D-Link® **DXE-820T**

## Dual Port 10GBASE-T RJ-45 PCI Express Adapter

# Manual

V1.00

# Index

**Information in this document is subject to change without notice.**

Documentation is provided as is without warranty of any kind, either express or implied, including any kind of implied or express warranty of non-infringement or the implied warranties of merchantability or fitness for a particular purpose.

Broadcom Corporation reserves the right to make changes without further notice to any products or data herein to improve reliability, function, or design. Information furnished by Broadcom Corporation is believed to be accurate and reliable. However, Broadcom Corporation does not assume any liability arising out of the application or use of this information, nor the application or use of any product or circuit described herein, neither does it convey any license under its patent rights or the rights of others.

Broadcom, the pulse logo, Connecting everything, the Connecting everything logo, NetXtreme, Ethernet@Wirespeed, LiveLink, and Smart Load Balancing are among the trademarks of Broadcom Corporation and/or its affiliates in the United States, certain other countries, and/or the EU. Microsoft and Windows are trademarks of Microsoft Corporation. Linux is a trademark of Linus Torvalds. Intel is a trademark of Intel Corporation. Magic Packet is a trademark of Advanced Micro Devices, Inc. Red Hat is a trademark of Red Hat, Inc. PCI Express is a trademark of PCI-SIG. Any other trademarks or trade names mentioned are the property of their respective owners.

# 1. Introduction

Thank you for choosing the D-Link DXE-820T, the value leader among 10 Gigabit Ethernet adapters for PCI Express Bus personal computers. The D-Link DXE-820T PCI Express 10 Gigabit Ethernet Adapter is a high performance adapter designed for the high-speed PCI Express Bus Architecture. The PCI Express 10 Gigabit Ethernet Adapter offers increased bandwidth, reliability, and more functionality than standard PCI network cards. It is specifically designed to allow throughput at rates up to 40 Gbps, thus eliminating the bottleneck that exists with current 32 and 64 bit PCI bus architecture

## 1.1. Functionality and Features

### 1.1.1. Functional Description

DXE-820T is a 10 GbE converged network interface controller (C-NIC) that can simultaneously perform accelerated data networking and storage networking on a standard Ethernet network. The C-NIC offers acceleration for popular protocols used in the data center, such as:

- Internet Small Computer Systems Interface (iSCSI) offload for accelerating network storage access featuring centralized boot functionality (iSCSI boot).

Enterprise networks that use multiple protocols and multiple network fabrics, benefit from the C-NICs ability to combine data communications, storage, and clustering over a single Ethernet fabric by boosting server CPU processing performance and memory utilization while alleviating I/O bottlenecks.

Using the Broadcom teaming software, you can split your network into virtual LANs (VLANs) as well as group multiple network adapters together into teams to provide network load balancing and fault tolerance functionality. See **Configuring Teaming** and **Broadcom Gigabit Ethernet Teaming Services** for detailed information about teaming. See **Virtual LANs**, for a description of VLANs. See **Configuring Teaming** for instructions on configuring teaming and creating VLANs on Windows operating systems

### 1.1.2. Features

The following is a list of DXE-820T features.

- NIC Partitioning
- Other performance features
  - TCP, IP, UDP checksum

  - TCP segmentation
  - Adaptive interrupts
  - Receive Side Scaling (RSS)
- Manageability
  - Statistics for SNMP MIB II, Ethernet-like MIB, and Ethernet MIB (IEEE Std 802.3z, Clause 30)
  - SMBus controller
  - IPMI support
- Advanced network features
  - Jumbo frames (up to 9 KB). The OS and the link partner must support jumbo frames.
  - Virtual LANs
  - IEEE Std 802.3ad Teaming
  - Smart Load Balancing Teaming
  - Flow Control (IEEE Std 802.3x)
  - LiveLink™ (supported in both the 32-bit and 64-bit Windows operating systems)
- Logical Link Control (IEEE Std 802.2)
- Layer-2 Priority Encoding (IEEE Std 802.1p)
- High-speed on-chip RISC processor
- Up to 3 classes of service (CoS)
- Integrated 96 KB frame buffer memory
- Quality of Service (QoS)
- Support for multicast addresses via 256 bits hashing hardware function
- JTAG support
- PCI Power Management Interface (v1.1)
- 64-bit BAR support
- iSCSI Boot support
- Virtualization
  - Microsoft
  - VMware

# 2. System Requirements

## 2.1. Hardware Requirements

- IA32- or EMT64-based computer that meets operating system requirements
- One open PCI Express slot. Depending on the PCI Express support on your adapter, the slot may be of type PCI Express 1.0a x1, PCI Express 1.0a x4, or PCI Express Gen2 x8.
- 128-MB RAM (minimum)

## 2.2. Operating System Requirements

### 2.2.1. Microsoft Windows

One of the following versions of Microsoft Windows:

- Windows Server 2008 family
- Windows Server 2008 R2 family
- Windows Server 2012 family

### 2.2.2. Linux

Although the adapter driver should work with many Linux kernel versions and distributions, it has only been tested on 2.4x kernels (starting from 2.4.24) and 2.6.x kernels. The driver may not compile on kernels older than 2.4.24. Testing is concentrated on i386 and x86_64 architectures. Only limited testing has been done on other architectures. Minor changes to some source files and Makefile may be needed on some kernels.

**Note**: *Support for the 2.4.21 kernels is provided in Red Hat Enterprise Linux 3.*

### 2.2.3. VMware ESX

- VMware ESX
- VMware ESX 3.5
- VMware ESX 4.0
- VMware ESX 4.1
- VMware ESXi 5.0
- VMware ESXi 5.1

## 2.3. Connecting the Network Cables

**Note***: DXE-820T supports Automatic MDI Crossover (MDIX), which eliminates the need for crossover cables when connecting machines back-to-back. A straight-through Category 5 cable allows the machines to communicate when* *connected directly together.*

1. Select an appropriate cable. Table 1 lists the copper cable requirements for connecting to 10/100/1000BASE-T and 10GBASE-T ports:

<table>
<tr><th colspan="4">Table 1: 10/100/1000BASE-T and 10GBASE-T Cable Specifications</th></tr>
<tr><td>Port Type</td><td>Connector</td><td>Media</td><td>Maximum Distance</td></tr>
<tr><td>10BASE-T</td><td>RJ-45</td><td>Category 3, 4, or 5 unshielded twisted pairs (UTP)</td><td>100m (328 ft.)</td></tr>
<tr><td>100/1000BASE-T[1]</td><td>RJ-45</td><td>Category 5[2] UTP</td><td>100m (328 ft.)</td></tr>
<tr><td>10GBASE-T</td><td>RJ-45</td><td>Category 6[3] UTP<br>Category 6A[3] UTP</td><td>50m (164 ft.)<br>100m (328 ft.)</td></tr>
<tr><td colspan="4">1: 1000BASE-T signaling requires four twisted pairs of Category 5 balanced cabling, as specified in ISO/IEC 11801:2002 and ANSI/EIA/TIA-568-B.<br>2: Category 5 is the minimum requirement. Category 5e and Category 6 are fully supported.<br>3: 10GBASE-T signaling requires four twisted pairs of Category 6 or Category 6A (augmented Category 6) balanced cabling, as specified in ISO/IEC 11801:2002 and ANSI/TIA/EIA-568-B.</td></tr>
</table>

2. Connect one end of the cable to the RJ-45 connector on the adapter.
3. Connect the other end of the cable to an RJ-45 Ethernet network port.

# 3. Installing Management Applications

The Broadcom Advanced Control Suite version 4 (BACS4) is a management application for configuring the DXE-820T, also known as Converge Network Adapters (CNAs). BACS4 software operates on Windows and Linux server and client operating systems. This chapter describes how to install the BACS4 management application.

There are two main components of the BACS4 utility: the provider component and the client software.

A provider is installed on a server, or "managed host", that contains one or more CNAs. The provider collects information on the CNAs and makes it available for retrieval from a management PC on which the client software is installed. The client software enables viewing information from the providers and configuring the CNAs. The BACS client software includes a graphical user interface (GUI) and a command line interface (CLI).

## 3.1. Installing the Broadcom Advanced Control Suite and Related Management Applications

### 3.1.1. Installing on a Windows System

The Broadcom Advanced Control Suite (BACS) software and related management applications can be installed from the installation CD or by using the silent install option.

**The following are installed when running the installer:**

- **Control Suite.** Broadcom Advanced Control Suite (BACS). If selected, a GUI and a CLI client will be installed.
- **BASP.** Broadcom Advanced Server Program. This is a Broadcom intermediate NDIS driver to configure VLAN, Team, Load Balancing etc.
- **SNMP.** The Simple Network Management Protocol subagent. This feature allows the SNMP manager to monitor the Broadcom network adapters.
- **CIM Provider.** Common Information Model provider. This component presents the network adapter information to WMI based management applications. Select this component on a host which has D-Link DXE-820T adapter installed and which you want to manage using the GUI client.

**Notes:**

- Ensure that the Broadcom network adapter(s) is physically installed in the system before installing BACS.
- Before you begin the installation, close all applications, windows, or dialog boxes.
- BASP is not available on Windows Small Business Server (SBS) 2008.

**Using the Installer**

To install the management applications:

1. Insert the installation CD into the CD or DVD drive.
2. On the installation CD, open the MgmtApps folder, select IA32 or x64, and then double-click Setup.exe to open the InstallShield Wizard.
3. Click **Next** to continue.
4. After you review the license agreement, click the box to accept the terms in the license agreement and then click **Next** to continue.
5. Select the features you want installed.
6. Click **Next**.
7. Click **Install**.
8. Click **Finish** to close the wizard.
9. After successful installation, you can start the GUI from Windows Start menu.

**Using Silent Installation**

**Notes:**

- All commands are case sensitive.
- User must "Run as Administrator" for Vista when using "msiexec" for "silent" install/uninstall(s).
- For detailed instructions and information about unattended installs, refer to the Silent.txt file in the MgmtApps folder.

**To perform a silent install (or upgrade) from within the installer source folder, type the following:**

```
setup /s /v/qn
```

If performing a silent upgrade, your system may reboot automatically. To suppress the reboot, type the following:

```
setup /s /v"/qn REBOOT=ReallySuppress"
```

**To perform a silent install and create a log file**

Type the following:

```
setup /s /v"/qn /L f:\ia32\1testlog.txt"
```

The 1testlog.txt log file will be created at f:\ia32.

**To perform a silent uninstall from any folder on the hard drive**

```
msiexec /x "{26E1BFB0-E87E-4696-9F89-B467F01F81E5}" /qn
```

**Notes:**

- The hexadecimal number above may differ from your current installer. Check the Key name corresponding with the Broadcom Advanced Control Suite (BACS) application in HKLM\Software\Microsoft\Windows\CurrentVersion\Uninstall for the correct hexadecimal number.
- After performing a silent uninstall, it is necessary to reboot the system before reinstalling this installer. If a reboot is not performed, BASP will not install correctly.

**To perform a silent install by feature on IA32 platforms**

Use ADDSOURCE to include any of the features listed below.

**Note:** *CHM32 or CHM64 installs the BACS help file and must be included when installing the BACS feature.*

```
setup /s /v"/qn ADDSOURCE=BACSi32,CHM32,BASPi32,SNMPi32,CIMi32"
```

**To perform a silent install by feature on AMD64/EM64T platforms**

Type the following:

```
setup /s /v"/qn ADDSOURCE=BACSa64,CHMa64,BASPa64,SNMPa64"
```

**To perform a silent install from within a batch file**

To perform a silent install from within a batch file type the following and wait for the install to complete before continuing with the next command line:

```
start /wait setup /s /w /v/qn
```

### 3.1.2. Installing on a Linux System

The Broadcom Advanced Control Suite (BACS) software can be installed on a Linux system using the Linux RPM package. This installation includes a BACS GUI and a CLI client.

**Before you begin:**

- Ensure that the Broadcom network adapter(s) is physically installed and the appropriate device driver for the NICis is installed on the system that will be managed by this utility.
- Ensure that the CIM provider is installed properly on the system that will be managed by this utility.
- For managing iSCSI on Linux hosts, ensure that the open-iscsi and sg utilities are installed on the Linux host.

**To install BACS**

1. Download the latest BACS management application RPM package.
2. Install the RPM package using the following command:

```
% rpm -i BACS-{version}.{arch}.rpm
```

**To Use BACS**

- To use the GUI, on XWindow, double-click the BACS4 desktop icon, or access the BACS program from the task bar under System Tools.
- To use BACS CLI, refer to the file BACSCLI_Readme.txt provided with the release files.

**To remove BACS**

To uninstall the RPM package, use the following command:

```
% rpm -e BACS
```

### 3.1.3. Managing Management Applications (Windows)

**Modifying the Management Application**

To modify the management applications:

In Control Panel, double-click **Add or Remove Programs**.

1. Click **Broadcom Management Programs** and then click **Change**.
2. Click **Next** to continue.
3. Click **Modify** to change program features.
4. Click **Next** to continue.
5. Click on an icon to change how a feature is installed.
6. Click **Next**.
7. Click **Install**.

8. Click **Finish** to close the wizard.
9. Reboot your system to complete the modification of the management applications.

### Repairing Management Applications

To repair the management applications:

1. In Control Panel, double-click **Add or Remove Programs**.
2. Click **Broadcom Management Programs**, and then click **Change**.
3. Click **Next** to continue.
4. Click **Repair** to repair errors in installed applications.
5. Click **Next** to continue.
6. Click **Install**.
7. Click **Finish** to close the wizard.

### Removing Management Applications

To remove all management applications:

1. In Control panel, double-click **Add or Remove Programs**.
2. Click **Broadcom Management Programs**, and then click **Remove**.
3. Reboot your system to complete the removal of management applications.

To remove the management application using the CLI:

Enter following command:

```
rpm -e BACS
```

# 4. Advanced Teaming Concepts

The concept of grouping multiple physical devices to provide fault tolerance and load balancing is not new. It has been around for years. Storage devices use RAID technology to group individual hard drives. Switch ports can be grouped together using technologies such as Cisco Gigabit EtherChannel, IEEE 802.3ad Link Aggregation, Bay Network Multilink Trunking, and Extreme Network Load Sharing. Network interfaces on servers can be grouped together into a team of physical ports called a virtual adapter.

## 4.1. Network Addressing

To understand how teaming works, it is important to understand how node communications work in an Ethernet network. This document is based on the assumption that the reader is familiar with the basics of IP and Ethernet network communications. The following information provides a high-level overview of the concepts of network addressing used in an Ethernet network. Every Ethernet network interface in a host platform, such as a computer system, requires a globally unique Layer 2 address and at least one globally unique Layer 3 address. Layer 2 is the Data Link Layer, and Layer 3 is the Network layer as defined in the OSI model. The Layer 2 address is assigned to the hardware and is often referred to as the MAC address or physical address. This address is pre-programmed at the factory and stored in NVRAM on a network interface card or on the system motherboard for an embedded LAN interface. The Layer 3 addresses are referred to as the protocol or logical address assigned to the software stack. IP and IPX are examples of Layer 3 protocols. In addition, Layer 4 (Transport Layer) uses port numbers for each network upper level protocol such as Telnet or FTP. These port numbers are used to differentiate traffic flows across applications. Layer 4 protocols such as TCP or UDP are most commonly used in today's networks. The combination of the IP address and the TCP port number is called a socket.

Ethernet devices communicate with other Ethernet devices using the MAC address, not the IP address. However, most applications work with a host name that is translated to an IP address by a Naming Service such as WINS and DNS. Therefore, a method of identifying the MAC address assigned to the IP address is required. The Address Resolution Protocol (ARP) for an IP network provides this mechanism. For IPX, the MAC address is part of the network address and ARP is not required. ARP is implemented using an ARP Request and ARP Reply frame. ARP Requests are typically sent to a broadcast address while the ARP Reply is typically sent as unicast traffic. A unicast address corresponds to a single MAC address or a single IP address. A broadcast address is sent to all devices on a network.

## 4.2. Teaming and Network Addresses

A team of adapters function as a single virtual network interface and do not appear any different to other network devices than a non-teamed adapter. A virtual network adapter advertises a single Layer 2 and one or more Layer 3 addresses. When the teaming driver initializes, it selects one MAC address from one of the physical adapters that make up the team to be the Team MAC address. This address is typically taken from the first adapter that gets initialized by the driver. When the system hosting the team receives an ARP request, it selects one MAC address from among the physical adapters in the team to use as the source MAC address in the ARP Reply. In Windows operating systems, the IPCONFIG /all command shows the IP and MAC address of the virtual adapter and not the individual physical adapters. The protocol IP address is assigned to the virtual network interface and not to the individual physical adapters.

For switch-independent teaming modes, all physical adapters that make up a virtual adapter must use the unique MAC address assigned to them when transmitting data. That is, the frames that are sent by each of the physical adapters in the team must use a unique MAC address to be IEEE compliant. It is important to note that ARP cache entries are not learned from received frames, but only from ARP requests and ARP replies.

## 4.3. Description of Teaming Types

There are three methods for classifying the supported teaming types:

- One is based on whether the switch port configuration must also match the adapter teaming type.
- The second is based on the functionality of the team, whether it supports load balancing and failover or just failover.
- The third is based on whether the Link Aggregation Control Protocol (LACP) is used or not.

Table 2 shows a summary of the teaming types and their classification.

