

Cisco UCS B-Series Best Network Configuration Practices

Last Updated: March 15, 2023

Contents

Overview.....	3
Cisco UCS Fabric Interconnects.....	3
Cisco UCS B-Series.....	4
Cisco UCS Virtual Interface Cards.....	6
Bandwidth on blade servers.....	6
Cisco UCS B480 M5 Blade Servers.....	12
Fibre Channel Policies.....	17
Performance testing.....	18
Resources.....	25

Overview

The Cisco Unified Computing System™ (Cisco UCS®) is a next-generation data center platform that unites computing, networking, storage access, and virtualization resources into a cohesive system designed to reduce Total Cost of Ownership (TCO) and increase business agility.

The Cisco UCS Virtual Interface Card (VIC) incorporates the Converged Network Adapter (CNA) technology and offers a comprehensive feature set, providing investment protection across multiple generations of the Cisco UCS blade and rack servers.

Audience

The audience for this document consists of system architects, system engineers, and any other technical staff who are responsible for planning and maintaining the Cisco UCS infrastructure. Although every effort has been made to make this document appeal to the widest possible audience, the document assumes that readers have an understanding of Cisco UCS hardware, terminology, and configurations.

Objective

This document addresses some of the frequently asked questions about the Cisco UCS VIC adapters and the recommended practices from a Fibre Channel perspective. Topics addressed include: bandwidth on blade servers with various configurations of adapters and fabric interconnects and port extenders and best practices for adapter policies and Quality of Service (QoS).

Cisco UCS Fabric Interconnects



Figure 1.
Cisco UCS 6200 Series Fabric Interconnect

The Cisco UCS 6200 Series Fabric Interconnects are a core part of the Cisco Unified Computing System, providing both network connectivity and management capabilities for the system (Figure 2). The Cisco UCS 6200 Series offers line-rate, low-latency, lossless 10 Gigabit Ethernet, Fibre Channel over Ethernet (FCoE), and Fibre Channel functions.

- Although Cisco supports this product, we are no longer selling it. The end-of-sale date was May 31,2019.
- The end-of-support date will be May 31, 2024.



Figure 2.
Cisco UCS 6300 Series Fabric Interconnect

The Cisco UCS 6300 Series Fabric Interconnects are a core part of Cisco UCS, providing both network connectivity and management capabilities for the system (Figure 2). The Cisco UCS 6300 Series offers line-rate, low-latency, lossless 10 and 40 Gigabit Ethernet, FCoE, and Fibre Channel functions.



Figure 3.
Cisco UCS 6400 Series Fabric Interconnects

The Cisco UCS 6400 Series offer line-rate, low-latency, lossless 10/25/40/100 Gigabit Ethernet, FCoE, and Fibre Channel functions (Figure 3).

Using a cut-through architecture, the Cisco UCS 6454 and 64108 support deterministic, low-latency, line-rate 10/25/40/50/100 Gigabit Ethernet ports, switching capacity of 3.82 Tbps for the 6454, 7.42 Tbps for the 64108, and 200-Gbps bandwidth between the 6454 Fabric Interconnect or 64108 and the I/O module (IOM) or 2408 Fabric Extender per 5108 blade chassis.

Cisco UCS B-Series



Figure 4.
Cisco UCS 5108 Blade Server Chassis with blade servers front and back

A Cisco UCS 5108 chassis has 32x 10-Gbps backplane Ethernet connections between an IOM and the 8 server bays, thus each server has 4x 10-Gbps backplane traces per IOM. This setup gives 1.2 TB of aggregate throughput through the chassis.

Various interface types on I/O modules follow:

- A Host Interface (HIF) is a backplane port on an IOM that connects to a blade-server. The 10-Gbps-KR connections per server can mux to 40-Gbps-KR4 in IOM 2304 and 2408. 40-Gbps-KR4 can mux with IOM 1340+PE/1440+PE or 1380/1440.
- A Network Interface (NIF) is an external-facing interface on an IOM that connects to a fabric interconnect with SFP+ cables of 10/25/40-Gbps speeds depending on the IOM.



Figure 5.
Cisco UCS 2208XP Series Fabric Extender

- The Cisco UCS 2208XP Fabric Extender has eight 10 Gigabit Ethernet, FCoE-capable, Enhanced Small Form-Factor Pluggable (SFP+) ports that connect the blade chassis to the fabric interconnect.
- Each Cisco UCS 2208XP has thirty-two 10 Gigabit Ethernet ports connected through the midplane to each half-width slot in the chassis.
- Typically configured in pairs for redundancy, two fabric extenders provide up to 160 Gbps of I/O to the chassis.



Figure 6.
Cisco UCS 2304 Fabric Extender

The Cisco UCS 2304 Fabric Extender has four 40 Gigabit Ethernet, FCoE-capable, Quad Small Form-Factor Pluggable (QSFP+) ports that connect the blade chassis to the fabric interconnect. Each Cisco UCS 2304 can provide one 40 Gigabit Ethernet port connected through the midplane to each half-width slot in the chassis, giving it a total of eight 40-Gbps interfaces to the compute node. Typically configured in pairs for redundancy, two fabric extenders provide up to 320 Gbps of I/O to the chassis.



Figure 7.
Cisco UCS 2408 Fabric Extender

The Cisco UCS 2408 Fabric Extender has eight 25 Gigabit Ethernet, FCoE-capable, Small Form-Factor Pluggable 28 (SFP28) ports that connect the blade chassis to the fabric interconnect. Each Cisco UCS 2408 provides 10 Gigabit Ethernet ports connected through the midplane to each half-width slot in the chassis, giving it a total of 32 10-Gbps interfaces to Cisco UCS blades.

Cisco UCS Virtual Interface Cards



Figure 8.
Cisco UCS 1400 Series VICs

The Cisco UCS 1400 Series VICs extend the network fabric directly to both servers and virtual machines so that you can use a single connectivity mechanism 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.

For a comprehensive look at the supported features, please refer to the [Cisco UCS Virtual Interface Card 1400 Series Data Sheet](#). The VIC 1440 also works with an optional VIC port extender (PE), which enables extra bandwidth on the VIC 1440 adapter.

Bandwidth on blade servers

The half-width (B200 M5) or full-width (B480 M5) blade server on a UCS 5108 chassis will see different throughputs and NIC/HBA (vNIC/vHBA in UCS) speeds per VIC adapter depending on a combination of VIC adapter and IOM.

Cisco UCS B200 M5 Blade Servers

Figure 9 shows the connectivity on a Cisco UCS 5108 chassis with an IOM 2408 connected to the Cisco 6454 Fabric Interconnect. The red lines are the 25-Gbps physical connections from the IOM to the fabric interconnect, and the blue lines show the 10-Gbps-KR backplane connections from the IOM toward the VIC on a Cisco B200 M5 Blade Server.




