

## Contents

<b>Overview</b> . . . . .	<b>2</b>
<b>Audience</b> . . . . .	<b>2</b>
<b>Test Environment</b> . . . . .	<b>2</b>
Cisco UCS . . . . .	2
fNIC Driver . . . . .	2
<b>fNIC Tunable Parameters</b> . . . . .	<b>2</b>
<b>Operating System Support</b> . . . . .	<b>2</b>
<b>Install fNIC Drivers</b> . . . . .	<b>3</b>
<b>Configuring fNIC Tunables</b> . . . . .	<b>3</b>
LUN Queue Depth . . . . .	3
Displaying fnic_max_qdepth Parameter Value . . . . .	4
Boot-Time Configuration . . . . .	5
Load-Time Configuration . . . . .	5
Run-Time Configuration . . . . .	5
IO Throttle Count . . . . .	6
UCSM Non-Standalone I/O Throttle Count Configuration . . . . .	7
Cisco Integrated Management Controller Standalone I/O Throttle Count Configuration . . . . .	8
<b>References</b> . . . . .	<b>8</b>

## Overview

Cisco Unified Computing System™ (Cisco UCS®) 2.1.2a introduces the fNIC Tunables feature, which provides the capability to tune the fNIC Queue Depth and I/O Throttle Count parameters of the Cisco UCS Virtual Interface Card (VIC) Fibre Channel network interface card (fNIC) driver in Linux and VMware ESX implementations. This guide provides an overview of these two parameters and the methodologies and syntax for modifying their values. Prior to Cisco UCS 2.1.2a, both of these parameters were statically set to industry best practice defaults, and these values are still recommended for the vast majority of architectures. The ability to tune these parameters provides flexibility for those customers whose architectures may require nondefault values. Modifying default values to specific solutions is beyond the scope of this guide.

## Audience

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

## Test Environment

### Cisco UCS

Cisco UCS Manager (UCSM) 2.1.2a

(2) Cisco UCS 6248UP 48-Port Fabric Interconnects

(2) Cisco UCS 2208XP I/O Modules

(1) Cisco UCS 5108 Blade Server Chassis

(1) Cisco UCS B200 M3 Blade Server with UCS VIC 1240 modular LAN on motherboard (mLOM)

(1) Cisco UCS B200 M2 Blade Server with UCS VIC M81KR

### fNIC Driver

1.5.0.45 (minimum version)

## fNIC Tunable Parameters

The two new fNIC tunable parameters introduced in UCSM 2.1.2a are Queue Depth and I/O Throttle Count. Definitions for each of these parameters are as follows:

- **Queue Depth:** The total number of I/O requests that can be outstanding on a per-logical unit number (LUN) basis.
- **I/O Throttle Count:** The total number of I/O requests that can be outstanding on a per-virtual host bus adapter (vHBA) basis.

## Operating System Support

The Queue Depth and I/O Throttle Count tunable parameters are supported on the following operating systems:

- Red Hat Linux 6.x (RHEL 6.x)
- Red Hat Linux 5.x (RHEL 5.x)
- SUSE Linux Enterprise 11 (SLES 11)
- XenServer (XS)
- VMWare ESX 5.x (ESX)
- **Not Supported:** VMWare ESX 4.x

## Install fNIC Drivers

This guide assumes that the UCS 2.1.2a or later fNIC drivers have been installed and are running. Please see the Cisco UCS Manager Install and Upgrade Guides, Virtual Interface Card Drivers section, for complete driver installation instructions: <http://www.cisco.com/c/en/us/support/servers-unified-computing/ucs-manager/products-installation-guides-list.html>

In addition, please review the Cisco UCS Hardware Compatibility List (HCL) to confirm hardware, operating system, and driver compatibility: [http://www.cisco.com/en/US/products/ps10477/prod\\_technical\\_reference\\_list.html](http://www.cisco.com/en/US/products/ps10477/prod_technical_reference_list.html)