**Table 2: Available Teaming Types**

| **Teaming Type** | **Switch-Dependent (Switch must support specific type of** | **Link Aggregation Control Protocol Support Required on the Switch** | **Load Balancing** | **Failover** |
|---|---|---|---|---|

| | | | | |
|---|---|---|---|---|
| | **team)** | | | |
| Smart Load Balancing and Failover (with two to eight load balance team members) | | | v | v |
| SLB (Auto-Fallback Disable) | | | v | V |
| Link Aggregation (802.3ad) | v | v | v | V |
| Generic Trunking (FEC/GEC)/802.3ad-Draft Static | v | | v | v |

### 4.3.1. Smart Load Balancing and Failover

The Smart Load Balancing™ and Failover type of team provides both load balancing and failover when configured for load balancing, and only failover when configured for fault tolerance. This type of team works with any Ethernet switch and requires no trunking configuration on the switch. The team advertises multiple MAC addresses and one or more IP addresses (when using secondary IP addresses). The team MAC address is selected from the list of load balance members. When the system receives an ARP request, the software-networking stack will always send an ARP Reply with the team MAC address. To begin the load balancing process, the teaming driver will modify this ARP Reply by changing the source MAC address to match one of the physical adapters.

Smart Load Balancing enables both transmit and receive load balancing based on the Layer 3/ Layer 4 IP address and TCP/UDP port number. In other words, the load balancing is not done at a byte or frame level but on a TCP/UDP session basis. This methodology is required to maintain in-order delivery of frames that belong to the same socket conversation. Load balancing is supported on 2 to 8 ports. These ports can include any combination of add-in adapters and LAN on Motherboard (LOM) devices. Transmit load balancing is achieved by creating a hashing table using the source and destination IP addresses and TCP/UDP port numbers. The same combination of source and destination IP addresses and TCP/UDP port numbers will generally yield the same hash index and therefore point to the same port in the team. When a port is selected to carry all the frames of a given socket, the unique MAC address of the physical adapter is included in the frame, and not the team MAC address. This is required to comply with the IEEE 802.3 standard. If two adapters transmit using the same MAC address, then a duplicate MAC address situation would occur that the switch could not handle.

**Note***: IPv6 addressed traffic will not be load balanced by SLB because ARP is not a feature of IPv6.*

Receive load balancing is achieved through an intermediate driver by sending gratuitous ARPs on a client-by-client basis using the unicast address of each client as the destination address of the ARP request (also known as a directed ARP). This is considered client load balancing and not traffic load balancing. When the intermediate driver detects a significant load imbalance between the physical adapters in an SLB team, it will generate G-ARPs in an effort to redistribute incoming frames. The intermediate driver (BASP) does not answer ARP requests; only the software protocol stack provides the required ARP Reply. It is important to understand that receive load balancing is a function of the number of clients that are connecting to the system through the team interface.

SLB receive load balancing attempts to load balance incoming traffic for client machines across physical ports in the team. It uses a modified gratuitous ARP to advertise a different MAC address for the team IP Address in the sender physical and protocol address. This G-ARP is unicast with the MAC and IP Address of a client machine in the target physical and protocol address respectively. This causes the target client to update its ARP cache with a new MAC address map to the team IP address. G-ARPs are not broadcast because this would cause all clients to send their traffic to the same port. As a result, the benefits achieved through client load balancing would be eliminated, and could cause out-of-order frame delivery. This receive load balancing scheme works as long as all clients and the teamed system are on the same subnet or broadcast domain.

When the clients and the system are on different subnets, and incoming traffic has to traverse a router, the received traffic destined for the system is not load balanced. The physical adapter that the intermediate driver has selected to carry the IP flow carries all of the traffic. When the router sends a frame to the team IP address, it broadcasts an ARP request (if not in the ARP cache). The server software stack generates an ARP reply with the team MAC address, but the intermediate driver modifies the ARP reply and sends it over a particular physical adapter, establishing the flow for that session.

The reason is that ARP is not a routable protocol. It does not have an IP header and therefore, is not sent to the router or default gateway. ARP is only a local subnet protocol. In addition, since the G-ARP is not a broadcast packet, the router will not process it and will not update its own ARP cache.

The only way that the router would process an ARP that is intended for another network device is if it has Proxy ARP enabled and the host has no default gateway.

This is very rare and not recommended for most applications.

Transmit traffic through a router will be load balanced as transmit load balancing is based on the source and destination IP address and TCP/UDP port number. Since routers do not alter the source and destination IP address, the load balancing algorithm works as intended.

Configuring routers for Hot Standby Routing Protocol (HSRP) does not allow for receive load balancing to occur in the adapter team. In general, HSRP allows for two routers to act as one router, advertising a virtual IP and virtual MAC address. One physical router is the active interface while the other is standby. Although HSRP can also load share nodes (using different default gateways on the host nodes) across multiple routers in HSRP groups, it always points to the primary MAC address of the team.

### 4.3.2. Generic Trunking

Generic Trunking is a switch-assisted teaming mode and requires configuring ports at both ends of the link: server interfaces and switch ports. This is often referred to as Cisco Fast EtherChannel or Gigabit EtherChannel. In addition, generic trunking supports similar implementations by other switch OEMs such as Extreme Networks Load Sharing and Bay Networks or IEEE 802.3ad Link Aggregation static mode. 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 intermediate driver controls load balancing and failover for outgoing traffic only, while incoming traffic is controlled by the switch firmware and hardware. As is the case for Smart Load Balancing, the BASP intermediate driver uses the IP/TCP/UDP source and destination addresses to load balance the transmit traffic from the server. Most switches implement an XOR hashing of the source and destination MAC address.

**Note**: *Generic Trunking is not supported on iSCSI offload adapters.*

### 4.3.3. Link Aggregation (IEEE 802.3ad LACP)

Link Aggregation is similar to Generic Trunking except that it uses the Link Aggregation Control Protocol to negotiate the ports that will make up the team. LACP must be enabled at both ends of the link for the team to be operational. If LACP is not available at both ends of the link, 802.3ad provides a manual aggregation that only requires both ends of the link to be in a link up state. Because manual aggregation provides for the activation of a member link without performing the LACP message exchanges, it should not be considered as reliable and robust as an LACP negotiated link. 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 removal of aggregate link members is provided by the marker protocol that can be optionally enabled for Link Aggregation Control Protocol (LACP) enabled aggregate links.

The Link Aggregation group advertises a single MAC address for all the ports in the trunk. The MAC address of the Aggregator can be the MAC addresses of one of the MACs that make up the group. LACP and marker protocols use a multicast destination address.

The Link Aggregation control function determines which links may be aggregated and 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 will not cause loss of connectivity. The traffic will simply be failed over to the remaining links in the trunk.

### 4.3.4. SLB (Auto-Fallback Disable)

This type of team is identical to the Smart Load Balance and Failover type of team, with the following exception—when the standby member is active, if a primary member comes back on line, the team continues using the standby member rather than switching back to the primary member. This type of team is supported only for situations in which the network cable is disconnected and reconnected to the network adapter. It is not supported for situations in which the adapter is removed/installed through Device Manager or Hot-Plug PCI.

If any primary adapter assigned to a team is disabled, the team functions as a Smart Load Balancing and Failover type of team in which auto-fallback occurs.

## 4.4. Teaming and Other Advanced Networking Properties

Before creating a team, adding or removing team members, or changing advanced settings of a team member, make sure each team member has been configured similarly. Settings to check include VLANs and QoS Packet Tagging, Jumbo Frames, and the various offloads.

A team does not necessarily inherit adapter properties; rather various properties depend on the specific capability. For instance, an example would be flow control, which is a physical adapter property and has nothing to do with BASP, and will be enabled on a particular adapter if the miniport driver for that adapter has flow control enabled.

**Note**: *All adapters on the team must support the property listed in Table 7 in order for the team to support the property.*

### 4.4.1. IEEE 802.1p QoS Tagging

The IEEE 802.1p standard includes a 3-bit field (supporting a maximum of 8 priority levels), which allows for traffic prioritization. The BASP intermediate driver does not support IEEE 802.1p QoS tagging.

### 4.4.2. Jumbo Frames

The use of jumbo frames was originally proposed by Alteon Networks, Inc. in 1998 and increased the maximum size of an Ethernet frame to a maximum size of 9000 bytes. Though never formally adopted by the IEEE 802.3 Working Group, support for jumbo frames has been implemented in D-Link network adapters. The BASP intermediate driver supports jumbo frames, provided that all of the physical adapters in the team also support jumbo frames and the same size is set on all adapters in the team.

### 4.4.3. IEEE 802.1Q VLANs

In 1998, the IEEE approved the 802.3ac standard, which defines frame format extensions to support Virtual Bridged Local Area Network tagging on Ethernet networks as specified in the IEEE 802.1Q specification. The VLAN protocol permits insertion of a tag into an Ethernet frame to identify the VLAN to which a frame belongs. If present, the 4-byte VLAN tag is inserted into the Ethernet frame between the source MAC address and the length/type field. The first 2-bytes of the VLAN tag consist of the IEEE 802.1Q tag type, whereas the second 2 bytes include a user priority field and the VLAN identifier (VID). Virtual LANs (VLANs) allow the user to split the physical LAN into logical subparts. Each defined VLAN behaves as its own separate network, with its traffic and broadcasts isolated from the others, thus increasing bandwidth efficiency within each logical group. VLANs also enable the administrator to enforce appropriate security and quality of service (QoS) policies. The BASP supports the creation of 64 VLANs per team or adapter: 63 tagged and 1 untagged. The operating system and system resources, however, limit the actual number of VLANs. VLAN support is provided according to IEEE 802.1Q and is supported in a teaming environment as well as on a single adapter. Note that VLANs are supported only with homogeneous teaming and not in a multivendor teaming environment. The BASP intermediate driver supports VLAN tagging. One or more VLANs may be bound to a single instance of the intermediate driver.

### 4.4.4. Preboot Execution Environment

The Preboot Execution Environment (PXE) allows a system to boot from an operating system image over the network. By definition, PXE is invoked before an operating system is loaded, so there is no opportunity for the BASP intermediate driver to load and enable a team. As a result, teaming is not supported as a PXE client, though a physical adapter that participates in a team when the operating system is loaded may be used as a PXE client. Whereas a teamed adapter cannot be used as a PXE client, it can be used for a PXE server, which provides operating system images to PXE clients using a combination of Dynamic Host Control Protocol (DHCP) and the Trivial File Transfer Protocol (TFTP). Both of these protocols operate over IP and are supported by all teaming modes.

# 5. NIC Partitioning

NIC partitioning (NPAR) divides a D-Link switch into multiple virtual NICs by having multiple PCI physical functions per port. Each PCI function is associated with a different virtual NIC. To the OS and the network, each physical function appears as a separate NIC port.

**Note**: *Link speed cannot be configured for 1 Gbps when NPAR is enabled.*

*The number of partitions for each port is from one to four; thus, making a dual-port NIC capable of a maximum of eight partitions. Each partition behaves as if it is an independent NIC port.*

On quad-port adapters:

- The 1G ports do not support NPAR.
- On 10G ports, only two functions per port are supported.

Benefits of a partitioned 10G NIC include:

- Reduced cabling and ports when used to replace many 1G NICs.
- Server segmentation with separate subnets/VLANs.
- High server availability with NIC failover and NIC link bandwidth aggregation.
- Server I/O virtualization with a virtual OS and monolithic OS support.
- Changes to the OS is not required.
- SLB type teaming is supported.

## 5.1. Supported Operating Systems for NIC Partitioning

The DXE-820T supports NIC partitioning on the following operating systems:

- Windows Server 2008 family
- Windows Server 2012 family
- Linux 64-bit, RHEL 5.5 and later, SLES11 SP1 and later
- VMware ESX, ESXi 4.1, ESXi 5.0, and ESXi 5.1.

**Note:** *32-bit Linux operating systems have a limited amount of memory space available for Kernel data structures. Therefore, it is recommended that only 64-bit Linux be used when configuring NPAR.*

## 5.2. Configuring for NIC Partitioning

When NIC partitioning is enabled on an adapter, you must explicitly configure storage offloads on a PF to use FCoE and iSCSI offload functionality on an adapter.

NIC partitioning can be configured using Broadcom's Comprehensive Configuration Management (CCM) utility.

**Note:** *In NPAR mode, SR-IOV cannot be enabled on any VF on which storage offload (FCoE or iSCSI) is configured. This does not apply to adapters in Single Function (SF)* *mode.*

To configure a NIC for partitioning using the CCM utility

1. Select the NIC from Device List.
2. From the Main Menu, select Device Hardware Configuration.
3. Change the Multi-Function Mode to NPAR.
4. Configure the NIC parameters for your configuration based on the options shown in Table 1.

Table 3 lists the configuration parameters available from the NIC Partitioning Configuration screen.

**Table 3: Configuration Options**

| **Parameter** | **Description** | **Options** |
|---|---|---|
| Flow Control | Configures the Flow Control mode for this port. | • Auto<br>• TX Flow Control<br>• RX Flow Control<br>• TX/RX Flow Control<br>• None |
| PF#0, PF#2, PF#4, PF#6 | Displays the physical function (PF) information regarding the partition(s) on port 0. Select to configure. | See Table 2 for configuration options. |
| PF#1, PF#3, PF#5, PF#7 | Displays the physical function (PF) information regarding the partition(s) on port 1. Select to configure. | See Table 2 for configuration options. |
| Reset Configuration to Default | Resets the NIC partition configuration to the factory default settings. | |

Table 4 describes the functions available from the **PF# X** screen.

**Table 4: Function Description**

| Function | Description | Option |
|---|---|---|
| Ethernet Protocol | Enables/disables Ethernet protocol. | Enable<br>Disable |
| iSCSI Offload Protocol | Enables/disables iSCSI protocol. | Enable<br>Disable |
| FCoE Offload protocol | Enables/disables FCoE protocol. | Enable<br>Disable |
| Bandwidth Weight | Configures the weight or importance of a particular function. There are four functions per port and the weight is used to arbitrate | The sum of all weights for the four functions are either |

| | | |
|---|---|---|
| | between the functions in case of congestion. | 0 or 100. |
| Maximum Bandwidth | Configures the maximum bandwidth (in percentage) of the physical port link. | |
| Network MAC Address | Displays the network MAC address. | |
| iSCSI MAC Address | Displays the iSCSI MAC address. | |
| FCoE FIP MAC Address | Displays the FCoE MAC address. | |
| FCoE WWPN | FCoE World Wide Port Name. | |
| FCoE WWNN | FCoE World Wide Node Name. | |
| **Note**: *Ensure that the Network MAC Address and the iSCSI MAC Address are not the same.* | | |

# 6. Using Broadcom Advanced Control Suite 4

Broadcom Advanced Control Suite (BACS) is an integrated utility that provides useful information about each network adapter that is installed in your system. BACS also enables you to perform detailed tests, diagnostics, and analyses on each adapter, as well as to view and modify property values and view traffic statistics for network objects. BACS operates on Windows and Linux operating systems.

Broadcom Advanced Server Program (BASP), which runs within Broadcom Advanced Control Suite, is used to configure teams for load balancing, fault tolerance, and virtual local area networks (VLANs). BASP functionality is available only on systems that use at least one Broadcom network adapter. BASP operates on Windows operating systems only.

**Note:** *Some features of BACS are relevant only to particular adapters or adapter families. Because a single instance of BACS can be used to communicate with multiple hosts and adapter types, this document describes all BACS features*.

## 6.1. Starting Broadcom Advanced Control Suite

In Control Panel, click Broadcom Control Suite 4, or click the BACS icon in the taskbar located at the bottom of the Windows desktop.

On Linux systems, you can double-click the BACS4 desktop icon, or access the BACS program from the task bar under System Tools. (If you are having difficulty launching BACS on a Linux system, see the related topic in Troubleshooting BACS.)

## 6.2. BACS Interface

The BACS interface is comprised of the following regions:

- Explorer View pane
- Context View selector
- Context View pane
- Menu bar
- Description pane

By default, the Explorer View pane is docked and pinned on the left side of the main window, the Context View pane on the right, the Context View selector below the menu bar, and the Description pane below the Context View pane. Drag the splitter between any two panes to vary the size of the panes.

### 6.2.1. Explorer View Pane

You can dock and pin the Explorer View pane on the left side, right side, top, or bottom of the main window.

The Explorer View pane lists the objects that can be viewed, analyzed, tested, or configured by BACS. When an item is selected in the Explorer View pane, the tabs showing the information and options that are available for the item appear in the Context View pane.

The organization of this panel is designed to present the manageable objects in the same hierarchical manner as drivers and its subcomponents. This simplifies the management of various elements of the converged network interface controller (C-NIC). The top level of the hierarchy is the Host container, which lists all hosts managed by BACS. Below the hosts are the installed network adapters, with the manageable elements, such as physical port, VBD, NDIS, FCoE, and iSCSI below the adapters.

The icon next to each device in the Explorer View pane shows its status. An icon next to a device name that appears normal means the device is connected and working.

- X. A red "X" that appears on the device's icon indicates the device is currently not connected to the network.
- Greyed out. A device icon that appears greyed out indicates the device is currently disabled.

### 6.2.2. Context View Selector

The Context View selector appears below the menu bar and includes the filter and tab categories. Although you can expand and collapse the categories that appear on tabs in the Context View pane, you can alternatively display a category by selecting the box next to the category name.