Figure 9.
Cisco UCS B200 M5 connectivity in Cisco UCS 5108 chassis with IOM 2408

Throughput for the B200 M5 server depends on which VICs are installed on the server and how many backplane connections the installation enables. Broadly speaking, the following examples explain the throughput to be expected from a blade server:

- B200 M5 with only VIC 1440 will have two 10-Gbps-KR connections enabled toward IOM 2208/2304/2408.
- B200 M5 with VIC 1440 + port extender will have four 10-Gbps-KR connections toward IOM 2304/2408.
- B200 M5 with VIC 1440 + 1480 will have four 10-Gbps-KR connections enabled toward IOM 2208/2304/2408.

Table 1 shows the throughput per B200 M5 with various 1440 and 1480 combinations. The throughput from the server is calculated across both the fabric interconnects in the Cisco UCS domain. Table 2 shows B200 M5 adapter and module combinations for virtual NIC/virtual Host Bus Adapter (vNIC/vHBA) speeds per adapter on **each** fabric interconnect.

Table 1. B200 M5 adapter and module combinations throughput per blade server with two CPUs installed for two fabric interconnects. > **Total aggregated bandwidth for 2x fabric interconnects /IO**

Fabric interconnect	2x FI 6400	2x FI 6300	2x FI 6200/6300/6400	2x FI 6200/6300/6400
I/O module	2x IOM 2408	2x IOM 2304	2x IOM 2208	2x IOM 2204
1340	40G	40G	40G	20G
1340 + PE	80G	80G	80G	40G
1340 + 1380	80G	80G	80G	40G
1440	40G	40G	40G	20G
1440 + PE	80G	80G	40G	20G
1440 + 1480	80G	80G	80G	40G

Table 2. B200 M5 adapter and module combinations for vNIC/vHBA speeds per adapter on each fabric interconnect > **Total aggregated bandwidth for 1x fabric interconnect**

Fabric interconnect	1x FI 6400/6536	OS vNIC speed	Max bandwidth single flow	Fabric interconnect	6300/6536 Fabric Interconnect	OS vNIC speed	Max bandwidth single flow
I/O module	1x IOM 2408			I/O module	1x IOM 2304		
1340	2x 10 Gbps	20 Gbps	10 Gbps	1340	2x 10 Gbps	20 Gbps	10 Gbps
1340 + PE	4x 10 Gbps	40 Gbps	10 Gbps	1340 + PE	40 Gbps	40 Gbps	40 Gbps
1340 + 1380	2x 10 Gbps + 2x 10 Gbps	20 Gbps 20 Gbps	10 Gbps 10 Gbps	1340 + 1380	2x 10 Gbps + 2x 10 Gbps	20 Gbps 20 Gbps	10 Gbps 10 Gbps
1440	2x 10 Gbps	20 Gbps	10 Gbps	1440	2x10 Gbps	20 Gbps	10 Gbps
1440 + PE	40 Gbps**	40 Gbps	25Gbps	1440 + PE	40 Gbps	40 Gbps	40 Gbps
1440 + 1480	2x 10 Gbps + 2x 10 Gbps	20 Gbps 20 Gbps	10 Gbps 10 Gbps	1440 + 1480	2x 10 Gbps + 2x 10 Gbps	20 Gbps 20 Gbps	10 Gbps 10 Gbps

** Supported from Cisco UCSM Release 4.1(2), single-flow maximum of 25 Gbps

Highlighted area is the suggested configuration and was used for testing with two 6400 Fabric Interconnects and two IOM 2408 Fabric Extenders.

Fabric interconnect	1x FI 6200/6300/6400	OS vNIC speed	Max bandwidth single flow	Fabric interconnect	1x FI 6200/6300/6400	OS vNIC speed	Max bandwidth single flow
I/O module	1x IOM 2208			I/O module	1x IOM 2204		
1340	2x 10 Gbps	20 Gbps	10 Gbps	1340	10 Gbps	10 Gbps	10 Gbps
1340 + PE	4x 10 Gbps	40 Gbps	10 Gbps	1340 + PE	2x 10 Gbps	20 Gbps	10 Gbps
1340 +	2x 10 Gbps +	20 Gbps	10 Gbps	1340 +	10 Gbps +	10 Gbps	10 Gbps
1380	2x 10 Gbps	20 Gbps	10 Gbps	1380	10 Gbps	10 Gbps	10 Gbps
1440	2x 10 Gbps	20 Gbps	10 Gbps	1440	10 Gbps	10 Gbps	10 Gbps
1440 + PE	2x 10 Gbps	20 Gbps	10 Gbps	1440 + PE	10 Gbps	10 Gbps	10 Gbps
1440 +	2x 10 Gbps +	20 Gbps	10 Gbps	1440 +	10 Gbps +	10 Gbps	10 Gbps
1480	2x 10 Gbps	20 Gbps	10 Gbps	1480	10 Gbps	10 Gbps	10 Gbps

One vHBA is defined for each fabric interconnect A and B.

Note: VMware ESXi will display lower speeds, but bandwidths listed are achievable.

Note: You cannot mix VIC 1440 cards with 1380 cards or VIC 1340 cards with 1480 cards.

In Table 2, 2x 10-Gbps means the vNIC will have an aggregate speed of 20 Gbps and a maximum single flow of 10 Gbps. A 40-Gbps vNIC on a VIC 1400 with IOM 2304 can have a single-flow maximum of 40 Gbps, whereas with IOM 2408, it would be a single-flow maximum of 25 Gbps and aggregate of 40 Gbps across multiple flows. Details about the 40-Gbps capability of VIC 1440/1480 with IOM 2408 and IOM 2304 are explained in the next section.

On a B200 M5, VIC 1440 and 1480 adapters can have 10, 20, or 40-Gbps vNICs. The vNIC speed seen on the server depends on the various VIC combinations installed on the blade server, as in the following:

- With a VIC 1440 on a B200 M5, the server will see 20Gbps (2x 10 Gbps) throughput with all IOMs except the IOM 2204 and fabric interconnect 6324.
- A VIC 1440 with a port expander on a B200 M5 will enable 40-Gbps-KR4 toward an IOM 2304 and IOM 2408. An IOM 2304 in Cisco UCS Manager (UCSM) Releases 4.0 and later supports this solution; meanwhile, support for 40-Gbps-KR4 with an IOM 2408 is available from Cisco UCSM Release 4.12 and later).

