



Mellanox Innova™ IPsec 4 Lx Ethernet Adapter Quick Start Guide

Rev: 1.2



NOTE:

THIS HARDWARE, SOFTWARE OR TEST SUITE PRODUCT (“PRODUCT(S)”) AND ITS RELATED DOCUMENTATION ARE PROVIDED BY MELLANOX TECHNOLOGIES “ASIS” WITH ALL FAULTS OF ANY KIND AND SOLELY FOR THE PURPOSE OF AIDING THE CUSTOMER IN TESTING APPLICATIONS THAT USE THE PRODUCTS IN DESIGNATED SOLUTIONS. THE CUSTOMER’S MANUFACTURING TEST ENVIRONMENT HAS NOT MET THE STANDARDS SET BY MELLANOX TECHNOLOGIES TO FULLY QUALIFY THE PRODUCT(S) AND/OR THE SYSTEM USING IT. THEREFORE, MELLANOX TECHNOLOGIES CANNOT AND DOES NOT GUARANTEE OR WARRANT THAT THE PRODUCTS WILL OPERATE WITH THE HIGHEST QUALITY. ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT ARE DISCLAIMED. IN NO EVENT SHALL MELLANOX BE LIABLE TO CUSTOMER OR ANY THIRD PARTIES FOR ANY DIRECT, INDIRECT, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES OF ANY KIND (INCLUDING, BUT NOT LIMITED TO, PAYMENT FOR PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY FROM THE USE OF THE PRODUCT(S) AND RELATED DOCUMENTATION EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.



Mellanox Technologies
350 Oakmead Parkway Suite 100
Sunnyvale, CA 94085
U.S.A.
www.mellanox.com
Tel: (408) 970-3400
Fax: (408) 970-3403

© Copyright 2017. Mellanox Technologies Ltd. All Rights Reserved.

Mellanox®, Mellanox logo, Accelio®, BridgeX®, CloudX logo, CompustorX®, Connect-IB®, ConnectX®, CoolBox®, CORE-Direct®, EZchip®, EZchip logo, EZappliance®, EZdesign®, EZdriver®, EZsystem®, GPUDirect®, InfiniHost®, InfiniBridge®, InfiniScale®, Kotura®, Kotura logo, Mellanox CloudRack®, Mellanox CloudXMellanox®, Mellanox Federal Systems®, Mellanox HostDirect®, Mellanox Multi-Host®, Mellanox Open Ethernet®, Mellanox OpenCloud®, Mellanox OpenCloud Logo®, Mellanox PeerDirect®, Mellanox ScalableHPC®, Mellanox StorageX®, Mellanox TuneX®, Mellanox Connect Accelerate Outperform logo, Mellanox Virtual Modular Switch®, MetroDX®, MetroX®, MLNX-OS®, NP-1c®, NP-2®, NP-3®, Open Ethernet logo, PhyX®, PlatformX®, PSIPHY®, SiPhy®, StoreX®, SwitchX®, Tlera®, Tlera logo, TestX®, TuneX®, The Generation of Open Ethernet logo, UFM®, Unbreakable Link®, Virtual Protocol Interconnect®, Voltaire® and Voltaire logo are registered trademarks of Mellanox Technologies, Ltd.

All other trademarks are property of their respective owners.

For the most updated list of Mellanox trademarks, visit <http://www.mellanox.com/page/trademarks>

Table of Contents

Table of Contents	3
Revision History	4
Chapter 1 Overview	5
1.1 Prerequisites	5
1.1.1 Tunnel Mode	5
1.1.2 Transport Mode	5
Chapter 2 Fast Installation and Update	7
2.1 Hardware Installation	7
2.2 Content of Innova IPsec Package	7
2.3 Software, Firmware and Tools Update	7
Chapter 3 Test Set-up Architecture	9
3.1 Tunnel Mode Set-up	9
3.2 Transport Mode Set-up	10
Chapter 4 Performance Test Procedure	11
4.1 Tunnel Mode	11
4.1.1 Loading Kernel and Driver	11
4.1.2 Configuring IP Addresses	11
4.1.3 IP Routing Topology	12
4.1.4 Setting up IPsec Tunnel	13
4.1.5 Running Throughput Test	14
4.2 Transport Mode	14
4.2.1 Loading Kernel and Driver	14
4.2.2 Configuring IP Addresses	15
4.2.3 Setting up IPsec SA	15
4.2.4 Running Throughput Test	16

Revision History

This document was printed on January 31, 2017.

