



OneCommandTM NIC Teaming and VLAN Manager

Version 1.0

for Windows Server 2003

Windows Server 2008

Windows Server 2008 R2

User Manual

Copyright © 2003-2010 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.

US patent notice is given for one or more of the following: 6226680, 6247060, 6334153, 6389479, 6393487, 6427171, 6427173, 6434620, 6591302, 6658480, 6697868, 6751665, 6757746, 6941386, 6965941, 6687758, 7042898, 7133940, 7124205, 7089326, 6938092, 6996070.

Emulex, 3333 Susan Street
Costa Mesa, CA 92626

Introduction.....	1
NIC Teaming	1
Why Teaming?.....	1
Teaming Types	1
Failover (FO).....	1
Smart Load Balancing (SLB).....	1
Generic Trunking.....	2
Link Aggregation Control Protocol (LACP)	2
Installation	3
Updating the Application	3
Uninstalling the Application	3
Configuration	4
Starting the Application.....	4
Configuring Teams and VLANs.....	4
Creating and Configuring a Team.....	6
Removing an Adapter During Team Creation	7
Primary and Secondary Adapters	7
Configuring a VLAN for an Adapter.....	7
Team-Adapter Configuration Display.....	11
Troubleshooting	15

Introduction

NIC Teaming

In order to increase throughput and bandwidth, and to increase link availability, you can configure multiple NIC cards to appear to the network as a single interface. This is referred to as NIC teaming, or multilink trunking.

NIC teaming allows you to group multiple NICs as a single virtual device. Depending on the teaming mode, one or more interfaces can be active. When multiple NICs are combined this way, the group is called a team.

Why Teaming?

NIC teaming has several advantages.

- Increased bandwidth
Two or more Ethernet connections are combined to share the load, thus increasing bandwidth.
- Load balancing
Link aggregation enables distribution of processing and communication across multiple links.
- Higher link availability
Prevents a single link failure from disturbing traffic flow.

Teaming Types

There are different types of teaming:

- Switch independent
 - Failover - If configured for fault tolerance, the system provides only failover.
 - Smart load balancing - If configured for load balancing, failover is included.
- Switch dependent
 - Generic trunking - Link aggregation static mode
 - Link Aggregation Control Protocol (LACP)

Failover (FO)

Failover is a basic feature of teaming. Though a team may have multiple members, only one member is active at a time. When the active team member disconnects (due to link down, disabled or any other reason) the failover mechanism selects another team member (which is in linkup state) and traffic continues.

All the team members use the same MAC address, the MAC address of the primary team member. When the primary team member reports linkup state, traffic fails over to the primary adapter, if the team is configured for auto failback.

Smart Load Balancing (SLB)

Load balancing works with any Ethernet switch. It requires no trunking configuration of the switch itself. The team of NICs advertise multiple MAC addresses and one or more IP addresses. The team MAC address is selected by the Virtual Team Adapter (VTA) from the list of load balancing members. When

the server receives an ARP request, the software-networking stack sends an ARP reply with the team MAC address. To begin the load balancing process, the NIC Teaming driver modifies this ARP reply by changing the source MAC address to match the MAC address of one of the physical adapters.

You can balance both transmit and receive loads to maintain in-order delivery of frames.

Generic Trunking

Generic Trunking, (also called IEEE 802.3ad Link Aggregation static mode) is a switch-assisted teaming mode that requires configuring the ports at both ends of the link, both the server and the switch ports.

In this mode, the team advertises one MAC Address and one IP Address when the protocol stack responds to ARP requests. In addition, each physical adapter in the team uses the same team MAC address when transmitting frames. This is possible since the switch at the other end of the link is aware of the teaming mode and will handle the use of a single MAC address by every port in the team. The forwarding table in the switch will reflect the trunk as a single virtual port.

In this teaming mode, the NIC Teaming driver controls load balancing and failover for outgoing traffic only, while incoming traffic is controlled by the switch firmware and hardware.

Link Aggregation Control Protocol (LACP)

IEEE 802.3ad LACP is similar to link aggregation static mode except that it uses the LACP to negotiate the ports that make up the team. The LACP must be enabled at both the server and the switch for the team to operate.

The LACP automatically determines which member links can be aggregated and then aggregates them. It provides for the controlled addition and removal of physical links for the link aggregation so that no frames are lost or duplicated. The marker protocol (that can be enabled for LACP) provides removal of aggregate link members.

The link aggregation group advertises a single MAC address for all the ports in the team. The MAC address of the team/aggregator can be the MAC addresses of one of the NICs in the group. The LACP and marker protocols use a multicast destination address.

The link aggregation control function determines which links may be aggregated. It then binds the ports to an aggregator function in the system and monitors conditions to determine if a change in the aggregation group is required. Link aggregation combines the individual capacity of multiple links to form a high performance virtual link. The failure or replacement of a link in an LACP trunk does not cause loss of connectivity. The traffic is simply failed over to the remaining links in the trunk.