**Filter View**

In a multiple-host environment using several C-NICs, there can be a large number of manageable elements per adapter that can be difficult and cumbersome to view, configure, and manage all elements. Use the filter to select a particular device function. Possible filter views include:

- All
- Team view
- NDIS view
- iSCSI view
- FCoE view
- iSCSI Target view
- FCoE Target view

### 6.2.3. Context View Pane

The Context View pane displays all the parameters that you can view for the object selected in the Explorer View pane. The parameters are grouped by tabs and categories, depending on the parameter type. The available tabs are Information, Configuration, Diagnostics, and Statistics. Because the BACS interface is context-sensitive, only the parameters that apply to the selected object can be viewed or configured in the Context View pane.

### 6.2.4. Menu Bar

The following appear on the menu bar, but because the menu items are context-sensitive, not all items will be available at all times:

**File menu**

- Team Save As: Saves the current team configurations to a file.
- Team Restore: Restores any saved team configuration from a file.

**Action menu**

- Remove Host: Removes the selected host.
- Refresh Host: Refreshes the selected host.

**View menu**

- Explorer View: Displays/hides the Explorer View pane.
- Tool Bar: Displays/hides the tool bar.
- Status Bar: Displays/hides the status bar.
- Broadcom Logo: Displays/hides the Broadcom Logo on BACS to optimize the maximum viewable space.

**Tools menu**

- Options: Used for configuring BACS preferences.

**Teams (Windows only)**

- Create Teams: Creates new teams with either the Teaming Wizard or in Advanced mode.
- Manage Teams: Manages existing teams with either the Teaming Wizard or in Advanced mode.

**iSCSI menu**

- Discovery Wizard: Locates targets and helps to configure the HBA.
- Manage Targets Wizard: Manages targets.
- Manage iSNS Servers: Manages Internet Storage Name Service (iSNS) servers to allow discovery, management, and configuration of iSCSI devices.
- Manage Discovery Portals: Manages iSCSI discovery portals.

**Discovery Wizard**

The Discovery Wizard is available from the iSCSI menu. Follow the prompts in the wizard to discover iSCSI targets via the SendTargets method or the Internet Storage Name Service (iSNS) server.

**Manage Targets Wizard**

The Manage Targets Wizard is available from the iSCSI menu. Follow the prompts in the wizard to add and remove targets, and to login or logout of a target.

**Manage iSNS Servers**

The Manage iSNS Servers window is available from the iSCSI menu. From this window, you can add or remove Internet Storage Name Service (iSNS) servers.

**Manage Discovery Portals**

The Manage Discovery Portals window is available from the iSCSI menu. From this window, you can add or remove iSCSI discovery portals.

**Boot Configuration Wizard**

The Boot Configuration Wizard is available by right-clicking a port. Follow the prompts in the wizard to configure the iSCSI boot parameters.

**Hardware and Resource Configuration Wizard**

The Hardware and Resource Configuration Wizard is used to configure properties for hardware resources. Follow the prompts in the wizard to configure hardware resources. You can preview the configuration before committing the changes.

### 6.2.5. Description Pane

The Description pane provides information, configuration instructions, and options for the selected parameter in the Context View pane.

## 6.3. Configuring Preferences in Windows

To enable or disable the BACS tray icon in Windows

On Windows systems, BACS places an icon in the Windows taskbar when the program is installed. Use the Options window to turn this icon on or off.

1. From the Tools menu, select Options.
2. Select or clear Enable BACSTray (the option is enabled by default).
3. Click OK.

**Setting the teaming mode in Windows**

1. From the Tools menu, select Options.
2. Select Expert Mode if you do not need the assistance of the teaming wizard to create teams; otherwise, select Wizard Mode.
3. Click OK.

**Setting the Explorer View refresh time in Windows**

1. From the Tools menu, select Options.
2. Select Auto to set the Explorer View refresh time to 5 seconds. Otherwise, select Custom and select a time, in seconds.
3. Click OK.

## 6.4. Connecting to a Host

You can add one or more Windows or Linux hosts to manage from BACS.

**To add a local host**

1. From the Action menu, click Add Host.
2. For both Windows and Linux hosts, do not change the default settings. The User name and Password are not required while connecting to the local host.
3. Select Persist if you want BACS to save the information for this host.
4. Click Ok. BACS can now be used to view information and manage the host.

**To add a remote host**

1. From the Action menu, click Add Host.
2. Type the remote host's name or IP address in the Host box.
3. Select the protocol from the Protocol list. The protocol options for Windows are WMI, WSMan, or Try All. The protocol options for Linux are CimXML, WSMan, or Try All. The Try All option forces the GUI client to try all options.
4. Select the HTTP scheme, or the HTTPS scheme for added security.
5. Type the Port Number value you used to configure the host, if it is different than the default value of 5985.
6. Type the User name and Password.
7. Select Persist if you want BACS to save the information for this host. The host will appear in the Explorer Pane whenever you reopen BACS, and you will not need to enter the host IP address or host name when connecting to the host. For security reasons, you must enter the User name and Password every time.
8. Click OK.

## 6.5. Managing the Host

At the host level, you can view host information and configure parameters from the following tabs:

- Information
- Configuration

**To view host information**

Select the host in the Explorer View pane, and then select the Information tab to view host-level information.

### 6.5.1. Information Tab: Host Information

**Host Name**

Displays the name of the host.

**OS Version Info**

Displays the operating system, including the version.

**Platform**

Displays the hardware architecture platform (for example, 32-bit or 64-bit)

### 6.5.2. Information Tab: iSCSI Initiator

The iSCSI Initiator section of the Information tab is available if iSCSI is enabled on the host.

**Name**

Displays the iSCSI initiator name in IQN format.

**Portal List**

Displays all iSCSI portal IP addresses configured on the selected host.

**Note**: *Some information may not be available for all D-Link network adapters.*

**To configure the host**

Select the host in the Explorer View pane, and then select the Configuration tab to configure host-level parameters.

### 6.5.3. Configuration Tab: System Management

**Chimney Offload State**

Enable or disable chimney offload at the host level, rather than at the device level, and then click Apply.

**Configuration Tab: iSCSI Initiator**

**Name**

The current IQN name is displayed. Click the IQN name to modify the host's iSCSI initiator name, and then click Apply.

## 6.6. Managing the Network Adapter

The installed network adapters appear one level below the host in the hierarchical tree in the Explorer View pane. At the adapter level, you can view information and configure parameters from the following tabs:

- Information
- Configuration

### 6.6.1. Viewing Adapter Information

Select the network adapter in the Explorer View pane, and then select the Information tab to view adapter-level information.

### 6.6.2. Viewing Resource Information

The Resources section of the Information tab displays information about connections and other essential functions for the selected network adapter.

**Note**: *Some information may not be available for all D-Link network adapters.*

**Information Tab: Resources**

**Bus Type**

The type of input/output (I/O) interconnect used by the adapter.

**Bridge**

The bridge type, which is the PCI-E to PCI-X bridge. This information is only available for D-Link DXE-820T network adapters.

**Bridge Lanes**

The number of PCI-E lanes connected to the bridge. This information is only available for D-Link DXE-820T network adapters.

**Bridge Speed**

The clock speed on PCI-E bus. This information is only available for D-Link DXE-820T network adapters.

**Slot Number**

The slot number on the system board occupied by the adapter. This item is not available for PCI Express type adapters.

**Bus Speed**

The bus clock signal frequency used by the adapter. This item is not available for PCI Express type adapters.

**Bus Width**

The number of bits that the bus can transfer at a single time to and from the adapter. This item is not available for PCI Express type adapters.

**Bus Number**

Indicates the number of the bus where the adapter is installed.

**Device Number**

The number assigned to the adapter by the operating system.

**Function Number**

The port number of the adapter. For a single-port adapter, the function number is 0. For a two-port adapter, the function number for the first port is 0, and the function number for the second port is 1.

**Interrupt Request**

The interrupt line number that is associated with the adapter. Valid numbers range from 2 to 25.

**Memory Address**

The memory mapped address that is assigned to the adapter. This value can never be 0.

**MSI Version**

This is the Message Signaled Interrupts (MSI) version being used. The option MSI corresponds to the PCI 2.2 specification that supports 32 messages and a single MSI address value. The option MSI-X corresponds to the PCI 3.0 specification that supports 2,048 messages and an independent message address for each message.

### 6.6.3. Viewing Hardware Information

The Hardware section of the Information tab displays information about the hardware settings for the selected network adapter.

**Note:** *Some information may not be available for all Broadcom network adapters.*

**Information Tab: Hardware**

**ASIC Version**

The chip version of the DXE-820T adapter (this information is not available for adapters made by others).

**Bootcode Version**

The version of the boot code. This information is only available for D-Link DXE-820T network adapters.

**Family Firmware Version**

The global firmware version that represents all firmware on the device.

**Management Firmware**

The firmware version installed on the system.

**Vendor ID**

The vendor ID.

**Device ID**

The adapter ID.

**Subsystem Vendor ID**

The subsystem vendor ID.

**Subsystem ID**

The subsystem ID.

**External PHY Firmware Version**

The external PHY firmware version.

### 6.6.4. Configuring Adapter Parameters

Select the network adapter in the Explorer View pane, and then select the Configuration tab to configure adapter-level parameters.

### 6.6.5. Configure Multi-function Parameters Using the Wizard

Click Configure to configure multi-function parameters.

### 6.6.6. Hardware and Resource Configuration Wizard: Introduction

The Hardware and Resource Configuration Wizard will help you modify device hardware configuration and resource configuration. Select a multi-function mode and then click Next.

**Multi-Function mode**

Displays the multi-function mode.

**Number of Partitions**

Displays the number of partitions. The value is 4 and cannot be changed.

### 6.6.7. Hardware and Resource Configuration Wizard: Port Configuration

Select a port to configure and then click Next.

**Flow Control**

The possible values are Auto, Tx Pause, Rx Pause, Tx/Rx pause, and Disable. The configuration is done at the port level and applies to all functions under the port. The flow control value is a default value for the port. The effective configuration can be different based on the switch port configuration and whether or not DCB/DCBX is enabled.

**Link Speed**

Configure the link speed. The default speed is 1Gb for 1Gb adapters and 10Gb for 10Gb adapters.

### 6.6.8. Hardware and Resource Configuration Wizard: Configure Resources

The following are configurable at the function level. Click Next after making changes.

When NIC partitioning is enabled (on DXE-820T adapters only), four functions are created under each port. The functions are numbered 0 to 7. All odd function numbers (1, 3, 5, 7) are created on one port and all even functions (0, 2, 4, 6) are created on the remaining port.

**Ethernet/NDIS**

Ethernet/NDIS capability is enabled when selected.

**iSCSI**

iSCSI functionality is enabled when selected.

**FCoE**

FCoE functionality is enabled when selected (NetXtreme II adapters only).

**Maximum Bandwidth (%)**

- The maximum bandwidth setting defines an upper threshold value, ensuring that this limit will not be exceeded during transmission. The valid range for this value is between 1 and 100. The maximum bandwidth value is defined as a percentage of the physical link speed.

- It is possible for the sum of all maximum bandwidth values across the four functions of a single port to exceed the physical link speed value of either 10 Gbps or 1 Gbps. This case is considered as oversubscription. In a case where oversubscription congestion occurs on transmit, the Relative Bandwidth Weight value comes into effect.
- The Maximum Bandwidth setting is only valid in the context of Tx, but not Rx.

**Relative Bandwidth Weight (%)**

- The relative bandwidth setting represents a weight or importance of a particular function. There are four functions per port and the weight is being used in order to arbitrate between the functions in case of congestion.
- The sum of all weights for the four functions on a single port are either 0 or 100.
- A value of 0 for all functions means that each function will be able to transmit at 25% of the physical link speed, not to exceed the Maximum Bandwidth setting
- A value for a function between 1 and 100 represent a percentage of the physical link speed and is used by an internal arbitration logic as an input value (weight). A higher value will cause this function to transmit relatively more data, compared to a function (on the same port) that has defined a lower value.

### 6.6.9. Hardware and Resource Configuration Wizard: Commit and Finish

Click Apply to commit changes to the system or click Cancel. Click Finish to save your changes and exit the wizard.

## 6.7. Managing Ethernet Controller (Port)

From BACS, you can group various traffic classes in to priority group and allocate bandwidth to each priority group.

When the Ethernet controller is selected in the Object Explorer panel, following four tabs will be displayed in the context view panel:

- Information Tab
- Configuration tab
- Statistics Tab
- Diagnostic Tab

### 6.7.1. Viewing Port Level Information

Selecting Ethernet controller in the object explorer will allow user to view various types of information at the port level.

1. Select PortX (where X is either 0 or 1) below Adapter in the object explorer.
2. Various components of the port will be displayed below port in the object explorer. You can click on the "+" icon near Port to expand or collapse the tree

below.

3. Select Information tab in the context view panel on the right side.

### 6.7.2. Viewing Vital Signs

The Vital Signs section of the Information tab has useful information about the network adapters that are installed in your system, such as the link status of the adapter and general network connectivity.

To view Vital Signs information for any installed network adapter, select the name of the adapter listed in the Explorer View pane, then click the Information tab.

**Notes:**

- Information about Broadcom network adapters may be more comprehensive than information about network adapters made by others.
- Some information may not be available for all Broadcom network adapters.

**MAC Address**

A physical MAC (media access control) address that is assigned to the adapter by the manufacturer. The physical address is never all 0s.

**Permanent MAC Address**

The unique hardware address assigned to the network adapter.

**iSCSI MAC Address**

If an iSCSI network adapter is loaded onto the system, this parameter will display the iSCSI MAC address.

**IPv4 DHCP**

The IP address is from a DHCP server if the value is Enable.

**IP Address**

The network address associated with the adapter. If the IP address is all 0s, the associated driver has not been bound with Internet Protocol (IP).

**IPv6 DHCP**

The IP address is from a DHCP server if the value is Enable.

**IPv6 IP Address**

The IPv6 network address associated with the adapter.

**IPv6 Scope Id**

Since local-use addresses can be reused, the Scope ID for link-local addresses specifies the link where the destination is located. The Scope ID for site-local addresses specifies the site where the destination is located. The Scope ID is relative to the sending host.

**IPv6 Flow Info**

The non-zero Flow Info is used to classify traffic flows. If Flow Info equals zero, then the packets are not a part of any flow.

**Default Gateway**

The default gateway value is the network address of the gateway that will be used by the management firmware for packets destined for hosts external to the local network segment.

**Link Status**

The status of the network link.

- Up. A link is established.
- Down. A link is not established.

**Duplex**

The adapter is operating in the indicated duplex mode.

**Speed**

The link speed of the adapter, in megabits per second.

**Offload Capabilities**

The offload capabilities supported by the adapter. This information is only available for D-Link DXE-820T network adapters.

- iSCSI. iSCSI offload for block-level transfer of data.
- LSO. Large Send Offload (LSO) prevents an upper level protocol such as TCP from breaking a large data packet into a series of smaller packets with headers appended to them.
- CO. Checksum Offload (CO) allows the TCP/IP/UDP checksums for send and receive traffic to be calculated by the adapter hardware rather than by the host CPU.

**LiveLink IP Address**

The network address of the LiveLink enabled adapter.

**Local Connection**

Identifies the module to which the blade server is attached.

- Chassis SW. Chassis switch module
- Chassis PHY. Pass-through module
- None. No modules attached

**BASP State**

Information about the status of the BASP application. This information is displayed only when there is a team (see Configuring Teaming).

### 6.7.3. Viewing NIC Partitioning Information

The NIC partitioning feature is available on D-Link DXE-820T network adapters only.

The NIC Partitioning section of the Information tab displays information about the partitions for the selected network adapter.

To view NIC Partitioning for any installed network adapter, click the name of the adapter listed in the Explorer View pane, then click the Information tab.

**Note**: *Some information may not be available for all Broadcom network adapters.*

*NIC partitioning divides a D-Link switch into multiple virtual NICs by having multiple PCI physical functions per port. Each PCI function is associated with a different virtual NIC. To the OS and the network, each physical function appears as a separate NIC port. For more information, see the NIC Partitioning topic in the D-Link DXE-820T Network Adapter User Guide.*

**Number of Partitions**

The number of partitions for the port. Each port can have from one to four partitions with each partition behaving as if it is an independent NIC port.

**Network MAC Address**

The MAC address of the port.

**iSCSI MAC Address**

If an iSCSI adapter is loaded onto the system, the iSCSI MAC address will appear.

**Flow Control**

The flow control setting of the port.

**Physical Link Speed**

The physical link speed of the port, either 1G or 10G.

**Relative Bandwidth Weight (%)**

- The relative bandwidth setting represents a weight or importance of a particular function. There are up to four functions per port. The weight is used to arbitrate between the functions in the event of congestion.
- The sum of all weights for the functions on a single port is either 0 or 100.
- A value of 0 for all functions means that each function will be able to transmit at 25% of the physical link speed, not to exceed the Maximum Bandwidth setting.
- A value for a function between 1 and 100 represent a percentage of the physical link speed and is used by an internal arbitration logic as an input value (weight). A higher value will cause this function to transmit relatively more data, compared to a function (on the same port) that has defined a lower value.

**Maximum Bandwidth (%)**

- The maximum bandwidth setting defines an upper threshold value, ensuring that this limit will not be exceeded during transmission. The valid range for this value is between 1 and 100. The maximum bandwidth value is defined as a percentage of the physical link speed.
- It is possible for the sum of all maximum bandwidth values across the four functions of a single port to exceed the physical link speed value of either 10 Gbps or 1 Gbps. This case is considered as oversubscription. In a case where oversubscription congestion occurs on transmit, the Relative Bandwidth Weight

value comes into effect.