Table 1 - Revision History Table

Date	Rev	Comments/Changes
January 2017	1.2	<ul style="list-style-type: none"> • Updated Section 1.1, “Prerequisites,” on page 5 • Added Chapter 2, “Fast Installation and Update” on page 7 • Updated Section 4.1.1, “Loading Kernel and Driver,” on page 11 • Updated Section 4.2.1, “Loading Kernel and Driver,” on page 14 • Updated Section 4.2.3, “Setting up IPsec SA,” on page 15
January 2017	1.1	<ul style="list-style-type: none"> • Added Transport Mode as supported IPsec operation mode • Updated Section 1.1, “Prerequisites,” on page 5 • Updated Section 4.1.1, “Loading Kernel and Driver,” on page 11
November 2016	1.0	First Release

1 Overview

This guide describes the network topology, server configuration and process of an IPsec performance test set-up intended to test the Innova IPsec performance, in order to achieve the expected performance results.

The document assumes the reader has read the *Mellanox Innova™ IPsec 4 Lx Ethernet Adapter Card User Manual* and is familiar with the Innova IPsec adapter high level architecture and configuration.

1.1 Prerequisites

1.1.1 Tunnel Mode

Table 2 lists the hardware, software and firmware that the set up requires:

Table 2 - Set-up Requirements

Requirements	Description
Hardware	<ul style="list-style-type: none"> Four servers Four ConnectX-4 adapter cards Two Innova IPsec adapter cards Three copper cables
FPGA Image Version	<ul style="list-style-type: none"> 1161 - 10GbE 1160 - 40GbE
Firmware Version	14.18.0354 - 10GbE/40GbE
MFT Version	mft-4.6.0-48 - 10GbE/40GbE
Driver/Kernel Components	<ul style="list-style-type: none"> kernel-4.9.0_rc6_37cc319_2017_01_22_16_06_10-1.x86_64.rpm kernel-devel-4.9.0_rc6_37cc319_2017_01_22_16_06_10-1.x86_64.rpm iproute2-4.9.0rc6_ipsec_offload-1.x86_64.
Platforms/Operating Systems	<ul style="list-style-type: none"> RHEL7.1 RHEL7.2

1.1.2 Transport Mode

Table 3 lists the hardware, software and firmware that the set up requires:

Table 3 - Set-up Requirements

Requirements	Description
Hardware	<ul style="list-style-type: none"> Two servers Two Innova IPsec adapter cards One copper cable
FPGA Image Version	<ul style="list-style-type: none"> 1161 - 10GbE 1160 - 40GbE

Table 3 - Set-up Requirements

Requirements	Description
Firmware Version	14.18.0354 - 10GbE/40GbE
MFT Version	mft-4.6.0-48 - 10GbE/40GbE
Driver/Kernel Components	<ul style="list-style-type: none">• kernel-4.9.0_rc6_37cc319_2017_01_22_16_06_10-1.x86_64.rpm• kernel-devel-4.9.0_rc6_37cc319_2017_01_22_16_06_10-1.x86_64.rpm• iproute2-4.9.0rc6_ipsec_offload-1.x86_64.
Platforms/Operating Systems	<ul style="list-style-type: none">• RHEL7.1• RHEL7.2

2 Fast Installation and Update

This section describes the required steps only for a fast installation of Innova IPsec adapter card. For further details, please refer to the *Mellanox Innova™ IPsec 4 Lx Ethernet Adapter Card User Manual*.

2.1 Hardware Installation

1. Shut down your system if active.
2. After shutting down the system, turn off power and unplug the cord.
3. Place the adapter in a standard PCI Express slot.

2.2 Content of Innova IPsec Package

Mellanox provides an Innova IPsec package which includes the following:

- FPGA image bin file
- Kernel RPM files
- Offload scripts (xfrm, iproute)
- MFT tar file
- Firmware bin files

2.3 Software, Firmware and Tools Update



Please make sure to follow the below steps in the same order. Updating the FPGA image must be performed first.

Step 1. To download the package, please refer to www.mellanox.com => Products => Programmable Adapter Cards => Innova IPsec => FW & SW.

➤ **To update the FPGA image:**

Step 1. Locate the FPGA image bin file in the Images folder.

