

Cisco UCS fNIC Tunables

September 2024

Contents

Overview	3
Audience	3
Test environments	3
fNIC tunable parameters	4
Install the fNIC and NFNIC drivers	4
Display the fNIC driver version	5
Configure the fNIC tunable parameters	5
LUN Queue Depth	5
Configuring I/O Throttle Count	12
For more information	17

Overview

Cisco Unified Computing System™ (Cisco UCS®) allows you to tune the Fibre Channel network interface card (fNIC) Logical Unit Number (LUN) Queue Depth and I/O Throttle Count parameters of the Cisco UCS Virtual Interface Card (VIC) fNIC driver in Linux, VMware ESX, and Microsoft Windows implementations. This document provides an overview of these two parameters and the methodologies and syntax for modifying their values. Both parameters are currently set to best-practice defaults, and these values are still recommended for most architectures. The capability to tune these parameters provides flexibility for those customers whose architectures require nondefault values.

Audience

This document is intended for Cisco® systems engineers and customers involved in systems administration and performance engineering on Cisco UCS Linux, VMware ESX, and Windows implementations. It assumes advanced knowledge and understanding of operating system configurations in the context of storage technologies.

Test environments

The test environments for the solution described in this document include the following components:

- Cisco UCS Manager Release 4.2.1(f)
 - Two Cisco UCS 6454 Fabric Interconnects
 - Two Cisco UCS 2408 I/O Modules
 - Cisco UCS 5108 Blade Server Chassis
 - Cisco UCS B200 M5 Blade Server with Cisco UCS VIC 1440 and Port Expander modular LAN on motherboard (mLOM)
 - Cisco UCS B200 M6 Blade Server with Cisco UCS VIC 1440 and 1480 mLOM
- Cisco Intersight™ Managed Mode Release 4.2.1(f)
 - Two Cisco UCS 6454 Fabric Interconnects
 - Two Cisco UCS 9108 Intelligent Fabric Modules
 - Cisco UCS X9508 Chassis
 - Cisco UCS X210 M6 Compute Node with Cisco UCS VIC 14425 and Port Expander mLOM
- Cisco Integrated Management Controller Release 4.2.1(f)
 - Cisco UCS C220 M6S with Cisco UCS VIC 1467 mLOM

The following fNIC drivers were tested:

- Citrix Xen 8.2 LTSR: fNIC Version 2.0.0.72
- Red Hat Enterprise Linux (RHEL) 7.9 and 8.5: fNIC Version 2.0.0.85
- SUSE Linux Enterprise Server (SLES) 12.5 and 15.3: fNIC Version 2.0.0.72
- VMware ESX 6.7U3 and 7.0U3: NFNIC Version 5.0.0.15
- Microsoft Windows Server 2022: fNIC Version 4.0.0.1

fNIC tunable parameters

The two fNIC tunable parameters are LUN Queue Depth and I/O Throttle Count. Definitions for each of these parameters are as follows:

- LUN Queue Depth: The total number of I/O requests that can be outstanding per LUN
- I/O Throttle Count: The total number of I/O requests that can be outstanding per virtual Host Bus Adapter (vHBA)

Install the fNIC and NFNIC drivers

This guide assumes that the supported fNIC and NFNIC drivers have been installed and are running. Refer to the Cisco UCS VIC installation and upgrade guides for complete driver installation instructions:

<https://www.cisco.com/c/en/us/support/servers-unified-computing/ucs-virtual-interface-card/products-installation-guides-list.html>.

In addition, review the Cisco UCS Hardware Compatibility List (HCL) to confirm hardware, operating system, and driver compatibility: <https://ucshcltool.cloudapps.cisco.com/public/>.

Note: Citrix XenServer uses Driver Update Disks (DuDs) to update hardware drivers using Citrix XE Command-Line Interface (CLI) commands such as **update-upload** and **update-apply**. Cisco provides Citrix with Cisco VIC fNIC driver RPMs. Citrix then uses these packages to create DuDs and posts these ISO files on Citrix.com. Cisco includes VIC fNIC drivers on the driver ISO files for each Cisco UCS release. However, you should install these RPM packages only in a test or development environment, and you should use the official Citrix DuD ISO file to update Cisco VIC fNIC drivers in the production environment.