- The Maximum Bandwidth setting is only valid in the context of Tx, but not Rx.

### 6.7.4. Testing the Network

The Network Test option on the Diagnostics tab lets you verify IP network connectivity. This test verifies if the driver is installed correctly and tests connectivity to a gateway or other specified IP address on the same subnet.

The network test uses TCP/IP to send ICMP packets to remote systems, then waits for a response. If a gateway is configured, the test automatically sends packets to that system. If a gateway is not configured or if the gateway is unreachable, the test prompts for a destination IP address.

**Notes:**

- The network test option is not available on adapters that are grouped into a team (see Configuring Teaming).
- This feature can be used with Windows Server managed hosts only. It is not available for hosts operating on Linux or other OSes. You can, however use BACS on a Linux client to connect to a Windows Server host and run the network test utility.

**To run the network test using the BACS GUI**

1. Click the name of the adapter to test in the Explorer View pane.
2. From the Select a test to run list, select Network Test.
3. To change the destination IP address, select IP address to ping, then click the browse button (...). In the Network Test window, enter a Destination IP address, then click OK.
4. Click Run.

The results of the network test are displayed in the Status field.

To run the network test using the BACS CLI

You can use the following CLI command to perform a network diagnostic test for the specified target. This command is available for NDIS and virtual adapters.

```
BACScli -t <target type> -f <target format> -i <target ID> networkdiag [-p <IP
address>]
```

Examples:

1. The following command runs the network test for the current selected NDIS adapter.

   ```
   BACScli -t NDIS -f mac -i 0010181a1b1c "networkdiag -p 192.168.1.5"
   ```
2. The following command runs the network test for the current selected virtual adapter. Since there is no IP address specified, BACScli will use gateway address for the test.

In Interactive mode, use the list <view> and select <idx> commands to select the desired target device. Use networkdiag [-p <IP address>] to run the network diagnostics test for the selected target.

Examples:

1. The following command runs the network test for the currently selected NDIS adapter.
   networkdiag -p 192.168.1.5
2. The following command runs the network test for the current selected virtual adapter.
   Networkdiag

### 6.7.5. Running Diagnostic Tests in Windows

The Diagnostic Tests option on the Diagnostics tab lets you check the state of the physical components on a Broadcom network adapter. You can trigger the tests manually, or choose to have BACS continuously perform them. If the test are performed continuously, then the number of passes and fails in the Result field for each test increments every time the tests are performed. For example, if a test is performed four times and there are no fails, the value in the Result field for that test is 4/0. However, if there were 3 passes and 1 fail, the value in the Result field is 3/1.

**Notes:**

- This feature can be used with Windows Server managed hosts only. It is not available for hosts operating on Linux or other OSes. You can, however use BACS on a Linux client to connect to a Windows Server host and run the diagnostic test utility.
- You must have administrator privileges to run diagnostic tests.
- The network connection is temporarily lost while these tests are running.
- Some tests are not supported on all DXE-820T adapters.

**To run the diagnostic tests once using the BACS GUI**

1. Click the name of the adapter to test in the Explorer View pane and select the Diagnostics tab.
2. From the Select a test to run list, select Diagnostic Tests.
3. Select the diagnostic tests you want to run. Click Select All to select all tests or Clear All to clear all test selections.
4. Select the number of times to run the tests from Number of loops.
5. Click Run test(s).
6. In the error message window that warns of the network connection being temporarily interrupted, click Yes. The results are displayed in the Result field for each test.

**Control Registers**

This test verifies the read and write capabilities of the network adapter registers by writing various values to the registers and verifying the results. The adapter driver uses these registers to perform network functions such as sending and receiving information. A test failure indicates that the adapter may not be working properly.

**MII Registers**

This test verifies the read and write capabilities of the registers of the physical layer (PHY). The physical layer is used to control the electrical signals on the wire and to configure network speeds such as 1000 Mbit/s.

**EEPROM**

This test verifies the content of the electrically erasable programmable read-only memory (EEPROM) by reading a portion of the EEPROM and computing the checksum. The test fails if the computed checksum is different from the checksum stored in the EEPROM. An EEPROM image upgrade does not require a code change for this test.

**Internal Memory**

This test verifies that the internal memory of the adapter is functioning properly. The test writes patterned values to the memory and reads back the results. The test fails if an erroneous value is read back. The adapter cannot function if its internal memory is not functioning properly.

**On-Chip CPU**

This test verifies the operation of the internal CPUs in the adapter.

**Interrupt**

This test verifies that the Network Device Driver Interface Specification (NDIS) driver is able to receive interrupts from the adapter.

**Loopback MAC and Loopback PHY**

These tests verify that the NDIS driver is able to send packets to and receive packets from the adapter.

**Test LED**

This test causes all of the port LEDs to blink 5 times for the purpose of identifying the adapter.

To run the diagnostic tests using the BACS CLI

You can use the following CLI command to run diagnostics tests on a specified target. This command is available for physical device ports only:

```
BACScli -t <target type> -f <target format> -i <target ID> "diag {[-c REG ] [-c MII ] [-c
EEP] [-c MEM] [-c CPU] [-c INT] [-c MACLB ] [-c PHYLB] [-c LED] | [-c ALL]} [-l <cnt> ]
[ -v <LEDIntv> ]"
```

Examples:

1. The following command displays all the diagnostics tests available for the current selected target.
   ```
   BACScli -t PHYPORTS -f bdf -i 01:00.00 "diag"
   ```
2. The following command runs the MII and LED tests for the selected target:
   ```
   BACScli -t PHYPORTS -f bdf -i 01:00.00 "diag -c MII -c LED"
   ```
3. The following command runs all the tests five times with an LED test interval of 8 ms for the selected target:
   ```
   BACScli -t PHYPORTS -f bdf -i 01:00.00 "diag -c all -l 5 -v 8"
   ```

In Interactive mode, use the list <view> and select <idx> commands to select the desired target device. Use the following command to run diagnostic tests for the selected target:

```
diag {[-c REG ] [-c MII ] [-c EEP] [-c MEM] [-c CPU] [-c INT] [-c MACLB ] [-c PHYLB] [-c LED] | [-c ALL]} [-l <cnt> ] [ -v <LEDIntv> ]
```

Examples:

1. The following command displays all the diagnostics tests available for the current selected target.
   ```
   diag
   ```
2. The following command runs the MII and LED test for the selected target.
   ```
   diag -c MII -c LED
   ```
3. The following command runs all the tests five times, with an LED test interval of 8 ms for the selected target.
   ```
   diag -c all -l 5 -v 8
   ```

## 6.8. Analyzing Cables in Windows

The Cable Analysis option on the Diagnostics tab lets you monitor the conditions of each wire pair in an Ethernet Category 5 cable connection within an Ethernet network. The analysis measures the cable quality and compares it against the IEEE 802.3ab specification for compliance.

**Notes:**

- This feature can be used with Windows Server managed hosts only. It is not available for hosts operating on Linux or other OSes. You can, however use BACS on a Linux client to connect to a Windows Server host and run the cable analysis utility.
- You must have administrator privileges to run the cable analysis test.
- The network connection is temporarily lost during an analysis.
- This option is not available for DXE-820T 10 GbE network adapters.
- This option is not available for all Broadcom network adapters.
- This option is available for D-Link DXE-820T VBD drivers.

**To run a cable analysis using BACS GUI**

1. Connect the cable to a port on a switch where the port is set to Auto and the Speed & Duplex driver settings are also set to Auto.
2. Click the name of the adapter to test in the Explorer View pane.

   **Note:** *For D-Link DXE-820T network adapters, select a VBD driver; for other adapters, select an NDIS driver.*
3. From the Select a test to run list, select Cable Analysis.
4. Click Run.
5. In the error message window that warns of the network connection being temporarily interrupted, click Yes.

**Distance**

The valid cable length in meters (except when the Noise result is returned).

**Status**

The result of the analysis for the indicated pair.

- Good. Good cable/PCB signal paths, but no gigabit link.
- Crossed. Pin short or crosstalk along two or more cable/PCB signal paths.
- Open. One or both pins are open for a twisted pair.
- Short. Two pins from the same twisted pair are shorted together.
- Noise. Persistent noise present (most likely caused by Forced 10/100).
- GB Link. Gigabit link is up and running.
- N/A. Algorithm failed to reach a conclusion.

**Link**

The link connection speed and mode.

**Status**

The status after the test is run, either completed or failed.

There are several factors that could have an effect on the test results:

- Link partner. Various switch and hub manufacturers implement different PHYs. Some PHYs are not IEEE compliant.
- Cable quality. Category 3, 4, 5, and 6 may affect the test results.
- Electrical interference. The testing environment may affect the test results.

**To run a cable analysis using BACS CLI**

You can use the following CLI commands to run cable analysis for the specified target. This command is available for physical device ports only.

```
BACScli -t <target type> -f <target format> -i <target ID> cablediag
```

Example:

1. The following command runs the cable diagnostics test for the current selected target.

   ```
   BACScli -t PHYPORTS -f bdf -i 01:00.00 "cablediag"
   ```

In Interactive mode, use the list <view> and select <idx> commands to select the desired target device. Use the cablediag command to run the cable analysis test for the selected target.

Example:

1. The following command runs the cable diagnostics test for the currently selected NDIS adapter.
   cablediag

## 6.9. Managing the LAN Device

The LAN function represents the Ethernet (NDIS) functionality available under the PCI Function. User can view current values of various NDIS driver parameters, configure NDIS driver parameters, view attached FCoE targets and LUN information by selecting FCoE object in object explorer panel.

The available tabs for the NDIS function are as follows:

At the NDIS level, you can view parameters, configure parameters, and run tests from the following tabs:

- Information
- Configuration
- Diagnostics
- Statistics

### 6.9.1. Viewing NDIS Information

Select the NDIS driver in the Explorer View pane, and then select the Information tab to view NDIS-level information.

**Notes:**

- Information about Broadcom network adapters may be more comprehensive than information about network adapters made by others.
- Some information may not be available for all Broadcom network adapters.

**Viewing Driver Information**

Information Tab: Driver Information

**Driver Status**

The status of the adapter driver.

- Loaded. Normal operating mode. The adapter driver has been loaded by the OS and is functioning.
- Not Loaded. The driver associated with the adapter has not been loaded by the OS.
- Information Not Available. The value is not obtainable from the driver that is associated with the adapter.

**Driver Name**

The file name of the adapter driver.

**Driver Version**

The current version of the adapter driver.

**Driver Date**

The creation date of the adapter driver.

**Information Tab: Vital Signs**

**IP Address**

The network address associated with the adapter. If the IP address is all 0s, the associated driver has not been bound with Internet Protocol (IP).

**IPv6 IP Address**

The IPv6 network address associated with the adapter.

**MAC Address**

A physical MAC (media access control) address that is assigned to the adapter by the manufacturer. The physical address is never all 0s.

**Permanent MAC Address**

The unique hardware address assigned to the network adapter.

**Offload Capabilities**

The offload capabilities supported by the adapter. This information is only available for D-Link DXE-820T network adapters.

- iSCSI. iSCSI offload for block-level transfer of data.
- LSO. Large Send Offload (LSO) prevents an upper level protocol such as TCP from breaking a large data packet into a series of smaller packets with headers appended to them.
- CO. Checksum Offload (CO) allows the TCP/IP/UDP checksums for send and receive traffic to be calculated by the adapter hardware rather than by the host CPU.

### 6.9.2. Configuring the NDIS Driver

Select the NDIS driver in the Explorer View pane, and then select the Configuration tab to configure NDIS-level parameters. After making changes, click Apply to confirm the changes to all properties. Click Reset to return the properties to their original values. Click Defaults to restore all settings to their default values.

**Notes:**

- Clicking Reset after clicking Defaults, but before clicking Apply, will purge all values.
- Apply must be clicked to make changes go into effect.
- Any changes to existing settings will be lost upon clicking Defaults.

**Notes:**

- You must have administrator privileges to change the values for a property.
- The list of available properties for your particular adapter may be different.
- Some properties may not be available for all Broadcom network adapters.
- If an adapter is included as a member of a team and you change any advanced property, then you must rebuild the team to ensure that the team's advanced properties are properly set.

**Configuration Tab: Advanced**

**Ethernet@Wirespeed**

Enables a Gigabit Ethernet adapter to establish a link at a lower speed when only two pairs of wires are available in the cabling plant. The default setting for this property is Enabled.

**Flow Control**

Enables or disables the receipt or transmission of PAUSE frames. PAUSE frames allow the network adapter and a switch to control the transmit rate. The side that is receiving the PAUSE frame momentarily stops transmitting. Enable flow control to reduce the number of packets lost.

**Note**: If Jumbo Packet is set to 5000 bytes or greater on network adapters that support 10 Gbps link speed, ensure that Flow Control is set to Auto to prevent the system performance from performing at less than optimal levels. This limitation exists on a per-port basis.

- Auto (default). Receive and transmit PAUSE frame functionality are optimized. This option indicates that the adapter automatically adjusts the flow control settings for optimal performance, and its purpose is not enabling auto negotiation of the flow control parameters.
- Disable. Receive and transmit PAUSE frame functionality are disabled.
- Rx Enabled. Receive PAUSE frame is enabled.
- Rx & Tx Enabled. Receive and transmit PAUSE frame are enabled.
- Tx Enabled. Transmit PAUSE frame is enabled.

**IPv4 Checksum Offload**

Normally, the checksum function is computed by the protocol stack. When you select one of the Checksum Offload property values (other than None), the checksum can be computed by the network adapter.

- Rx Enabled. Enables receive TCP/IP/UDP checksum offload.
- Tx Enabled. Enables transmit TCP/IP/UDP checksum offload.
- Tx/Rx Enabled (default). Enables transmit and receive TCP/IP/UDP checksum offload.
- None. Disables checksum offload.

**IPv4 Large Send Offload**

Normally, the TCP segmentation is done by the protocol stack. When you enable the Large Send Offload property, the TCP segmentation can be done by the network adapter. The default setting for this property is Enabled. This property is only available for D-Link DXE-820T network adapters.

**IPv6 Checksum Offload**

Normally, the checksum function is computed by the protocol stack. When you select one of the Checksum Offload property values (other than None), the checksum can be computed by the network adapter.

- Rx Enabled. Enables receive TCP/IP/UDP checksum offload.
- Tx Enabled. Enables transmit TCP/IP/UDP checksum offload.
- Tx/Rx Enabled (default). Enables transmit and receive TCP/IP/UDP checksum offload.
- None. Disables checksum offload.

**IPv6 Large Send Offload**

Normally, the TCP segmentation is done by the protocol stack. When you enable the Large Send Offload property, the TCP segmentation can be done by the network adapter. The default setting for this property is Enabled. This property is only available for D-Link DXE-820T network adapters.

**Jumbo Packet**

Enables the network adapter to transmit and receive oversized Ethernet frames that are greater than 1514 bytes, but less than or equal to 9000 bytes in length (9600 bytes for network adapters that operate at 10 Gbps). This property requires the presence of a switch that is able to process jumbo frames. This property is only available for DXE-820T network adapters.

Frame size is set at 1500 bytes by default. To increase the size of the received frames, raise the byte quantity in 500-byte increments.

**Note**: *If Jumbo Packet is set to 5000 bytes or greater on network adapters that support 10 Gbps link speed, ensure that Flow Control is set to Auto to prevent the system performance from performing at less than optimal levels. This limitation exists on a per-port basis.*

*If SR-IOV is enabled on a virtual function (VF) on the adapter, ensure that the same jumbo packet settings is configured on both the VF and the Microsoft synthetic adapter. You can configure these values using Windows Device Manager > Advanced properties.*

*If there is a mismatch in the values, the SRIOV function will be shown the Degraded state in Hyper-V > Networking Status.*

**Locally Administered Address**

The Locally Administered Address is a user-defined MAC address that is used in place of the MAC address originally assigned to the network adapter. Every adapter in the network must have its own unique MAC address. This locally administered address consists of a 12-digit hexadecimal number.

- Value. Assigns a unique node address for the adapter.
- Not Present (default). Uses the factory-assigned node address on the adapter.

The appropriate assigned ranges and exceptions for the locally administered address include the following:

- The range is 00:00:00:00:00:01 to FF:FF:FF:FF:FF:FD.
- Do not use a multicast address (least significant bit of the high byte = 1).
- Do not use all 0s or all Fs.

**Receive Side Scaling**

Allows configuring network load balancing across multiple CPUs. The default setting for this property is Enabled.

**Switch Configuration.**

Allows configuring of the connected switch for the network adapters.

**Note:** Switch Configuration only applies to blade configurations.

- SW_Config_10G (default). Sets the switch speed to 10 Gbit/s.
- SW_Config_1G. Sets the switch speed to 1 Gbit/s.

**Speed & Duplex**

The Speed & Duplex property sets the connection speed and mode to that of the network. Note that Full-Duplex mode allows the adapter to transmit and receive network data simultaneously.

- 10 Mb Full. Sets the speed at 10 Mbit/s and the mode to Full-Duplex.
- 10 Mb Half. Sets the speed at 10 Mbit/s and the mode to Half-Duplex.
- 100 Mb Full. Sets the speed at 100 Mbit/s and the mode to Full-Duplex.
- 100 Mb Half. Sets the speed at 100 Mbit/s and the mode to Half-Duplex.
- 1 Gb Full. Sets the speed at 1000 Mb Full-Duplex mode only. Not available for 1 Gb ports.
- 2.5 Gb Full. Sets the speed at 2.5
- 10 GB Full. Sets the speed to 10 Gbit/s and the mode to Full-Duplex. Not available for 1 Gb ports.
- Auto (default). Sets the speed and mode for optimum network connection (recommended).