Step 2. Burn the FPGA image:

```
mlx_fpga -d /dev/mst/mt4117_pciconf0_fpga_rdma burn <fpga_image.bin>
```

➤ **To install the most updated kernel:**

Step 1. Locate the RPM files in the Kernel folder:

- rpm -i kernel-<kernel_version>.rpm
- rpm -i kernel-devel-<kernel_version>.rpm

Step 2. Reboot your system and select the relevant kernel to load from the grub menu.

➤ **To update MFT:**

Step 1. Untar the MFT tar file.

Step 2. Install MFT by running:

```
install.sh
```

Step 3. Start MFT:

Step a. `modprobe mlx_accel_tools`

Step b. `mst start -with_fpga`

Step c. `mst statusstatus`

```
apps-13:~ # mlx_accel_tools
apps-13:~ # mst start --with_fpga
apps-13:~ # mst status
MST modules:
-----
MST PCI module is not loaded
MST PCI configuration module is not loaded
MST devices:
-----
No MST devices were found nor MST modules were loaded.
You may need to run 'mst start' to load MST modules.
FPGA devices:
-----
/dev/mst/mt4117_pciconf0_fpga_i2c
/dev/mst/mt4117_pciconf1_fpga_rdmaa
```

- a. It is recommended to use the RDMA device as it uses the fast path to the FPGA. I2C is used for recovery purposes when RDMA is not functional.

➤ ***To burn the latest Firmware:***

Step 1. Locate the firmware bin file in FW folder:

Step 2. Burn the firmware:

```
mlxburn -d /dev/mst/mt4117_pciconf0 -i <fw.bin>
```

Step 3. To load the firmware:

```
mlxfwreset -d /dev/mst/mt4117_pciconf0 reset -y
```

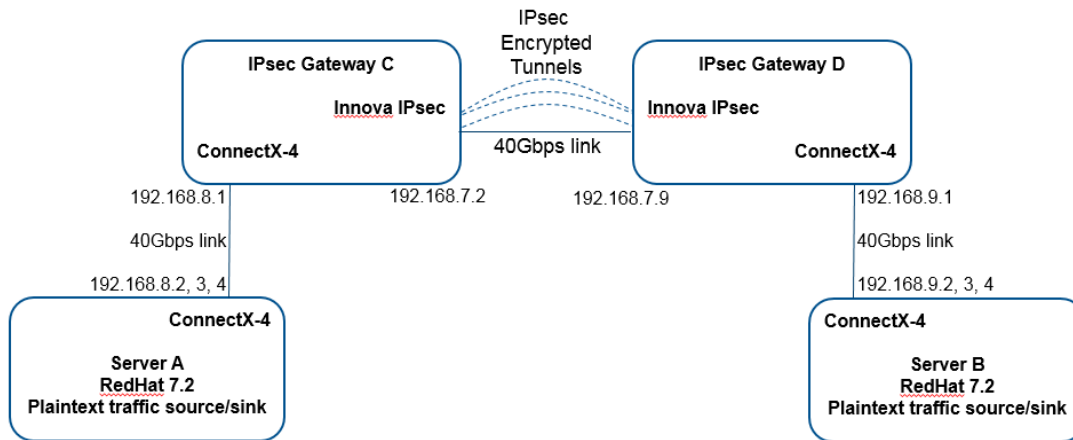

3 Test Set-up Architecture

3.1 Tunnel Mode Set-up

Two of the four servers - Server A and Server B - are used for generating and receiving packets. A ConnectX-4 adapter card supporting 10/40GbE should be installed on both server A and B. The other two servers - Gateway C and Gateway D - are used for IP forwarding (encryption and decryption). Each gateway server should have one ConnectX-4 and one Innova IPsec adapter card installed. The IPsec configuration is set to the following:

- Encryption algorithm: AES-GCM 128/256-bit key, and 128-bit ICV
- IPsec operation mode: Tunnel mode
- IPsec protocol: ESP
- IP version: IPv4

For test procedure details, see [Chapter 4, “Performance Test Procedure”](#) on page 11.




3.2 Transport Mode Set-up

Server A and Server B are used for generating and receiving packets. An Innova IPsec adapter card should be installed in both Server A and Server B. The IPsec configuration is set to the following:

- Encryption algorithm: AES-GCM 128/256-bit key, and 128-bit ICV
- IPsec operation mode: Transport mode
- IPsec protocol: ESP
- IP version: IPv4

For test procedure details, see [Chapter 4, “Performance Test Procedure”](#) on page 11.



4 Performance Test Procedure

4.1 Tunnel Mode

4.1.1 Loading Kernel and Driver

The following steps should be performed on both gateway servers C and D.

