

Achieve Optimal Network Throughput on the Cisco UCS Virtual Interface Card 1457



November 2018

Executive summary

This document describes the network I/O performance characteristics of the Cisco UCS® C240 M5 Rack Server using the Cisco UCS Virtual Interface Card (VIC) 1457, a 25-Gbps modular LAN-on-motherboard (mLOM) card based on the Cisco UCS VIC 1400 platform. The document discusses how to achieve optimal throughput on the Cisco UCS C240 M5 server in a variety of scenarios. The goal of this document is to demonstrate to customers that, with this Cisco Unified Computing System™ (Cisco UCS) server platform, they are investing in an enterprise-class server platform, and that they can achieve their desired level of network throughput by implementing the configuration presented here.

All performance data presented in this document was obtained using the iperf3 testing tool, with analysis based on the aggregate throughput of all the streams used to saturate the 25-Gbps links. The primary focus of this document is on how to achieve optimal throughput regardless of the environment.

Introduction

The Cisco Unified Computing System, or Cisco UCS, is a next-generation data center platform that unites computing, networking, storage access, and virtualization resources in a cohesive system designed to reduce total cost of ownership (TCO) and increase business agility. The system

integrates a low-latency, lossless 10-, 25-, and 40-Gigabit Ethernet network fabric with enterprise-class blade and rack x86-architecture servers. The system is an integrated, scalable, multichassis platform in which all resources participate in a unified management domain.

In fact, there are as many different network environments as there are customers, and no single solution can address them all. The Cisco UCS VIC 1457 is based on the most recent generation of the Cisco UCS VIC 1400 platform, which supports 25-Gbps throughput. This additional throughput option introduces a new set of challenges in a variety of situations. This document discusses how to meet those challenges and achieve optimal throughput.

Audience

This document assumes that readers have some familiarity with data center trends, server and storage technology, and the Cisco UCS product line. Readers should also have a working knowledge of Cisco Integrated Management Controller (IMC) web user interface, because this guide does not offer the start-to-finish configuration details usually found in more in-depth guides.

Network I/O challenges

Data center bandwidth requirements are growing at a fast rate, and with the latest-generation processors, high-performing nonvolatile memory express (NVMe) solid-state disks (SSDs), graphics processing units (GPUs), applications are demanding high bandwidth and faster I/O processing.

High-performance computing (HPC), distributed applications, artificial intelligence (AI), and machine learning need high-performing networks and bandwidth. With the increased demand for bandwidth and faster I/O operations, with server network I/O exceeding 10 Gbps, a new set of challenges has arisen. To achieve the maximum desired throughput with higher bandwidth, you must address two factors in particular:

- **Multiple application threads:** As network I/O increases, higher CPU frequencies are needed to accommodate the I/O load. If a higher frequency is not possible, you will need more CPU cores to process all the requests. To address this requirement, you need an application that can take advantage of today's multithreaded environments.
- **Need for reduced request load:** With network throughput of 25 Gbps, a tremendous amount of processing is required to service all the requests. You can reduce the load of those requests in two ways. The typical maximum transmission unit (MTU) value for most networks for many years has been 1500. However, by increasing the frame size so that the system uses an oversized or jumbo frame, you can reduce the processing load on the network adapter application-specific integrated circuit (ASIC) and the CPU in the system. Alternatively, you can increase the number of queues to achieve a more distributed method of receiving packets than with only a single queue.

Scope of this document

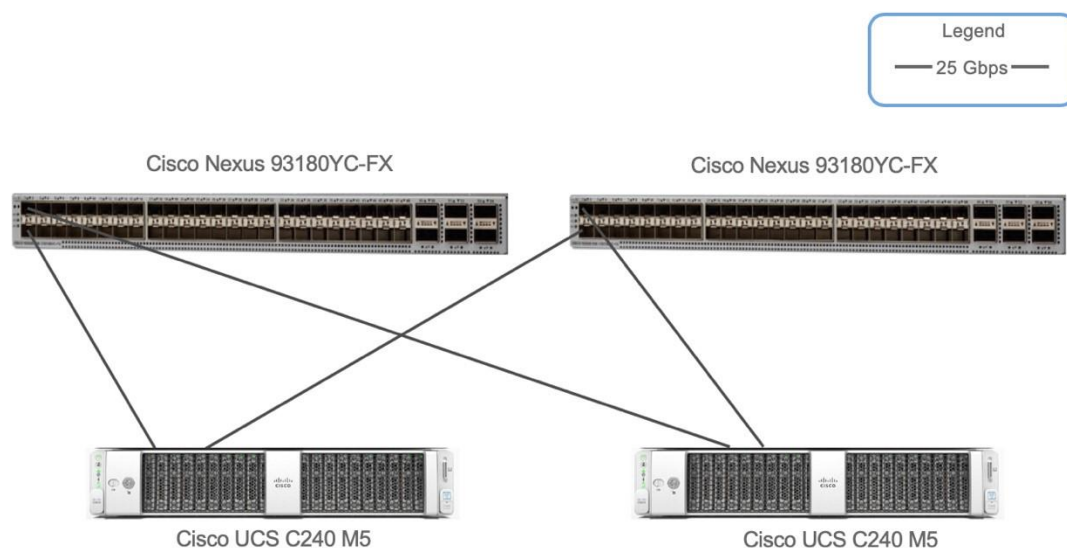
This document is limited to the discussion of testing and throughput measurement for the 25-Gbps link. All network performance testing was conducted between two Cisco UCS C240 M5 Rack Servers, and bandwidth was measured by connecting a 25-Gbps link from Cisco UCS VIC 1457 to a Cisco Nexus® switch. Default and optimal performance configuration and performance results were captured for MTU settings of both 1500 and 9000.