Port-extender support with the VIC 1440 on M5 blade servers is available with the following fabric interconnect and IOM combinations:

- 6300 Series Fabric Interconnect + IOM 2304
- 6400 Series Fabric Interconnect + IOM 2408 from Cisco UCSM Release 4.1(2)

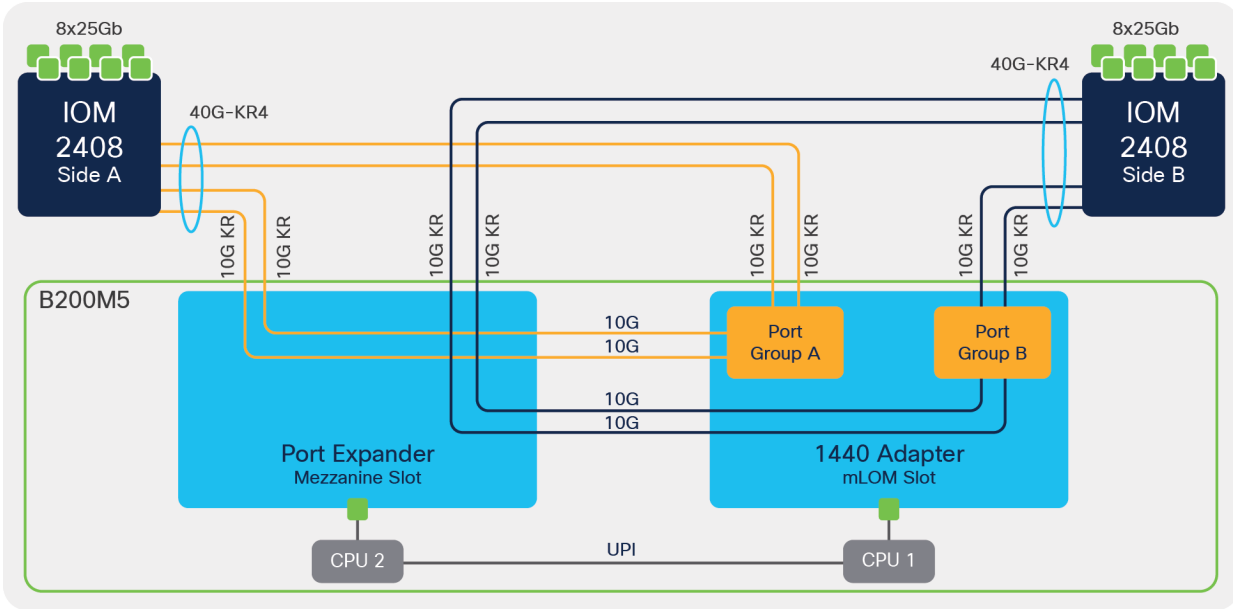


Figure 10.
B200 M5 with VIC 1440 + port expander and IOM

On the Cisco UCS 5108 chassis a B200 M5 or B480 M5 server with a VIC 1440 + port extender will enable a 40-Gbps-KR4 interface toward the IOM. In this case, the servers will see a vNIC bandwidth of 40 Gbps and each vNIC will support a maximum single flow of 40 Gbps. Figure 11 depicts the backplane connections that are enabled on a Cisco UCS 5108 chassis with a VIC 1440 and port extender in a B200-M5 chassis.

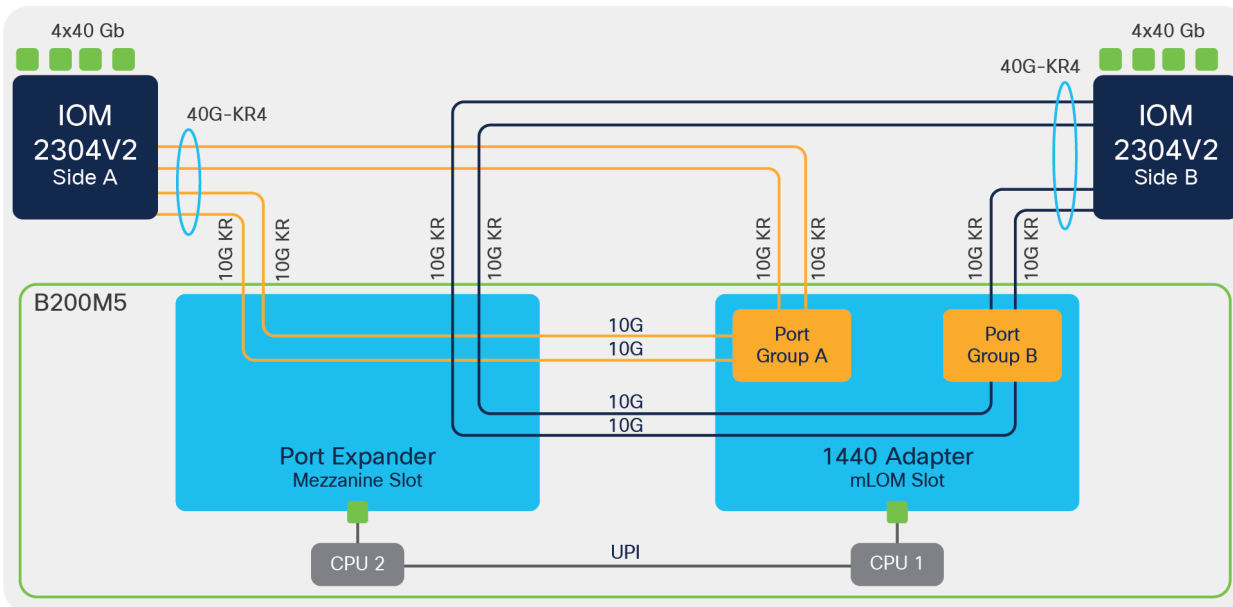


Figure 11.
B200 M5 with VIC 1440 + port extender and IOM 2304

With a port extender, the vNIC on the VIC 1440 will see a speed of 40 Gbps due to the 40-Gbps-KR4 interface that is enabled toward the IOM 2304 as shown in Figure 11.

Without a port expander, the vNIC on the 1440 will see a speed of only 20 Gbps, or 2x 10 Gbps, and the vNIC will have an aggregate bandwidth of 20 Gbps with a maximum single-flow bandwidth of 10 Gbps.

Similarly, the vNIC on the VIC 1480 with a 6300 + IOM 2304 will have a 2x 10-Gbps (20-Gbps) or native 40-Gbps speed, depending on the server on which it is placed.

- On the B200 M5 with a VIC 1440 + 1480 adapters, vNICs on both will see a speed of 2x 10 Gbps (20 Gbps).
- When a VIC 1480 is installed in slot 3 of a B480 M5 Blade Server, the vNIC on the 1480 will see a native 40-Gbps speed after the 40-Gbps-KR4 interface toward the IOM 2304 is enabled.

On a Cisco UCS 5108 chassis with IOM 2408 and a B200 M5 or B480 M5 server, 40-Gbps-KR4 for the VIC 1440 and 1480 are enabled from Cisco UCSM Release 4.1(2a).

Please note that these interfaces are 40-Gbps-KR4 interfaces from the VIC 1440 or 1480 toward the IOM 2408 and not 4x 10-Gbps hardware port-channel interfaces.

The 40-Gbps-KR4 is enabled for the VIC 1440 + port extender on the B200 M5 and B480 M5. On the VIC 1480, the 40-Gbps-KR4 is enabled when it is installed on slot 3 of a B480 M5.