  **Notes:**

  - Auto is the recommended setting. This setting allows the network adapter to dynamically detect the line speed of the network. Whenever the network capability changes, the network adapter automatically detects and adjusts to the new line speed and duplex mode. A speed of 1 Gbit/s is enabled by selecting Auto, when that speed is supported.

  - 1 Gb Full Auto must be attached to a link partner that is also capable of a 1 Gb connection. Since the connection is limited to a 1 Gb connection only, the Ethernet@Wirespeed feature will be disabled. If the link partner supports a 1 Gb connection only, management traffic (IPMI or UMP) in the absence of an operating system may also be affected.
  - 10 Mb Half and 100 Mb Half settings force the network adapter to connect to the network in Half-Duplex mode. Note that the network adapter may not function if the network is not configured to operate at the same mode.
  - 10 Mb Full and 100 Mb Full settings force the network adapter to connect to the network in Full-Duplex mode. The network adapter may not function if the network is not configured to operate at the same mode.

**Speed & Duplex (SerDes)**

- 1 Gb Full. Forces the speed to 1 Gb Full based on a matching setting for its link partner.
- Auto (default). Sets the speed to auto-negotiate with its link partner at the highest matching speed.
- Auto with 1Gb Fallback Full. Sets the speed to auto-negotiate with its link partner, but if the attached link partner is forced at 1 Gbit/s, it will fall back to this mode.
- Hardware Default. Sets the speed to negotiate according to the setting specified by the manufacturer (see manufacturer documentation for more information).

**Note**: *The following properties pertain to Windows Vista operating systems.*

**TCP/UDP Checksum Offload (IPv4)**

Allows configuring checksum offload for the IPv4 protocol.

- Disable. Disables checksum offload.
- Rx Enabled. Enables receive TCP/IP/UDP checksum offload.
- Tx Enabled. Enables transmit TCP/IP/UDP checksum offload.
- TX & Rx Enabled (default). Enables transmit and receive TCP/IP/UDP checksum offload.

**Priority & VLAN**

Allows enabling both the prioritization of network traffic and VLAN tagging. VLAN tagging only occurs when the VLAN ID setting is configured with a value other than 0 (zero).

- Priority & VLAN Enabled (default). Allows for packet prioritization and VLAN tagging.
- Priority & VLAN Disabled. Prevents packet prioritization and VLAN tagging.
- Priority Enabled. Allows packet prioritization only.

- VLAN Enabled. Allows VLAN tagging only.

**Note:** *If an intermediate driver is managing the network adapter for VLAN tagging, the Priority & VLAN Disabled and Priority Enabled settings should not be used. Use the Priority & VLAN Enabled setting and change the VLAN ID to 0 (zero).*

**VLAN ID**

Enables VLAN tagging and configures the VLAN ID when Priority & VLAN Enabled is selected as the Priority & VLAN setting. The range for the VLAN ID is 1 to 4094 and must match the VLAN tag value on the connected switch. A value of 0 (default) in this field disables VLAN tagging.

Risk Assessment of VLAN Tagging through the NDIS Miniport Driver

Broadcom's NDIS 6.0 miniport driver provides the means to allow a system containing a DXE-820T adapter to connect to a tagged VLAN. On Windows XP systems, this support was only provided through the use of an intermediate driver (e.g., Broadcom Advanced Server Program - BASP). Unlike BASP, however, the NDIS 6 driver's support for VLAN participation is only for a single VLAN ID.

Also unlike BASP, the NDIS 6.0 driver only provides VLAN tagging of the outbound packet, but does not provide filtering of incoming packets based on VLAN ID membership. This is the default behavior of all miniport drivers. While the lack of filtering packets based on VLAN membership may present a security issue, the following provides a risk assessment based on this driver limitation for an IPv4 network:

A properly configured network that has multiple VLANs should maintain separate IP segments for each VLAN. This is necessary since outbound traffic relies on the routing table to identify which adapter (virtual or physical) to pass traffic through and does not determine which adapter based on VLAN membership.

Since support for VLAN tagging on Broadcom's NDIS 6.0 driver is limited to transmit (Tx) traffic only, there is a risk of inbound traffic (Rx) from a different VLAN being passed up to the operating system. However, based on the premise of a properly configured network above, the IP segmentation and/or the switch VLAN configuration may provide additional filtration to limit the risk.

In a back-to-back connection scenario, two computers on the same IP segment may be able to communicate regardless of their VLAN configuration since no filtration of VLAN membership is occurring. However, this scenario assumes that the security may already be breached since this connection type is not typical in a VLAN environment.

If the risk above is not desirable and filtering of VLAN ID membership is required, then support through an intermediate driver would be necessary.

**iSCSI Crash Dump**

Crash dump is used to collect information on adapters that were booted remotely using iSCSI. To enable crash dump, set to Enable and reboot the system. If you perform an upgrade of the device drivers, re-enable iSCSI Crash Dump. If iSCSI Boot is configured to boot in the HBA path, then this parameter cannot be changed.

**Interrupt Moderation**

Enables interrupt moderation, which limits the rate of interrupt to the CPU during packet transmission and packet reception. The disabled option allows one interrupt for every packet transmission and packet reception. Enable is the default option.

**Number of RSS Queues**

Allows configuring RSS queues. For 1 Gbps network adapters, the RSS queue options are Auto (default), 2, 4, and 8. For 10 Gbps network adapters, the RSS queue options are Auto (default), 2, 4, 8, and 16.

**Receive Buffers**

The number of receive buffers. Receive buffers are data segments that allow the network adapter to allocate receive packets to memory. For 1 Gbps adapters, the range of valid receive buffers is 50 to 5000 in increments of 1 with 750 receive buffers as the default value.

**Receive Buffers (0=Auto)**

The number of receive buffers. Receive buffers are data segments that allow the network adapter to allocate receive packets to memory. For 10 Gbps adapters, the range of valid receive buffers is 0 to 3000 in increments of 50 with 0 receive buffers as the default value.

**Transmit Buffers (0=Auto)**

The number of transmit buffers. Transmit buffers are data segments that allow the network adapter to monitor transmit packets in the system memory. The range of valid transmit buffers is 0 to 5000 in increments of 1 with 250 transmit buffers as the default value.

**Pause on Exhausted Host Ring**

For BCM57711 and BCM57712 network adapters, there are two possible scenarios that can trigger pause frames to be generated: a host ring buffer is exhausted or the on-chip buffers are depleted. With RSS enabled inside the system, it is possible to achieve better Ethernet throughput if no pause frames are being generated in a case where a host ring buffer (of multiple RSS rings) is exhausted. The default is Disabled.

**Quality of Service**

Enables Quality of Service (QoS) to provide different priorities to different applications.

**Recv Segment Coalescing (IPv4)**

Enable Receive Segment Coalescing (IPv4). Receive Segment Coalescing is an offload technology that reduces CPU utilization for network processing on the receive side by offloading tasks from the CPU to a network adapter.

**Recv Segment Coalescing (IPv6)**

Enable Receive Segment Coalescing (IPv6). Receive Segment Coalescing is an offload technology that reduces CPU utilization for network processing on the receive side by offloading tasks from the CPU to a network adapter.

## 6.10. Viewing Resource Information

The Resources section of the Information tab displays information about connections and other essential functions for the selected network adapter.

**Note**: *Some information may not be available for all Broadcom network adapters.*

**Information Tab: Resources**

**Bus Type**

The type of input/output (I/O) interconnect used by the adapter.

### 6.10.1. Configuring System Settings

System Management on the Configurations tab allow you to view and change the values of the available properties for the system. The potentially available properties and their respective settings are described below.

### 6.10.2. Viewing Statistics

The information provided on the Statistics tab allows you to view traffic statistics for both Broadcom network adapters and network adapters made by others. Statistical information and coverage are more comprehensive for DXE-820T adapters.

To view Statistics information for any installed network adapter, click the name of the adapter listed in the Explorer View pane, then click the Statistics tab.

If any of the sections described below is not visible, then from the Context View tab on the right side of the window, select Statistics and then select the name of the missing section.

Click Refresh to get the most recent values for each statistic. Click Reset to change all values to zero for the current BACS session.

**Notes:**

- Team statistics are not compiled for a Broadcom network adapter if it is disabled.
- Some statistics may not be available for all Broadcom network adapters.

**General Statistics**

General Statistics show the transmitted and received statistics to and from the adapter.

**Frames Tx OK**

A count of the frames that were successfully transmitted. This counter is incremented when the transmit status is reported as Transmit OK.

**Frames Rx OK**

A count of the frames that were successfully received. This does not include frames received with frame-too-long, frame check sequence (FCS), length, or alignment errors, nor frames lost due to internal MAC sublayer errors. This counter is incremented when the receive status is reported as Receive OK.

**Directed Frames Tx**

A count of directed data frames that were successfully transmitted.

**Multicast Frames Tx**

A count of frames that were successfully transmitted (as indicated by the status value Transmit OK) to a group destination address other than a broadcast address.

**Broadcast Frames Tx**

A count of frames that were successfully transmitted (as indicated by the transmit status Transmit OK) to the broadcast address. Frames transmitted to multicast addresses are not broadcast frames and are excluded.

**Directed Frames Rx**

A count of directed data frames that were successfully received.

**Multicast Frames Rx**

A count of frames that were successfully received and are directed to an active nonbroadcast group address. This does not include frames received with frame-too-long, FCS, length, or alignment errors, nor frames lost because of internal MAC sublayer errors. This counter is incremented as indicated by the Receive OK status.

**Broadcast Frames Rx**

A count of frames that were successfully received and are directed to a broadcast group address. This count does not include frames received with frame-too-long, FCS, length, or alignment errors, nor frames lost because of internal MAC sublayer errors. This counter is incremented as indicated by the Receive OK status.

**Frames Rx with CRC Error**

The number of frames received with CRC errors.

**Initiator Login Statistics**

iSCSI login enables a connection for iSCSI use between the initiator and the target and is used to authenticate parties, negotiate the session's parameters, open security association protocol, and mark the connection as belonging to an iSCSI session.

**Login Accept Responses**

The number of login requests accepted by the target.

**Login other failed Responses**

The number of login requests that were not accepted by the target.

**Login Redirect Responses**

The number of responses that required further action by the initiator.

**Login Authentication Failed Responses**

The number of login requests that failed due to party authentication failure.

**Login target authentication failure**

The number of instances where the login could not authenticate the target.

**Login target negotiation failure**

The number of instances where the login could not negotiate the sessions parameters.

**Normal logout command PDU**

The number of normal logout commands issued by the initiator to remove a connection from a session or to close a session.

**Other logout command PDU**

The number of logout commands issued by the initiator for reasons other than to remove a connection from a session or to close a session.

**Local Initiator login failures**

The number of login failures likely caused by the initiator.

**Initiator Instance Statistics**

The statistics in this area pertain to all sessions.

**Session digest errors**

The number of sessions with errors due to an invalid payload or header.

**Session connection timeout error**

The number of sessions that were terminated due to any of the many timeout errors.

**Session format error**

The number of sessions with errors due to inconsistent fields, reserved fields not 0, non-existent LUN, etc.

**Sessions failed**

The number of failed sessions.

**Custom**

Custom statistics.

**Total Offload iSCSI Connections**

The total number of offloaded iSCSI connections.

**Session Statistics**

The statistics in this area only pertain to the named session.

**Session Name**

The name used for the session between the initiator and the target.

**Session Id**

The identifier used for the session between the initiator and the target.

**Bytes sent**

The number of bytes sent for the named session.

**Bytes received**

The number of bytes received for the named session.

**PDU sent**

The number of iSCSI PDUs sent for the named session.

**PDU received**

The number of iSCSI PDUs received for the named session.

**Digest errors**

The number of errors due to an invalid payload or header for the named session.

**Connection Timeout errors**

The number of connection timeout errors for the named session.

**Format errors**

The number of errors due to inconsistent fields, reserved fields not 0, non-existent LUN, etc. for the named session.

**IEEE 802.3 Statistics**

**Frames Rx with Alignment Error**

A count of the frames that were not an integral number of octets in length and do not pass the FCS check. This counter is incremented when the receive status is reported as Alignment Error.

**Frames Tx with one Collision**

A count of the frames that were involved in a single collision and were subsequently transmitted successfully. This counter is incremented when the result of a transmission is reported as Transmit OK, and the attempt value is 2.

**Frames Tx with more than one Collision**

A count of the frames that were involved in more than one collision and were subsequently transmitted successfully. This counter is incremented when the transmit status is reported as Transmit OK, and the value of the attempts variable is greater than 2 and less than or equal to the attempt limit.

**Frames Tx after Deferral**

A count of the frames that were delayed being transmitted on the first attempt because the medium was busy. The frames involved in any collision are not counted.

**Custom Statistics**

**Note**: *Custom statistics are available only for an enabled Broadcom network adapter.*

**Out of Recv. Buffer**

The number of times the adapter ran out of Receive Buffer Descriptors. This information is only available for D-Link DXE-820T network adapters.

**Frames size less than 64-byte with bad FCS**

The number of frames with a size less than 64 bytes with bad FCS.

**MAC Rx w/ Pause Command and Length = 0**

MAC control frames with the pause command and a length equal to 0.

**MAC Rx w/ Pause Command and Length greater than 0**

MAC control frames with the pause command and a length greater than 0.

**MAC Rx w/ no Pause Command**

MAC control frames with no pause command.

**MAC Sent X-on**

MAC Transmit with X-on was on.

**MAC Sent X-off**

MAC Transmit with X-on was off.

**Large Send Offload Transmit Requests**

The number of times the adapter was requested to transmit a packet performing TCP segmentation.

**Total Offload TCP Connections**

The total number of offloaded TCP connections.

**SR-IOV Switch Statistics**

This area shows the statistics for SR-IOV switches.

**Num of Active VFs**

This shows the number of active Virtual Functions (VF).

### 6.10.3. Viewing Resource Reservations

**Notes:**

- Resource Reservation information is only available for DXE-820T adapters and VBD drivers.
- Not all offload technologies are available with all adapters.
- Resource Reservation information is not available in BACS on Linux systems.
-

### 6.10.4. Configuring the IP Address for iSCSI Offload

For iSCSI-booted adapters, the Configurations tab is not available and you will not be able to perform this procedure.

**To set the IP address of the iSCSI HBA for iSCSI offload**

The iSCSI Management section of the Configurations tab allows you to set the IP address of the iSCSI HBA when using iSCSI protocol to offload network processing from the CPU to the Broadcom network adapter.

1. Click the name of the iSCSI device in the SCSI controller section of the Explorer View pane.

2. Depending on the protocol you will be using, for IPv4 DHCP or IPv6 DHCP, select Enable (not available for iSCSI booted adapters) to set the IP address dynamically using a DHCP server. Or select Disable to set the IP address using a static IP address. Enter the IP Address, Subnet Mask, and Default Gateway.
3. Configure the VLAN ID for the iSCSI HBA by entering a number for VLAN ID. The value must be between 1 and 4094.
4. After the configurations are complete, click Apply to save the settings or click Reset to revert back to the previous settings.

## 6.11. Configuring Teaming

BACS does not support teaming on Linux systems. Linux provides a similar built-in functionality called Channel Bonding. Refer to the Linux OS documentation for more information.

The teaming function allows you to group any available network adapters together to function as a team. Teaming is a method of creating a virtual NIC (a group of multiple adapters that functions as a single adapter). The benefit of this approach is that it enables load balancing and failover. Teaming is done through the Broadcom Advanced Server Program (BASP) software. For a comprehensive description of the technology and implementation considerations of the teaming software, refer to the "Broadcom Gigabit Ethernet Teaming Services" section of your Broadcom network adapter user guide.

Teaming can be accomplished by either of the following methods:

- Using the Broadcom Teaming Wizard
- Using Expert Mode

**Notes**:

- For further information regarding teaming protocols, see "Teaming" in your Broadcom network adapter user guide.
- If you do not enable LiveLink™ when configuring teams, disabling Spanning Tree Protocol (STP) at the switch is recommended. This minimizes the downtime due to spanning tree loop determination when failing over. LiveLink mitigates such issues.
- BASP is available only if a system has one or more Broadcom network adapters installed.
- Large Send Offload (LSO), and Checksum Offload properties are enabled for a team only when all of the members support and are configured for the feature.
- To physically remove a teamed NIC from a system, you must first delete the NIC from the team. Not doing this before shutting down the system could result in breaking the team on a subsequent reboot, which may result in unexpected team behavior.

- If an adapter is included as a member of a team and you change any advanced property, then you must rebuild the team to ensure that the team's advanced properties are properly set.
- You must have administrator privileges to create or modify a team.
- The load balance algorithm in a team environment in which members are connected at different speeds favors members connected with a Gigabit Ethernet link over members connected at lower speed links (100 Mbps or 10 Mbps) until a threshold is met. This is normal behavior.

### 6.11.1. Team Types

You can create four types of load balance teams:

- Smart Load Balance and Failover
- Link Aggregation (802.3ad)
- Generic Trunking (FEC/GEC)/802.3ad-Draft Static
- SLB (Auto-Fallback Disable) – The Auto-Fallback Disable feature is configured for Smart Load Balance and Failover type teams in the Teaming Wizard.