## Configuring fNIC Tunables

### LUN Queue Depth

**Description:** The total number of I/O requests that can be outstanding on a per-LUN basis.

**Parameter name:** fnic\_max\_qdepth

**Default value:** 32

#### Configuration capabilities:

- Boot Time
  - » Requires reboot for changes to take effect.
  - » Changes are persistent across reboots.
- Load Time
  - » Disruptive to SAN attached storage.
  - » Requires stopping SAN I/O, removing fNIC module dependencies, and removing and reloading the fNIC module.
  - » Not possible with the Boot from SAN configuration.
  - » Changes are nonpersistent across reboots.
- Run Time
  - » Nondisruptive.
  - » Changes apply only to newly discovered LUNs after modification.
  - » Changes are nonpersistent 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 5.x, 6.x - Local Boot	X	X	X
RHEL 5.x, 6.x - Boot from SAN	X		X
SLES 11 - Local Boot	X	X	X
SLES 11 - Boot from SAN	X		X
XS - Local Boot	X	X	X
XS - Boot from SAN	X		X
ESX 5.x - Local Boot	X		
ESX 5.x - Boot from SAN	X		

### Displaying fnic\_max\_qdepth Parameter Value

The current and postconfiguration value of the fnic\_max\_qdepth parameter can be displayed using the following commands.

#### RHEL 6.x / SLES11 / XS

- To display the current fnic\_max\_qdepth, run the following command:

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

- To display the current fnic\_max\_qdepth 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 command-line interface (CLI) (install lsscsi if not already installed):

```
# lsscsi -l
```

#### ESX 5.x

- From the ESX CLI, run the following command (the fnic\_max\_qdepth parameter will not be listed until it has been explicitly set per the configuration instructions below):

```
# cat /etc/vmware/esx.conf |grep fnic
```

**Results:**

```
/vmkernel/module/fnic/options = fnic_max_qdepth=128
```

- From the ESX CLI, run the following esxcli command (the fnic\_max\_qdepth value will be empty until the fnic\_max\_qdepth parameter has been explicitly set per the configuration instructions below):

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

**Results:**

Name	Type	Value	Description
-----	----	-----	-----
fnic_max_qdepth	uint 128		Queue depth to report for each LUN

## Boot-Time Configuration

Example commands for setting `fnic_max_qdepth` to 128:

### RHEL 6.x / SLES11 / XS

1. In `/etc/grub.conf`, append `fnic.fnic_max_qdepth=<value>` to the kernel line:

```
kernel /vmlinuz-2.6.32-358.el6.x86_64 ro root=/dev/mapper/vg_craigsrh-lv_root rd_NO_LUKS LANG=en_US.UTF-8
rd_NO_MD SYSFONT=latacyrheb-sun16 crashkernel=auto rd_LVM_LV=vg_craigsrh/lv_swap rd_LVM_LV=vg_craigsrh/
lv_root KEYBOARDTYPE=pc KEYTABLE=us rd_NO_DM rhgb quiet fnic.fnic_max_qdepth=128
```

2. Reboot for the new `fnic_max_qdepth` to take effect.

### RHEL 5.x

1. In `/etc/modprobe.conf`, add an “options `fnic`” line to set `fnic_max_qdepth` at boot time:

```
options fnic fnic_log_level=4 fnic_trace_max_pages=128 fnic_max_qdepth=128
```

2. Rebuild `initrd`:

```
mkinitrd -v -f /boot/initrd-`uname -r`.img `uname -r`
```

3. Reboot for the new `fnic_max_qdepth` to take effect.

### ESX 5.x

1. Set `fnic_max_qdepth` by running the following command in ESX 5 CLI:

```
# esxcli system module parameters set -p fnic_max_qdepth=128 -m fnic
```

2. Reboot for the new `fnic_max_qdepth` to take effect.

## Load-Time Configuration

### RHEL / SLES11 / XS

1. Run one of the following commands at the CLI:

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

or

```
# insmod fnic.ko fnic_max_qdepth=128
```

## Run-Time Configuration

### RHEL / SLES11 / XS

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

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



## IO Throttle Count

**Description:** The total number of I/O requests that can be outstanding on a per-vHBA basis.

**Parameter name:** I/O Throttle Count

**Parameter values:**

**Note:** Prior to the UCSM 2.1.2a release, Cisco UCSM displayed an I/O Throttle Count parameter in the non-standalone Linux and ESX Adapter Policies, and in the standalone Cisco Integrated Management Controller vHBA Properties field. This parameter was programmatically ignored; modifying it had no effect on the actual adapter setting, which remained statically set. As of the Cisco UCSM 2.1.2a release, this value is now honored.

**Note:** Some values' actual parameter settings are different than what is displayed in the UCSM and Cisco Integrated Management Controller (IMC) GUIs. These differences are implemented to support backward compatibility and forward consistency.