Note that the same settings as those used to achieve optimal throughput on the Cisco UCS C240 M5 server are applicable to other Cisco UCS C-Series Rack Servers and S-Series Storage Servers. However, test results for those servers are not presented in this document.

Infrastructure components

Figure 1 shows the Cisco UCS components used in the tests described in this document.

Figure 1. Cisco UCS components: Two Cisco Nexus 93180YC-FX Switches and two Cisco C240 M5 Rack Servers



Hardware components

Cisco UCS C240 M5

The Cisco UCS C240 M5 Rack Server is a 2-socket, 2-rack-unit (2RU) server offering industry-leading performance and expandability. It supports a wide range of storage and I/O-intensive infrastructure workloads, from big data and analytics to collaboration. Cisco UCS C-Series Rack Servers can be deployed as standalone servers or as part of a Cisco UCS managed environment to take advantage of Cisco's standards-based unified computing innovations that help reduce customers' TCO and increase their business agility.

In response to ever-increasing computing and data-intensive real-time workloads, the enterprise-class Cisco UCS C240 M5 server extends the capabilities of the Cisco UCS portfolio in a 2RU form factor. It incorporates the Intel® Xeon® Scalable processors, supporting up to 20 percent more cores per socket, twice the memory capacity, and five times more NVMe PCI Express (PCIe) SSDs than the previous generation of servers. These improvements deliver significant performance and efficiency gains that will improve your application performance.

The C240 M5 delivers outstanding levels of storage expandability with exceptional performance, with:

- The latest Intel Xeon Scalable CPUs, with up to 28 cores per socket
- Up to 24 DDR4 DIMMs for improved performance
- Up to 26 hot-swappable small-form-factor (SFF) 2.5-inch drives, including 2 rear hot-swappable SFF drives (up to 10 support NVMe PCIe SSDs on the NVMe-optimized chassis version), or 12 large-form-factor (LFF) 3.5-inch drives plus 2 rear hot-swappable SFF drives
- Support for a 12-Gbps SAS modular RAID controller in a dedicated slot, leaving the remaining PCIe 3.0 slots available for other expansion cards
- mLOM slot that can be used to install a Cisco UCS VIC without consuming a PCIe slot, supporting dual 10, 25, or 40-Gbps network connectivity
- Dual embedded Intel x550 10GBASE-T LAN-on-motherboard (LOM) ports
- Modular M.2 or Secure Digital (SD) cards that can be used for boot
- C240 M5 servers can be deployed as standalone servers or in a Cisco UCS managed environment. When used in combination with Cisco UCS Manager, the C240 M5 brings the power and automation of unified computing to enterprise applications, including Cisco® SingleConnect technology, drastically reducing switching and cabling requirements.

Table 1 lists the specifications for the Cisco UCS C240 M5 Rack Server.

Table 1. Part numbers for Cisco UCS C220 M5 high-density SFF and LFF base rack server models