Installation

The OneCommand NIC Teaming and Multiple VLAN Manager is installed automatically when the Windows driver is installed.

If the Windows driver installation process does not include rebooting the computer, you must reboot manually.

Updating the Application

If you want to update the NIC Teaming manager to preserve the existing configuration, run `elx_octeamupdate.exe` after you install the newer version.

Uninstalling the Application

If you want to remove all team interfaces and VLAN interfaces, run `elx_octeamuninstall.exe` before you uninstall the package.

Configuration

A team of adapters functions as a single virtual network interface and appears the same as a non-teamed adapter to other network devices.

A protocol address such as an IP address is usually assigned to the physical adapter. However, when the OneCommand NIC Teaming and Multiple VLAN Manager is installed, the protocol address is assigned to the Team Adapter and not to the physical adapters that make up the team.

The IPCONFIG /all command shows the IP and MAC addresses of the virtual adapter and not of the individual physical adapters.

Starting the Application

Click the **OC NIC Teaming Manager** shortcut to start the application.



Figure 1: The OC NIC Teaming Manager shortcut

Configuring Teams and VLANs

You configure teams and VLANs using the NIC Teaming and VLAN Manager dialog box.

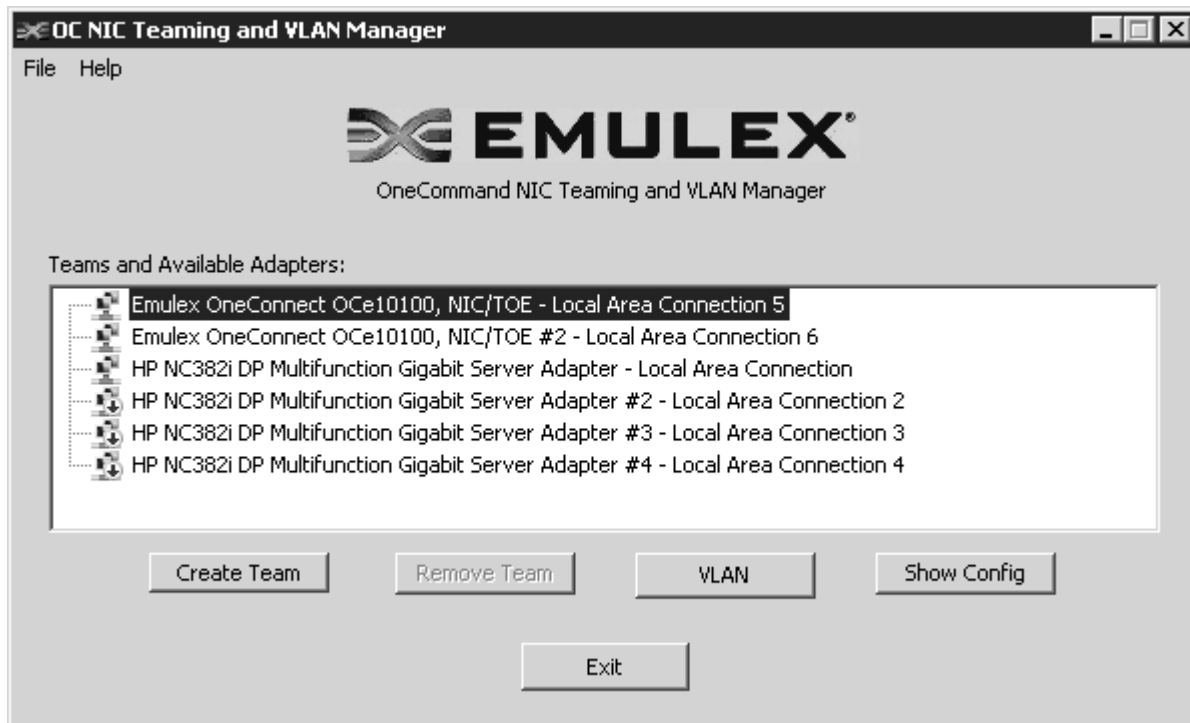


Figure 2: The NIC Teaming and VLAN Manager dialog box

NIC Teaming and VLAN Manager Field Definitions

- Available Adapters - Displays all the physical adapters as well as Team adapters configured, on the system as a tree view. For example:



NIC Teaming and VLAN Manager buttons

- Create Team - Click to create a new team. The team configuration dialog appears.
- Remove Team - Click to delete a team configured earlier. If no teams have been configured, the button is unavailable. Select the earlier configured team. Click **Remove Team**.
- VLAN - Click to configure VLAN for the selected physical or team adapter:
- Show Config - Click to view adapter configuration. Select the adapter from the list of Available Adapters to view its configuration information.
- Exit - Click to exit the application.