Step 4. Make sure the correct kernel version is running:

```
[root@gateway-c ~]# uname -r
kernel-4.9.0_rc6_3602b52_2016_12_06_16_08_22
```

The running kernel should be the version you installed from the Innova IPsec software package. See [Section 1.1, “Prerequisites,” on page 5](#).

Step 5. Load the Innova IPsec driver:

```
[root@gateway-c ~]# modprobe mlx_ipsec
```

Step 6. Make sure the drivers are loaded:

```
[root@gateway-c ~]# lsmod | grep mlx
mlx_ipsec          32768  0
mlx_accel_core    32768  1 mlx_ipsec
mlx5_ib           159744  1 mlx_accel_core
mlx5_core         319488  2 mlx5_ib,mlx_accel_core
```

The above four drivers should be running.

Step 7. Verify no error messages in kernel log:

```
[root@gateway-c ~]# dmesg | tail
[ 5983.745124] mlx_accel_hw_dev_add_one called for 0000:00:08.0
[ 5983.745627] infiniband mlx5_0: mlx_accel_ib_dev_add_one called
[ 5983.746909] mlx5_core 0000:00:08.0: FPGA image is ready
[ 5983.864881] mlx_accel_core_client_register called for mlx_ipsec
[ 5983.866169] mlx_accel_core_conn_create called for mlx5_0-0000:00:08.0
[ 5983.910555] mlx_accel_core_connect called for mlx5_0-0000:00:08.0
[ 5983.959818] net ens8: mlx_ipsec added on device mlx5_0-0000:00:08.0
```

4.1.2 Configuring IP Addresses

Step 1. Configure IP addresses for the six interfaces on the four servers:

Note: Replace interface names in the commands below as appropriate. See [Chapter 3, “Test Set-up Architecture”](#) on page 9.

```
[root@server-a ~]# ifconfig ens1 192.168.8.2/24 up
[root@server-a ~]# ip addr add 192.168.8.3/24 dev ens1
[root@server-a ~]# ip addr add 192.168.8.4/24 dev ens1
[root@gateway-c ~]# ifconfig ens1 192.168.8.1/24 up
[root@gateway-c ~]# ifconfig ens2 192.168.7.2/24 up
[root@gateway-d ~]# ifconfig ens2 192.168.7.9/24 up
[root@gateway-d ~]# ifconfig ens1 192.168.9.1/24 up
[root@server-b ~]# ifconfig ens1 192.168.9.2/24 up
[root@server-b ~]# ip addr add 192.168.9.3/24 dev ens1
[root@server-b ~]# ip addr add 192.168.9.4/24 dev ens1
```

Step 2. Verify connectivity and link speed for all six interfaces on all four servers. You may use `ethtool` to query the link state and/or ping to test connectivity. Example:

```
[root@server-a ~]# ping 192.168.8.1
```

4.1.3 IP Routing Topology

Step 1. Turn on IP forwarding on both gateway servers:

```
[root@gateway-c ~]# echo 1 > /proc/sys/net/ipv4/ip_forward
[root@gateway-d ~]# echo 1 > /proc/sys/net/ipv4/ip_forward
```

Step 2. Add routes between all three subnets:

```
[root@server-a ~]# ip route add 192.168.7.0/24 via 192.168.8.1
[root@server-a ~]# ip route add 192.168.9.0/24 via 192.168.8.1
[root@gateway-c ~]# ip route add 192.168.9.0/24 via 192.168.7.9
[root@gateway-d ~]# ip route add 192.168.8.0/24 via 192.168.7.2
[root@server-b ~]# ip route add 192.168.7.0/24 via 192.168.9.1
[root@server-b ~]# ip route add 192.168.8.0/24 via 192.168.9.1
```

Step 3. Verify forwarding connectivity between the two servers. At this point, traffic is neither encrypted nor encapsulated:

```
[root@server-a ~]# ping 192.168.9.2
[root@server-b ~]# ping 192.168.8.2
```

Step 4. Check the IPsec counters before processing offloaded traffic on both gateways:

```
[root@gateway-c ~]# ethtool -S ens2 | grep ipsec
ipsec_dec_in_packets: 0
ipsec_dec_out_packets: 0
ipsec_dec_bypass_packets: 6
ipsec_enc_in_packets: 0
ipsec_enc_out_packets: 0
ipsec_enc_bypass_packets: 6
ipsec_dec_drop_packets: 0
ipsec_dec_auth_fail_packets: 0
ipsec_enc_drop_packets: 0
ipsec_add_sa_success: 0
ipsec_add_sa_fail: 0
ipsec_del_sa_success: 0
ipsec_del_sa_fail: 0
ipsec_cmd_drop: 0
[root@gateway-d ~]# ethtool -S ens2 | grep ipsec
...
```