  **Note**: *DXE-820T network adapters with iSCSI enabled is supported only in an SLB team type.*

**Smart Load Balance and Failover**

In this type of team, a standby member handles the traffic if all of the load balance members fail (a failover event). All load balance members have to fail before the standby member takes over. When one or more of the load balance members is restored (fallback), the restored team member(s) resumes the handling of the traffic. The LiveLink feature is supported for this type of team.

**Link Aggregation (802.3ad)**

In this type of team, you can dynamically configure the network adapters that have been selected to participate in a given team. If the link partner is not correctly configured for IEEE 802.3ad link configuration, errors are detected and noted. All adapters in the team are configured to receive packets for the same MAC address. The outbound load balancing scheme is determined by the BASP driver. The link partner of the team determines the load balancing scheme for inbound packets. In this mode, at least one of the link partners must be in active mode.

**Generic Trunking (FEC/GEC)/802.3ad-Draft Static**

This type of team is very similar to the link aggregation type, in that all adapters in the team must be configured to receive packets for the same MAC address. This mode does not provide link aggregation control protocol (LACP) or marker protocol support. This mode supports a variety of environments where the link partners are statically configured to support a proprietary trunking mechanism. Trunking supports load balancing and failover for both outbound and inbound traffic.

**SLB (Auto-Fallback Disable)**

This team is identical to Smart Load Balance and Failover, with the following exception: when the standby member is active, if a primary member comes back online, the team continues using the standby member rather than switching back to the primary member. This type of team is supported only for situations in which the network cable is disconnected and reconnected to the network adapter. It is not supported for situations in which the adapter is removed/installed through Device Manager or Hot-Plug PCI. If any primary adapter assigned to a team is disabled, the team functions as a Smart Load Balancing and Failover type of team in which auto-fallback occurs. The LiveLink feature is supported for this type of team.

### 6.11.2. Standby Team Member and Auto-Fallback Disable Mode

You can designate one team member in an SLB type of team to be the standby member. The standby member does not actively send and receive normal network traffic while other adapters on the team are active. If all of the active adapters on the team fail or are disconnected, the standby member takes over the handling of the network activities.

In Auto-Fallback Disable mode, if a load balance member returns on line, the team continues using the standby member rather than switching back to using the load balance member. Consequently, the adapter that was initially designated a load balance member remains in an inactive state and becomes the new standby member.

### 6.11.3. LiveLink

LiveLink is a feature of BASP that is available for the Smart Load Balancing (SLB) and SLB (Auto-Fallback Disable) type of teaming. The purpose of LiveLink is to detect link loss beyond the switch and to route traffic only through team members that have a live link.

### 6.11.4. Using the Broadcom Teaming Wizard

You can use the Broadcom Teaming Wizard to create a team, configure an existing team if a team has already been created, or create a VLAN.

1. Create or edit a team:

   To create a new team, select Create a Team from the Team menu, or right-click one of the devices in the "Unassigned Adapters" section and select Create a Team. This option is not available if there are no devices listed in the "Unassigned Adapters" sections, which means all adapters are already assigned to teams.

To configure an existing team, right-click one of the teams in the list and select Edit Team. This option is only available if a team has already been created and is listed in the Team Management pane.

**Note**: *If you prefer to work without the wizard for now, click Expert Mode. If you want to always use Expert Mode to create a team, select Default to Expert Mode on next start. See Using Expert Mode.*

2. To continue using the wizard, click Next.

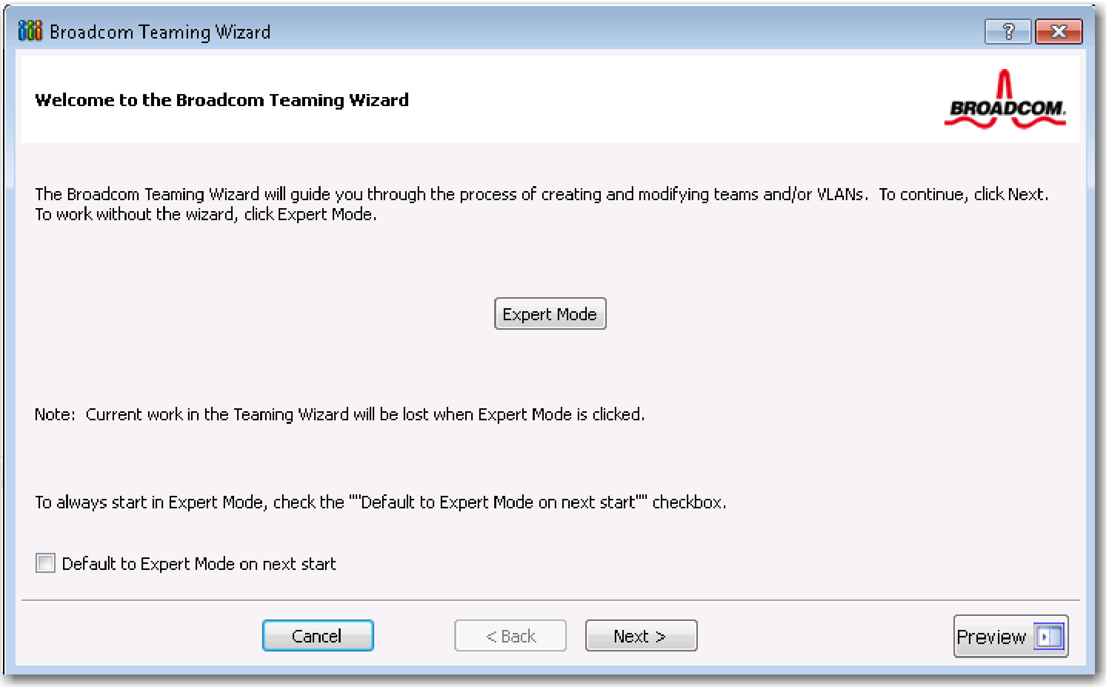

Broadcom Teaming Wizard

Welcome to the Broadcom Teaming Wizard

The Broadcom Teaming Wizard will guide you through the process of creating and modifying teams and/or VLANs. To continue, click Next. To work without the wizard, click Expert Mode.

Note: Current work in the Teaming Wizard will be lost when Expert Mode is clicked.

To always start in Expert Mode, check the ""Default to Expert Mode on next start"" checkbox.

□ Default to Expert Mode on next start

3. Type the team name and then click Next. If you want to review or change any of your settings, click Back. Click Cancel to discard your settings and exit the wizard.

**Note:** *The team name cannot exceed 39 characters, cannot begin with spaces, and cannot contain any of the following characters*: & \ / : * ? < > |

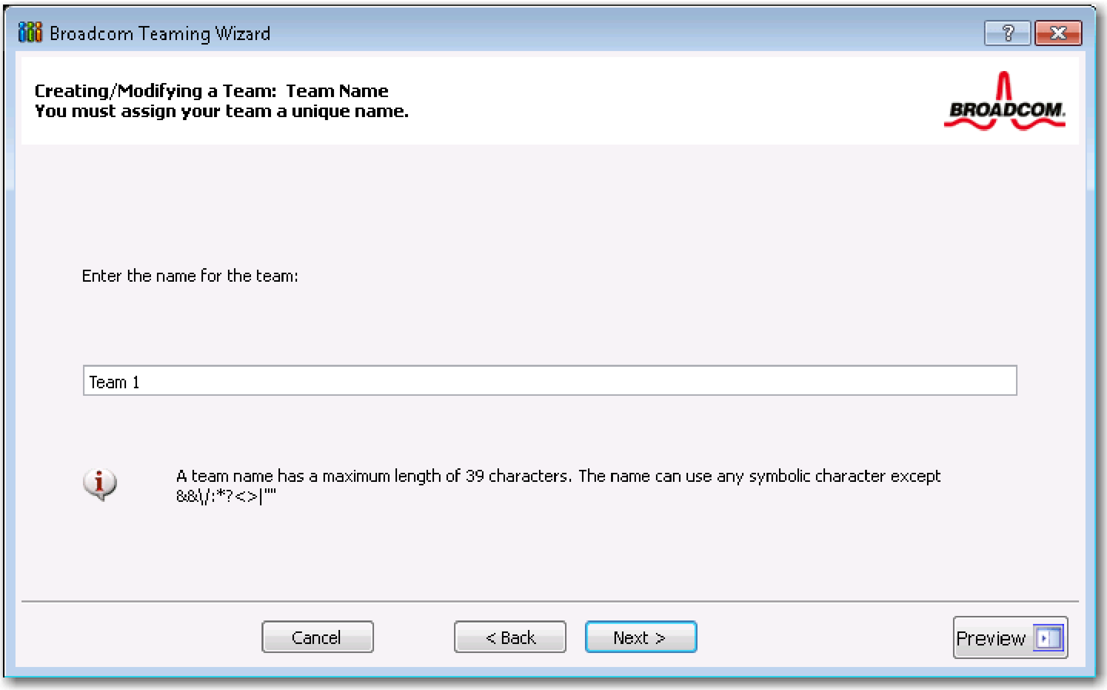

Creating/Modifying a Team: Team Name
You must assign your team a unique name.

Enter the name for the team:

Team 1

A team name has a maximum length of 39 characters. The name can use any symbolic character except &&\/:*?<>|""

4. Select the type of team you want to create.
5. Select Enable Hyper-V Mode if you want to enable Windows virtualization services.
6. If the team type is an SLB type team, click Next. If the team type is not an SLB type team, then a dialog box appears. Verify that the network switch connected to the team members is configured correctly for the team type, click OK, and continue.

   **Note**: *Network adapters with iSCSI enabled are supported only in an SLB team type. To continue with the creation of non-SLB team types, first disable iSCSI by deselecting iSCSI Offload Engine from the Resource Reservations area of the Configurations tab.*

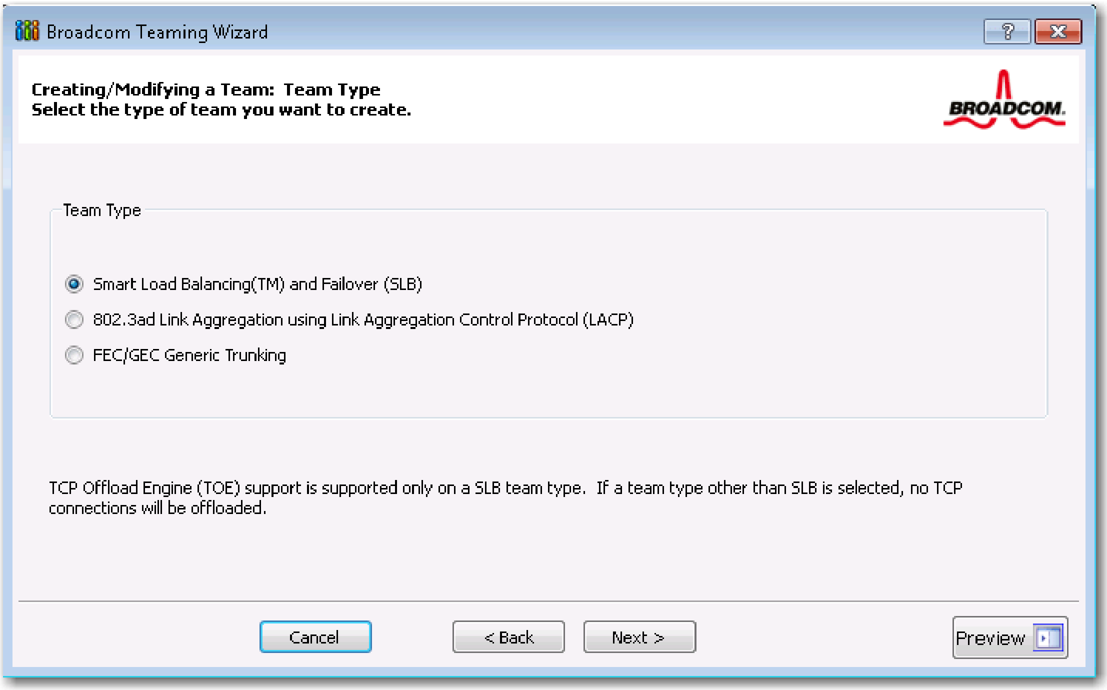

Creating/Modifying a Team: Team Type
Select the type of team you want to create.

Team Type

● Smart Load Balancing(TM) and Failover (SLB)
○ 802.3ad Link Aggregation using Link Aggregation Control Protocol (LACP)
○ FEC/GEC Generic Trunking

TCP Offload Engine (TOE) support is supported only on a SLB team type. If a team type other than SLB is selected, no TCP connections will be offloaded.

7. From the Available Adapters list, click the adapter you want to add to the team and then click Add. Remove team members from the Team Members list by clicking the adapter and then clicking Remove. Click Next.

   **Note:** *There must be at least one Broadcom network adapter assigned to the team.*

   The Large Send Offload (LSO) and Checksum Offload (CO) columns indicate if the LSO, Jumbo MTU, and/or the CO properties are supported for the adapter. The LSO, Jumbo MTU, and CO properties are enabled for a team only when all of the members support and are configured for the feature. If this is the case, then the team offload capabilities appear on the bottom of the screen.

   Adding a network adapter to a team where its driver is disabled may negatively affect the offloading capabilities of the team. This may have an impact on the team's performance. Therefore, it is recommended that only driver-enabled network adapters be added as members to a team.

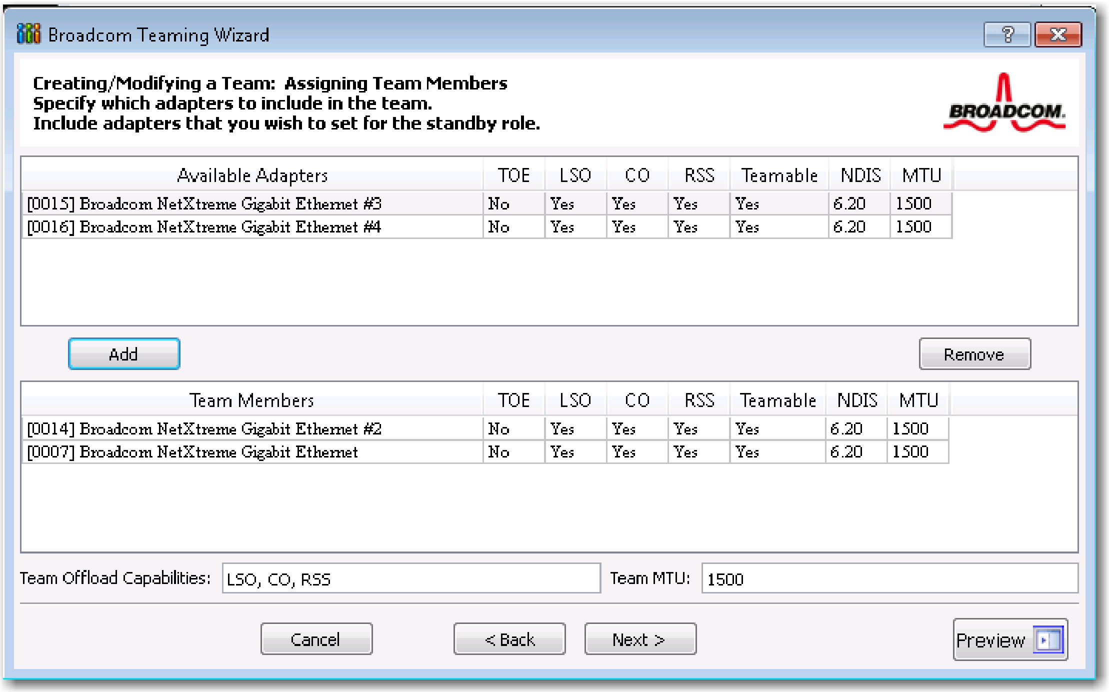

Broadcom Teaming Wizard

**Creating/Modifying a Team: Assigning Team Members**
**Specify which adapters to include in the team.**
**Include adapters that you wish to set for the standby role.**

| Available Adapters | TOE | LSO | CO | RSS | Teamable | NDIS | MTU |
|---|---|---|---|---|---|---|---|
| [0015] Broadcom NetXtreme Gigabit Ethernet #3 | No | Yes | Yes | Yes | Yes | 6.20 | 1500 |
| [0016] Broadcom NetXtreme Gigabit Ethernet #4 | No | Yes | Yes | Yes | Yes | 6.20 | 1500 |

| Team Members | TOE | LSO | CO | RSS | Teamable | NDIS | MTU |
|---|---|---|---|---|---|---|---|
| [0014] Broadcom NetXtreme Gigabit Ethernet #2 | No | Yes | Yes | Yes | Yes | 6.20 | 1500 |
| [0007] Broadcom NetXtreme Gigabit Ethernet | No | Yes | Yes | Yes | Yes | 6.20 | 1500 |

Team Offload Capabilities: LSO, CO, RSS

Team MTU: 1500

8. If you want to designate one of the adapters as a standby member (optional), select Use the following member as a standby member, then choose the standby member from the list of adapters.
9. The Auto-Fallback Disable mode feature allows the team to continue using the standby member rather than switching back to the primary member if the primary member comes back online. To enable this feature, select Enable Auto-Fallback Disable mode. Click Next.