Creating and Configuring a Team

Use the Create Team dialog box to create a team.

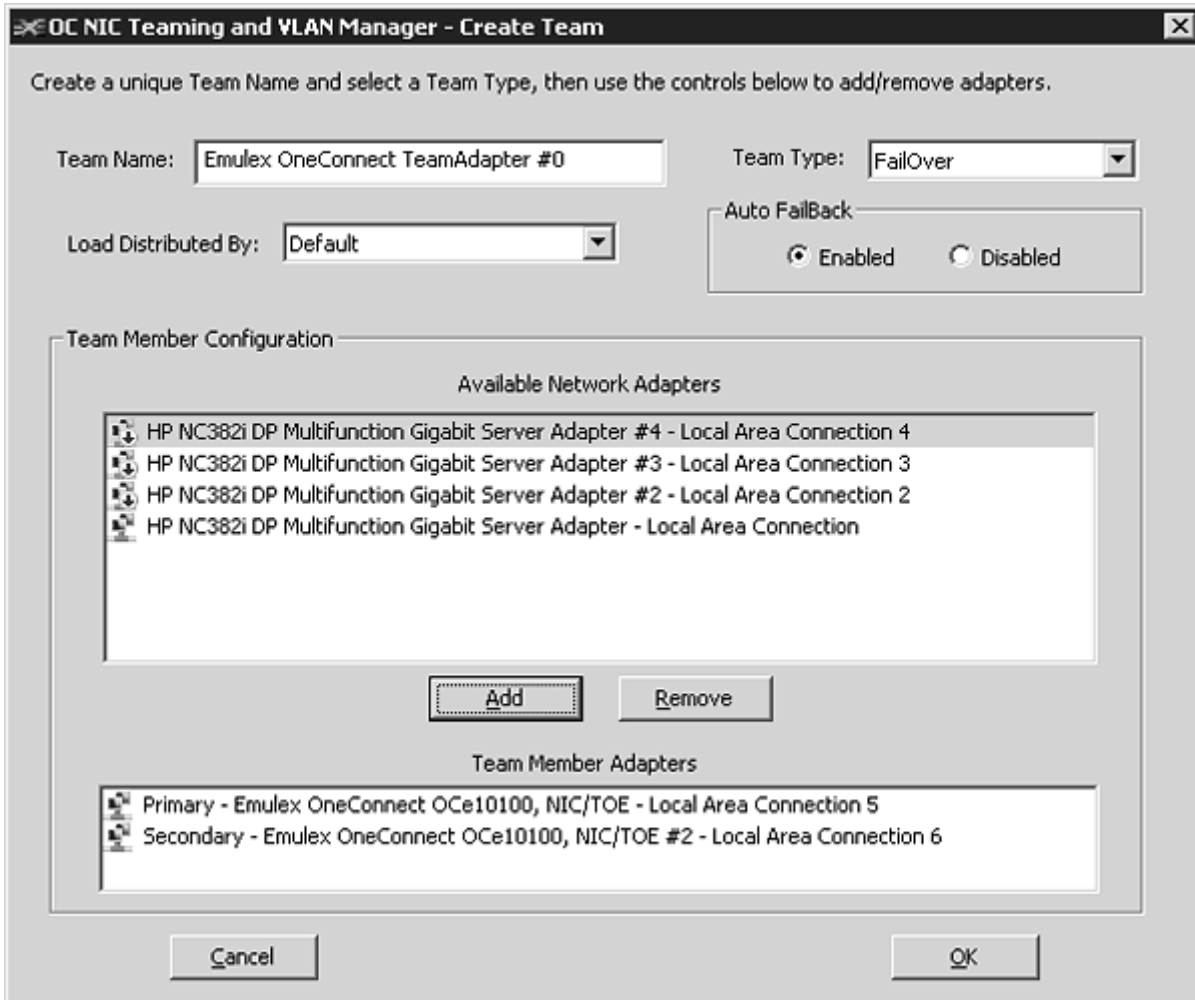


Figure 3: The Create Team dialog box

To create and configure a new team:

1. Enter a new team name.
2. Choose a Team Type. Select Failover > Load balancing > 802.3ad static > LACP (Default: Failover).
3. If Team Type is not Failover, select the criteria for Load Distributed By:
 - a. Default - This is the default option.
 - b. Destination MAC Address.
 - c. Destination IP Address.
4. Enable or Disable Auto Failback. The default is Enable.
5. The Available Network Adapters area lists adapters that are not members of any team. To add an adapter to the team, select the adapter from the Available Network Adapters list and click **Add**. The adapter appears in the Team Member Adapters list.

Note: Every team must include at least one Emulex adapter.

