

Contents

[Introduction](#)

[Background Information](#)

[How is a Packet Processed by a Switch](#)

[Padding Modified with Tagged VLANs when Traffic traverses N9K](#)

[Solution](#)

Introduction

This document describes how to troubleshoot the corrupted ethernet packet on Cisco Nexus 9000 when a padding information is corrupted or malformed.

Contributed by Sivakumar Sukumar, Cisco TAC Engineer.

Background Information

The minimal size of an Ethernet frame is 64 bytes, no matter VLAN tag is

The minimal Ethernet payload size is:

- 46 bytes if the VLAN tag is absent.
- 42 bytes if the VLAN tag is present.

You can verify this fact:

- On Wikipedia, section **Payload**: https://en.wikipedia.org/wiki/Ethernet_frame
- On the IEEE 802.3 standard (http://people.ee.duke.edu/~mbrooke/EE164.02/Spring_2004/group_2/index_files/8023.pdf), w 39, and the elements of a tagged MAC frame is defined on page 43 section 3.5.

The minimal size of an ethernet packet is 64 bytes, no matter VLAN header is present there or not. The server is allowed to send a 64 bytes long packet that contains a VLAN, which you should accept and process correctly.

Note: This behaviour is correctly handled by a catalyst 4500x. Not by a Nexus 9k.

How is a Packet Processed by a Switch

Step 1. Receive a **VALID** 64 bytes Ethernet frame.

Step 2. Remove the Frame Check Sequence (FCS), so the packet becomes 60 bytes long.

Step 3. Remove the VLAN tag, so the packet becomes 56 bytes long.

Step 4. Add padding to make the packet 60 bytes long.

Step 5. It adds the FCS, making the packet 64 bytes long.

Padding should not get modified when a packet goes through cut-through switch.

Padding Modified with Tagged VLANs when Traffic traverses N9K

Instead of padding with zeroes, the packet is padded with garbish characters, in most of the cases it has no impact because checksums are not modified and so nobody uses these datas. However, if customers have a special usage and need to recompute checksums, these garbish datas leads to corruption of checksums in the end (I believe other appliances, like NAT/load-balancers might see the issue too)

Device is a N9K 93120TX (was initially detected on a 9372TX though), version is latest NXOS 7.0(3)I2(2a)

Use Linux hosts with directly connected hardware to the N9K (no virtualization of any kind) here (1000base-T links)

Use this configuration :

You can use netcat to send and receive packets.

As shown in the image, it sends Side (VLAN 100 tagged), port e1/59 on the switch

```
6: eth1.100@eth1: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 9000 qdisc noqueue state UP group default
link/ether 44:a8:42:2c:5f:c4 brd ff:ff:ff:ff:ff:ff
inet 10.1.1.1/24 brd 10.1.1.255 scope global eth1.100
    valid_lft forever preferred_lft forever
inet6 fe80::46a8:42ff:fe2c:5fc4/64 scope link
    valid_lft forever preferred_lft forever
```

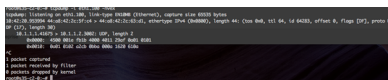
```
root@s35-c2-0:~# nc 10.1.1.2 3002 -u
a
^C
root@s35-c2-0:~# █
```

It receives Side (VLAN 100 tagged), port e1/60 on the switch, as shown in the image:


```
7: eth1.100@eth1: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 9000 qdisc noqueue state UP group default
link/ether 44:a8:42:2c:63:d1 brd ff:ff:ff:ff:ff:ff
inet 10.1.1.2/24 brd 10.1.1.255 scope global eth1.100
    valid_lft forever preferred_lft forever
inet6 fe80::46a8:42ff:fe2c:63d1/64 scope link
    valid_lft forever preferred_lft forever
```

```
root@s35-c2:~# nc -l -u -p 3002
a
^C
root@s35-c2:~# █
```