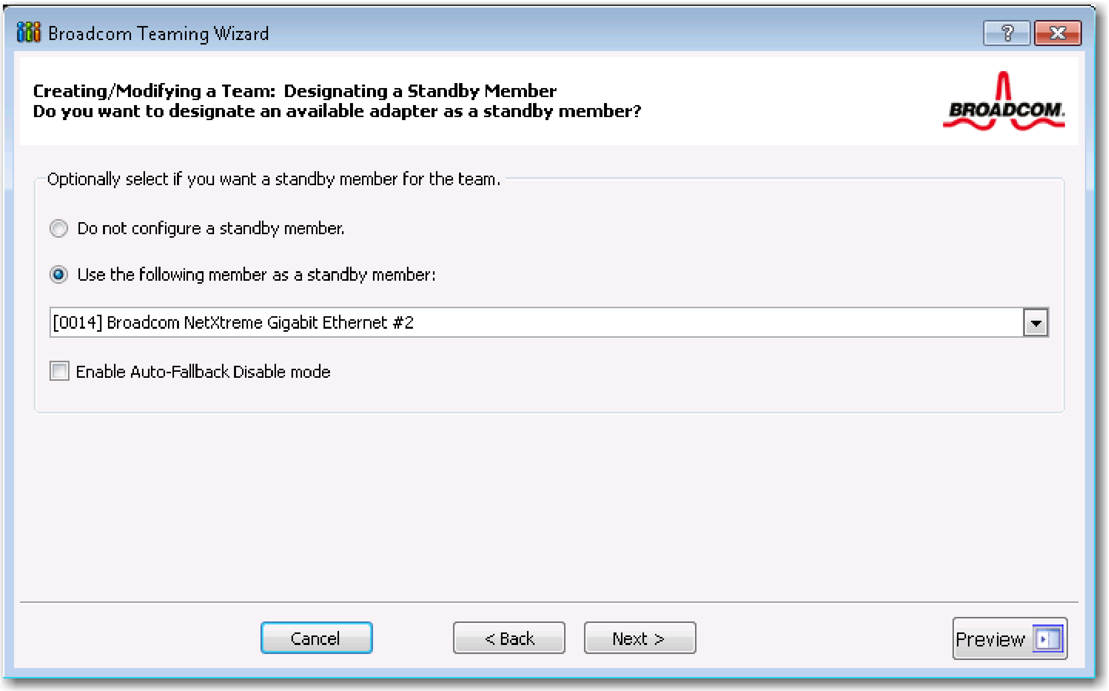

Broadcom Teaming Wizard

**Creating/Modifying a Team: Designating a Standby Member**
**Do you want to designate an available adapter as a standby member?**

Optionally select if you want a standby member for the team.

○ Do not configure a standby member.

● Use the following member as a standby member:

[0014] Broadcom NetXtreme Gigabit Ethernet #2

□ Enable Auto-Fallback Disable mode

10. If you want to configure LiveLink, select Yes, otherwise select No, then click

Next.

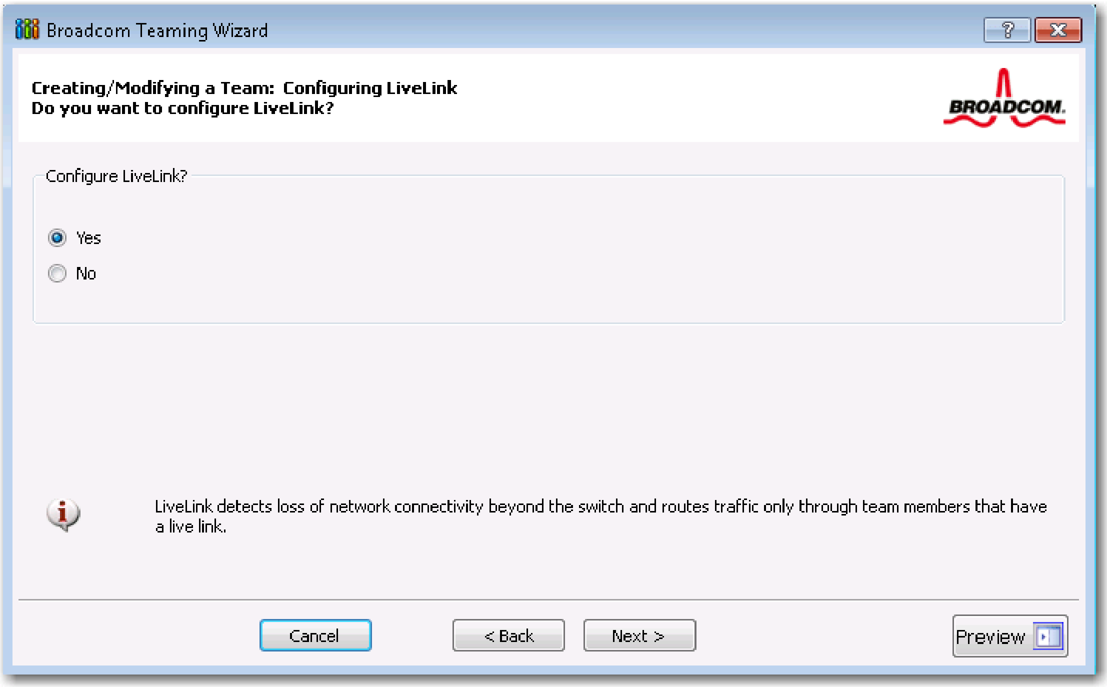

Broadcom Teaming Wizard

**Creating/Modifying a Team: Configuring LiveLink**
**Do you want to configure LiveLink?**

Configure LiveLink?

● Yes
○ No

LiveLink detects loss of network connectivity beyond the switch and routes traffic only through team members that have a live link.

11. Select the probe interval (the number of seconds between each retransmission of a link packet to the probe target) and the maximum number of probe retries (the number of consecutively missed responses from a probe target before a failover is triggered).
12. Set the Probe VLAN ID to allow for connectivity with probe targets residing on a tagged VLAN. The number set must match the VLAN ID of the probe targets as well as the port(s) on the switch to which the team is connected.

    **Note**: *Each LiveLink enabled team can only communicate with Probe Targets on a single VLAN. Also, VLAN ID 0 is equivalent to an untagged network. If the Probe VLAN ID is set to a value other than 0, then a VLAN must be created with an identical VLAN tag value (see Step 18.).*
13. Click the probe target at the top of the list, click Edit Target IP Address, type the target IP address in the IP Address box for one or all probe targets, and then click OK. Click Next.

    **Note**: *Only the first probe target is required. You can specify up to three additional probe targets to serve as backups by assigning IP addresses to the other probe targets.*
14. Select a listed team member, click Edit Member IP Address, and then type the member IP address in the IP Address box. Repeat for all listed team members and then click OK. Click Next.

    **Note**: *All of the member IP addresses must be in the same subnet as the subnet of the probe targets.*

15. If you want to create a VLAN on the team, select Add VLAN, or if you want to change the settings of an existing VLAN, select Edit VLAN, then click Next. If you do not want to create or edit a VLAN, select Skip Manage VLAN, then click Next, and continue with the wizard from the Finish screen (see Step 20. of this procedure).

    VLANs enable you to add multiple virtual adapters that are on different subnets. The benefit of this is that your system can have one network adapter that can belong to multiple subnets.

    **Note:** *VLANs can only be created when all team members are DXE-820T adapters.*

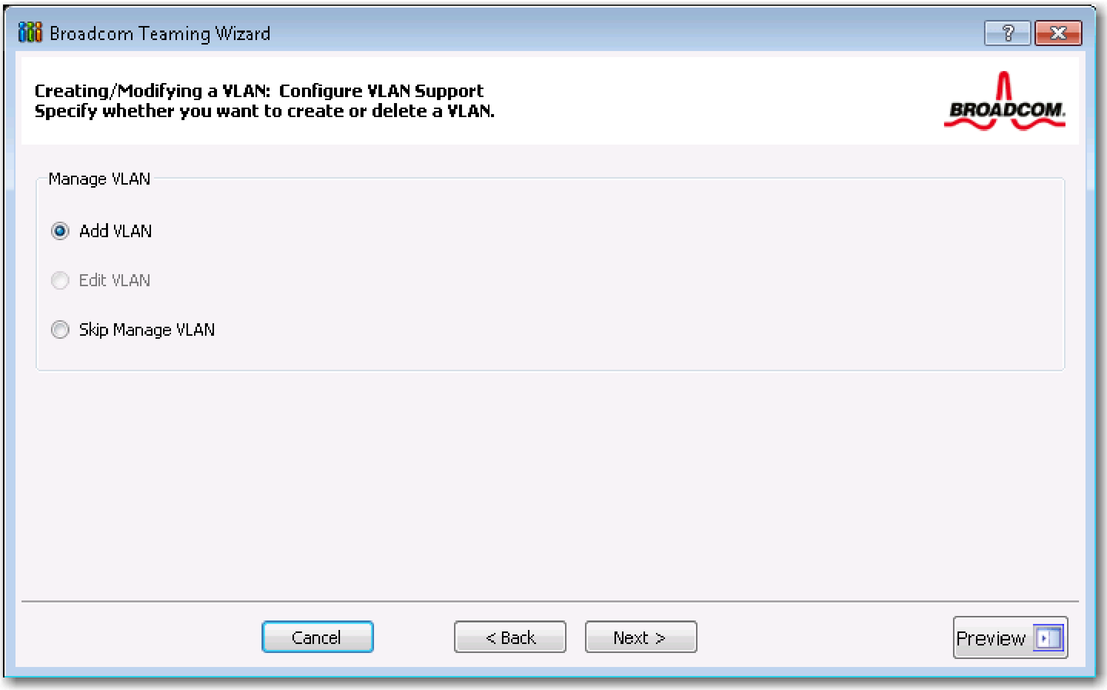

Broadcom Teaming Wizard

**Creating/Modifying a VLAN: Configure VLAN Support**
**Specify whether you want to create or delete a VLAN.**

Manage VLAN

● Add VLAN

○ Edit VLAN

○ Skip Manage VLAN

16. Type the VLAN name and then click Next.

    **Note**: *The team name cannot exceed 39 characters, cannot begin with spaces, and cannot contain any of the following characters*: & \ / : * ? < > |

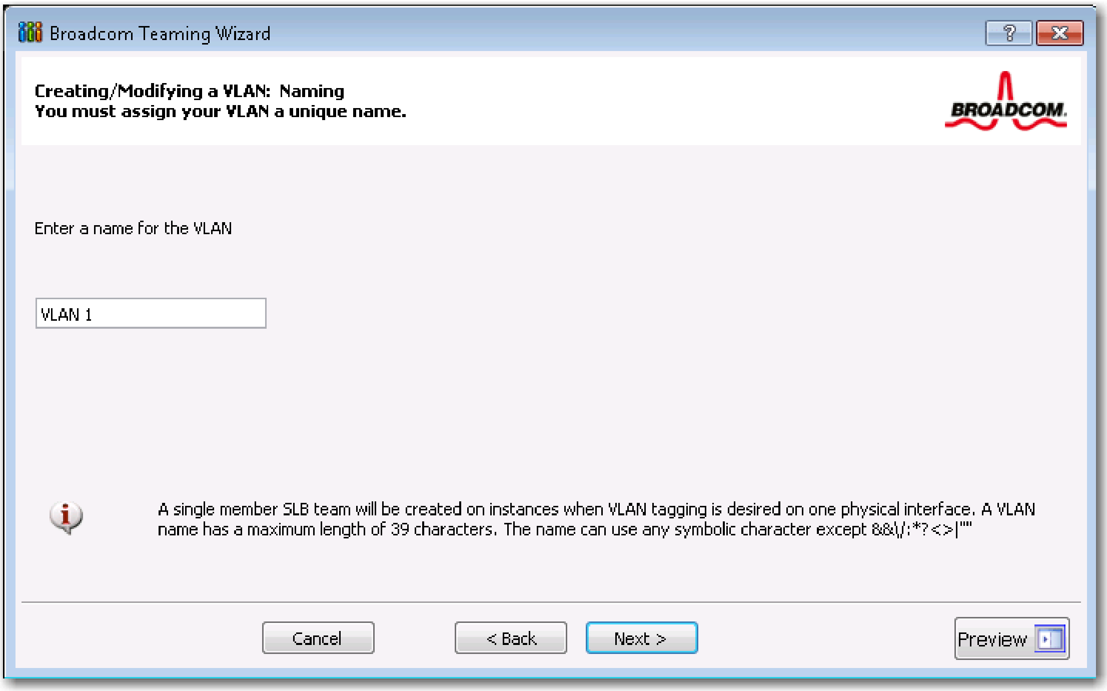

**Creating/Modifying a VLAN: Naming**
**You must assign your VLAN a unique name.**

Enter a name for the VLAN

VLAN 1

A single member SLB team will be created on instances when VLAN tagging is desired on one physical interface. A VLAN name has a maximum length of 39 characters. The name can use any symbolic character except &&\/:*?<>|""

17. To tag the VLAN, select Tagged and then click Next. Otherwise, click Untagged, click Next, and continue with the wizard to add additional VLANs.

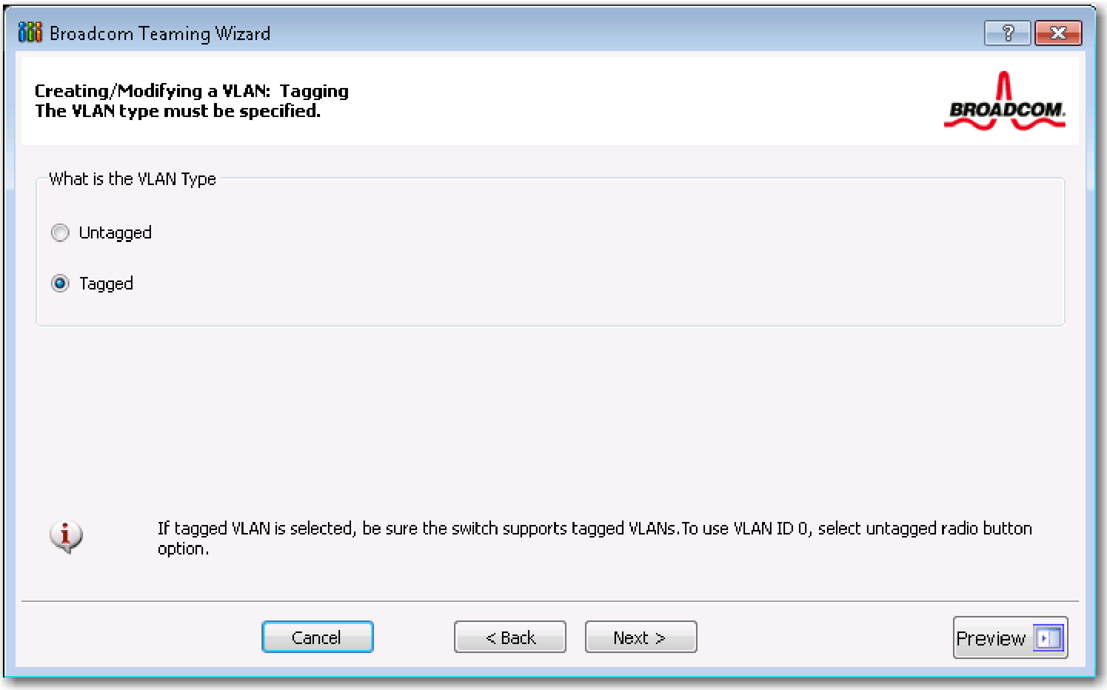

**Creating/Modifying a VLAN: Tagging**
**The VLAN type must be specified.**

What is the VLAN Type

○ Untagged

● Tagged

If tagged VLAN is selected, be sure the switch supports tagged VLANs.To use VLAN ID 0, select untagged radio button option.

18. Type the VLAN tag value and then click Next. The value must be between 1 and 4094.

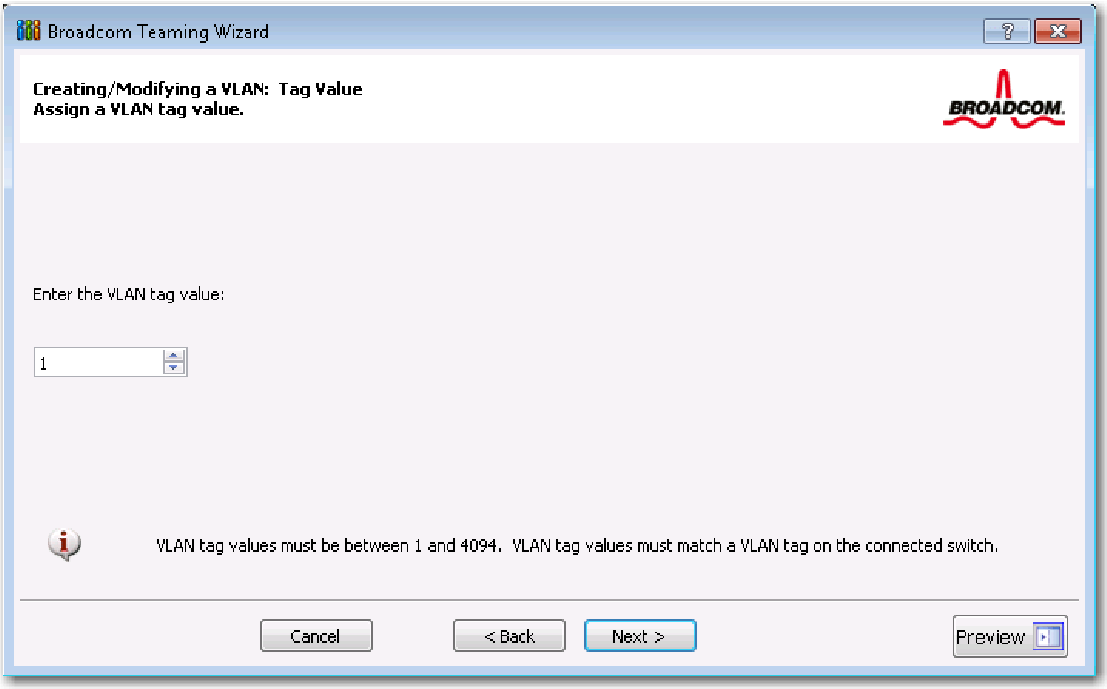

**Creating/Modifying a VLAN: Tag Value**
**Assign a VLAN tag value.**

Enter the VLAN tag value:

1

VLAN tag values must be between 1 and 4094. VLAN tag values must match a VLAN tag on the connected switch.