6. The Team Member Adapters area lists adapters that are members of the team.
7. Click **OK** to create a new team.

Removing an Adapter During Team Creation

To remove an adapter during team creation:

1. From the Create Team dialog box, select the adapter you want to remove from Team Member Adapters list.
2. Click **Remove**. The adapter is removed from the Team Member Adapters list.
3. Click **OK**.

Primary and Secondary Adapters

The first Emulex adapter you add to a team is always the primary adapter. To change the primary adapter, remove the primary adapter from the team and then add the adapter you want to be the primary adapter to the team.

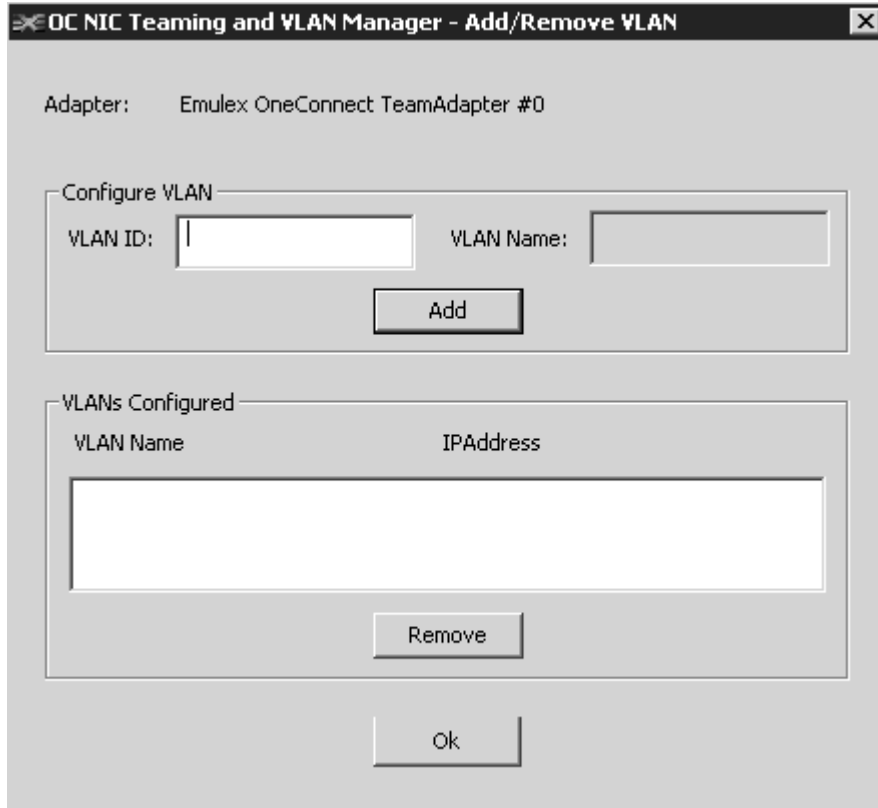
Note: If the adapter you want to be the primary adapter is already a team member, you must remove it and add it to the team.

Configuring a VLAN for an Adapter

To configure a VLAN for a physical or team adapter:

1. In the NIC Teaming and VLAN Manager dialog box, see Figure 2, select the physical/team adapter in Available Adapters to which you want to add a VLAN.

2. Click **VLAN**. The Add/Remove VLAN dialog appears. Figure 4 shows the Add/Remove VLAN dialog-box.



OC NIC Teaming and VLAN Manager - Add/Remove VLAN

Adapter: Emulex OneConnect TeamAdapter #0

Configure VLAN

VLAN ID: VLAN Name:

Add

VLANs Configured

VLAN Name	IP Address

Remove

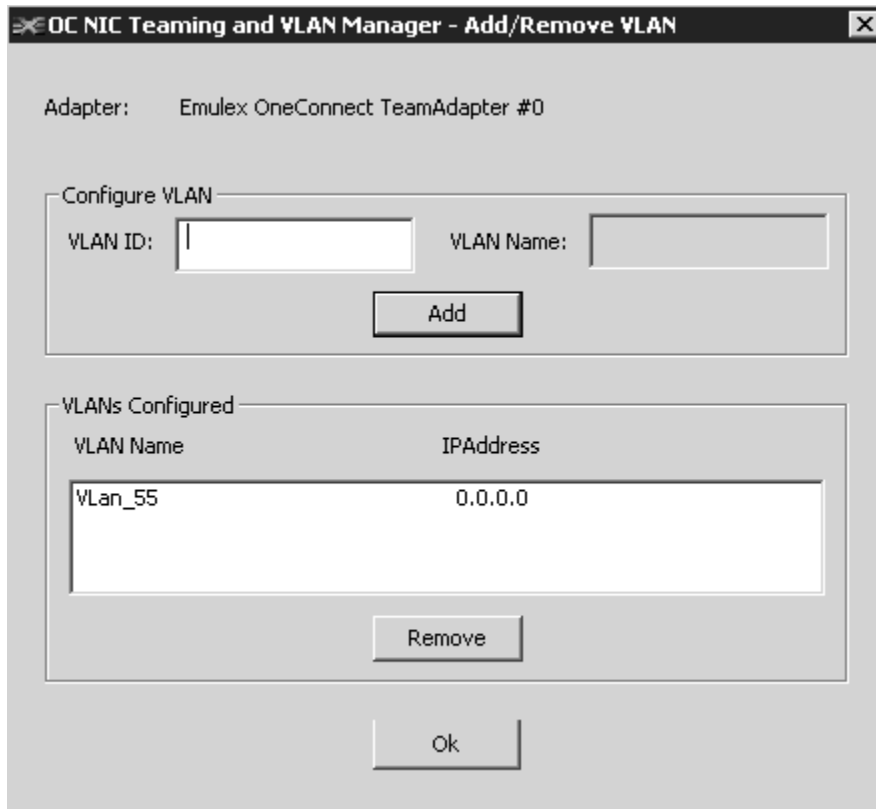
Ok

Figure 4: Add/Remove VLAN dialog box

3. Enter a VLAN ID. Enter a VLAN tag (value: 1-4096). The VLAN Name shows the VLAN Name in the format Vlan_<VLAN ID>.
4. Click **Add** to add the VLAN to the adapter. You can create multiple VLANs for an adapter. The VLANs Configured list shows the list of all VLANs configured for the adapter.

Note: To delete a VLAN configured for the adapter, select the VLAN from the VLANs Configured list and click **Remove**.

Figure 5 shows the newly created VLAN with the list of currently configured VLANs for the adapter.



Adapter: Emulex OneConnect TeamAdapter #0

Configure VLAN

VLAN ID: VLAN Name:

Add

VLANs Configured

VLAN Name	IPAddress
Vlan_55	0.0.0.0

Remove

Ok

Figure 5: Newly created VLAN screen

- Click **Ok**. The VLAN is added to the list of configured VLANs.

Figure 6 shows the VLAN in the list.

OC NIC Teaming and VLAN Manager - Add/Remove VLAN

Adapter: Emulex OneConnect TeamAdapter #0

Configure VLAN

VLAN ID: VLAN Name:

Add

VLANs Configured

VLAN Name	IPAddress
Vlan_54	169.254.131.104
Vlan_55	169.254.250.25

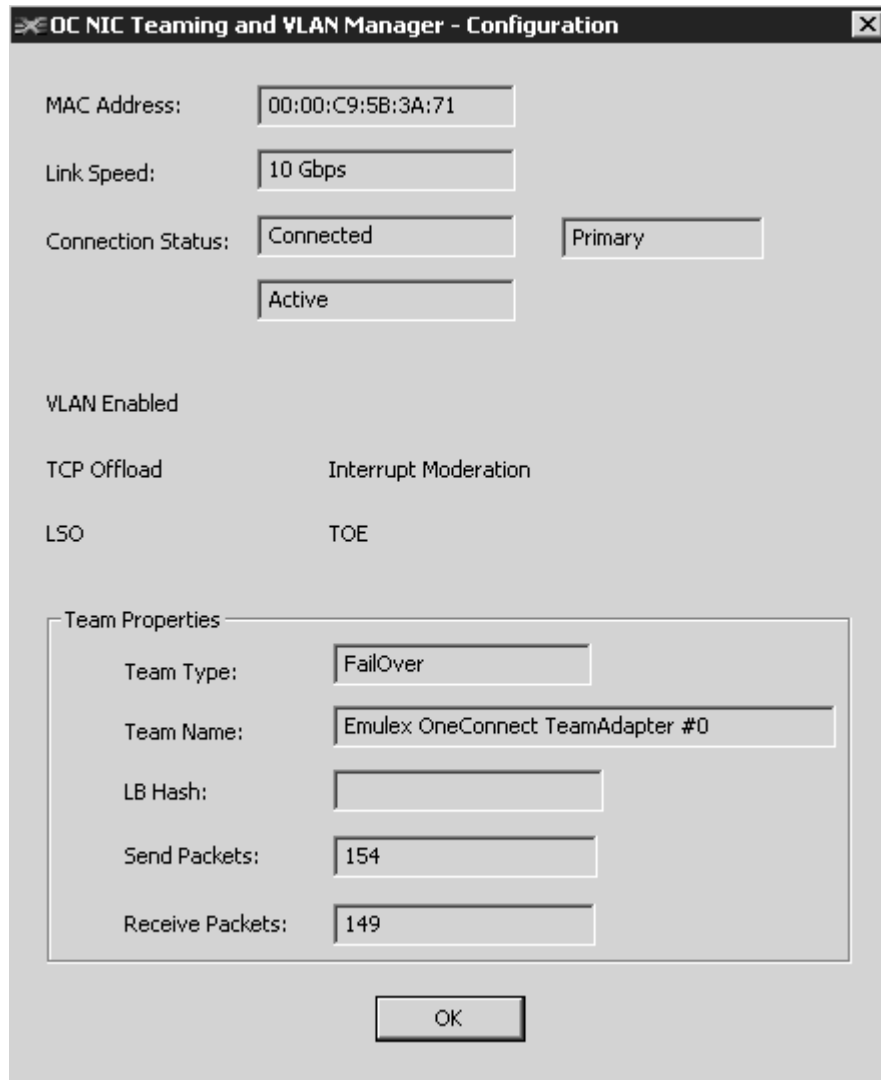
Remove

Ok

Figure 6: Configured VLANs

Team-Adapter Configuration Display

The Teaming and VLAN Manager Configuration window is read-only and shows the current configuration of the adapter.



OC NIC Teaming and VLAN Manager - Configuration

MAC Address: 00:00:C9:5B:3A:71

Link Speed: 10 Gbps

Connection Status: Connected Primary

Active

VLAN Enabled

TCP Offload Interrupt Moderation

LSO TOE

Team Properties

Team Type: FailOver

Team Name: Emulex OneConnect TeamAdapter #0

LB Hash:

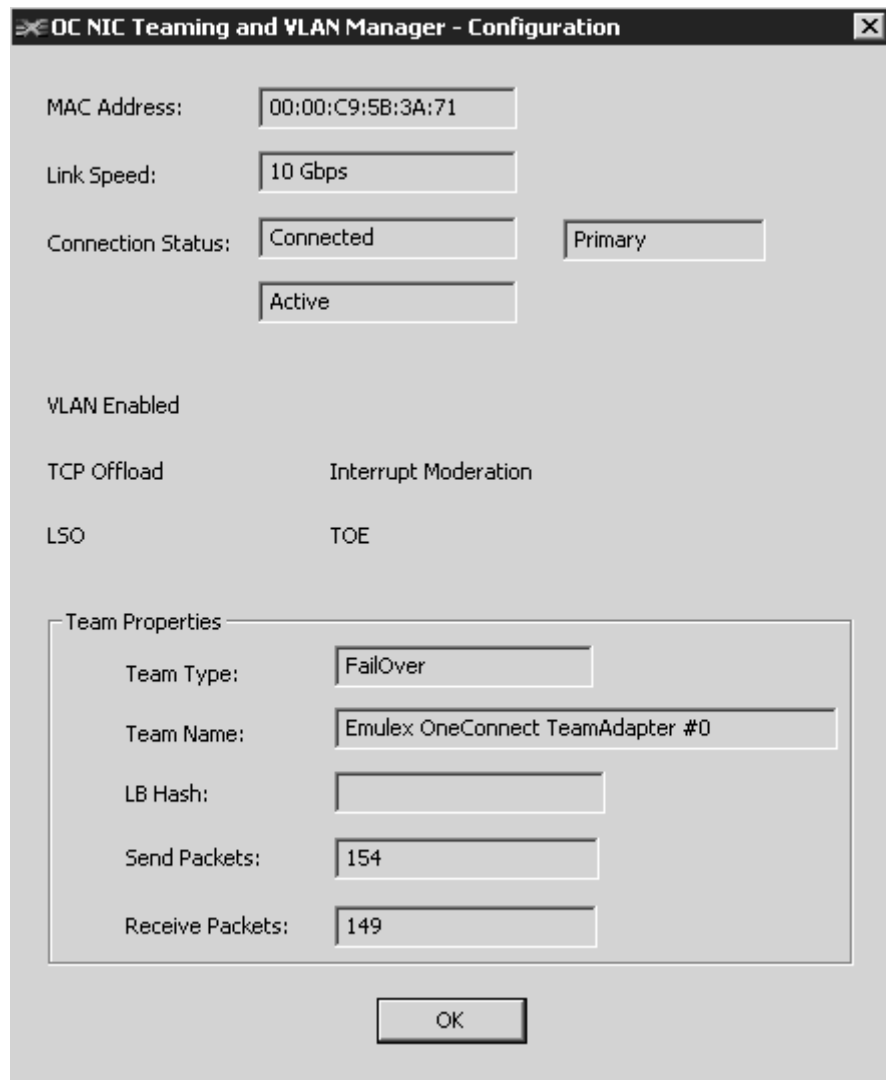
Send Packets: 154

Receive Packets: 149

OK

Figure 7: The Teaming and VLAN Manager Configuration window

Figure 8 shows the Teaming and VLAN Manager Configuration display with the Primary adapter Active.



