

ASR 5000 Series ICMP Packet Generation from the CLI and Fragmentation Identification



Document ID: 119210

Contributed by Dave Damerjian, Cisco TAC Engineer.

Jul 27, 2015

Contents

Introduction

Problem

Solution

IP Fragmentation

Introduction

This document describes the byte counts reported by the show port commands and the ping command when pings are executed in the CLI on the Aggregation Services Router (ASR) 5000 Series platform. It also demonstrates the effects of fragmentation when the packets sent are greater than the configured maximum transmission unit (MTU) in the interface. This is good background information to have when you troubleshoot user-plane issues with ping or even in general for packets that pass through the chassis. Also, experimentation with pings on a node is a great way to confirm the concepts explained in this document.

Problem

When you specify the size of an Internet Control Message Protocol (ICMP) packet, the size refers to the raw payload that gets wrapped up into the packet. It does NOT include the ICMP header (8 bytes) OR the IP header (20 bytes). Also notable is the Ethernet header (14 bytes = destination MAC (6) + source MAC (6) + Ethernet frame type (2)), the VLAN tag (4 bytes), and the trailing Ethernet Frame Check Sequence (FCS, 4 bytes), the later of which is NOT to be displayed in a Wireshark trace.

Solution

When you view the output from the **show port [npu | datalink] counters** command, the math that can be applied is shown in this example. Match the colors in order to obtain a straightforward understanding. This exercise works on a port with little to no traffic as it allows a ping to be sent without any other traffic to cloud the resultant output.

Payload size: 56 bytes (which is also the default for this command)

Payload + ICMP header: 64

Payload + ICMP header + IP header: 84

Payload + ICMP header + IP header + Ethernet header + VLAN tag: 102

Payload (56) + ICMP header (8) + IP header (20) + Ethernet header (14) + VLAN tag (4) + FCS (4): 106

```

context Ctx
  interface 21/1 broadcast
    ip address 10.193.82.118 255.255.255.0
    ip mtu 1500
  #exit

```

```

port ethernet 21/1
  no shutdown
  vlan 30
  no shutdown
  bind interface 21/1 Ctx
  #exit
#exit

```

[Ctx]ASR5000> **show ip arp**

Flags codes:

I - Incomplete, R - Reachable, M - Permanent, S - Stale,
 D - Delay, P - Probe, F - Failed

Address	Link Type	Link Address	Flags	Mask	Interface
10.193.82.1	ether	00:00:0C:07:AC:1E	R		21/1

[Ctx]ASR5000> **show ip route**

"*" indicates the Best or Used route. S indicates Stale.

Destination	Nexthop	Protocol	Prec	Cost	Interface
*0.0.0.0/0	10.193.82.1	static	1	0	21/1

[Ctx]ASR5000> **clear port data counters 21/1**

Saturday April 12 14:06:21 UTC 2014

[Ctx]ASR5000> **clear port npu count 21/1**

Saturday April 12 14:06:21 UTC 2014

[Ctx]ASR5000> **ping 10.193.82.1 count 1 size 56**

Saturday April 12 14:06:21 UTC 2014

PING 10.193.82.1 (10.193.82.1) 56(84) bytes of data.

64 bytes from 10.193.82.1: icmp_seq=1 ttl=255 time=0.957 ms

--- 10.193.82.1 ping statistics ---

1 packets transmitted, 1 received, 0% packet loss, time 0ms

rtt min/avg/max/mdev = 0.957/0.957/0.957/0.000 ms

Note that the datalink counters include the FCS bytes, while both datalink and npu counters include the ethernet frame and vlan tags.