Note: To avoid initial transient packet drop for a flow on the IOM 2408 due to this bandwidth mismatch of 40 Gbps from the server toward the IOM 2408 and the 25 Gigabit Ethernet connection between the IOM 2408 and the 6400 Fabric Interface, you can rate-limit the vNICs to 25 Gbps. You can achieve vNIC rate limiting through a QoS policy applied to the vNIC (Refer to Figure 12.).

Figure 12.
vNIC rate-limiting configuration

Like the VIC 1440 adapter, the vNIC on the 1480 adapter with the 6400 Fabric Interconnect and IOM 2408 can have 2x 10-Gbps (20-Gbps) or native 40-Gbps speed, depending on the server on which it is installed.

- On the B200 M5 with the 1440 + 1480 adapters, vNICs on both the 1440 and 1480 will see 2x 10-Gbps (20-Gbps) speed.
- When the 1480 is installed in slot 3 on the B480 M5, vNICs on the 1480 will see native 40-Gbps speed.

The 25-Gbps rate-limiting recommendation is applicable for vNICs on the 1480 when a 40-Gbps-KR4 interface with IOM 2408 is enabled from Cisco UCSM Release 4.1(2).

Before Cisco UCSM Release 4.1(2), there was no support for the port extender with the 1440 adapter and IOM 2408, and both the 1440 and 1480 supported only 2x 10 Gbps (20 Gbps) per vNIC on each of these adapters with the IOM 2408.

Cisco UCS B480 M5 Blade Servers

Figure 13 shows the connectivity on a Cisco UCS 5108 chassis with an IOM 2408 connected to a 6454 Fiber Interconnect. The red lines are the 25-Gbps physical connections from the IOM to the fiber interconnect, and the blue lines show the 10-Gbps-KR backplane connections from the IOM toward the VIC on a B480 M5 server.




Figure 13.
UCS B480 M5 connectivity in UCS 5108 chassis with IOM 2408

On a B480 M5, VIC 1440 and 1480 adapters can have 10-, 20-, or 40-Gbps vNICs. The vNIC speed seen on the server depends on the various VIC combinations installed on the blade server, as in the following:

- With VIC 1440 on a B480 M5, the server will see 20-Gbps (2x10) throughput with all IOMs except the IOM 2204 and Inteconnect interconnect 6324.
- The VIC 1440 with a port extender on the B480 M5 will enable 40-Gbps-KR4 toward the IOM 2304 and IOM 2408. The IOM 2304 in Cisco UCS Manager (UCSM) Releases 4.0 and later supports this solution; support for 40-Gbps-KR4 with the IOM 2408 is available starting with Cisco UCSM Release 4.1(2).

- The VIC 1480 in slot 3 of B480 M5 can have 40-Gbps throughput. The IOM 2304 in Cisco UCSM Releases 4.0 and later supports this solution. With the IOM 2408, 20-Gbps (2x10 Gbps) was supported initially and 40-Gbps supported starting at Cisco UCSM Release 4.1(2).
- The VIC 1480 adapter in slot 3 of B480 M5 can have 40-Gbps throughput. It is supported by the IOM 2304 in Cisco UCSM Releases 4.0 and later. With the IOM 2408, 20 Gbps (2x 10 Gbps) was supported initially and 40 Gbps is supported starting with Cisco UCSM Release 4.1(2). Tables 3 and 4 show the Cisco UCS B480 M5 adapter and module combinations throughput per blade server with **four** CPUs installed for **two** fabric interconnects and vNIC/vHBA speed per adapter on **each** fabric interconnect, respectively.

Table 3. B480 M5 adapter and module combinations throughput per blade server with **four** CPUs installed for **two** fabric interconnects. > **Total aggregated bandwidth for 2x fabric interconnects I/O**

Fabric Interconnect	2x FI 6400	2x FI 6300	2x FI 6200/6300/6400	2x FI 6200/6300/6400
I/OM	2x IOM 2408	2x IOM 2304	2x IOM 2208	2x IOM 2204
1340	40G	40G	40G	20G
1340 + PE	80G	80G	80G	80G
1340 + 1380 (slot3 in B480)	120G	120G	120G	60G
1340 + PE + 1380	160G	160G	160G	80G
1340 + 1380 + 1380	160G	160G	160G	80G
1440	40G	40G	40G	20G
1440 + PE	80G	80G	N/A	N/A
1440 + 1480 (slot3 in B480)	120G	120G	80G	40G
1440 + PE + 1480	160G	160G	N/A	N/A
1440 + 1480 + 1480	160G	160G	120G	60G

*N/A = Not supported

Table 4. B480 M5 adapter and module combinations vNIC/vHBA speed per adapter on **each** fabric interconnect
 > **Total aggregated bandwidth for 1x fabric interconnect**

Fabric interconnect	1x FI 6400/6536	OS vNIC speed	Max bandwidth single flow	Fabric interconnect	1x FI 6300/6536	OS vNIC speed	Max bandwidth single flow
I/OM	1x IOM 2408			I/OM	1x IOM 2304		
1340	2x 10 Gbps	20 Gbps	10 Gbps	1340	2x 10 Gbps	20 Gbps	10 Gbps
1340 + PE	4x 10 Gbps	40 Gbps	10 Gbps	1340 + PE	40 Gbps	40 Gbps	40 Gbps
1340 + 1380 (slot3 in B480)	2x 10 Gbps + 40 Gbps	20 Gbps 40 Gbps	10 Gbps 10 Gbps	1340 + 1380 (slot3 in B480)	2x 10 Gbps+ 40 Gbps	20 Gbps 40 Gbps	10 Gbps 40 Gbps
1340 + PE + 1380	40 Gbps + 40 Gbps	40 Gbps	10 Gbps	1340 + PE + 1380	40 Gbps + 40 Gbps	40 Gbps 40 Gbps	40 Gbps 40 Gbps
1340 + 1380 + 1380	2x 10 Gbps + 2x 10 Gbps + 40 Gbps	20 Gbps 20 Gbps 40 Gbps	10 Gbps 10 Gbps 10 Gbps	1340 + 1380 + 1380	2x 10 Gbps + 2x 10 Gbps + 40 Gbps	20 Gbps 20 Gbps 40 Gbps	10 Gbps 10 Gbps 40 Gbps
1440	2x 10 Gbps	20 Gbps	10 Gbps	1440	2x10 Gbps	20 Gbps	10 Gbps
1440 + PE	40 Gbps**	40 Gbps	25 Gbps	1440 + PE	40 Gbps	40 Gbps	40 Gbps
1440 + 1480 (slot3 in B480)	2x 10 Gbps + 40 Gbps**	20 Gbps 40 Gbps	10 Gbps 25 Gbps	1440 + 1480 (slot3 in B480)	2x 10 Gbps + 40 Gbps	20 Gbps 40 Gbps	10 Gbps 40 Gbps
1440 + PE + 1480	40 Gbps** + 40 Gbps**	40 Gbps 40 Gbps	25 Gbps 25 Gbps	1440 + PE + 1480	40 Gbps + 40 Gbps	40 Gbps 40 Gbps	40 Gbps 40 Gbps
1440 + 1480 + 1480	2x 10 Gbps + 2x 10 Gbps + 40 Gbps**	20 Gbps 20 Gbps 40 Gbps	10 Gbps 10 Gbps 25 Gbps	1440 + 1480 + 1480	2x 10 Gbps + 2x 10 Gbps + 40 Gbps	20 Gbps 20 Gbps 40 Gbps	10 Gbps 10 Gbps 40 Gbps

** Supported from Cisco UCSM Release 4.1(2), single-flow maximum of 25 Gbps

Highlighted area is the suggested configuration; it was used for testing with 2x 6400 Fabric Interconnects and 2x IOM 2408.