Part number	Description
Form factor	2RU rack server
Processors	Intel Xeon Scalable processors (1 or 2)
Memory	24 DDR4 DIMM slots: 8, 16, 32, 64, and 128 GB and up to 2666 MHz
PCIe expansion	6 PCIe 3.0 slots plus 1 dedicated 12-Gbps RAID controller slot and 1 dedicated mLOM slot
RAID controllers	<ul style="list-style-type: none"> • Internal controllers: Cisco 12-Gbps Modular RAID Controller (PCIe 3.0) with 2- or 4-GB Flash-Backed Write Cache (FBWC), providing enterprise-class data protection for up to 26 SAS and SATA hard-disk drives (HDDs), SSDs, or NVMe PCIe SSDs; or Cisco 12-Gbps modular SAS host bus adapter (HBA) • External controller: Cisco 12-Gbps 9400-8e SAS HBA
Internal storage	Backplane options: <ul style="list-style-type: none"> • Up to 26 x 2.5-inch SAS and SATA HDDs and SSDs and up to 4 NVMe PCIe drives • Up to 10 x 2.5-inch NVMe PCIe and 16 SAS and SATA HDDs and SSDs • Up to 12 x 3.5-inch SAS and SATA HDDs and SSDs, and 2 rear 2.5-inch HDDs and SSDs and up to 4 NVMe PCIe drives
Embedded network interface cards (NICs)	Dual 10GBASE-T Intel x550 Ethernet ports
mLOM	Dedicated mLOM slot that can flexibly accommodate 1-, 10-, 25-, 40-, and

Part number	Description
	100-Gbps adapters
Power supplies	Hot-pluggable, redundant 770W AC, 1050W AC, 1050W DC, and 1600W AC
Other storage	<ul style="list-style-type: none"> • Dual internal Cisco FlexFlash SD cards (32, 64, and 128 GB) for installing an operating system or hypervisor • Support for RAID 0 mirroring between SD cards • Dedicated Baseboard Management Controller (BMC) MicroSD card (32 GB) for server utilities • Dual M.2 SATA SSD or NVMe
Management	<ul style="list-style-type: none"> • Cisco Intersight™ • Cisco Integrated Management Controller (IMC) • Cisco IMC Supervisor • Cisco UCS Manager • Cisco UCS Central Software • Cisco UCS Director • Cisco UCS Performance Manager
Rack options	Cisco ball-bearing rail kit with optional reversible cable management farm
Hardware and software interoperability	See the Cisco Hardware and Software Interoperability List for a complete list of supported operating systems peripheral options
Operating systems	Microsoft Windows Server 2012, Microsoft Windows 2016, Red Hat Enterprise Linux (RHEL), SUSE Linux, or VMware vSphere

Cisco UCS VIC 1457

The Cisco UCS VIC 1400 platform extends the network fabric directly to both servers and virtual machines so that a single


connectivity mechanism can be used to connect both physical and virtual servers with the same level of visibility and control. Cisco VICs provide complete programmability of the Cisco UCS I/O infrastructure, with the number and type of I/O interfaces configurable on demand with a zero-touch model.

Cisco VICs support Cisco SingleConnect technology, which provides an easy, intelligent, and efficient way to connect and manage computing in your data center. Cisco SingleConnect unifies LAN, SAN, and systems management into one simplified link for rack servers, blade servers, and virtual machines. This technology reduces the number of network adapters, cables, and switches needed and radically simplifies the network, reducing complexity. Cisco VICs can support 256 PCIe virtual devices, either virtual NICs (vNICs) or virtual HBAs (vHBAs), a high rate of I/O operations per second (IOPS), support for lossless Ethernet, and 10- and 40-Gbps connection to servers. The PCIe 3.0 x16 interface helps ensure optimal bandwidth to the host for network-intensive applications, with a redundant path to the fabric interconnect. Cisco VICs support NIC teaming, with fabric failover for increased reliability and availability. In addition, it provides a policy-based, stateless, agile server infrastructure for your data center.

The VIC 1400 series is designed exclusively for the M5 generation of Cisco UCS B-Series Blade Servers, C-Series Rack Servers, and S-Series Storage Servers. The adapters are capable of supporting 10-, 25, and 40-Gigabit Ethernet and Fibre Channel over Ethernet (FCoE). The VIC incorporates Cisco's next-generation converged network adapter (CNA) technology and offers a comprehensive feature set, providing investment protection for future feature software releases. In addition, the VIC supports Cisco Data Center Virtual Machine Fabric Extender (VM-FEX) technology. This technology extends the Cisco UCS fabric interconnect ports to virtual machines, simplifying server virtualization deployment.

The Cisco UCS VIC 1457 (Figure 2) is a quad-port Small Form-Factor Pluggable (SFP28) mLOM card designed for the M5 generation of Cisco UCS C-Series Rack Servers. The card supports 10- and 25-Gbps Ethernet and FCoE. The card can present PCIe standards-compliant interfaces to the host, and these can be dynamically configured as either NICs or HBAs.

Figure 2. Cisco UCS VIC 1457



The Cisco UCS VIC 1400 platform provides the following features and benefits:

- **Stateless and agile platform:** The personality of the card is determined dynamically at boot time using the service profile associated with the server. The number, type (NIC or HBA), identity (MAC address and World Wide Name [WWN]), failover policy, bandwidth, and quality-of-service (QoS) policies of the PCIe interfaces are all determined using the service profile. The capability to define, create, and use interfaces on demand provides a stateless and agile server infrastructure.
- **Network interface virtualization:** Each PCIe interface created on the VIC is associated with an interface on the Cisco UCS fabric interconnect, providing complete network separation for each virtual cable between a PCIe device on the VIC and the interface on the fabric interconnect.

