



Universal Multichannel Reference Guide

Revision 4.6

EMULEX CONFIDENTIAL. Copyright © 2012–2013 Emulex. All rights reserved worldwide. No part of this document may be reproduced by any means or translated to any electronic medium without the prior written consent of Emulex.

Information furnished by Emulex is believed to be accurate and reliable. However, no responsibility is assumed by Emulex for its use; or for any infringements of patents or other rights of third parties which may result from its use. No license is granted by implication or otherwise under any patent, copyright or related rights of Emulex.

Emulex, the Emulex logo, AutoPilot Installer, AutoPilot Manager, BlockGuard, Connectivity Continuum, Convergenomics, Emulex Connect, Emulex Secure, EZPilot, FibreSpy, HBAnyware, InSpeed, LightPulse, MultiPulse, OneCommand, OneConnect, One Network. One Company., SBOD, SLI, and VEngine are trademarks of Emulex. All other brand or product names referenced herein are trademarks or registered trademarks of their respective companies or organizations.

Emulex provides this manual "as is" without any warranty of any kind, either expressed or implied, including but not limited to the implied warranties of merchantability or fitness for a particular purpose. Emulex may make improvements and changes to the product described in this manual at any time and without any notice. Emulex assumes no responsibility for its use, nor for any infringements of patents or other rights of third parties that may result. Periodic changes are made to information contained herein; although these changes will be incorporated into new editions of this manual, Emulex disclaims any undertaking to give notice of such changes.

Emulex, 3333 Susan Street

Costa Mesa, CA 92626

Table of Contents

List of Tables	4
1. Introduction	5
Abbreviations	5
2. Universal Multichannel	7
Overview	7
Considerations	7
Capabilities	7
Physical Functions	8
Assigning LPVIDs	9
Overriding LPVIDs with VLAN IDs	10
.....	10
3. Configuration	11
Recommended Configuration Guidelines	11
Configuration Options	12
LPVID Assignment	12
Bandwidth Assignments	12
Configurable Parameters	12
High Availability with UMC	13
Windows	13
Linux	13
VMware	14
Solaris	14
Considerations	15
Windows NIC Teaming Considerations	15
UMC NIC Channels on the Same Physical Port Cannot Be Placed in the Same Team	15
Teaming is Supported on Multiple Channels	15
Windows Deployment Considerations	15
Using the LPVID	15
Windows Device Manager VLAN ID	16
Multiple VLAN Configuration in Windows	16
4. Troubleshooting	17

List of Tables

Table 2-1	UMC Port Mapping - NIC Only.....	8
Table 2-2	UMC Port Mapping - NIC and Storage.....	8
Table 2-3	UMC Port Mapping for High Availability.....	9
Table 3-1	UMC Configurable Parameters	12
Table 4-1	UMC Troubleshooting	17

1. Introduction

This manual describes universal multichannel (UMC) and how to configure it for supported Emulex adapters. UMC provides powerful port partitioning and traffic management capabilities to optimize bandwidth allocation. With UMC, system administrators can partition a dual-port OCe11100-series adapter into up to eight Peripheral Component Interconnect Express (PCIe) functions (logical ports or channels) with assigned bandwidth that can be integrated into both physical and virtual servers. This is particularly beneficial for virtualized servers where individual functions can be assigned for virtual machine (VM) migration, system management, and input/output (I/O) intensive applications running in VMs.

Abbreviations

BIOS	basic input/output system
CLP	command line protocol
CLI	command line interface
DCB	Data Center Bridging
FCoE	Fibre Channel over Ethernet
Gb/s	gigabits per second
GbE	gigabit Ethernet
GUI	graphic user interface
IEEE	Institute of Electrical and Electronics Engineers
I/O	input/output
IP	internet protocol
iSCSI	Internet Small Computer System Interface
LACP	Link Aggregation Control Protocol
LAN	local area network
LPVID	logical port (channel) VLAN ID
MAC	media access control
Mb/s	megabits per second
MSI	message signaled interrupts
MTU	maximum transmission unit
NIC	network interface card (or controller)
NFS	network file system
OCM	OneCommand™ Manager
OS	operating system
PCI	Peripheral Component Interconnect
PCIe	Peripheral Component Interconnect Express
PF	physical function