Display the fNIC driver version

Use the following commands to display the current fNIC driver version:

- ESX 6.7 and 7.x:

```
# vmkload_mod -s nfnic |grep Version
```

Or:

```
# esxcli software vib list |grep nfnic
```

- RHEL 7.9 and 8.x, SLES 12.5 and 15.x, and XS 8.x

```
# /sbin/modinfo fnic
```

Alternatively, use `systool` (install the Linux package `sysfsutils` if it is not already installed); run the following command:

```
# systool -vm fnic
```

- Windows

From Microsoft Windows PowerShell, run the following command:

```
> Get-WmiObject Win32_PnPSignedDriver | select devicename, driverversion | where  
{$_ .devicename -like "*Storport*"}
```

Alternatively, in the Windows GUI, navigate to Computer Management > Device Manager > Storage Controllers, right-click Cisco VIC-FCoE Storport Miniport, open the Driver tab, and review the displayed driver details.

Configure the fNIC tunable parameters

The two configurable fNIC parameters are LUN Queue Depth and I/O Throttle Count.

LUN Queue Depth

Description: The total number of I/O requests that can be outstanding per LUN

Note: The Cisco VIC LUN Queue Depth parameter listed in the Cisco Integrated Management Controller (IMC), Cisco UCS Manager, and Cisco Intersight Fibre Channel adapter policies is applicable only to the Windows operating system. All other operating systems require you to modify the LUN Queue Depth parameter at the operating system command line.

Parameter name: `fnic_max_qdepth`

Default value (with updated asynchronous driver):

- RHEL 7.9 and 8.x: 256
- SLES 12.5 and 15.x: 256
- ESX 6.7 and 7.x: 32
- Windows: 20 to 255 (dynamic)
- XS 8.x 256

Configuration capabilities:

- Boot time
 - A reboot is required for changes to take effect.
 - Changes are persistent across reboots.
- Load time
 - Configuration is disruptive to SAN-attached storage.
 - Configuration requires you to stop SAN I/O, remove fNIC module dependencies, and remove and then reload the fNIC module.
 - Configuration is not possible with the boot-from-SAN configuration.
 - Changes are not persistent across reboots.
- Run time
 - Configuration is nondisruptive.
 - Changes apply only to newly discovered LUNs after modification.
 - Changes are not persistent across reboots.

Table 1 lists the boot configuration capabilities of the supported operating systems.

Table 1. Local-boot and boot-from-SAN configuration capabilities by operating system (X signifies available capability)

	Boot time	Load time	Run time
RHEL: Local boot	X	X	X
RHEL: Boot from SAN	X		X
SLES: Local boot	X	X	X
SLES: Boot from SAN	X		X
XS: Local boot	X	X	X
XS: Boot from SAN	X		X
ESX: Local boot	X		
ESX: Boot from SAN	X		
Windows Server: Local boot	X		
Windows Server: Boot from SAN	X		

Display the fnic_max_qdepth parameter value

You can display the current and post configuration values of the fnic_max_qdepth parameter using the commands shown here.

ESX 6.7 and 7.x

From the ESX CLI, run the following command (the fnic_max_qdepth parameter value will not be listed until it has been explicitly set according to the configuration instructions here):

```
# esxcli system module parameters list -m nfnic
```

RHEL 7.9 and 8.x, SLES 12.5 and 15.x, and XS 8.x

To display the current fnic_max_qdepth value, run the following command:

```
# cat /sys/module/fnic/parameters/fnic_max_qdepth
```

Alternatively, use systool (install the Linux package sysfsutils if it is not already installed) and run the following command:

```
# systool -vm fnic
```

To display the current fnic_max_qdepth value on a per-LUN basis, which is relevant if the value was changed using the run-time configuration for newly discovered LUNs, run the following command at the CLI (install lsscsi if it is not already installed):

```
# lsscsi -l (Not available for XS 8.x)
```

Windows

You can display Cisco VIC fNIC parameters using the Cisco command-line tool for Windows called fctool.exe. This tool has traditionally been available only from the Cisco Technical Assistance Center (TAC) as a troubleshooting tool. Cisco currently does not support this tool, but you can download and use it at your own risk at the following link: <https://community.cisco.com/t5/unified-computing-system/cisco-vic-fnic-fctool-utility-for-windows/ta-p/4663692>

