to Ethernet packets of up to 9000 bytes in size. Supervisor III and IV can handle packets up to a maximum size of 9198 bytes. This value includes the 802.1q tag or ISL VLAN tag, but does not include the Ethernet header and CRC trailer. Thus, the maximum Ethernet frame size, including the Ethernet header/trailer, is $9198 + 18 = 9216$ bytes.

Note: There is a mismatch in the maximum supportable packet size between Catalyst 4000s and Catalyst 6000s. Catalyst 6000s can support packets whose packet size is as large as 9216 bytes, which means it supports a maximum total Ethernet frame size of $9216 + 18 = 9234$ bytes.

Jumbo frames are used in situations where certain applications would benefit from using a large frame size (for example, Network File System (NFS)) for better throughput.

Baby Giants

Software Support

Baby giant support has been available since Cisco IOS software release 12.1(12c)EW for the Supervisor III and IV.

Hardware Support

The baby giant feature is supported on all modules on the Catalyst 4000/4500 with the **exception** of the following two modules:

- WS-X4418-GB module (ports 3-18 only)
- WS-X4412-2GB-TX. (ports 1-12 only)

The frames are dropped if forwarded to these ports.

Configuration

To enable baby giants, issue the **system mtu global config** command, as shown below.

```
4507(config)#system mtu ?
<1500-1552> MTU size in bytes

4507(config)#system mtu 1552
Global Ethernet MTU is set to 1552 bytes.
Note: this is the Ethernet payload size, not the total
Ethernet frame size, which includes the Ethernet
header/trailer
```

Note: The baby giant configuration applies to all interfaces that support this feature. You cannot enable the baby giant configuration per interface.

Note: If you enable jumbo frame support on a specific interface, it will override the baby giant global configuration for that interface.

Verification

The **show system mtu** command will display the globally configured MTU, as shown below.

```
Switch#show system mtu
Global Ethernet MTU is 1552 bytes.
```

The **show interfaces <interface-id> mtu** command gives the configured MTU. MTU reflects the existing configured baby giant or jumbo frame, as shown below.

```
Switch#sh interfaces fastEthernet 4/1 mtu

Port      Name          MTU
Fa4/1     Fa4/1         1552
```

The **show interface <interface-id>** command shows the globally configured baby giant value, as shown below.

```
Switch#sh int fas 4/1
FastEthernet4/1 is up, line protocol is down (notconnect)
Hardware is Fast Ethernet Port, address is 0009.e845.633f (bia 0009.e845.633f)
MTU 1552 bytes, BW 100000 Kbit, DLY 100 usec,
reliability 255/255, txload 1/255, rxload 1/255
Encapsulation ARPA, loopback not set
Keepalive set (10 sec)
Auto-duplex, Auto-speed
input flow-control is off, output flow-control is off
ARP type: ARPA, ARP Timeout 04:00:00
Last input never, output never, output hang never
Last clearing of "show interface" counters never
Input queue: 0/75/0/0 (size/max/drops/flushes); Total output drops: 0
Queueing strategy: fifo
Output queue: 0/40 (size/max)
5 minute input rate 0 bits/sec, 0 packets/sec
5 minute output rate 0 bits/sec, 0 packets/sec
L3 in Switched: ucast: 0 pkt, 0 bytes - mcast: 0 pkt, 0 bytes
L3 out Switched: ucast: 0 pkt, 0 bytes - mcast: 0 pkt, 0 bytes
0 packets input, 0 bytes, 0 no buffer
Received 0 broadcasts (0 IP multicast)
0 runts, 0 giants, 0 throttles
0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored
0 input packets with dribble condition detected
```

```
0 packets output, 0 bytes, 0 underruns
0 output errors, 0 collisions, 0 interface resets
0 babbles, 0 late collision, 0 deferred
0 lost carrier, 0 no carrier
0 output buffer failures, 0 output buffers swapped out
```

The **show interfaces <interface-id> counters all** command provides statistics for jumbo frames, as shown below.

```
sup3# sh interfaces gigabitEthernet 1/1 counters all
```