Cisco Nexus 93180YC-FX Switch

The Cisco Nexus 93180YC-FX Switch (Figure 3) is a 1RU switch with latency of less than 1 microsecond that supports 3.6 Tbps of bandwidth. The 48 downlink ports on the 93180YC-FX can be configured to work as 1-, 10-, or 25-Gbps Ethernet or FCoE ports or as 16- or 32-Gbps Fibre Channel ports, converging primary storage and computing servers, back-end storage, and policybased switching on the leaf node. The six uplink ports can be configured as 40- and 100-

Gbps Ethernet or FCoE ports, offering flexible migration options. The switch has FC-FEC and RS-FEC enabled for 25-Gbps support.

Figure 3. Cisco Nexus 93180YC-FX Switch



Table 2 lists the Cisco Nexus 9300-FX platform switch specifications.

Table 2. Cisco Nexus 9300-FX platform switch specifications

Feature	Cisco Nexus 93180YC-FX
Ports	48 x 10/25-Gbps and 6 x 40/100-Gbps QSFP28 ports
Downlink supported speeds	1/10/25-Gbps Ethernet 8/16/32-Gbps Fibre Channel
CPU	6 cores
System memory	64 GB
SSD drive	128 GB
System buffer	40 MB
Management ports	1 RJ-45 port L1 and L2 ports are unused
USB ports	1
RS-232 serial ports	1
Power supplies (up to 2)	500W AC, 930W DC, or 1200W HVAC/HVDC
Typical power (AC/DC)	260W
Maximum power (AC/DC)	425W
Input voltage (AC)	100 to 240V
Input voltage (high-voltage AC [HVAC])	200 to 277V
Input voltage (DC)	-48 to -60V
Input voltage (high-voltage DC [HVDC])	-240 to -380V
Frequency (AC)	50 to 60 Hz

Feature	Cisco Nexus 93180YC-FX
Fans	4
Airflow	Port-side intake and exhaust
Physical dimensions (H x W x D)	1.72 x 17.3 x 22.5 in. (4.4 x 43.9 x 57.1 cm)
Acoustics	57 dBA at 40% fan speed, 68.9 dBA at 70% fan speed, and 77.4 dBA at 100% fan speed
RoHS compliance	Yes
Mean time between failure (MTBF)	238,470 hours

Infrastructure configuration

Both Cisco UCS C240 M5 servers are configured in standalone mode (without using Cisco UCS Manager) and they are connected to a pair of Cisco Nexus 93180YC-FX Switches. Port 0 from each VIC is connected to port 1 of Cisco Nexus 93180YC-FX Switch A, and port 3 from each VIC is connected to port 1 of Cisco Nexus 93180YC-FX Switch B.

Table 3 and 4 shows the Hardware configuration and software version used for the testing

Table 3. Hardware configuration

Component	Configuration
2 x Cisco UCS C240 M5	Each with 2 x Intel Xeon Scalable 6132 CPU 12 x 32 GB-2R-DDR4-2666MHz UCS VIC 1457

Table 4. Software version

Component	Version
CIMC	4.0(1a)
Operating System	Redhat 7.5
enic driver	3.0.107.37-492.52
enic firmware	5.0(1c)

Three infrastructure configuration states were captured in the results:

- **State 1:** The first configuration state is as close to a default configuration as possible. The Cisco Nexus is configured with a standard 1500-MTU frame size. The adapter profiles on each Cisco UCS C240 M5 are configured with the default adapter profile using the Cisco IMC

- **State 2:** The second configuration state applies jumbo frames to the vNIC and Cisco Nexus 93180YC-FX Switches.
- **State 3:** The final configuration state applies jumbo frames to the Cisco Nexus 93180YC-FX Switches, and an optimal adapter profile is applied to the network adapter on the vNIC.

Workload Characterization

The tools most commonly used to test network performance are iperf and iperf3. The classic tool is iperf. It has been in use for many years. The more recent version, iperf3, is fundamentally different from the original iperf. It benefits from a more compact code base and a more modern implementation. iperf3 supports multiple streams, and all streams used by its process are single threaded. It supports the tuning of various parameters related to timing, buffers, and protocols (TCP, User Datagram Protocol [UDP], and Stream Control Transmission Protocol [SCTP] with IPv4 and IPv6). For each test it reports the bandwidth, loss, and other parameters.