The fctool.exe tool requires the Microsoft Visual C++ redistributable VCRUNTIME140.dll dynamic link library, which is available from Microsoft within the vc_redist.x64.exe package. When fctool.exe is run, if there is a missing DLL file error, please download and install the latest Microsoft Visual C++ Redistributable from this link: <https://docs.microsoft.com/en-US/cpp/windows/latest-supported-vc-redist?view=msvc-170>

To display the fNIC LUN Queue Depth value, at a Windows CMD prompt from which fctool.exe is accessible, run the following command to list the fNICs and their port numbers:

```
> fctool -list
```

```
fnic [04] --> vnic [16] maca [0025b5fa0008]
  node_wnn 20000025b5ff0008 port_wnn 20000025b5fa0008
  maxdatafieldsize 2112 edtov 2000 ratov 10000
  drv fnic2k12.sys (fre) version 4.0.0.1
  svc fnic2k12
  PCI bus 27 slot 32
```

```
fnic [05] --> vnic [17] maca [0025b5fb0008]
node_wwn 20000025b5ff0008 port_wwn 20000025b5fb0008
maxdatafieldsize 2112 edtov 2000 ratov 10000
drv fnic2k12.sys (fre) version 4.0.0.1
svc fnic2k12
PCI bus 27 slot 64
```

The port number for fnic[04] is 4 and the port number for fnic[05] is 5. The number in the bracket following fnic denotes the port number. To display the LUN Queue Depth value of the fNIC at port 4, run the following command:

```
> fctool -p 4 -res
fnic --> fnic resources
node_wwn: 20000025b5ff0008
port_wwn: 20000025b5fa0008
.
.
LUN Queue Depth: 128
```

Boot-time configuration

The example here shows commands for setting fnic_max_qdepth to 128. (The value 128 is only an example; refer to your storage array manufacturer documentation and best practices for specific values.)

RHEL 7.9 and 8.x, and SLES 12.5 and 15.x

Follow these steps:

1. Create or edit the file `/etc/modprobe.d/fnic.conf`
2. Add the following line:

```
options fnic fnic_max_qdepth=128
```

3. Save `/etc/modprobe.d/fnic.conf`
4. Rebuild initramfs by running one of the following commands:

```
# dracut -f -v
```

Or:

```
# dracut -v -f /boot/initramfs-`uname -r`.img `uname -r` (Not available for SLES 15.1)
```

5. Reboot to make the change take effect.

ESX 6.7 and 7.x

At the ESX CLI, enter the following command:

```
# esxcli system module parameters set -m nfnic -p lun_queue_depth_per_path=128
```


XS 8.x

Citrix XenServer uses DuDs to update hardware drivers using Citrix XE CLI commands such as `update-upload` and `update-apply`. Cisco provides Citrix with Cisco VIC fNIC driver RPM packages. Citrix then uses these packages to create DuDs and posts these ISO files on Citrix.com. Cisco includes VIC fNIC drivers on the driver ISO files for each Cisco UCS release. However, you should install these RPM packages only in a test or development environment, and you should use the official Citrix DuD ISO file to update Cisco VIC fNIC drivers in the production environment.

1. Create or edit the file `/etc/modprobe.d/fnic.conf`

2. Add the following line:

```
options fnic fnic_max_qdepth=128
```

3. Save `/etc/modprobe.d/fnic.conf`

4. Rebuild initramfs by running the following command:

```
# dracut --include --force /etc/modprobe.d/fnic.conf /etc/modprobe.d/fnic.conf  
/boot/initrd-`uname -r`.img
```

Windows

See Cisco Intersight, Cisco UCS Manager, or Cisco IMC LUN Queue Depth configuration for Microsoft Windows sections below.

Load-time configuration

The example here shows commands for setting `fnic_max_qdepth` to 128. (The value 128 is only an example; refer to your storage array manufacturer documentation and best practices for specific values.)

RHEL 7.9 and 8.x, SLES 12.5 and 15.x, and XS 8.x

Follow these steps:

1. Unload the fNIC driver:

```
# modprobe -r fnic
```

2. Load the fNIC driver with the modified `fnic_max_qdepth` parameter:

```
# modprobe fnic fnic_max_qdepth=128
```

Run-time configuration

The example here shows commands for setting `fnic_max_qdepth` to 128. (The value 128 is only an example; refer to your storage array manufacturer documentation and best practices for specific values.)