Port	InBytes	InUcastPkts	InMcastPkts	InBcastPkts
Gil/1	0	0	0	0
Port	OutBytes	OutUcastPkts	OutMcastPkts	OutBcastPkts
Gil/1	0	0	0	0
Port	InPkts 64	OutPkts 64	InPkts 65-127	OutPkts 65-127
Gil/1	0	0	0	0
Port	InPkts 128-255	OutPkts 128-255	InPkts 256-511	OutPkts 256-511
Gil/1	0	0	0	0
Port	InPkts 512-1023	OutPkts 512-1023		
Gil/1	0	0		
Port	InPkts 1024-1518	OutPkts 1024-1518	InPkts 1519-1548	OutPkts 1519-1548
Gil/1	0	0	0	0
Port	InPkts 1549-9216	OutPkts 1549-9216		
Gil/1	0	0		

<output truncated>

Caveats

Baby giants are supported, however they are counted as oversized giant packets. The error counter is incremented in the output of the **show interface** command on the following line cards:

- WS-X4504-FX-MT
- WS-X4232-RJ-XX
- WS-X4148-FX-MT
- WS-X4148-RJ21
- WS-X4148-RJ21
- WS-X4232-GB-RJ (Ports 3-34)
- WS-X4124-FXMT
- WS-X4148-RJ
- WS-X4148-RJ
- WS-X4148-RJV

Jumbo Frames

Software Support

Jumbo frame support has been available since Cisco IOS software release 12.1(13)EW for the Supervisor III and IV.

Hardware Support

Jumbo frames are supported only on non-blocking Gigabit ports. The following is a list of Gigabit modules and their specific ports supporting jumbo frames:

- both Supervisor uplink ports on Supervisor III (WS-X4013) and Supervisor IV (WS-X4014)
- WS-X4306-GB
- WS-X4232-GB-RJ (ports 1-2 only)
- WS-X4418-GB (ports 1-2 only)
- WS-X4412-2GB-TX (ports 13-14 only)

Configuration

To configure jumbo frame support, issue the **mtu <mtu-size>** interface configuration command, as shown below.

```
sup3#conf t
Enter configuration commands, one per line. End with CNTL/Z.
sup3(config)#interface gigabitEthernet 1/1
sup3(config-if)#mtu ?
<1500-9198> MTU size in bytes

sup3(config-if)#mtu 9198
sup3(config-if)#end
```

Jumbo frame support can be enabled on the following types of interfaces:

- port-channel interface
- Switched Virtual Interface (SVI)
- physical interface (L2/ Layer 3 (L3))

Verification

The **show interfaces <interface-id> mtu** command gives the configured interface level jumbo frame configuration, as shown below.

```
sup3#sh interfaces gigabitEthernet 1/1 mtu

Port      Name           MTU
Gi1/1    9198
```

The **show interface <interface-id>** command provides the configured MTU for the specific interface.

Note: Jumbo frame interface configuration will supersede the global MTU configuration. In the output below, the system MTU is configured for 1552, however interface Gigabit Ethernet 1/1 is configured for jumbo frame support of 9198 bytes.

```
sup3#show interfaces gigabitEthernet 1/1
GigabitEthernet1/1 is up, line protocol is down (notconnect)
  Hardware is Gigabit Ethernet Port, address is 0004.9a80.a400 (bia 0004.9a80.a400)
  MTU 9198 bytes, BW 1000000 Kbit, DLY 10 usec,
    reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation ARPA, loopback not set
  Keepalive set (10 sec)
  Auto-duplex, Auto-speed
  input flow-control is off, output flow-control is off
  ARP type: ARPA, ARP Timeout 04:00:00
  Last input never, output never, output hang never
```

```

Last clearing of "show interface" counters never
Input queue: 0/2000/0/0 (size/max/drops/flushes); Total output drops: 0
Queueing strategy: fifo
Output queue: 0/40 (size/max)
5 minute input rate 0 bits/sec, 0 packets/sec
5 minute output rate 0 bits/sec, 0 packets/sec
  0 packets input, 0 bytes, 0 no buffer
    Received 0 broadcasts (0 multicast)
      0 runts, 0 giants, 0 throttles
    0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored
    0 input packets with dribble condition detected
  0 packets output, 0 bytes, 0 underruns
    0 output errors, 0 collisions, 0 interface resets
    0 babbles, 0 late collision, 0 deferred
    0 lost carrier, 0 no carrier
    0 output buffer failures, 0 output buffers swapped out