Fabric interconnect	1x FI 6200/6300/6400	OS vNIC speed	Max bandwidth single flow	Fabric interconnect	1x FI 6200/6300/6400	OS vNIC speed	Max bandwidth single flow
I/OM	1x IOM 2208			I/OM	1x IOM 2204		
1340	2x 10 Gbps	20 Gbps	10 Gbps	1340	10 Gbps	10 Gbps	10 Gbps
1340 + PE	4x 10 Gbps	40 Gbps	10 Gbps	1340 + PE	2x 10 Gbps	20 Gbps	10 Gbps
1340 + 1380 (slot3 in B480)	2x 10 Gbps + 4x 10 Gbps	20 Gbps 40 Gbps	10 Gbps 10 Gbps	1340 + 1380 (slot3 in B480)	10 Gbps + 2x 10 Gbps	10 Gbps 20 Gbps	10 Gbps 10 Gbps
1340 + PE + 1380	4x 10 Gbps + 4x 10 Gbps	40 Gbps 40 Gbps	10 Gbps 10 Gbps	1340 + PE + 1380	2x 10 Gbps + 2x 10 Gbps	20 Gbps 20 Gbps	10 Gbps 10 Gbps
1340 + 1380 + 1380	2x 10 Gbps + 2x 10 Gbps + 40 Gbps	20 Gbps 20 Gbps 40 Gbps	10 Gbps 10 Gbps 10 Gbps	1340 + 1380 + 1380	1x 10 Gbps + 1x 10 Gbps 2x 10 Gbps	10 Gbps 10 Gbps 20 Gbps	10 Gbps 10 Gbps 10 Gbps
1440	2x 10 Gbps	20 Gbps	10 Gbps	1440	10 Gbps	10 Gbps	10 Gbps
1440 + PE	2x 10 Gbps	20 Gbps	10 Gbps	1440 + PE	10 Gbps	10 Gbps	10 Gbps
1440 + 1480 (slot3 in B480)	2x 10 Gbps + 2x 10 Gbps	20 Gbps 20 Gbps	10 Gbps 10 Gbps	1440 + 1480 (slot3 in B480)	10 Gbps + 10 Gbps	10 Gbps 10 Gbps	10 Gbps 10 Gbps
1440 + PE + 1480	2x 10 Gbps + 2x 10 Gbps	20 Gbps 20 Gbps	10 Gbps 10 Gbps	1440 + PE + 1480	10 Gbps + 10 Gbps	10 Gbps 10 Gbps	10 Gbps 10 Gbps
1440 + 1480 + 1480	2x 10 Gbps + 2x 10 Gbps + 2x 10 Gbps	20 Gbps 20 Gbps 20 Gbps	10 Gbps 10 Gbps 10 Gbps	1440 + 1480 + 1480	10 Gbps + 10 Gbps + 10 Gbps	10 Gbps 10 Gbps 10 Gbps	10 Gbps 10 Gbps 10 Gbps

One vHBA is defined for each fabric interconnect side A and side B.

Note: ESXi will display lower speeds, but bandwidth listed in the tables above are achievable.

Note: You cannot mix VIC 1440 cards with 1380 cards or VIC 1340 cards with 1480 cards.

On the Cisco UCS 5108 chassis, the B200 M5 or B480 M5 server with a VIC 1440 + port extender will enable 40-Gbps-KR4 interface toward the IOM. In this case the servers will see vNIC bandwidth of 40 Gbps and each vNIC will support a maximum single flow of 40 Gbps. Figure 14 depicts the backplane connections that are enabled on a Cisco UCS 5108 chassis with a VIC 1440 adapter and port extender in a B200 M5 chassis.

Fibre Channel Policies

Three policies are associated with the Fibre Channel adapter in a Cisco UCS server:

1. Fibre Channel Adapter
2. Fibre Channel Network
3. Storage-area network (SAN) connectivity

The SAN connectivity policy combines three policies to configure and place the vHBAs on the server. The Fibre Channel Network Policy controls the VLAN/VSAN mapping and has no effect on Fibre channel throughput.

The Fibre Channel Adapter policy offers the ability to set various options that can affect Fibre Channel throughput.

Many of the settings in the Fibre Channel Adapter policy are used to set Fibre Channel login (flogi) values, as well as determine the number of logical unit numbers (LUNs) to which the vHBA can attach. Figure 16 shows the Fibre Channel Adapter policy configuration.

SAN / Policies / root / Fibre Channel Adapter Policies / FC Adapter Policy default

General Events

Actions

Delete

Show Policy Usage

Use Global

Properties

Name : **default**

Description : default adapter policy

Owner : **Local**

Resources

Transmit Queues : **1**

Ring Size : [64-128]

Receive Queues : **1**

Ring Size : [64-2048]

I/O Queues : [1-64]

Ring Size : [64-512]

Options

FCP Error Recovery : Disabled Enabled

Flogi Retries : [0-infinite]

Flogi Timeout (ms) : [1000-255000]

Plogi Retries : [0-255]

Plogi Timeout (ms) : [1000-255000]

Error Detect Timeout (ms) : [0-255000]

Port Down Timeout (ms) : [0-240000]

IO Retry Timeout (seconds) : [1-59]

Port Down IO Retry : [0-255]

Link Down Timeout (ms) : [0-240000]

Resource Allocation Timeout (ms) : [0-255000]

IO Throttle Count : [256-1024]

Max LUNs Per Target : [1-1024]

LUN Queue Depth : [1-254]

Interrupt Mode : MSI X MSI IN Tx

vHBA Type : FC Initiator FC Target FC NVME Initiator FC NVME Target

Figure 16.
Fibre Channel Adapter policy configuration

Resources

- Transmit ring size: This value does not affect adapter performance. The default value is 64.
- Receive ring size: This value does not affect adapter performance. The default value is 64.
- I/O queues ring size: The number of descriptors in each I/O queue, this value can affect the performance of the adapter, so we recommend that you do not change the default value which is in the above screenshot.