RHEL 7.9 and 8.x, SLES 12.5 and 15.x, and XS 8.x

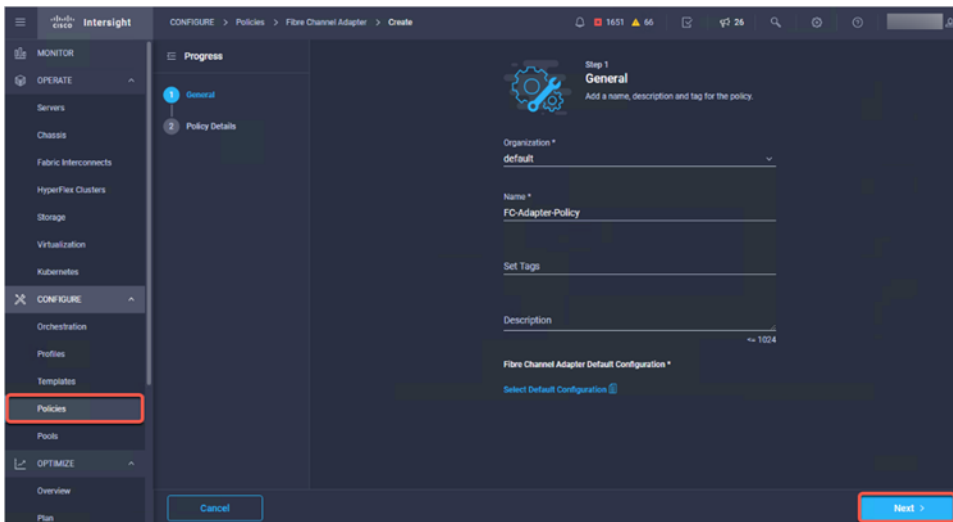
Configure `fnic_max_qdepth` with `sysfs` entry with the following command at the CLI:

```
# echo 128 > /sys/module/fnic/parameters/fnic_max_qdepth
```

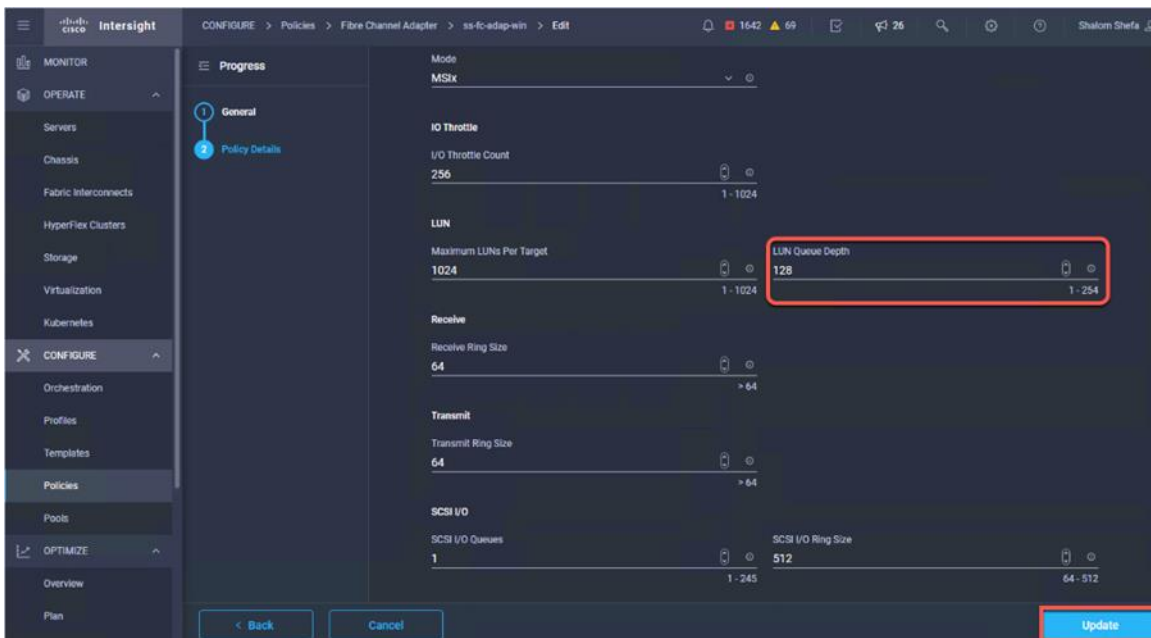
Cisco Intersight LUN Queue Depth configuration for Microsoft Windows

You can configure the LUN Queue Depth parameter in the Cisco Intersight platform through <https://www.intersight.com>. The Cisco Intersight Fibre Channel adapter policy LUN Queue Depth parameter is applicable to Windows only.