UCSM GUI / CIMC GUI configurable range = 1 to 1024

UCSM GUI default (non-standalone) = 16 displayed / 2048 actual parameter setting

CIMC GUI default (standalone) = 512 displayed / 512 actual parameter setting

UCSM GUI / CIMC GUI < 256 displayed / 256 actual parameter setting

UCSM GUI / CIMC GUI = 16 displayed / 2048 actual parameter setting

**Configuration capabilities:** Boot time only.

## UCSM Non-Standalone I/O Throttle Count Configuration

The I/O Throttle Count for non-standalone servers is configured through the UCSM GUI or equivalent UCSM XML commands. The I/O Throttle Count parameter is configurable in the Linux and VMWare FC Adapter Policies.

To change the I/O Throttle Count parameter, in the UCSM navigation tree, click the Servers tab, then expand the Policies and Adapter Policies in the navigation tree. Click the FC Adapter Policy Linux or FC Adapter Policy VMWare, and then, in the main window, expand the Options drop-down. Configure the I/O Throttle Count field per the list of parameter values above, and then click the Save Changes button.

The screenshot displays the UCSM GUI configuration for the FC Adapter Policy Linux. The navigation tree on the left shows the path: Servers > Policies > Adapter Policies > FC Adapter Policy Linux. The main window displays the configuration for this policy, with the 'Options' section expanded. The 'IO Throttle Count' is set to 256. 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), Port Down IO Retry (30), Link Down Timeout (30000), Resource Allocation Timeout (10000), and Max LUNs Per Target (256). The 'Interrupt Mode' is set to MSI X.

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]
Port Down IO Retry	30	[0-255]
Link Down Timeout (ms)	30000	[0-240000]
Resource Allocation Timeout (ms)	10000	
IO Throttle Count	256	[1-1024]
Max LUNs Per Target	256	[1-1024]
Interrupt Mode	MSI X	MSI, IN Tx

## Cisco Integrated Management Controller Standalone I/O Throttle Count Configuration

The I/O Throttle Count for standalone servers is configured in the Cisco Integrated Management Controller GUI.

To change the I/O Throttle Count parameter, log in to the CIMC, navigate to the Server tab, click the Summary link and then click Power on Server. Once the server has been powered on, click the Inventory link, select the Cisco VIC Adapters tab, choose the Adapter Card, and then select the vHBAs tab. Click fc0 or fc1, then click the Properties button. The vHBA Properties window will display. Scroll down to the Fibre Channel Port section, configure the I/O Throttle Count field per the parameter values listed above, then click the Save Changes button.

The screenshot displays the Cisco Integrated Management Controller (CIMC) GUI. The top navigation bar shows the CIMC Hostname as 'ucs-c240-m3' and the user logged in as 'admin@10.155.160.168'. The main interface is divided into several sections:

- Overall Server Status:** Shows a 'Good' status.
- Adapter Cards:** A table lists the installed adapter cards. One card is highlighted: PCI Slot 2, Product Name UCS VIC 1225, Serial Number FCH162974Z0.
- Adapter Card 2:** A sub-section for the selected card, showing tabs for General, vNICs, VM FEXs, and vHBAs. The vHBAs tab is active.
- Host Fibre Channel Interfaces:** A table lists the vHBAs. Two are shown: 'fc0' and 'fc1', both with WWPNs starting with 20:00:30:F7:0D:9A:0C:5F and 20:00:30:F7:0D:9A:0C:60 respectively.
- vHBA Properties Dialog:** A modal window for configuring the selected vHBA. The 'I/O Throttle Count' field is highlighted with a red circle and set to 1024. Other fields include 'Port Down Timeout' (10000 ms), 'Fibre Channel Interrupt' (MSIX), 'Fibre Channel Port' (LUNs per Target: 256), 'Fibre Channel Port FLOGI' (Retries: INFINITE, Timeout: 2000 ms), and 'Fibre Channel Port PLOGI'.

## References

Virtual Interface Card Drivers Installation Guides

<http://www.cisco.com/c/en/us/support/servers-unified-computing/ucs-manager/products-installation-guides-list.html>

UCSM Managed and Standalone Compatibility Documents

[http://www.cisco.com/en/US/products/ps10477/prod\\_technical\\_reference\\_list.html](http://www.cisco.com/en/US/products/ps10477/prod_technical_reference_list.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](http://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](http://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)

C07-730811-00 1/14