Options

- I/O throttle count: This count is the maximum number of data or control I/O operations that can be pending in the vHBA at one time. If this value is exceeded, the additional I/O operations wait in the queue until the number of pending I/O operations decreases and the additional operations can be processed. The default is 256. We recommend you consult your storage array documentation for the optimal value for this parameter.
- LUN queue depth: This parameter is the number of commands that the HBA can send and receive in a single transmission per LUN. The default LUN queue depth is 20.
- Interrupt mode: This mode describes the method used to send interrupts to the operating system from the driver. Message Signal Interrupt (MSI)-x mode is the recommended option.

Note: VMware recommends that you use Payment Card Industry (PCI) functions for VMDirectPath I/O support for MSI or MSI-X interrupts. Please refer to [vSphere VMDirectPath I/O and Dynamic DirectPath I/O: Requirements for Platforms and Devices \(2142307\)](#)

Performance testing

Hardware

- 1x 5108 chassis
- 2x 2408 IOMs
- 2x 6454 Fabric Interconnects
- 8x 25-Gbps cables to each IOM and fabric interconnect
- 1x B480 M5 with VIC 1440 adapter + port extender + VIC 1480 adapter
- NetApp connected through Fibre Channel

Software

- Windows Server 2022 Datacenter
- Red Hat Enterprise Linux 8.1
- iPerf 2

Figures 17 and 18 give information about the firmware Version 4.2(3b) in Cisco UCS Manager used for testing and the VIC 1440, port extender, and VIC 1480 in Cisco UCS Manager.

Equipment / Chassis / Chassis 1 / Servers / Server 3								
General Inventory Virtual Machines Installed Firmware CIMC Sessions SEL Logs VIF Paths Health Diagnostics Faults Events FSM Statistics Temperatures Power								
+ - Advanced Filter Export Print Update Firmware Activate Firmware Capability Catalog								
Name	Model	Package Version	Running Version	Startup Version	Backup Version	Update Status	Activate Status	
▼ Adapters								
Adapter 1	Cisco UCS VIC 1440	4.2(3b)B	5.2(3c)	5.2(3c)	5.2(2e)	Ready	Ready	
Adapter 3	Cisco UCS VIC 1480	4.2(3b)B	5.2(3c)	5.2(3c)	5.2(2e)	Ready	Ready	
Persistent Memory								
BIOS	Cisco UCS B480 M5 4 Socket Bla...	4.2(3b)B	B480M5.4.2.3b.0.1016222226	B480M5.4.2.3b.0.1016222226	B480M5.4.2.2d.0.0809221351	Ready	Ready	
Board Controller	Cisco UCS B480 M5 4 Socket Bla...	4.2(3b)B	19.0	19.0	N/A	N/A	Ready	
CIMC Controller	Cisco UCS B480 M5 4 Socket Bla...	4.2(3b)B	4.2(3b)	4.2(3b)	4.2(2e)	Ready	Ready	
▶ Storage Controller PCH 1	Lewisburg SSATA Controller [AHC...				N/A	N/A		
▶ Storage Controller SAS 1	Cisco FlexStorage 12G SAS RAID ...	4.2(3b)B	24.21.0-0156[6.36.00.3]NA	24.21.0-0156[6.36.00.3]NA	N/A	N/A	Ready	
▶ Storage Controller SAS 2	Cisco FlexStorage PCIe SSD/Pass...				N/A	N/A		

Figure 17.
Firmware Version 4.2(3b) in Cisco UCS manager used for testing

Equipment / Chassis / Chassis 1 / Servers / Server 3 / Adapters / Adapter 1

General DCE Interfaces NICs HBAs iSCSI vNICs Faults Events

Fault Summary

0 0 0 0

Status

Overall Status: **Operable**

Actions

Update Firmware
Activate Firmware
View POST Results

Properties

ID : 1
Product Name : **Cisco UCS VIC 1440**
Vendor : Cisco Systems Inc PID : UCSB-MLOM-40G-04
Revision : 0 Serial : FCH22437XJY
PCI Slot : N/A Integrated : No

Part Details

Port Expander Card

ID : 2
Product Name : **Cisco UCS Port Expander Card**
Vendor : Cisco Systems Inc PID : UCSB-MLOM-PT-01
Revision : 0 Serial : FCH22337QJY

Part Details

Firmware

Running Version : **5.2(3c)**
Package Version : **4.2(3b)B**
Backup Version : **5.2(2e)**
Update Status : **Ready**
Startup Version : **5.2(3c)**
Activate Status : **Ready**

Equipment / Chassis / Chassis 1 / Servers / Server 3 / Adapters / Adapter 3

General DCE Interfaces NICs HBAs iSCSI vNICs Faults Events

Fault Summary

0 6 0 0

Status

Overall Status: **Operable**

Actions

Update Firmware
Activate Firmware
View POST Results

Properties

ID : 3
Product Name : **Cisco UCS VIC 1480**
Vendor : Cisco Systems Inc PID : UCSB-VIC-M84-4P
Revision : 0 Serial : FCH22337UT1
PCI Slot : N/A Integrated : No

Part Details

Firmware

Running Version : **5.2(3c)**
Package Version : **4.2(3b)B**
Backup Version : **5.2(2e)**
Update Status : **Ready**
Startup Version : **5.2(3c)**
Activate Status : **Ready**

Figure 18.
Cisco VIC 1440, port extender, and 1480 in Cisco UCS manager

Policies / root / Adapter Policies / Eth Adapter Policy Win-HPN

General Events

Actions: Delete, Show Policy Usage, Use Global

Properties: Name: Win-HPN, Description: Recommended adapter settings for Windows high p, Owner: Local

Resources:

Pooled: Disabled Enabled

Transmit Queues: 1 [1-1000]

Ring Size: 256 [64-4096]

Receive Queues: 4 [1-1000]

Ring Size: 512 [64-4096]

Completion Queues: 5 [1-2000]

Interrupts: 512 [1-1024]

Options:

Transmit Checksum Offload: Disabled Enabled

Receive Checksum Offload: Disabled Enabled

TCP Segmentation Offload: Disabled Enabled

TCP Large Receive Offload: Disabled Enabled

Receive Side Scaling (RSS): Disabled Enabled

Accelerated Receive Flow Steering: Disabled Enabled

Network Virtualization using Generic Routing Encapsulation: Disabled Enabled

Virtual Extensible LAN: Disabled Enabled

GENEVE: Disabled Enabled

AzureStack-Host QoS: Disabled Enabled

Fallback Timeout (Seconds): 5 [0-600]

Interrupt Mode: MSI X MSI IN Tx

Interrupt Coalescing Type: Min Idle

Interrupt Timer (us): 125 [0-65535]

RoCE: Disabled Enabled

Advance Filter: Disabled Enabled

Interrupt Scaling: Disabled Enabled

Policies / root / Adapter Policies / FC Adapter Policy Windows

General Events

Actions: Delete, Show Policy Usage, Use Global

Properties: Name: Windows, Description: Recommended adapter settings for Windows, Owner: Local

Resources:

Transmit Queues: 1

Ring Size: 64 [64-128]