PXE	Preboot eXecution Environment
QoS	quality of service
SCSI	Small Computer System Interface
SR-IOV	single root I/O virtualization
TCP	transmission control protocol
TOE	TCP Offload Engine
UCNA	universal converged network adapter
UEFI	Unified Extensible Firmware Interface
UMC	universal multichannel
VLAN	virtual local area network
VM	virtual machine
vNIC	virtual NIC
vSwitch	virtual switch

2. Universal Multichannel

Overview

UMC enables up to four PCIe functions per port for OCe11100-series adapters.

- For OCe11100-series-N Ethernet adapters, each port can support up to four channels.
- For OCe11100-series-I storage adapters, each port can support up to three NIC channels and an iSCSI function.
- For OCe11100-series-F storage adapters, each port can support up to three NIC channels and an iSCSI or FCoE function.

Note: UMC is only supported on Emulex OneConnect OCe11100-series UCNAs. However, it is possible to use older OCe10100-series adapters with a UMC-enabled OCe11100-series adapter in the same server.

Bandwidths for each function can be specified as a percentage of the full 10 Gb/s for the adapter port.

Most servers are currently deployed with multiple 1GbE physical connections. Typically these additional ports are used to support virtual servers and high availability, and to provide bandwidth needed for I/O-intensive applications. UMC provides a similar capability for 10GbE networking using individually configurable partitions of the 10GbE port. With UMC, data centers can save on costs for cabling, adapters, switches, and power.

Considerations

- If UMC support is enabled, SR-IOV support must be disabled.
- The valid LPVID range is 2–4094.
- You cannot run LACP on UMC NICs.
- Make sure you correctly configure your switch to support UMC. You must configure the switch ports with the LPVIDs (4 per port) on the switch. Additionally, if you configure additional VLAN IDs on a channel, these VLAN IDs must also be configured on the switch port.

Capabilities

- Switch-agnostic support (works with any 10GbE switch)
- Creation of four PFs per physical port, with the mapping of PFs to physical ports as follows:
 - PF0, PF2, PF4, PF6 → Port0
 - PF1, PF3, PF5, PF7 → Port1
- Separate transmit and receive queues for each channel

- Channel isolation with unique VLAN assignments (based on the IEEE 802.1Q VLAN standard; each channel has its own independent broadcast and multicast domain)

Note: A corresponding VLAN configuration is required in the switch.

- Optimized virtual server deployments, which enable allocation of separate PFs for VM migration, console management, iSCSI and NFS storage, and network traffic for individual VMs

Physical Functions

When UMC is enabled on an OCe11100-series dual-port network adapter, each 10GbE port is partitioned into four isolated PFs (channels). A total of eight PFs are available on an OCe11100-series dual-port adapter.

Although support for UMC channels is switch-agnostic, a switch that supports DCB is required for FCoE and is also required when iSCSI is used with DCB. UMC is supported with powerful traffic management and provisioning capabilities such as dynamic rate control, priorities, MAC configuration, and VLAN assignment.

With UMC, physical functions (PF0–PF7) are presented to an operating system or hypervisor as eight independent adapters. UMC channels are presented to the operating system or hypervisor as a physical port with a separate MAC address and assigned bandwidth. Port mapping examples are shown in Table 2-1 and Table 2-2. In Table 2-1, all of the physical functions are configured as NICs. In Table 2-2, one of the physical functions for each port is configured for storage.