[Ctx]ASR5000> **show port datalink count 21/1**

Saturday April 12 14:06:25 UTC 2014

Counters for port 21/1:

Line Card Gigabit Ethernet Port

Rx Counter	Data	Tx Counter	Data
RX Unicast frames	112	TX Unicast frames	1
RX Multicast frames	3	TX Multicast frames	0
RX Broadcast frames	9	TX Broadcast frames	0
RX Size 64 frames	0	TX Size 64 frames	0
RX Size 65 .. 127 fr	14	TX Size 65 .. 127 fr	1
RX Size 128 .. 255 fr	0	TX Size 128 .. 255 fr	0
RX Size 256 .. 511 fr	110	TX Size 256 .. 511 fr	0
RX Size 512 .. 1023 fr	0	TX Size 512 .. 1023 fr	0
RX Size 1024 .. 1518 fr	0	TX Size 1024 .. 1518 fr	0
RX Size > 1518 frames	0	TX Size > 1518 frames	0
RX Bytes OK	43966	TX Bytes OK	106

```

RX Bytes BAD          0 | TX Bytes BAD          0
RX SHORT OK           0 | TX PAUSE                0
RX SHORT CRC          0 | TX ERR                  0
RX OVF                0 |
RX NORM CRC           0 |
RX LONG OK            0 |
RX LONG CRC           0 |
RX PAUSE              0 |
RX FALS CRS           0 |
RX SYM ERR            0 |
-----+-----

```

[Ctx]ASR5000> show port npu count 21/1

Saturday April 12 14:06:25 UTC 2014

Counters for port 21/1

Counter	Rx Frames	Rx Bytes	Tx Frames	Tx Bytes
Unicast	1	102	1	102
Multicast	3	202	0	0
Broadcast	8	512	0	0
IPv4 unicast	1	102	1	102
IPv4 non-unicast	1	66	0	0
IPv6 unicast	0	0	0	0
IPv6 non-unicast	0	0	0	0
Fragments received	0	0	n/a	n/a
Packets reassembled	0	0	n/a	n/a
Fragments to kernel	0	0	n/a	n/a
HW error	0	0	n/a	n/a
Port non-operational	0	0	0	0
SRC MAC is multicast	0	0	n/a	n/a
Unknown VLAN tag	0	0	n/a	n/a
Other protocols	2	136	n/a	n/a
Not IPv4	0	0	n/a	n/a
Bad IPv4 header	0	0	n/a	n/a
IPv4 MRU exceeded	0	0	n/a	n/a
TCP tiny fragment	0	0	0	0
No ACL match	0	0	0	0
Filtered by ACL	0	0	0	0
TTL expired	0	0	n/a	n/a
Flow lookup twice	0	0	n/a	n/a
Unknown IPv4 class	0	0	n/a	n/a
Too short: IP	0	0	n/a	n/a
Too short: ICMP	0	0	0	0
Too short: IGMP	0	0	0	0
Too short: TCP	0	0	0	0
Too short: UDP	0	0	0	0
Too short: IPIP	0	0	n/a	n/a
Too short: GRE	0	0	n/a	n/a
Too short: GRE key	0	0	n/a	n/a
Don't frag discards	n/a	n/a	0	0
Fragment packets	n/a	n/a	0	0
Fragment fragments	n/a	n/a	0	0
IPv4VlanMap dropped	0	0	n/a	n/a
IPSec NATT keep alive	0	0	n/a	n/a
MPLS Flow not found	0	0	n/a	n/a
MPLS unicast	0	0	0	0
Size < 17	0	0	0	0
Size 17 .. 64	8	512	0	0
Size 65 .. 127	4	304	1	102
Size 128 .. 255	0	0	0	0
Size 256 .. 511	0	0	0	0
Size 512 .. 1023	0	0	0	0
Size 1024 .. 2047	0	0	0	0
Size 2048 .. 4095	0	0	0	0
Size 4096 .. 4500	0	0	0	0
Size > 4500	0	0	0	0