Receive Queues: 1

Ring Size: 64 [64-2048]

I/O Queues: 1 [1-64]

Ring Size: 512 [64-512]

Options:

FCP Error Recovery: Disabled Enabled

Flogi Retries: 8 [0-Infinite]

Flogi Timeout (ms): 4000 [1000-255000]

Plugi Retries: 8 [0-255]

Plugi Timeout (ms): 20000 [1000-255000]

Error Detect Timeout (ms): 2000

Port Down Timeout (ms): 30000 [0-240000]

IO Retry Timeout (seconds): 5 [1-59]

Port Down IO Retry: 30 [0-255]

Link Down Timeout (ms): 30000 [0-240000]

Resource Allocation Timeout (ms): 10000

IO Throttle Count: 256 [256-1024]

Max LUNs Per Target: 1024 [1-1024]

LUN Queue Depth: 20 [1-254]

Interrupt Mode: MSI X MSI IN Tx

vHBA Type: FC Initiator FC Target FC NVME Initiator FC NVME Target

Figure 19. Ethernet vNIC and Fibre Channel adapter policy used for Windows.

Note: Each vHBA and vNIC should have a policy attached to it according to the OS that is running. Failure to select an adapter policy will result in slower performance speeds.

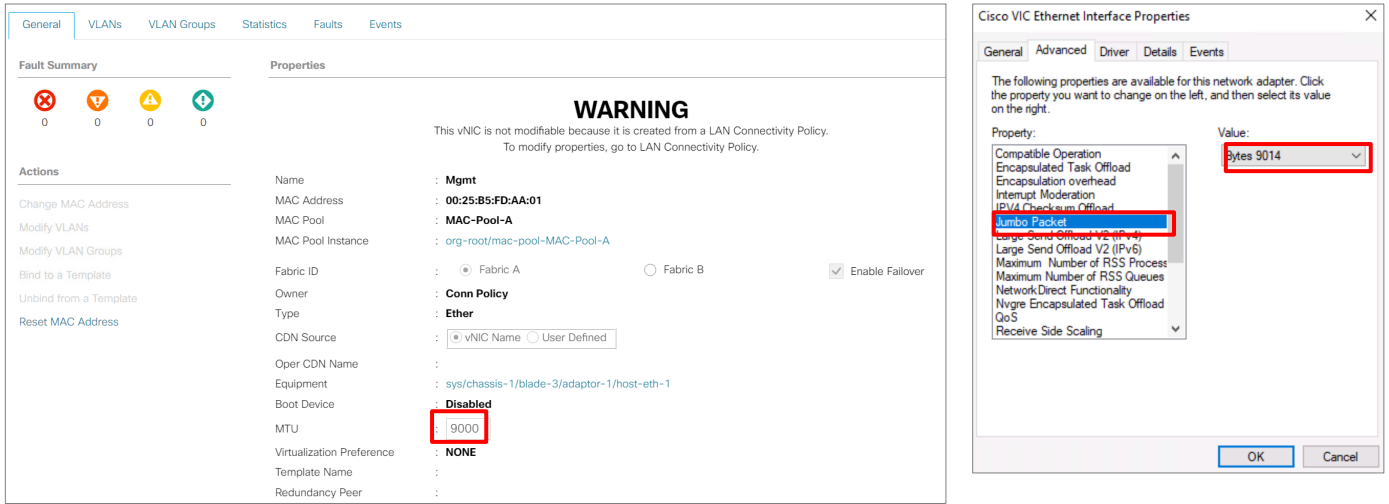


Figure 20. MTU on the adapter policy changed to 9000 and Windows OS jumbo packet to 9014

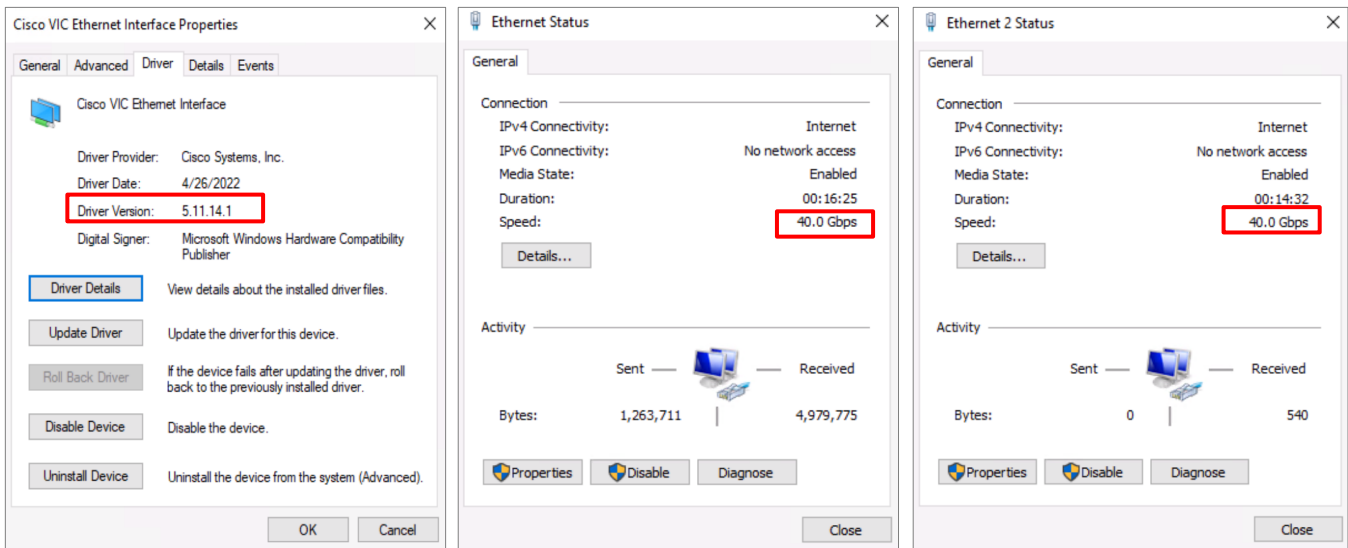


Figure 21. VIC driver used in Windows from Release 4.2.3a, VIC 1440 + port expander speed in Windows, and VIC 1480 speed in Windows