Table 2-1 UMC Port Mapping - NIC Only

Channel Number	Channel Type	Port 0 Physical Function	Port 1 Physical Function
0	NIC	PF0	PF1
1	NIC	PF2	PF3
2	NIC	PF4	PF5
3	NIC	PF6	PF7

Table 2-2 UMC Port Mapping - NIC and Storage

Channel Number	Channel Type	Port 0 Physical Function	Port 1 Physical Function
0	NIC	PF0	PF1
1	Storage (iSCSI or FCoE)	PF2	PF3
2	NIC	PF4	PF5
3	NIC	PF6	PF7

Assigning LPVIDs

The LPVID is the default VLAN ID (2-4094) that identifies the NIC's network channel. It is automatically inserted on untagged transmit packets, and automatically removed on receive packets. Effectively, the LPVID converts untagged packets into VLAN tagged packets without any operating system knowledge.

The LPVID for a UMC NIC is configured using the PXESelect utility, the UEFI BIOS utility, or the OneCommand Manager application CLI.

Note: An LPVID is required for a UMC physical function with a NIC function type (UMC NIC). A VLAN ID is optional for a UMC physical function with a storage function type (iSCSI or FCoE). You can configure an iSCSI VLAN ID through the iSCSISelect utility or OneCommand Manager application. You can configure an FCoE VLAN ID through the OneCommand Manager application.

For existing network environments, assign LPVIDs that correspond to VLANs previously configured on the switches in your network. The switch port should be set to trunking mode to carry traffic for multiple LPVIDs. The switch port must allow all LPVIDs that are assigned to UMC NICs for the adapter port.

UMC NICs work like physical adapter ports within the network. VLAN membership and trunking-mode between multiple switches are left to the discretion of the system administrator, but they must include *all* VLAN IDs and LPVIDs in order for traffic to be accepted by other systems in the network.

Note: Each UMC NIC channel must be configured with a unique LPVID within the physical port.

The same LPVID can be assigned to UMC NICs on both ports of the adapter for redundancy and high availability as shown in Table 2-3.

Table 2-3 UMC Port Mapping for High Availability

Port 0 Function Number	LPVID	Port 1 Physical Function	LPVID
PF0	2	PF1	2
PF2	5	PF3	5
PF4	9	PF5	9
PF6	12	PF7	12

See "High Availability with UMC" on page 13 for additional information on configuring UMC for high availability on various operating systems.

Overriding LPVIDs with VLAN IDs

LPVID values assigned to the NIC channels can be overridden using the operating system or hypervisor utilities to assign a VLAN ID to the NIC driver instance running on a UMC channel. The NIC driver sends the VLAN ID to the adapter each time it loads, which overrides the LPVID configured on that channel.

If the operating system for that PCI function has set up a VLAN ID:

- For transmit packets, the operating system-configured VLAN ID takes precedence over the LPVID.
- For receive packets, the operating system-configured VLAN ID and LPVID-tagged packets are both received.

If the operating system has not set up any VLAN IDs, then the LPVID is inserted into all packets.

Caution: LPVIDs and user-configured VLAN IDs must be different values.

3. Configuration

You must determine your network requirements such as VLANs, bandwidth, trunk ports, and load balance prior to configuring UMC. Converged traffic such as FCoE or iSCSI over DCB requires specialized network switches. If your configuration requirement is for multiple NICs only, you can use a standard 10GbE Layer 2 switch.

Complete the following steps before configuring UMC:

1. Install the Emulex OCE11100-series adapter.
2. Install the appropriate drivers.
3. Install the OneCommand Manager application or, for VMware systems, install the required CIM provider and OneCommand Manager for VMware vCenter version 1.5 or later.
4. Install the appropriate adapter firmware. The UFI firmware image contains both the firmware and boot code. A server reboot is required after the firmware update.

You can configure UMC using the following Emulex utilities:

- PXESelect utility – see the *Boot for NIC, iSCSI, and FCoE Protocols User Manual* for more information.
- UEFI BIOS utility – see the *Boot for NIC, iSCSI, and FCoE Protocols User Manual* for more information.
- OneCommand Manager application – see the *OneCommand Manager Application User Manual* or *OneCommand Manager CLI User Manual* for more information.
- OneCommand Manager for VMware vCenter version 1.5 or later – see the *OneCommand Manager for VMware vCenter User Manual* for more information.