IP Fragmentation

IP packet fragmentation occurs when the size of the packet, which includes the IP header, but not any Layer 2 data such as source/destination MAC, VLAN ID, or FSC, is broken into multiple segments in order to comply with the "ip mtu" setting for the interface (default 1500). In this example, with the MTU set to the default size of 1500, an ICMP payload of size 1472 attempted to be sent will not be fragmented (1472 + ICMP header (8) + IP header (20) = 1500), but anything larger than that will be fragmented. When fragmentation occurs, there are two packet fragments counted in the port npu and datalink counters, one in each appropriate-sized bucket. With the completion of a flood ping of 500 packets, 1000 fragments are sent; 500 for the large size and 500 for the small size. This example was run on an ASR 5500 (compared to the previous output example on an ASR 5000), but it works similarly on both platforms.

```
[ECS]ASR500> show ip interface
Thursday July 16 00:31:39 UTC 2015
Intf Name:      5/29-ECS
Intf Type:      Broadcast
Description:
VRF:            None
IP State:       UP (Bound to 5/29 vlan id 31, 802.1P prior 0, ifIndex 85786626)
IP Address:     10.213.137.105      Subnet Mask:    255.255.255.0
Bcast Address: 10.213.137.255      MTU:            1500
Resoln Type:    ARP              ARP timeout:    60 secs
L3 monitor LC-port switchover: Disabled
Number of Secondary Addresses: 0

port ethernet 5/29
no shutdown
vlan 31
no shutdown
bind interface 5/29-ECS ECS
#exit
#exit
```

First, establish a baseline for the counters. Specify the VLAN ID for the NPU counters (in order to confirm that all the pings have the VLAN ID appended).

```
[ECS]ASR5500> clear port datalink counters 5/29
Friday July 17 23:53:46 UTC 2015
```

```
[ECS]ASR5500> show port datalink counters 5/29
Friday July 17 23:53:46 UTC 2015
rCounters for port 5/29:
Line Card 10 Gigabit Ethernet Port
```

Rx Counter	Data	Tx Counter	Data
RX Bytes	406	TX Bytes	0
RX Unicast frames	0	TX Unicast frames	0
RX Multicast frames	4	TX Multicast frames	0
RX Broadcast frames	1	TX Broadcast frames	0
RX Size 64 frames	0	TX Size 64 frames	0
RX Size 65 .. 127 fr	5	TX Size 65 .. 127 fr	0
RX Size 128 .. 255 fr	0	TX Size 128 .. 255 fr	0
RX Size 256 .. 511 fr	0	TX Size 256 .. 511 fr	0
RX Size 512 .. 1023 fr	0	TX Size 512 .. 1023 fr	0
RX Size 1024 .. 1518 fr	0	TX Size 1024 .. 1518 fr	0
RX Size 1519 .. 1522 fr	0	TX Size 1519 .. 1522 fr	0

```
[ECS]ASR5500> clear port npu counters 5/29 vlan 31
Friday July 17 23:53:47 UTC 2015
```

```
[ECS]ASR5500> show port npu counters 5/29 vlan 31
```

Friday July 17 23:53:47 UTC 2015
 Counters for port 5/29 vlan id 31:

Counter	Rx Frames	Rx Bytes	Tx Frames	Tx Bytes
Unicast	0	0	0	0
Multicast	3	196	0	0
Broadcast	1	64	0	0
IPv4 unicast	0	0	0	0
IPv4 non-unicast	3	196	0	0
IPv6 unicast	0	0	0	0
IPv6 non-unicast	0	0	0	0
Fragments received	0	0	n/a	n/a
Packets reassembled	0	0	n/a	n/a
Fragments to kernel	0	0	n/a	n/a
HW error	0	0	n/a	n/a
Port non-operational	539	38520	0	0
SRC MAC is multicast	0	0	n/a	n/a
Unknown VLAN tag	0	0	n/a	n/a
Other protocols	2	128	n/a	n/a

Pings of size **1500** bytes which include the IP header are sent out unfragmented. The total size at the datalink layer (as it leaves the port) is:

$$1472 \text{ (payload)} + 8 \text{ (ICMP header)} + 20 \text{ (IP header)} + 14 \text{ (MAC source/dest)} + 4 \text{ (VLAN ID)} + 4 \text{ (FSC)} = 1522$$

```
[ECS]ASR5500> ping 10.213.137.1 size 1472 flood
Friday July 17 23:53:48 UTC 2015
PING 10.213.137.1 (10.213.137.1) 1472(1500) bytes of data.
```

```
--- 10.213.137.1 ping statistics ---
500 packets transmitted, 500 received, 0% packet loss, time 422ms
rtt min/avg/max/mdev = 0.405/0.800/0.994/0.143 ms, ipg/ewma 0.845/0.802 ms
```

```
[ECS]ASR5500> show port datalink counters 5/29
```

```
Friday July 17 23:53:54 UTC 2015
Counters for port 5/29:
```

```
Line Card 10 Gigabit Ethernet Port
```