Test configuration

The results presented here represent all three states detailed in the preceding “Infrastructure configuration” section. Performance was measured using iperf3 to examine 25-Gbps throughput under a variety of conditions.

State 1: Default configuration

In this configuration, the default adapter profile is used with an MTU of 1500. The default configuration is shown here (Figure 4). With an MTU of 1500 and four parallel streams, the throughput was 21.7 Gbps, and with eight parallel streams, the throughput was 23.5 Gbps.

Figure 4. Default configuration

▼ Ethernet Interrupt

Interrupt Count: (1 - 514)Interrupt Mode: ▼

▼ Ethernet Receive Queue

Receive Queue Count: (1 - 256)Receive Queue Ring Size: (64 - 4096)

▼ Ethernet Transmit Queue

Transmit Queue Count: (1 - 256)Transmit Queue Ring Size: (64 - 4096)

▼ Completion Queue

Completion Queue Count: (1 - 512)Completion Queue Ring Size:

▼ TCP Offload

Enable Large Receive: ☒ Enable TCP Rx Offload Checksum Valid... ☒Enable TCP Segmentation Offload: ☒ TCP Tx Offload Checksum Generation: ☒

▼ Receive Side Scaling

Enable TCP Receive Side Scaling: ☒ Enable TCP-IPv6 RSS: ☒Enable IPv4 RSS: ☒ Enable IPv6 Extension RSS: ☐Enable TCP-IPv4 RSS: ☒ Enable TCP-IPv6 Extension RSS: ☐Enable IPv6 RSS: ☒**Performance results**

Figure 5 shows the interface configuration, and Figure 6 shows the configuration results.

Figure 5. Interface configuration

```

eno6: flags=4163<UP,BROADCAST,RUNNING,MULTICAST> mtu 1500
    inet 172.102.2.171 netmask 255.255.0.0 broadcast 172.102.255.255
    inet6 fe80::5c57:58fd:7743:e75c prefixlen 64 scopeid 0x20<link>
    inet6 2001::31e7:45d2:3ee6:a659 prefixlen 64 scopeid 0x0<global>
    ether 00:5d:73:1c:a4:6d txqueuelen 1000 (Ethernet)
    RX packets 942 bytes 95544 (93.3 KiB)
    RX errors 0 dropped 0 overruns 0 frame 0
    TX packets 84 bytes 13274 (12.9 KiB)
    TX errors 0 dropped 0 overruns 0 carrier 0 collisions 0

```

Figure 6. Results

```
[root@localhost ~]# iperf3 -c 172.102.2.171 -t 60 -P 8
Connecting to host 172.102.2.171, port 5201
[ 4] local 172.102.1.149 port 59046 connected to 172.102.2.171 port 5201
[ 6] local 172.102.1.149 port 59048 connected to 172.102.2.171 port 5201
[ 8] local 172.102.1.149 port 59050 connected to 172.102.2.171 port 5201
[10] local 172.102.1.149 port 59052 connected to 172.102.2.171 port 5201
[12] local 172.102.1.149 port 59054 connected to 172.102.2.171 port 5201
[14] local 172.102.1.149 port 59056 connected to 172.102.2.171 port 5201
[16] local 172.102.1.149 port 59058 connected to 172.102.2.171 port 5201
[18] local 172.102.1.149 port 59060 connected to 172.102.2.171 port 5201
[ ID] Interval           Transfer     Bandwidth   Retr  Cwnd
[ 4]  0.00-1.00   sec    356 MBytes  2.98 Gbits/sec    0   581 KBytes
[ 6]  0.00-1.00   sec    354 MBytes  2.97 Gbits/sec    0   566 KBytes
[ 8]  0.00-1.00   sec    351 MBytes  2.94 Gbits/sec    0   573 KBytes
[10]  0.00-1.00   sec    351 MBytes  2.94 Gbits/sec    0   576 KBytes
[12]  0.00-1.00   sec    352 MBytes  2.95 Gbits/sec    0   573 KBytes
[14]  0.00-1.00   sec    351 MBytes  2.94 Gbits/sec    0   574 KBytes
[16]  0.00-1.00   sec    351 MBytes  2.94 Gbits/sec    0   573 KBytes
[18]  0.00-1.00   sec    350 MBytes  2.93 Gbits/sec    0   577 KBytes
[SUM] 0.00-1.00   sec    2.75 GBytes 23.6 Gbits/sec    0
-----
[ 4]  1.00-2.00   sec    350 MBytes  2.94 Gbits/sec    0   587 KBytes
[ 6]  1.00-2.00   sec    351 MBytes  2.94 Gbits/sec    0   581 KBytes
[ 8]  1.00-2.00   sec    350 MBytes  2.94 Gbits/sec    0   576 KBytes
[10]  1.00-2.00   sec    350 MBytes  2.93 Gbits/sec    0   583 KBytes
[12]  1.00-2.00   sec    350 MBytes  2.93 Gbits/sec    0   585 KBytes
[14]  1.00-2.00   sec    349 MBytes  2.93 Gbits/sec    0   581 KBytes
[16]  1.00-2.00   sec    349 MBytes  2.93 Gbits/sec    0   580 KBytes
[18]  1.00-2.00   sec    350 MBytes  2.93 Gbits/sec    0   587 KBytes
[SUM] 1.00-2.00   sec    2.73 GBytes 23.5 Gbits/sec    0
-----
[ 4]  2.00-3.00   sec    351 MBytes  2.94 Gbits/sec    0   587 KBytes
[ 6]  2.00-3.00   sec    351 MBytes  2.94 Gbits/sec    0   581 KBytes
[ 8]  2.00-3.00   sec    350 MBytes  2.93 Gbits/sec    0   576 KBytes
[10]  2.00-3.00   sec    350 MBytes  2.94 Gbits/sec    0   588 KBytes
[12]  2.00-3.00   sec    350 MBytes  2.93 Gbits/sec    0   585 KBytes
[14]  2.00-3.00   sec    350 MBytes  2.93 Gbits/sec    0   581 KBytes
[16]  2.00-3.00   sec    350 MBytes  2.94 Gbits/sec    0   580 KBytes
[18]  2.00-3.00   sec    350 MBytes  2.93 Gbits/sec    0   587 KBytes
[SUM] 2.00-3.00   sec    2.74 GBytes 23.5 Gbits/sec    0
```