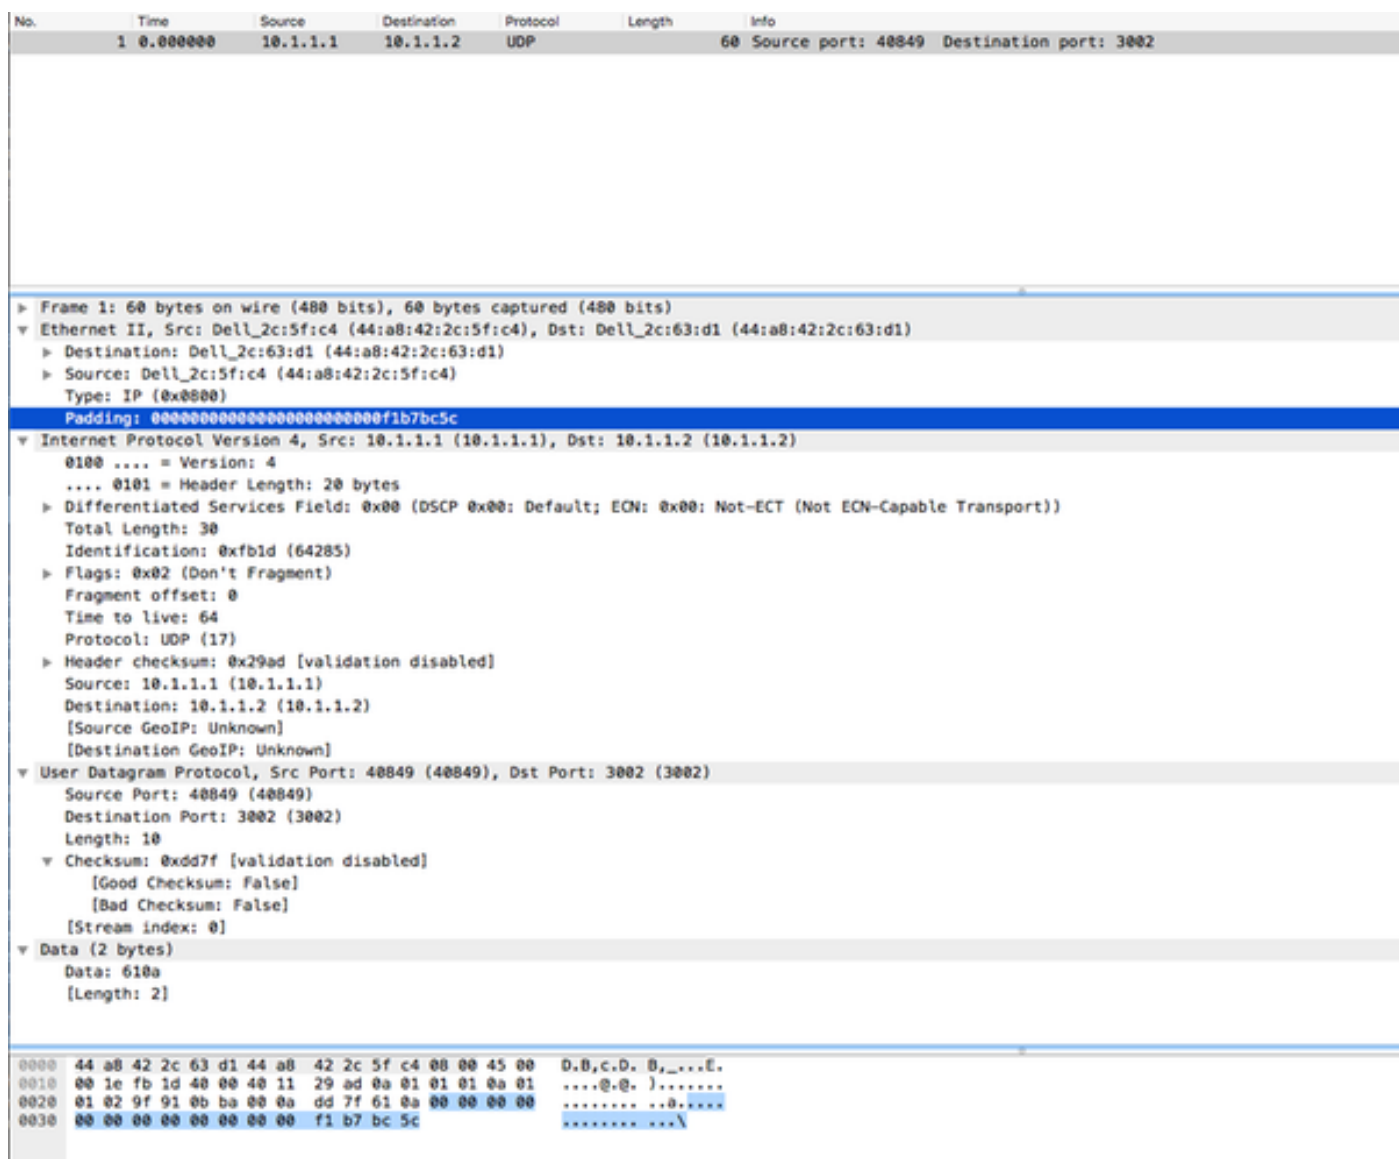
As shown in the image, the packet is transmitted.



The packet is received, as shown in the image:



As shown in the image, the wrong padding is highlighted.



This is also displayed with a packet analyzer (another packet, data is different than previous screenshots but test and bug is identical),

Solution

The work around is to [buffer-boost](#) on interface where we have this server connected.

C9396PX-1(config)# int et 1/7

C9396PX-1(config-if)# no buffer-boost