Rx Counter	Data	Tx Counter	Data
RX Bytes	771008	TX Bytes	765656
RX Unicast frames	536	TX Unicast frames	524
RX Multicast frames	45	TX Multicast frames	0
RX Broadcast frames	20	TX Broadcast frames	0
RX Size 64 frames	0	TX Size 64 frames	0
RX Size 65 .. 127 fr	75	TX Size 65 .. 127 fr	0
RX Size 128 .. 255 fr	24	TX Size 128 .. 255 fr	18
RX Size 256 .. 511 fr	2	TX Size 256 .. 511 fr	6
RX Size 512 .. 1023 fr	0	TX Size 512 .. 1023 fr	0
RX Size 1024 .. 1518 fr	0	TX Size 1024 .. 1518 fr	0
RX Size 1519 .. 1522 fr	500	TX Size 1519 .. 1522 fr	500
RX OverSize frames	0	TX OverSize frames	0
RX UnderSize frames	0	TX UnderSize frames	0
RX ExceedMaxSize frames	0		
RX Fragment frames	0	TX Fragment frames	0
RX Jabber frames	0	TX Jabber frames	0
RX Control frames	0	TX Control frames	0
RX Pause frames	0	TX Pause frames	0
RX FCS Error frames	0	TX FCS Error frames	0
RX Length Error frames	0	TX Length Error frames	0
RX Code Error frames	0		
RX ExMaxSize Err frames	0		


```

RX Size 1024 .. 1518 fr          0 | TX Size 1024 .. 1518 fr          0
RX Size 1519 .. 1522 fr        500 | TX Size 1519 .. 1522 fr      1000
RX ExceedMaxSize frames         0
RX Fragment frames              0 | TX Fragment frames              0
RX Jabber frames                0 | TX Jabber frames                0
RX Control frames               0 | TX Control frames               0
RX Pause frames                 0 | TX Pause frames                 0
RX FCS Error frames             0 | TX FCS Error frames             0
RX Length Error frames          0 | TX Length Error frames          0
RX Code Error frames            0
RX ExMaxSize Err frames         0
-----+-----

```

[ECS]ASR5500> **show port npu counters 5/29 vlan 31**

Friday July 17 23:54:06 UTC 2015

Counters for port 5/29 vlan id 31:

Counter	Rx Frames	Rx Bytes	Tx Frames	Tx Bytes
Unicast	554	766984	1562	1549040
Multicast	94	6962	0	0
Broadcast	53	4294	0	0
IPv4 unicast	607	771278	1562	1549040
IPv4 non-unicast	73	4904	0	0
IPv6 unicast	0	0	0	0
IPv6 non-unicast	21	2058	0	0
Fragments received	0	0	n/a	n/a
Packets reassembled	0	0	n/a	n/a
Fragments to kernel	0	0	n/a	n/a
HW error	0	0	n/a	n/a
Port non-operational	25146	1805666	0	0
SRC MAC is multicast	0	0	n/a	n/a
Unknown VLAN tag	0	0	n/a	n/a
Other protocols	68	4742	n/a	n/a
Not IPv4	21	2058	n/a	n/a
...				
Size 0 .. 63	0	0	501	19546
Size 64 .. 127	188	14154	2	140
Size 128 .. 255	60	8520	45	6450
Size 256 .. 511	0	0	15	4950
Size 512 .. 1023	0	0	0	0
Size 1024 .. 2047	500	759000	1000	1518000
Size 2048 .. 4095	0	0	0	0
Size 4096 .. 8191	0	0	0	0
Size >= 8192	0	0	0	0