State 2: Jumbo frames enabled

The Ethernet adapter and Cisco Nexus 93180YC-FX Switches is configured with an MTU of 9000 with the default adapter profile. By using a more efficient transmission unit, the adapter can achieve a maximum throughput of 24.7 Gbps without the need to increase the number of queues and interrupts. If you can modify the existing network to support jumbo frames, or if jumbo frames are already set, this change is relatively simple to make. However, many customers cannot or do not want to accommodate an MTU of 9000 in their data center environments.

To configure jumbo frames in the vNICs, follow these steps:

1. From the Cisco UCS C240 M5 server's IMC home screen, select the top-right button, which looks like this:



2. In the menu that appears, expand the Networking section and Adapter Card.
3. Select the vNICs tab.
4. In the left tree beneath the vNICs section, select the Ethernet adapter to be modified: most likely, eth0 or eth1, or both.
5. Set the MTU to 9000.

▼ vNIC Properties

▼ General

Name: <input type="text" value="eth0"/>	VLAN Mode: <input type="text" value="Trunk"/>
CDN: <input type="text" value="VIC-1-eth0"/>	Rate Limit: <input checked="" type="radio"/> OFF
MTU: <input type="text" value="9000"/> (1500 - 9000)	<input type="radio"/> <input type="text" value=""/>
Uplink Port: <input type="text" value="0"/>	Channel Number: <input type="text" value="N/A"/> (1 - 1000)
MAC Address: <input type="radio"/> Auto	PCI Link: <input type="text" value="0"/> (0 - 1)
<input checked="" type="radio"/> <input type="text" value="00:5D:73:1C:A4:6C"/>	Enable NVGRE: <input type="checkbox"/>
Class of Service: <input type="text" value="0"/> (0 - 6)	Enable VXLAN: <input type="checkbox"/>
Trust Host CoS: <input type="checkbox"/>	Advanced Filter: <input type="checkbox"/>
PCI Order: <input type="text" value="0"/> (0 - 7)	Port Profile: <input type="text" value="N/A"/>
Default VLAN: <input checked="" type="radio"/> None	Enable PXE Boot: <input type="checkbox"/>
<input type="radio"/> <input type="text" value=""/>	Enable VMQ: <input type="checkbox"/>
	Enable aRFS: <input type="checkbox"/>
	Enable Uplink Failover: <input type="checkbox"/>
	Fallback Timeout: <input type="text" value="N/A"/> (0 - 600)

6. Save the changes and reboot the server to make the changes take effect.

Performance results

Figure 7 shows the interface configuration, and Figure 8 shows the results.