```

The **show interfaces <interface-id> counters all** command provides statistics for jumbo frames, as shown below.

```

sup3# sh interfaces gigabitEthernet 1/1 counters all

Port          InBytes      InUcastPkts  InMcastPkts  InBcastPkts
Gi1/1          0             0             0             0

Port          OutBytes      OutUcastPkts  OutMcastPkts  OutBcastPkts
Gi1/1          0             0             0             0

Port          InPkts 64      OutPkts 64      InPkts 65-127  OutPkts 65-127
Gi1/1          0             0             0             0

Port          InPkts 128-255  OutPkts 128-255  InPkts 256-511  OutPkts 256-511
Gi1/1          0             0             0             0

Port          InPkts 512-1023  OutPkts 512-1023
Gi1/1          0             0

Port          InPkts 1024-1518  OutPkts 1024-1518  InPkts 1519-1548  OutPkts 1519-1548
Gi1/1          0             0             0             0

Port          InPkts 1549-9216  OutPkts 1549-9216
Gi1/1          0             0

```

<output truncated>

The **show system mtu** command shows the configured baby giant value, if configured. Jumbo frame support is per-interface, as shown below.

```

sup3# sh system mtu
Global Ethernet MTU is 1552 bytes.

```

Consideration when configuring Jumbo Frames on SVI

Make sure all interfaces in a VLAN are configured for jumbo frames before configuring jumbo frame support on an SVI. A packet's MTU is not checked on the ingress side of an SVI. However, it is checked on the egress side of an SVI. If the packet's MTU is larger than the egress SVI's MTU, the packet is fragmented by software (if the DF bit is not set), which results in poor performance. Software fragmentation only happens for L3 switching. When a packet is forwarded to an L3 port or an SVI with a smaller MTU, software fragmentation will occur.

In the output below, you can see that by issuing the **show vlan mtu** command, a mismatch has occurred for VLAN 1. The port Gig 4/1 in VLAN 1 can support only 1500 bytes, and therefore cannot fully support jumbo frames for that VLAN. Packets destined to these such ports, which do not support jumbo MTU, may get dropped for L2 switching. The packet continues to be forwarded if it is destined to Gig 1/1 or any non-blocking port in that VLAN.

It is recommended that the MTU of an SVI should always be smaller than the smallest MTU among all the switch ports in the VLAN. However, this is not enforced in software.

```
sup3# sh vlan mtu
```

VLAN	SVI_MTU	MinMTU(port)	MaxMTU(port)	MTU_Mismatch
1	9198 (TooBig)	1500 (Gi4/1)	9198 (Gi1/1)	Yes
2	1552	1552	1552	No
17	1552	1552	1552	No

Consideration of Configuring Jumbo Frames with Port-Channel

Jumbo frames can be enabled on interfaces configured for port-channel protocols. The following are some of the guidelines or restrictions:

- All the ports in a port-channel must have the same MTU.
- Changing the MTU of a port-channel interface changes the MTU of all member ports.
- If the MTU of a member port cannot be changed to the new value due to the member port being the blocking port, the port-channel is suspended .
- A port cannot join a port-channel if it has a different MTU than the others in the existing port-channel.
- If an individual member port's MTU is changed, the port is suspended.

Related Information

- [Configuring Jumbo/Giant Frame Support on Catalyst Switches](#)
- [Technical Support – Cisco Systems](#)

[Contacts & Feedback](#) | [Help](#) | [Site Map](#)

© 2013 – 2014 Cisco Systems, Inc. All rights reserved. [Terms & Conditions](#) | [Privacy Statement](#) | [Cookie Policy](#) | [Trademarks of Cisco Systems, Inc.](#)

Updated: Mar 24, 2005

Document ID: 29805