4.1.4 Setting up IPsec Tunnel

Step 1. Set up the IPsec tunnels between the two gateway servers using the script provided with the Innova IPsec software package. Replace gateway-d with the management DNS name or IP address of Gateway D server:

```
[root@gateway-c ~]# ./xfrm-offload-tunnel.sh -both -256 192.168.7.2 ens2 192.168.7.9
ens2 gateway-d 192.168.8.2 192.168.9.2
[root@gateway-c ~]# ./xfrm-offload-tunnel.sh -both -a -256 192.168.7.2 ens2 192.168.7.9
ens2 gateway-d 192.168.8.3 192.168.9.3
[root@gateway-c ~]# ./xfrm-offload-tunnel.sh -both -a -256 192.168.7.2 ens2 192.168.7.9
ens2 gateway-d 192.168.8.4 192.168.9.4
```

Step 2. Verify forwarding connectivity between the two servers. At this point, traffic should be encapsulated and encrypted:

```
[root@server-a ~]# ping 192.168.9.2
[root@server-b ~]# ping 192.168.8.2
```

Step 3. Check the IPsec counters after processing the offloaded traffic on both gateways:

```
[root@gateway-c ~]# ethtool -S ens2 | grep ipsec
ipsec_dec_in_packets: 8
ipsec_dec_out_packets: 8
ipsec_dec_bypass_packets: 6
ipsec_enc_in_packets: 8
ipsec_enc_out_packets: 8
ipsec_enc_bypass_packets: 6
ipsec_dec_drop_packets: 0
ipsec_dec_auth_fail_packets: 0
ipsec_enc_drop_packets: 0
ipsec_add_sa_success: 0
ipsec_add_sa_fail: 0
ipsec_del_sa_success: 0
ipsec_del_sa_fail: 0
ipsec_cmd_drop: 0
[root@gateway-d ~]# ethtool -S ens2 | grep ipsec
```

4.1.5 Running Throughput Test

Step 1. Run iperf3 servers on server B machine:

```
[root@server-b ~]# iperf3 -s -B 192.168.9.2 &
[root@server-b ~]# iperf3 -s -B 192.168.9.3 &
[root@server-b ~]# iperf3 -s -B 192.168.9.4 &
```

Step 2. Run iperf3 clients on server A machine:

```
[root@server-a ~]# iperf3 -c 192.168.9.2 -B 192.168.8.2 -t 1000 -i 1 &
[root@server-a ~]# iperf3 -c 192.168.9.3 -B 192.168.8.3 -t 1000 -i 1 &
[root@server-a ~]# iperf3 -c 192.168.9.4 -B 192.168.8.4 -t 1000 -i 1 &
```

Note the throughput achieved.

Step 3. While traffic is flowing, check the total traffic on both gateways:

```
[root@gateway-c ~]# nload ens2
[root@gateway-d ~]# nload ens2
```

Step 4. Check CPU usage on both gateways:

```
[root@gateway-c ~]# top
[root@gateway-d ~]# top
```

Press 'l' inside 'top' user interface to see individual CPU usage. Note the usage pattern on each gateway server is different.

4.2 Transport Mode

4.2.1 Loading Kernel and Driver

The following steps should be performed on both servers A and B.

Step 1. Make sure the correct kernel version is running:

```
[root@server-a ~]# uname -r
```

```
kernel-4.9.0_rc6_3602b52_2016_12_06_16_08_22
```

The running kernel should be the version you installed from the Innova IPsec software package. See [Section 1.1, “Prerequisites,” on page 5](#).

Step 2. Load the Innova IPsec driver:

```
[root@server-a ~]# modprobe mlx_ipsec
```

Step 3. Make sure the drivers are loaded:

```
[root@server-a ~]# lsmod | grep mlx
mlx_ipsec                32768  0
mlx_accel_core           32768  1 mlx_ipsec
mlx5_ib                  159744  1 mlx_accel_core
mlx5_core                 319488  2 mlx5_ib,mlx_accel_core
```

The above four drivers should be running.