OC NIC Teaming and VLAN Manager - Configuration

MAC Address: 00:00:C9:5B:3A:71

Link Speed: 10 Gbps

Connection Status: Connected Primary

Active

VLAN Enabled

TCP Offload Interrupt Moderation

LSO TOE

Team Properties

Team Type: FailOver

Team Name: Emulex OneConnect TeamAdapter #0

LB Hash:

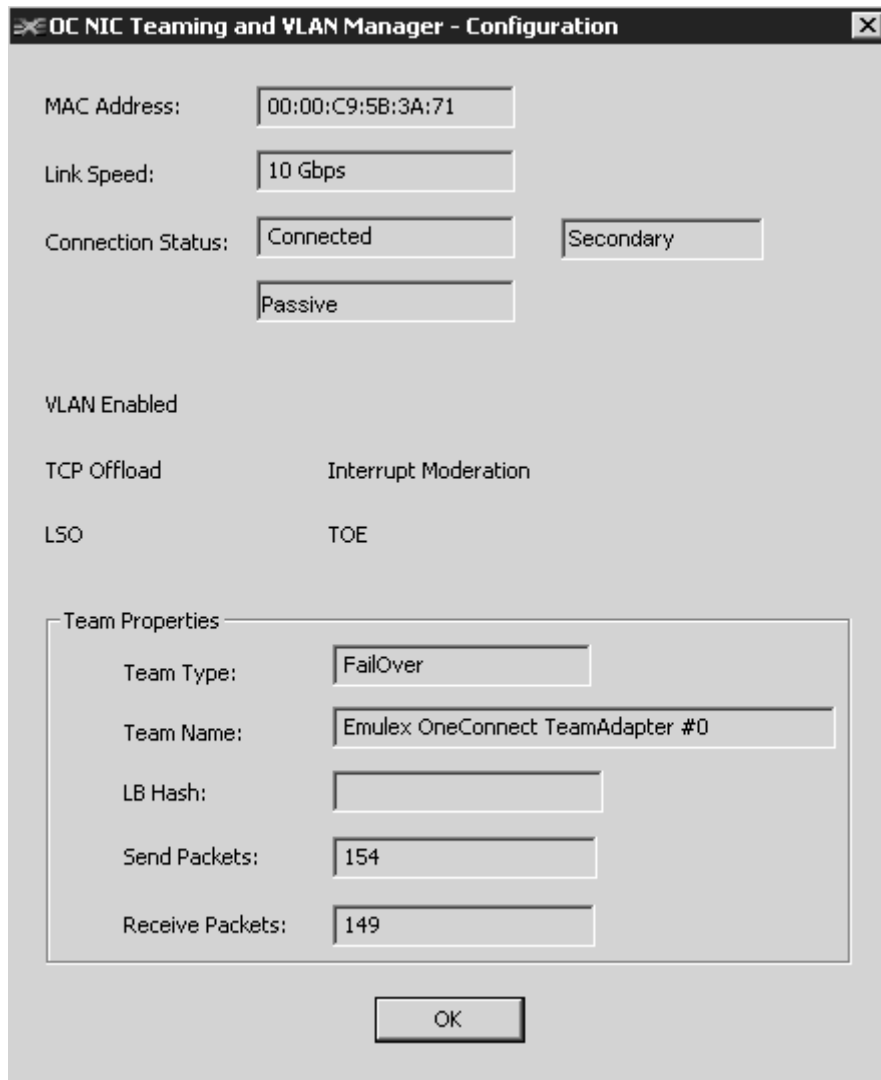
Send Packets: 154

Receive Packets: 149

OK

Figure 8: Team Member - Active

Figure 9 shows the Teaming and VLAN Manager Configuration display with the Secondary Team member Passive.



OC NIC Teaming and VLAN Manager - Configuration

MAC Address: 00:00:C9:5B:3A:71

Link Speed: 10 Gbps

Connection Status: Connected Secondary Passive

VLAN Enabled

TCP Offload Interrupt Moderation

LSO TOE

Team Properties

Team Type: FailOver

Team Name: Emulex OneConnect TeamAdapter #0

LB Hash:

Send Packets: 154

Receive Packets: 149

OK

Figure 9: Team Member - Passive

Team Member States

Each team member can be in one of the following states as shown in the Connection Status box:

- **Added:** When a Team member is successfully bound to the virtual adapter it becomes a member of team.
- **Connecting:** Once the Team member is Added, the virtual adapter will query it for the state of the physical adapter. At this point, the team member is in the **Connecting** state.
- **Connected:** When the virtual team adapter gets the status of the physical adapter, depending on the status returned, it enters either the **Connected** or the **Disconnected** state.
- **Active:** If LACP is set for the team, the LACP handshake starts and LACP adds the link to the team, and the adapter's status changes to **Active**. Otherwise, it remains in **Standby**. In non-LACP mode, after connecting, if you have not set the adapter to **Standby**, via the GUI, the adapter becomes **Active**.

- **Disconnected:** When the virtual team adapter gets the link-down status indication, either through LACP or because you disable or remove the link, it is disconnected and it releases the IP address. It is no longer a member of the team
- **Deleted:** If the adapter is **Disconnected**, you can delete it. If you delete a team member or the whole team, the virtual team adapter enters this state. It is no longer a member of the team.