1. Log in to Cisco Intersight and choose Configure > Policies > Create Policy. In the Select Policy Type panel, click Fibre Channel Adapter. Then click Start. Configure the name and other fields on the General configuration screen. Then click Next.



2. Configure the LUN Queue Depth field with the desired value. Then click the Create button (or the Update button if you are editing an existing policy).



3. Add the Fibre Channel adapter policy to the SAN connectivity policy of the Cisco UCS server profile and deploy or redeploy the server profile.

Cisco UCS Manager LUN Queue Depth configuration for Microsoft Windows

The example here shows how to set the UCS Manager FC Adapter Policy Windows LUN Queue Depth parameter to 128. (The value 128 is only an example; refer to your storage array manufacturer documentation and best practices for specific values.)

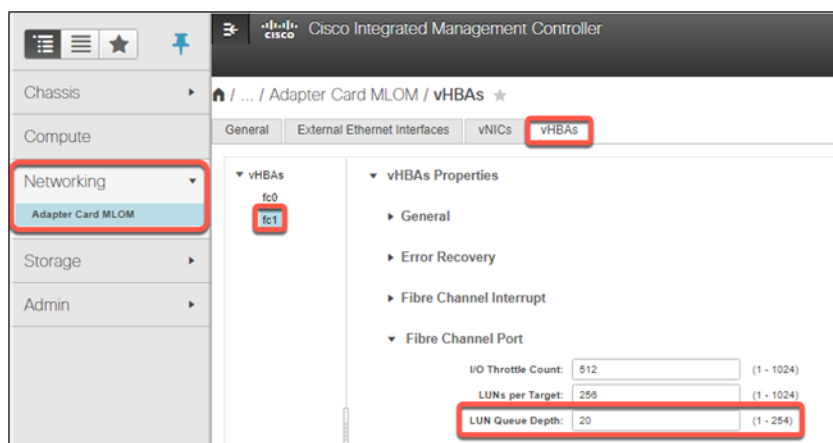
The screenshot displays the Cisco UCS Manager web interface. The breadcrumb navigation at the top reads: Servers / Policies / root / Adapter Policies / FC Adapter Policy Windows. The left-hand navigation pane shows a list of policies, with "FC Adapter Policy Windows" selected and highlighted by a red box. The main content area is divided into "General" and "Events" tabs, with "General" active. It shows the "Owner" as "Local". Below this are sections for "Resources" and "Options". The "Options" section contains various configuration parameters. The "LUN Queue Depth" parameter is highlighted with a red box and is currently set to 20, with a range of [1-254] indicated to its right. Other parameters include FCP Error Recovery (Disabled), Flogi Retries (8), Flogi Timeout (4000), Plogi Retries (8), Plogi Timeout (20000), Error Detect Timeout (2000), Port Down Timeout (30000), IO Retry Timeout (5), Port Down IO Retry (30), Link Down Timeout (30000), Resource Allocation Timeout (10000), IO Throttle Count (256), Max LUNs Per Target (1024), Interrupt Mode (MSI X), and vHBA Type (FC Initiator).

Parameter	Value	Range
FCP Error Recovery	Disabled	Enabled
Flogi Retries	8	[0-infinite]
Flogi Timeout (ms)	4000	[1000-255000]
Plogi Retries	8	[0-255]
Plogi 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

Cisco Integrated Management Controller standalone LUN Queue Depth configuration for Microsoft Windows

You can configure the LUN Queue Depth parameter for standalone servers in the Cisco IMC GUI. This parameter is applicable to Windows only.