Recommended Configuration Guidelines

Emulex recommends following these guidelines to ensure that UMC is configured properly:

- Set the LPVID on each NIC channel. Set the VLAN ID on switches and any remote servers (if UMC is not enabled on the remote server).
- Do not set the operating system VLAN ID to the same value as the LPVID when UMC is enabled. (If you must have both the operating system VLAN ID and LPVID set when UMC is enabled, you must ensure that the LPVID and operating system VLAN ID are not the same value).
- Configure the maximum bandwidth setting to take advantage of excess bandwidth in your 10Gb network or storage solution.
- Configure the minimum bandwidth setting for an FCoE channel to ensure adequate storage traffic speeds for your network usage model. The minimum bandwidth setting is guaranteed bandwidth.
- If UMC support is enabled, SR-IOV support must be disabled.
- With UMC enabled, any VLAN IDs to be used by the operating system or applications must also be configured in the NIC driver.
- You cannot run LACP on UMC NICs.

- Use the OneCommand Manager application help file or the CLI (HBACMD) help command, or refer to the appropriate user manuals, for information on the required UMC settings and syntax.
- Make sure you correctly configure your switch to support UMC. You must configure the switch ports with the LPVIDs (4 per port) on the switch. Additionally, if you configure additional VLAN IDs on a channel, these VLAN IDs must also be configured on the switch port.
- Use the Emulex utilities to verify that your UMC settings are complete, including OneCommand Manager application, OneCommand Manager for VMware vCenter, and the PXESelect and uEFI BIOS Option ROMS utilities during system boot.
- Reboot or reset the server when enabling or disabling UMC. The Emulex management utilities will prompt you for a reboot. After enabling UMC and rebooting, additional reboots are not necessary when making changes to the UMC channel (for example, configuring the LPVID).

Configuration Options

LPVID Assignment

You must assign an LPVID to all enabled NIC channels. The LPVID should be a value between 2 and 4094, and the LPVIDs must be unique across all channels on an adapter port.

Bandwidth Assignments

You must assign minimum and maximum bandwidths to all enabled channels. Bandwidths are assigned in percentages. The total of all the minimum bandwidths for all enabled channels on an adapter port must add up to 100%. The maximum bandwidth for each channel must be between the channel's minimum bandwidth and 100%.

Note: When the minimum bandwidth is set to 0, the logical link state for the channel is disabled.

Configurable Parameters

Table 3-1 summarizes all configurable parameters with UMC enabled or disabled.

Table 3-1 UMC Configurable Parameters

Configurable Parameter	UMC Disabled	UMC Enabled
Quality of service (QoS)	Based on physical port	Based on channel
Link State	Mimics physical link state	Configurable through bandwidth assignment
Link speed reported to operating system	Port speed	Based on QoS value

Table 3-1 UMC Configurable Parameters (Continued)

Configurable Parameter	UMC Disabled	UMC Enabled
Port speed	10 Gb/s	1-10 Gb/s
Configuration Entity	Not applicable	BIOS, CLP, Switch, OneCommand Manager application
PXE/iSCSI/FCoE Boot	Yes	Yes - first primary function per port
Personality change	NIC, iSCSI+NIC, FCoE+NIC	NIC, iSCSI+NIC, FCoE+NIC

High Availability with UMC

UMC NICs can be configured for high availability with the same procedures that are used for physical ports. The following sections provide information to configure high availability for Windows, Linux, VMware, and Solaris.

Note: UMC NICs do not support link aggregation or port trunking using IEEE 802.3ad LACP.

Windows

Note: The following conditions also apply when Hyper-V is used.

Use the OneCommand NIC Teaming and VLAN Manager utility to team UMC NICs with other UMC NICs or other physical adapters. Make sure the selected UMC NICs for teaming are not associated with the same physical port. This ensures that traffic fails over to the secondary port and switch in the event of a link failure.

Note: If you are creating teams with the NIC Teaming Driver utility when Hyper-V is in use, you should not set the VLAN IDs with the utility when UMC is enabled. Instead, you should set the VLAN IDs in the driver's Device Manager Property page and within the Hyper-V switch.