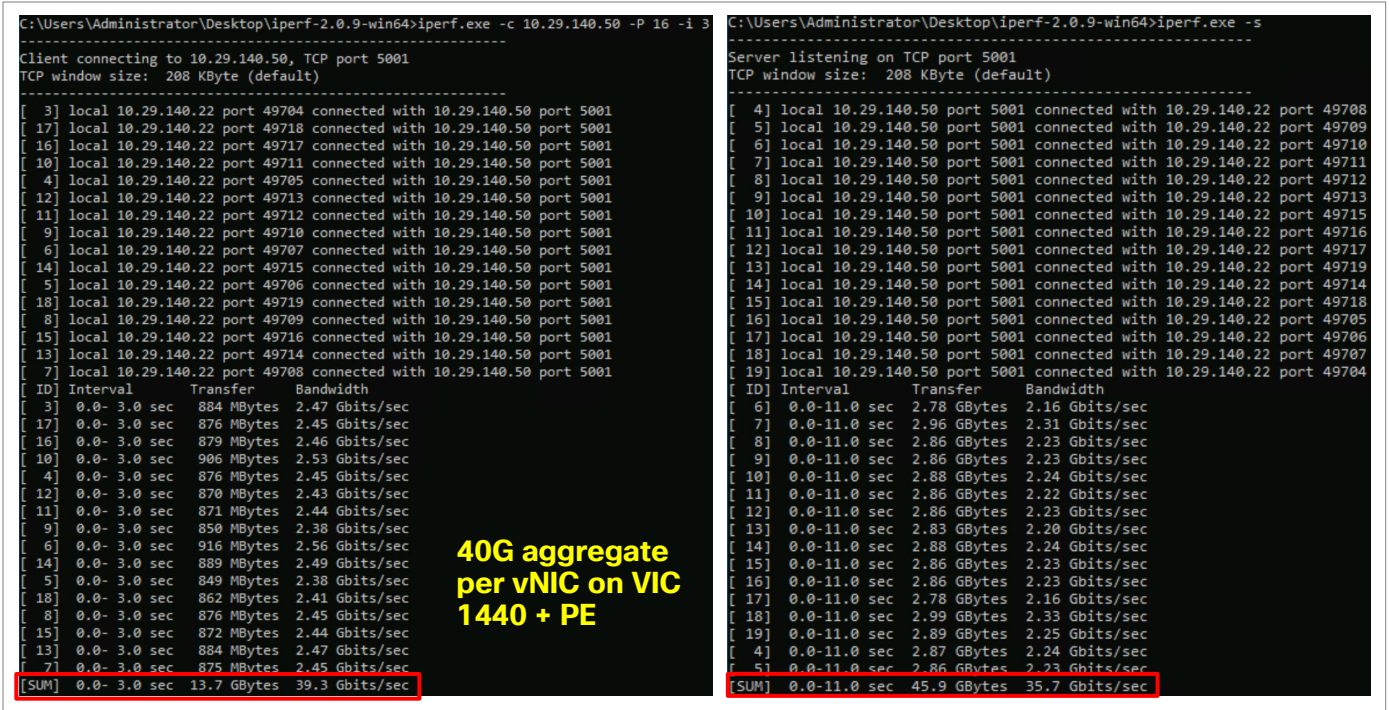


Figure 22.
iperf results from command line

Inside the same chassis with the vLAN subnet using iPerf bare metal-to-bare metal, the B480 M5 can receive 40 Gbps from another server with the same vNIC configuration from the previous figures. Both servers were running Windows Server 2022 Datacenter.

Note: If you are running iPerf in a single stream, the B480 M5 will show a maximum speed of only 10 Gbps. You should run the test in multiple streams to push maximum bandwidth. We recommend that you use this command: iperf.exe -c X.X.X.X -P 16 -i 3. (-P=parallel, 16=connections, -i=interval, 3=seconds).

Note: To achieve maximum bandwidth, it is best to test bare metal-to-bare metal because virtual machine-to-virtual machine tests will have some bandwidth restrictions and yield slower speeds.

Policies / root / Adapter Policies / Eth Adapter Policy Linux

General Events

Actions

Delete

Show Policy Usage

Use Global

Properties

Name: **Linux**

Description: Recommended adapter settings for linux

Owner: Local

Resources

Pooled: Disabled Enabled

Transmit Queues: 1 [1-1000]

Ring Size: 256 [64-4096]

Receive Queues: 1 [1-1000]

Ring Size: 512 [64-4096]

Completion Queues: 2 [1-2000]

Interrupts: 4 [1-1024]

Options

Transmit Checksum Offload: Disabled Enabled

Receive Checksum Offload: Disabled Enabled

TCP Segmentation Offload: Disabled Enabled

TCP Large Receive Offload: Disabled Enabled

Receive Side Scaling (RSS): Disabled Enabled

Accelerated Receive Flow Steering: Disabled Enabled

Network Virtualization using Generic Routing Encapsulation: Disabled Enabled

Virtual Extensible LAN: Disabled Enabled

GENEVE: Disabled Enabled

AzureStack-Host QoS: Disabled Enabled

Failback Timeout (Seconds): 5 [0-600]

Interrupt Mode: MSI X MSI IN Tx

Interrupt Coalescing Type: Min Idle

Interrupt Timer (us): 125 [0-65535]

RoCE: Disabled Enabled

Advance Filter: Disabled Enabled

Interrupt Scaling: Disabled Enabled

Policies / root / Adapter Policies / FC Adapter Policy Linux

General Events

Actions

Delete

Show Policy Usage

Use Global

Properties

Name: **Linux**

Description: Recommended adapter settings for Linux

Owner: Local

Resources

Transmit Queues: 1

Ring Size: 64 [64-128]

Receive Queues: 1

Ring Size: 64 [64-2048]

I/O Queues: 1 [1-64]

Ring Size: 512 [64-512]

Options

FCP Error Recovery: Disabled Enabled

Flogi Retries: 8 [0-infinite]

Flogi Timeout (ms): 4000 [1000-255000]

Plugi Retries: 8 [0-255]

Plugi Timeout (ms): 20000 [1000-255000]

Error Detect Timeout (ms): 2000

Port Down Timeout (ms): 30000 [0-240000]

IO Retry Timeout (seconds): 5 [1-50]

Port Down IO Retry: 30 [0-255]

Link Down Timeout (ms): 30000 [0-240000]

Resource Allocation Timeout (ms): 10000

IO Throttle Count: 256 [256-1024]

Max LUNs Per Target: 1024 [1-1024]

LUN Queue Depth: 20 [1-254]

Interrupt Mode: MSI X MSI IN Tx

vHBA Type: FC Initiator FC Target FC NVME Initiator FC NVME Target

Figure 23. Ethernet vNIC and Fibre Channel Adapter policy used for Linux

Note: Each vHBA and vNIC should have a policy attached to it according to the OS you are running. Failure to select an adapter policy will result in slower performance speeds.

Figure 24. MTU on the adapter policy changed to 9000 and Red Hat OS jumbo packet to 9000

Figure 25. VIC driver used in Red Hat from Release 4.2.3a, VIC 1440 + port extender (top) and VIC 1480 (bottom) speed 40 Gbps in Red Hat

Specification Sheets

- [Cisco UCS B200 M5 Blade Server](#)
- [Cisco UCS B480 M5 Blade Server](#)
- [Cisco UCS 6400 Series Fabric Interconnects](#)

White Papers

- [Cisco UCS VIC 1400 Series Best Practices in Ethernet Fabric White Paper](#)
- [Performance Tuning Guide for Cisco UCS M5 Servers](#)

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 <https://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: <https://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)