To change the LUN Queue Depth parameter, log in to the IMC, click the Networking panel, and then click the adapter card on which the vHBAs are located. In the main panel, click the vHBA tab, select the vHBA, and then expand vHBA Properties > Fibre Channel Port. Set the LUN Queue Depth parameter to the desired value for each vHBA. Click Save Changes after each vHBA has been modified. Reboot to make the changes take effect.



Configuring I/O Throttle Count

Description: The total number of I/O requests that can be outstanding per virtual HBA (vHBA)

Parameter name: I/O Throttle Count

Parameter values:

- Cisco UCS Intersight (IMM):
 - Configurable range = 1 to 1024
 - Linux: Default = 256
 - VMware: Default = 256
 - Windows: Default = 256
- Cisco UCS Manager:
 - Configurable range = 256 to 1024
 - Linux: Default = 256
 - VMware: Default = 256
 - Windows: Default = 256
- Cisco Integrated Management Controller:
 - Configurable range = 1 to 1024
 - Default = 512

Configuration capabilities: Boot time only

Display the I/O Throttle Count parameter value

You can display the current and post configuration values of the I/O Throttle Count parameter using the commands shown here.

ESX 6.7 and 7.x

Run this command:

```
# cat /var/log/vmkernel.log |grep throttle
```

RHEL 7.9 and 8.x

Run one of the following commands:

```
# journalctl |grep throttle
```

Or:

```
# cat /var/log/messages |grep throttle
```

SLES 12.5 and 15.x

Run one of the following commands:

```
# journalctl |grep throttle
```

Or:

```
# cat /var/log/messages |grep throttle (if the syslog daemon is installed)
```

XS 8.x

Enter this command:

```
# journalctl |grep throttle
```

Windows

You can display Cisco VIC fNIC parameters using the Cisco command-line tool for Windows called fctool.exe. This tool has traditionally been available only from the Cisco TAC as a troubleshooting tool. Cisco currently does not support this tool, but you can download and use it at your own risk at the following link:

<https://community.cisco.com/t5/unified-computing-system/cisco-vic-fnic-fctool-utility-for-windows/ta-p/4663692>

The fctool.exe tool requires the Microsoft Visual C++ redistributable VCRUNTIME140.dll dynamic link library, which is available from Microsoft within the vc_redist.x64.exe package. When fctool.exe is run, if there is a missing DLL file error, please download and install the latest Microsoft Visual C++ Redistributable from this link: <https://docs.microsoft.com/en-US/cpp/windows/latest-supported-vc-redist?view=msvc-170>

To display the fNIC I/O Throttle Count value, at a Windows CMD prompt from which fctool.exe is accessible, run the following command to list the fNICs and their port numbers:

```
> fctool -list
fnic [04] --> vnic [16] maca [0025b5fa0008]
    node_wwn 20000025b5ff0008 port_wwn 20000025b5fa0008
    maxdatafieldsize 2112 edtov 2000 ratov 10000
    drv fnic2k12.sys (fre) version 4.0.0.1
    svc fnic2k12
    PCI bus 27 slot 32

fnic [05] --> vnic [17] maca [0025b5fb0008]
    node_wwn 20000025b5ff0008 port_wwn 20000025b5fb0008
    maxdatafieldsize 2112 edtov 2000 ratov 10000
    drv fnic2k12.sys (fre) version 4.0.0.1
    svc fnic2k12
    PCI bus 27 slot 64
```

The port number for fnic[04] is 4 and the port number for fnic[05] is 5. The number in the bracket following fnic denotes the port number. To display the I/O Throttle Count value of the fNIC at port 4, run the following command:

> fctool -p 4 -res

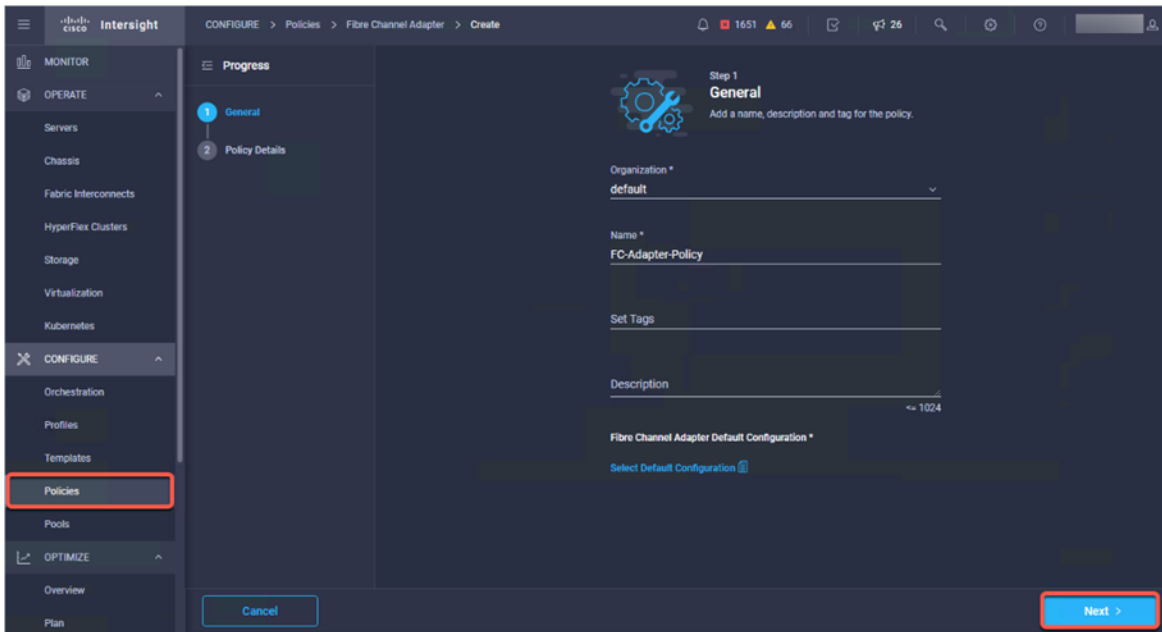
```
fnic --> fnic resources
    node_wwn: 20000025b5ff0008
    port_wwn: 20000025b5fa0008

throttle cnt: 256
```

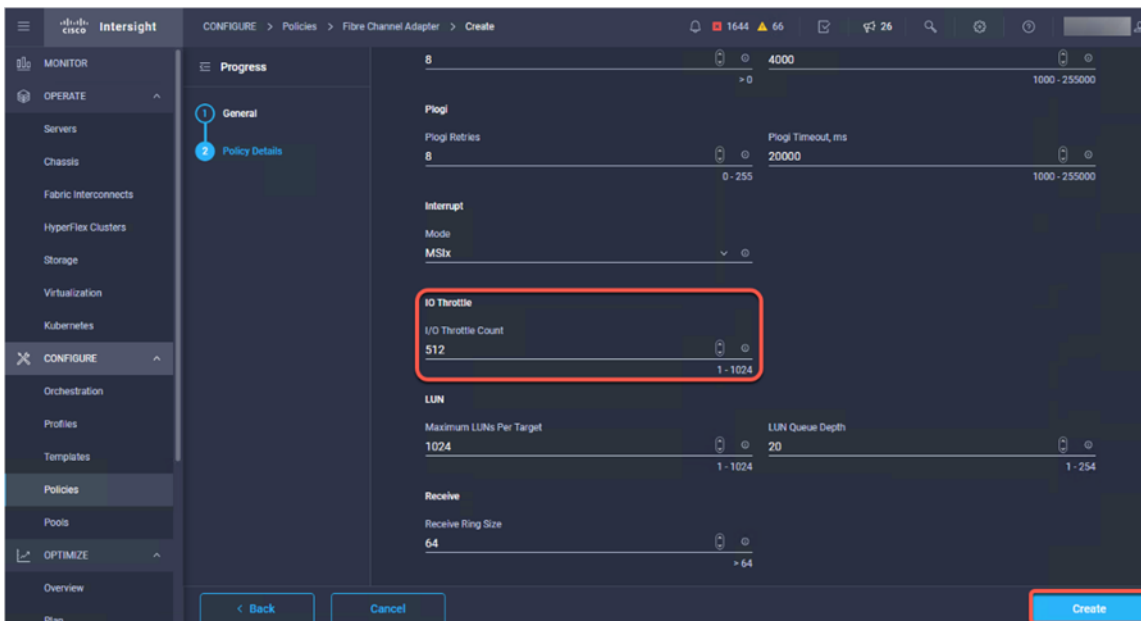
Cisco Intersight I/O Throttle Count configuration

You can configure the I/O Throttle Count parameter in the Cisco Intersight platform through <https://www.intersight.com>. The I/O Throttle Count parameter is applicable to Linux, VMware, and Windows through the Fibre Channel adapter policy.