Step 4. Verify no error messages in kernel log:

```
[root@server-a ~]# dmesg | tail
[ 5983.745124] mlx_accel_hw_dev_add_one called for 0000:00:08.0
[ 5983.745627] infiniband mlx5_0: mlx_accel_ib_dev_add_one called
[ 5983.746909] mlx5_core 0000:00:08.0: FPGA image is ready
[ 5983.864881] mlx_accel_core_client_register called for mlx_ipsec
[ 5983.866169] mlx_accel_core_conn_create called for mlx5_0-0000:00:08.0
[ 5983.910555] mlx_accel_core_connect called for mlx5_0-0000:00:08.0
[ 5983.959818] net ens8: mlx_ipsec added on device mlx5_0-0000:00:08.0
```

4.2.2 Configuring IP Addresses

Step 1. Configure IP addresses for the two interfaces on the two servers:

Note: Replace interface names in the commands below as appropriate. See [Chapter 3, “Test Set-up Architecture” on page 9](#).

```
[root@server-a ~]# ifconfig ens1 192.168.8.12/24 up
[root@server-a ~]# ip addr add 192.168.8.13/24 dev ens1
[root@server-a ~]# ip addr add 192.168.8.14/24 dev ens1
[root@server-b ~]# ifconfig ens1 192.168.8.22/24 up
[root@server-b ~]# ip addr add 192.168.8.23/24 dev ens1
[root@server-b ~]# ip addr add 192.168.8.24/24 dev ens1
```

Step 2. Verify connectivity and link speed for both interfaces. You may use `ethtool` to query the link state and/or ping to test connectivity. Example:

```
[root@server-a ~]# ping 192.168.8.22
```

4.2.3 Setting up IPsec SA

Step 1. Set up the IPsec tunnels between the two servers using the script provided with the Innova IPsec software package.

```
[root@server-a ~]# ./xfrm-offload-transport.sh -both 192.168.8.12 ens1 192.168.8.22 ens1 server-b
[root@server-a ~]# ./xfrm-offload-transport.sh -both -a 192.168.8.13 ens1 192.168.8.23 ens1 server-b
[root@server-a ~]# ./xfrm-offload-transport.sh -both -a 192.168.8.14 ens1 192.168.8.24 ens1 server-b
```

Step 2. Verify connectivity between the two servers. At this point, traffic should be encrypted:

```
[root@server-a ~]# ping 192.168.8.22  
[root@server-b ~]# ping 192.168.8.12
```

Step 3. Check the IPsec counters after processing the offloaded traffic on both servers:

```
[root@server-a ~]# ethtool -S ens1 | grep ipsec  
ipsec_dec_in_packets: 8  
ipsec_dec_out_packets: 8  
ipsec_dec_bypass_packets: 6  
ipsec_enc_in_packets: 8  
ipsec_enc_out_packets: 8  
ipsec_enc_bypass_packets: 6  
ipsec_dec_drop_packets: 0  
ipsec_dec_auth_fail_packets: 0  
ipsec_enc_drop_packets: 0  
ipsec_add_sa_success: 0  
ipsec_add_sa_fail: 0  
ipsec_del_sa_success: 0  
ipsec_del_sa_fail: 0  
ipsec_cmd_drop: 0  
[root@server-b ~]# ethtool -S ens1 | grep ipsec
```

4.2.4 Running Throughput Test

Step 1. Run iperf3 servers on server B machine:

```
[root@server-b ~]# iperf3 -s -B 192.168.8.22 &  
[root@server-b ~]# iperf3 -s -B 192.168.8.23 &  
[root@server-b ~]# iperf3 -s -B 192.168.8.24 &
```

Step 2. Run iperf3 clients on server A machine:

```
[root@server-a ~]# iperf3 -c 192.168.8.22 -B 192.168.8.12 -t 1000 -i 1 &  
[root@server-a ~]# iperf3 -c 192.168.8.23 -B 192.168.8.13 -t 1000 -i 1 &  
[root@server-a ~]# iperf3 -c 192.168.8.24 -B 192.168.8.14 -t 1000 -i 1 &
```

Note the throughput achieved.

Step 3. While traffic is flowing, check the total traffic on both servers:

```
[root@server-a ~]# nload ens1  
[root@server-b ~]# nload ens1
```

Step 4. Check CPU usage on both servers:

```
[root@server-a ~]# top  
[root@server-b ~]# top
```

Press '1' inside 'top' user interface to see individual CPU usage.