Figure 7. Interface configuration

```
eno6: flags=4163<UP,BROADCAST,RUNNING,MULTICAST> mtu 9000
    inet 172.102.2.171 netmask 255.255.0.0 broadcast 172.102.255.255
    inet6 fe80::5c57:58fd:7743:e75c prefixlen 64 scopeid 0x20<link>
    inet6 2001::31e7:45d2:3ee6:a659 prefixlen 64 scopeid 0x0<global>
    ether 00:5d:73:1c:a4:6d txqueuelen 1000 (Ethernet)
    RX packets 290 bytes 43692 (42.6 KiB)
    RX errors 0 dropped 0 overruns 0 frame 0
    TX packets 84 bytes 13383 (13.0 KiB)
    TX errors 0 dropped 0 overruns 0 carrier 0 collisions 0
```


Figure 8. Results

```
[root@localhost ~]# iperf3 -c 172.102.2.171 -t 60 -P 4
Connecting to host 172.102.2.171, port 5201
[ 4] local 172.102.1.149 port 38630 connected to 172.102.2.171 port 5201
[ 6] local 172.102.1.149 port 38632 connected to 172.102.2.171 port 5201
[ 8] local 172.102.1.149 port 38634 connected to 172.102.2.171 port 5201
[10] local 172.102.1.149 port 38636 connected to 172.102.2.171 port 5201
```

[ID]	Interval		Transfer	Bandwidth	Retr	Cwnd
[4]	0.00-1.00	sec	741 MBytes	6.21 Gbits/sec	0	673 KBytes
[6]	0.00-1.00	sec	738 MBytes	6.19 Gbits/sec	0	673 KBytes
[8]	0.00-1.00	sec	737 MBytes	6.18 Gbits/sec	0	673 KBytes
[10]	0.00-1.00	sec	741 MBytes	6.22 Gbits/sec	0	699 KBytes
[SUM]	0.00-1.00	sec	2.89 GBytes	24.8 Gbits/sec	0	

[4]	1.00-2.00	sec	739 MBytes	6.20 Gbits/sec	0	673 KBytes
[6]	1.00-2.00	sec	737 MBytes	6.18 Gbits/sec	0	699 KBytes
[8]	1.00-2.00	sec	738 MBytes	6.19 Gbits/sec	0	690 KBytes
[10]	1.00-2.00	sec	738 MBytes	6.19 Gbits/sec	0	760 KBytes
[SUM]	1.00-2.00	sec	2.88 GBytes	24.8 Gbits/sec	0	

[4]	2.00-3.00	sec	737 MBytes	6.18 Gbits/sec	0	690 KBytes
[6]	2.00-3.00	sec	738 MBytes	6.19 Gbits/sec	0	717 KBytes
[8]	2.00-3.00	sec	738 MBytes	6.19 Gbits/sec	0	751 KBytes
[10]	2.00-3.00	sec	737 MBytes	6.18 Gbits/sec	0	760 KBytes
[SUM]	2.00-3.00	sec	2.88 GBytes	24.7 Gbits/sec	0	

[4]	3.00-4.00	sec	737 MBytes	6.18 Gbits/sec	0	708 KBytes
[6]	3.00-4.00	sec	738 MBytes	6.19 Gbits/sec	0	717 KBytes
[8]	3.00-4.00	sec	735 MBytes	6.16 Gbits/sec	0	786 KBytes
[10]	3.00-4.00	sec	740 MBytes	6.21 Gbits/sec	0	760 KBytes
[SUM]	3.00-4.00	sec	2.88 GBytes	24.7 Gbits/sec	0	

State 3: Jumbo frames enabled and optimal adapter profile applied

The Ethernet adapter (vNIC) is configured with an MTU of 9000 and an optimized adapter profile. By applying both settings, maximum throughput should be observed in almost all circumstances, which was observed in the tests to be 24.7 Gbps.

To apply an adapter profile, follow these steps:

1. From the Cisco UCS 240 M5 server's IMC home screen, select the top-right button, which looks like this:



2. In the menu that appears, expand the Networking section and Adapter Card.
3. Select the vNICs tab.
4. In the left tree beneath the vNICs section, select the Ethernet adapter to be modified: most likely, eth0 or eth1, or both.

5. Set the MTU to 9000.
6. Provide a name for the Ethernet adapter policy. Change the following fields:
 - Resources
 - Transmit Queues: 8
 - Ring Size: 4096
 - Receive Queues: 8
 - Ring Size: 4096
 - Completion Queues: 16
 - Interrupts: 32
 - Options
 - Receive Side Scaling (RSS): Enabled

▼ Ethernet Interrupt

Interrupt Count: (1 - 514) Coalescing Time: (0-65535us)

Interrupt Mode: Coalescing Type:

▼ Ethernet Receive Queue

Receive Queue Count: (1 - 256)

Receive Queue Ring Size: (64 - 4096)

▼ Ethernet Transmit Queue

Transmit Queue Count: (1 - 256)

Transmit Queue Ring Size: (64 - 4096)

▼ Completion Queue

Completion Queue Count: (1 - 512)

Completion Queue Ring Size:

▼ TCP Offload

Enable Large Receive: ☒ Enable TCP Rx Offload Checksum Validation: ☒

Enable TCP Segmentation Offload: ☒ TCP Tx Offload Checksum Generation: ☒

▼ Receive Side Scaling

Enable TCP Receive Side Scaling: ☒ Enable TCP-IPv6 RSS: ☒

Enable IPv4 RSS: ☒ Enable IPv6 Extension RSS: ☐

Enable TCP-IPv4 RSS: ☒ Enable TCP-IPv6 Extension RSS: ☐

7. Save the changes and reboot the server.

Performance results

Figure 9 shows the interface configuration, and Figure 10 shows the configuration results.