1. Log in to Cisco Intersight and choose Configure > Policies > Create Policy. In the Select Policy Type panel, click Fibre Channel Adapter. Then click Start. Configure the name and other fields on the General configuration screen. Then click Next.



2. Configure the I/O Throttle Count field with the desired value. Then click the Create button (or the Update button if you are editing an existing policy).



3. Add the Fibre Channel adapter policy to the SAN connectivity policy of the Cisco UCS server profile and deploy or redeploy the server profile.

Cisco UCS Manager I/O Throttle Count configuration

You can configure the I/O Throttle Count parameter for Cisco UCS managed servers through the Cisco UCS Manager GUI or equivalent Cisco UCS Manager XML commands. The I/O Throttle Count parameter is configurable in the Linux, VMware, and Windows Fibre Channel adapter policies.

To change the I/O Throttle Count parameter, in the Cisco UCS Manager navigation tree, click Servers and then expand Policies and Adapter Policies in the navigation tree. Click the Linux, VMware, or Windows FC Adapter Policy and then, in the main window, expand the Options drop-down menu. Configure the I/O Throttle Count field with the desired value and then click Save Changes.

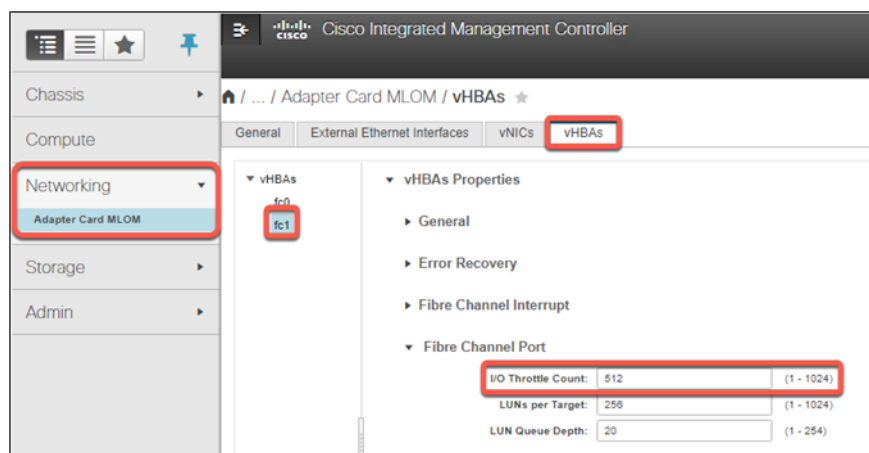
The screenshot displays the Cisco UCS Manager interface. The navigation tree on the left shows the path: Servers > Policies > root > Adapter Policies > FC Adapter Policy Linux. The main window shows the configuration for this policy. The 'Options' section is expanded, and the 'IO Throttle Count' parameter is highlighted with a red box. The value is set to 256, with a range of [256-1024].

Parameter	Value	Range
IO Throttle Count	256	[256-1024]
Max LUNs Per Target	1024	[1-1024]

Cisco Integrated Management Controller standalone I/O Throttle Count configuration

You can configure the I/O Throttle Count parameter for standalone servers in the Cisco IMC GUI. This parameter is applicable to Linux, VMware, and Windows.

To change the I/O Throttle Count parameter, log in to the IMC, click the Networking panel, and then click the adapter card on which the vHBAs are located. In the main panel, click the vHBA tab, select the vHBA, and then expand vHBA Properties > Fibre Channel Port. Set the I/O Throttle parameter to the desired value for each vHBA. Click Save Changes after each vHBA has been modified. Reboot to make the changes take effect.



For more information

Consult the following resources for additional information:

- Cisco VIC driver installation guides: <https://www.cisco.com/c/en/us/support/servers-unified-computing/ucs-virtual-interface-card/products-installation-guides-list.html>.
- Cisco UCS Hardware Compatibility List (HCL): <https://ucshcltool.cloudapps.cisco.com/public/>

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)