Use the OneCommand Manager application to identify and select the correct physical function, IP address, or MAC for the physical ports. Using the OneCommand NIC Teaming and VLAN Manager, select the adapter in the navigation pane (tree view). The MAC addresses of all UMC NICs and other physical NICs in the server appear on the General tab in the Connection Properties group box. Refer to the *OneCommand Teaming and VLAN Manager User Manual* for more information on NIC teaming with Windows.

Linux

Use the OneCommand Manager application to identify MAC addresses and physical PCIe function numbers. Use the `ipconfig` command to configure IP addresses for all UMC NICs. Use the Linux bonding driver to team UMC NICs with other UMC NICs or other physical adapters.

Note: See the documentation included with the Linux distribution for information on configuring Linux bonding.

Best practices:

- Verify the UMC-adapter physical port association, PCIe function, and MAC address prior to bonding UMC NICs.
- Ensure that all bonded UMC NICs are not partitioned on the same physical port. This provides redundancy in the event of a physical link failure.

VMware

From the vSphere Client, use Network Adapters to show the list of UMC NICs that can be provisioned as uplinks for vSphere standard or distributed switches. UMC NICs can be used for individual VMs, VMotion, console management, and other network traffic types, such as NAS.

VMware recognizes a UMC NIC as an available vmnic that can be configured using a vSwitch to handle network traffic. As a best practice, configure UMC NICs across physical ports using equal bandwidth to ensure consistency across active and standby or aggregated link paths. Also, ensure that the vmnics selected for redundant paths are not on the same physical port.

Solaris

Use the OneCommand Manager application to identify MAC addresses and physical PCIe function numbers. Use the `ipconfig` command to configure IP addresses for all UMC NICs.

Use the OneCommand Manager application to identify physical PCIe function numbers and determine their association with the physical ports using the port mapping outlined in “Physical Functions” on page 8. Using that information, the interfaces should be selected from the `/etc/path_to_inst` file. The interfaces belong to different physical ports and are used to form the members of the aggregation group.

The aggregation group is created using the `'dladm create-aggr'` command. For more information on this command, refer to the `'dladm'` man page or the “Oracle Solaris Administration: Network Interfaces and Network Virtualization” manual.

Caution: Using UMC in a Solaris hypervisor environment is not advised, if the UMC interface is going to be part of the hypervisor virtual switch. This configuration may cause performance issues.

Considerations

The following considerations should be understood when using UMC with the specific operating systems or environments.

Windows NIC Teaming Considerations

UMC NIC Channels on the Same Physical Port Cannot Be Placed in the Same Team

Each of the channel NIC ports is presented to the operating system as a separate and unique NIC port. However, this can create an issue when the NIC channels on the same physical port are placed in the same team.

While teaming can provide increased reliability and increased bandwidth, this is not the case in this circumstance. Since the channels are assigned to the same physical port, reliability is not increased because it is still the same as a single physical port and not equal to two ports combined. In addition, bandwidth increases are not actually realized. Greater throughput could be achieved by assigning the combined bandwidths of the channels to a single channel as the overhead of teaming would not be a factor in that case.

Teaming is Supported on Multiple Channels

Teaming is supported on multiple channels under UMC. However, LACP configuration is not supported when UMC is enabled.

Windows Deployment Considerations

The following sections describe how to configure UMC effectively for Windows systems.

Using the LPVID

The LPVID is a convenient method to assign a single VLAN tag to a UMC NIC, without the host even recognizing the VLAN usage. This tag must be configured according to the network topology, such as using a separate UMC NIC for each class of traffic: network storage, backup, virtual machines, clients, and so on. The mapping of the LPVID to an operating system NIC device requires noting the PCI function number with the corresponding LPVID, and comparing that to the function number listed in the Windows Device Manager.