Figure 9. Interface configuration

```

eno6: flags=4163<UP,BROADCAST,RUNNING,MULTICAST> mtu 9000
    inet 172.102.2.171 netmask 255.255.0.0 broadcast 172.102.255.255
    inet6 fe80::5c57:58fd:7743:e75c prefixlen 64 scopeid 0x20<link>
    inet6 2001::31e7:45d2:3ee6:a659 prefixlen 64 scopeid 0x0<global>
    ether 00:5d:73:1c:a4:6d txqueuelen 1000 (Ethernet)
    RX packets 290 bytes 43692 (42.6 KiB)
    RX errors 0 dropped 0 overruns 0 frame 0
    TX packets 84 bytes 13383 (13.0 KiB)
    TX errors 0 dropped 0 overruns 0 carrier 0 collisions 0

```

Figure 10. Results

```

[root@localhost ~]# iperf3 -c 172.102.2.171 -t 60 -P 4
Connecting to host 172.102.2.171, port 5201
[ 4] local 172.102.1.149 port 38630 connected to 172.102.2.171 port 5201
[ 6] local 172.102.1.149 port 38632 connected to 172.102.2.171 port 5201
[ 8] local 172.102.1.149 port 38634 connected to 172.102.2.171 port 5201
[10] local 172.102.1.149 port 38636 connected to 172.102.2.171 port 5201
[ ID] Interval           Transfer     Bandwidth       Retr   Cwnd
[ 4]  0.00-1.00   sec    741 MBytes  6.21 Gbits/sec    0   673 KBytes
[ 6]  0.00-1.00   sec    738 MBytes  6.19 Gbits/sec    0   673 KBytes
[ 8]  0.00-1.00   sec    737 MBytes  6.18 Gbits/sec    0   673 KBytes
[10]  0.00-1.00   sec    741 MBytes  6.22 Gbits/sec    0   699 KBytes
[SUM]  0.00-1.00   sec    2.89 GBytes 24.8 Gbits/sec    0
-----
[ 4]  1.00-2.00   sec    739 MBytes  6.20 Gbits/sec    0   673 KBytes
[ 6]  1.00-2.00   sec    737 MBytes  6.18 Gbits/sec    0   699 KBytes
[ 8]  1.00-2.00   sec    738 MBytes  6.19 Gbits/sec    0   690 KBytes
[10]  1.00-2.00   sec    738 MBytes  6.19 Gbits/sec    0   760 KBytes
[SUM]  1.00-2.00   sec    2.88 GBytes 24.8 Gbits/sec    0
-----
[ 4]  2.00-3.00   sec    737 MBytes  6.18 Gbits/sec    0   690 KBytes
[ 6]  2.00-3.00   sec    738 MBytes  6.19 Gbits/sec    0   717 KBytes
[ 8]  2.00-3.00   sec    738 MBytes  6.19 Gbits/sec    0   751 KBytes
[10]  2.00-3.00   sec    737 MBytes  6.18 Gbits/sec    0   760 KBytes
[SUM]  2.00-3.00   sec    2.88 GBytes 24.7 Gbits/sec    0
-----
[ 4]  3.00-4.00   sec    737 MBytes  6.18 Gbits/sec    0   708 KBytes
[ 6]  3.00-4.00   sec    738 MBytes  6.19 Gbits/sec    0   717 KBytes
[ 8]  3.00-4.00   sec    735 MBytes  6.16 Gbits/sec    0   786 KBytes
[10]  3.00-4.00   sec    740 MBytes  6.21 Gbits/sec    0   760 KBytes
[SUM]  3.00-4.00   sec    2.88 GBytes 24.7 Gbits/sec    0

```

For more information

For additional information, see:

- <https://www.cisco.com/c/en/us/products/collateral/servers-unified-computing/ucs-c-series-rack-servers/datasheet-c78-739279.html>
- <https://www.cisco.com/c/en/us/products/collateral/interfaces-modules/unified-computing-system-adapters/datasheet-c78-741130.html>

- <https://www.cisco.com/c/en/us/products/switches/nexus-93180yc-fx-switch/index.html>

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)

Printed in USA

C11-741550-00 11/18