Troubleshooting

There are several circumstances in which your system may operate in an unexpected manner. The Troubleshooting section provides information regarding unusual situations.

NDIS

DESCRIPTION: The team members are incompatible. The Virtual Team driver is NDIS 6.0 and team member (underlying miniport driver) is NDIS 5.0 / NDIS 5.1.

SEVERITY: Information

ACTION: None. NDIS takes care of compatibility issues.

Header Data Split

DESCRIPTION: Incompatible adapters. The team member (underlying miniport driver) supports the Header Data split feature.

SEVERITY: Information

ACTION: None. NDIS 6.1 onwards supports Header Data Split. Windows 2008 drivers can handle Header Data Split. If there are any old-NDIS drivers in the network stack, not supporting Header Data Split, then NDIS dynamically turns it OFF.

RSS

DESCRIPTION: RSS (Receive Side Scaling) adversely affects the Virtual Team Miniport Driver (IM driver).

SEVERITY: Information

ACTION: None: In Non 802.3ad mode or 802.3ad static mode, the driver is working as a pass-through driver. So RSS does not affect it. It creates ARP requests and receives ARP replies, but if the OOB fields in the packet are set to 0 (for RSS hash), there is no problem.

In 802.3ad LACP mode creates LACP packets, driver can use the information from OIDs to assign connection on the same CPU that miniport and protocol driver are using. The driver disables RSS by intercepting the OIDS and reports RSS not supported.

Failover

DESCRIPTION: RSS Failover does not work.

SEVERITY: Major

LOG: Teaming members use different hash values.

ACTION: None. A Virtual Teamed adapter sets the RSS capabilities to be similar for all the teaming members using OIDs. The RSS capabilities of the first adapter in a team is used as a base and the RSS capabilities of the subsequent adapters will be compared to it. If different adapters advertise different capabilities, only common capabilities are supported by the Team driver.

Offload features

DESCRIPTION: Adapters cannot be teamed because they support different offloading features.

SEVERITY: Information

ACTION: None. The minimum matching capabilities, among the adapters, becomes the offload capabilities.

SEE ALSO: Table 1, "Offload Configuration," on page 16

Table 1: Offload Configuration

Offload Capabilities	Available Configuration	Accepted Configuration
1. Checksum Offload	Ipv4, ipv6, TX,RX, TCP,UDP,IP,TCPOPTION, IPOPTION	Based on common factors of TCP/UDP/ IP checksum.
2. Large Send Offload	Ipv4,ipv6 Maxoffloadsize, Minsegmentcount	Based on common factors and minimum of mxoffloadsize and min segment count of all the teamed adapters.
3. TCP Connection offload	Ip4, ipv6 SupportSack TcpConnectionOffloadCapacity	Based on common factors of ipv4/ipv6/ support sack and minimum of TcpConnectionOffloadCapacity.
4. RSS	Base CPU Hash Information Indirection Table Size Indirection Table Offset Secret Key size Secret Key offset	After querying the team adapters, pass the common information to the stack and the pass data from the stack to all the team members (key, key size).

Jumbo frames

DESCRIPTION: Team adapters do not support Jumbo frames. because they support different offloading features.

SEVERITY: Information

ACTION: Different adapters support different frame sizes. Some support 1514 and 9032 only; some support 2048, 4096 and so forth. The Virtual Team driver considers the minimum packet size reported by `OID_GEN_MAXIMUM_FRAME_SIZE` and uses it as its MTU. Show Config shows whether the Team adapter supports LSO or not.

VLANs

DESCRIPTION: The adapter supports VLANs, different from the VLAN ID of the virtual team adapter.

SEVERITY: Information

ACTION: If the adapter is configured for VLAN and if you create a VLAN on top of the adapter, data transmission will not work properly.

Multiple VLAN configuration

DESCRIPTION: Must configure multiple VLAN IDs for each virtual team adapter.

SEVERITY: Information

ACTION: Configure an Intermediate Driver (IM) on the top of the virtual team adapter and create VLANs on the virtual team adapter.

Promiscuous mode

DESCRIPTION: Teamed physical adapter is not in promiscuous mode.

SEVERITY: Information

ACTION: None. The teaming/VLAN GUI tries to set the current MAC address for the team member. The driver reads the current MAC address of the adapter. If team MAC address is different from the current MAC address of the adapter, it sets filter type to be Promiscuous mode.

IPv6

DESCRIPTION: Network uses IPv6 - driver incompatible.

SEVERITY: Information

ACTION: None: Windows operating systems, from VISTA onwards, use only one network stack, so when the driver is bound to the TCP/IP stack, it can receive IPV4/IPV6 packets and it can handle both the packets.

The driver supports IPV6.