Teaming can be used with the LPVID. In this configuration, the LPVID is entirely hidden from the teaming software. Individual UMC NICs that form a team must all be configured with the same LPVID. By definition, that requires the functions be from separate Ethernet ports. The VLAN ID should not be configured anywhere in the teaming software or operating system.

Windows Device Manager VLAN ID

If you use VLAN IDs in the operating system, you must use different VLAN IDs than the LPVIDs, by specifying them within the NIC driver's property pages accessed through the Windows Device Manager. This does not require you to enter the PXESelect BIOS. See "Multiple VLAN Configuration in Windows" on page 16 for more information.

This VLAN ID replaces the LPVID for all transmit packets, but the function still receives both the LPVID packets and the new VLAN ID-configured packets. The VLAN ID configured in the Device Manager is visible to the operating system, unlike the LPVID. Therefore, receive packets may be filtered based on VLAN ID in teaming drivers or Hyper-V, so it may be necessary to configure these software components with the same VLAN ID in the following locations:

- Device Manager property page for each UMC NIC – This configures the VLAN ID.
- Teaming Software (if applicable) – VLAN IDs should only be defined within teaming when Hyper-V is not being used. In this case, the VLAN IDs must match the VLAN IDs defined in the Device Manager.

When Hyper-V is being used, do not define the VLAN IDs within teaming. They must be defined in the Device Manager and within the Hyper-V switch. The team should use its default native VLAN ID. This enables the team to run in VLAN promiscuous mode, which allows packets to pass unchanged whether they have a VLAN tag or not.

- Hyper-V Manager (if applicable) – The receive packets have a VLAN tag that is filtered by the virtual switch unless Hyper-V is configured for the same VLAN value. All Hyper-V interfaces must have the same VLAN ID.

The VLAN ID must be different than all LPVID values for the given Ethernet port.

Multiple VLAN Configuration in Windows

The Emulex NIC's VLAN ID parameter in Device Manager may contain a space-delimited list of VLAN IDs, or a VLAN ID range separated with a dash.

For example, the VLAN ID string "3-6 45 47" will filter on all VLANs 3, 4, 5, 6, 45, and 47. When multiple VLANs are configured, the driver is configured to receive all of these VLANs.

Note: If Hyper-V is in use, you should define the VLAN IDs in the Device Manager and within the Hyper-V switch. Do not define the VLAN IDs within the teaming utility.
If Hyper-V is not in use, the same VLAN IDs must be configured in the teaming utility.

Each NIC channel on a physical port should be configured with a unique set of VLANs. This allows the hardware to efficiently direct receive frames to the correct PCI function, even when the port is running in MAC promiscuous mode.

When UMC mode is enabled, untagged packets will not work. All traffic must have a VLAN assigned, whether explicitly in the host operating system or using the LPVID.

4. Troubleshooting

This section includes UMC troubleshooting information.

Table 4-1 UMC Troubleshooting

Issue	Answer/Solution
No network traffic on a particular channel.	<ul style="list-style-type: none">• LPVID value is set to 0 or an invalid value<ul style="list-style-type: none">◦ If the LPVID is set to 0, the LPVID has not been set. Enter a valid value for the LPVID. See the <i>Boot for NIC, iSCSI, and FCoE Protocols User Manual</i> for information on configuring the LPVID.◦ Verify that the LPVID and operating system VLAN ID are not the same value. They must be different values when UMC is enabled.◦ An invalid LPVID can send all traffic to the wrong VLAN. See the <i>Boot for NIC, iSCSI, and FCoE Protocols User Manual, OneCommand Manager Application User Manual, or OneCommand Manager CLI User Manual</i> for information on configuring the LPVID.• Admin logical link is down<ul style="list-style-type: none">◦ Indicates that the admin logical link is not set properly, or the minimum bandwidth on the channel has been set to 0. See the <i>Boot for NIC, iSCSI, and FCoE Protocols User Manual, OneCommand Manager Application User Manual, or OneCommand Manager CLI User Manual</i> for information on configuring the admin logical link.• Invalid switch configuration<ul style="list-style-type: none">◦ Verify that the switch configuration is set to allow VLANs.