19. Select Yes to add or manage another VLAN and then click Next. Repeat until you do not want to add or manage any additional VLANs.
    **Note**: *You can define up to 64 VLANs per team (63 VLANs that are tagged and 1 VLAN that is not tagged). Adding several VLANS may slow down the reaction time of the Windows interface due to memory and processor time usage for each VLAN. The degree to which Windows performance may suffer depends on system configuration.*

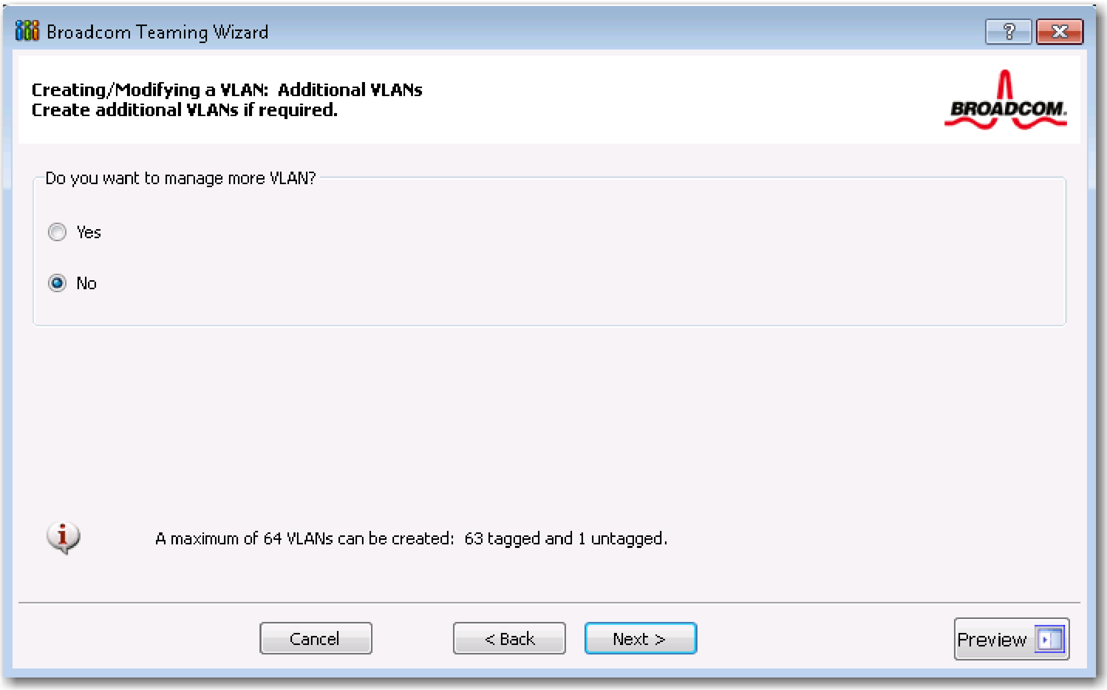

**Creating/Modifying a VLAN: Additional VLANs**
**Create additional VLANs if required.**

Do you want to manage more VLAN?

○ Yes

● No

A maximum of 64 VLANs can be created: 63 tagged and 1 untagged.

20. To apply and commit the changes to the team, select Commit changes to system

and Exit the wizard. To apply your changes but continue using the wizard, select Save changes and continue to manage more teams. Click Finish.

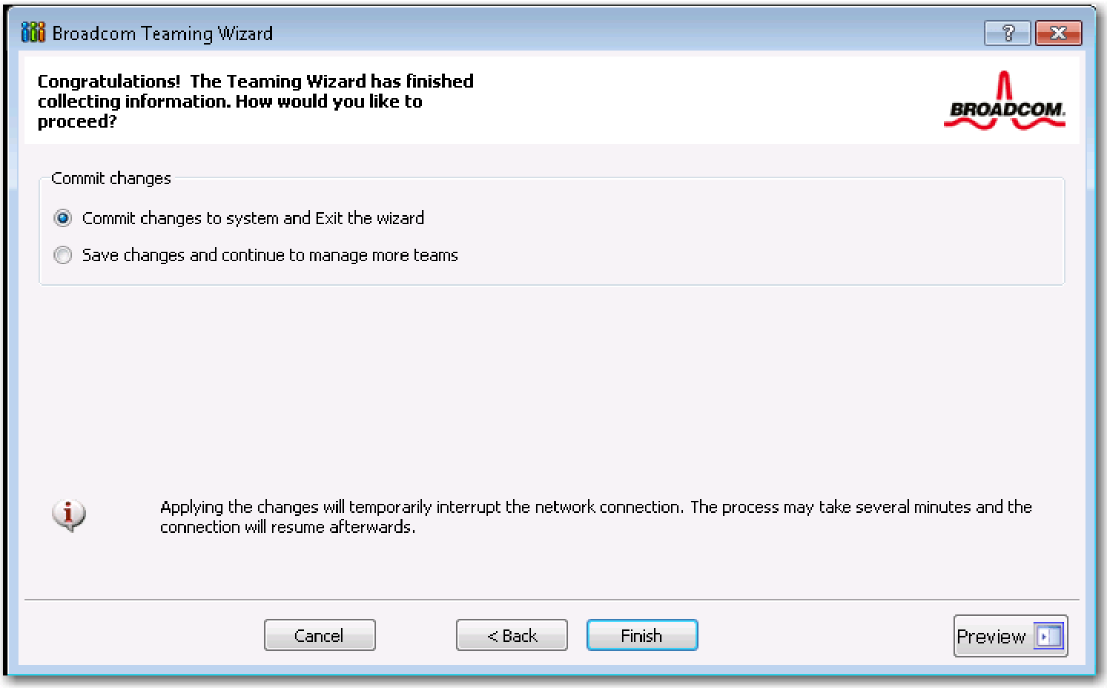

**Congratulations! The Teaming Wizard has finished collecting information. How would you like to proceed?**

Commit changes

● Commit changes to system and Exit the wizard

○ Save changes and continue to manage more teams

Applying the changes will temporarily interrupt the network connection. The process may take several minutes and the connection will resume afterwards.

**Note:** *At any point in the Broadcom Teaming Wizard procedure, click Preview to get a visual representation of what the team will look like before committing any changes.*

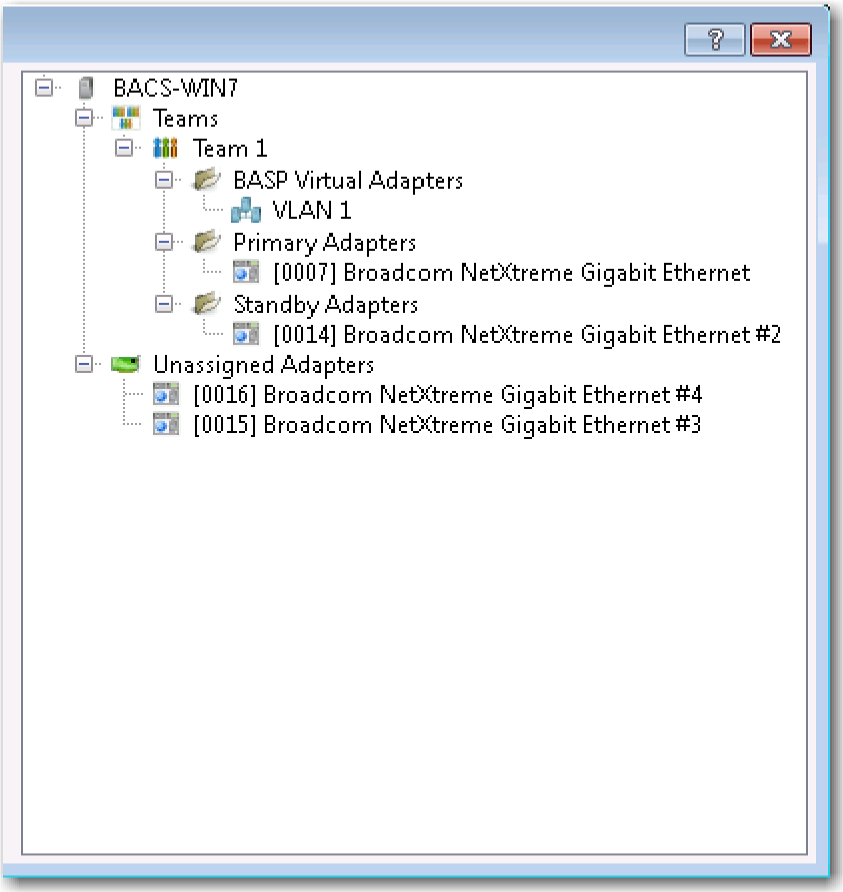

```
BACS-WIN7
  Teams
    Team 1
      BASP Virtual Adapters
        VLAN 1
      Primary Adapters
        [0007] Broadcom NetXtreme Gigabit Ethernet
      Standby Adapters
        [0014] Broadcom NetXtreme Gigabit Ethernet #2
  Unassigned Adapters
    [0016] Broadcom NetXtreme Gigabit Ethernet #4
    [0015] Broadcom NetXtreme Gigabit Ethernet #3
```

21. Click the team name in the Team Management pane to view the team's properties in the Information tab, transfer and receive data in the Statistics tab, and team customization options in the Configurations tab.

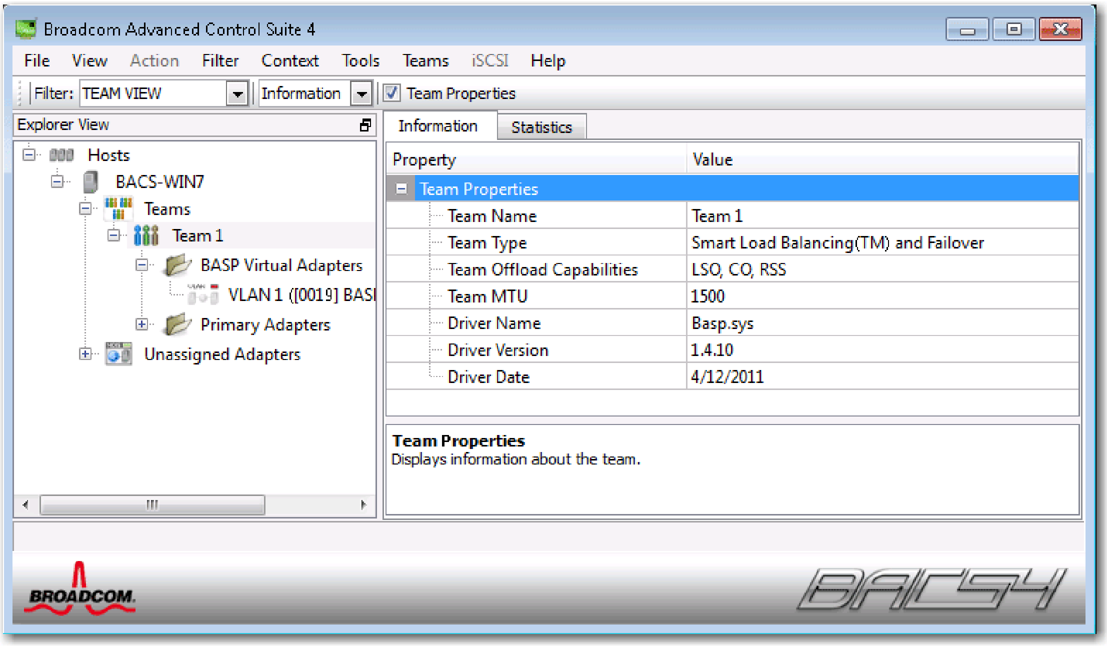

| Property | Value |
|---|---|
| ⊟ Team Properties | |
| Team Name | Team 1 |
| Team Type | Smart Load Balancing(TM) and Failover |
| Team Offload Capabilities | LSO, CO, RSS |
| Team MTU | 1500 |
| Driver Name | Basp.sys |
| Driver Version | 1.4.10 |
| Driver Date | 4/12/2011 |

**Team Properties**
Displays information about the team.

### 6.11.5. Using Expert Mode

Use Expert Mode to create a team, modify a team, add a VLAN, and configure LiveLink for a Smart Load Balance and Failover and SLB (Auto-Fallback Disable) team. To create a team using the wizard, see Using the Broadcom Teaming Wizard.

To set the default Teaming Mode, select Options from the Tools menu. In the Options window, click the General tab, then select Expert Mode or Wizard Mode (the default is Wizard Mode).

#### Creating a Team

**Note**: *Enabling Dynamic Host Configuration Protocol (DHCP) is not recommended for members of an SLB type of team.*

1. From the Teams menu, select Create a Team, or right-click one of the devices in the "Unassigned Adapters" section and select Create a Team. This option is not available if there are no devices listed in the "Unassigned Adapters" sections, which means all adapters are already assigned to teams.
2. Click Expert Mode.

   **Note:** *If you want to always use Expert Mode to create a team, click Default to use Expert Mode on your next startup.*
3. Click the Create Team tab.

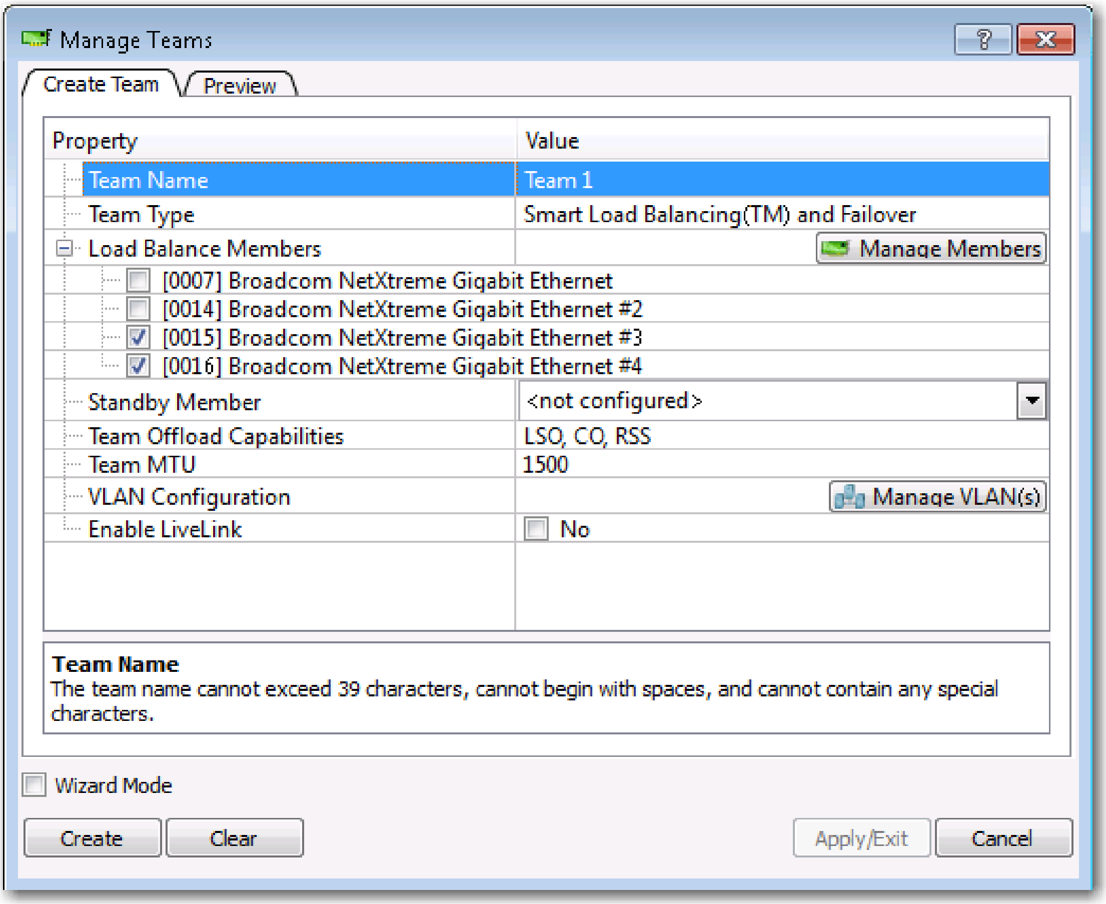

| Property | Value |
|---|---|
| Team Name | Team 1 |
| Team Type | Smart Load Balancing(TM) and Failover |
| Load Balance Members | Manage Members |
| □ [0007] Broadcom NetXtreme Gigabit Ethernet | |
| □ [0014] Broadcom NetXtreme Gigabit Ethernet #2 | |
| ☑ [0015] Broadcom NetXtreme Gigabit Ethernet #3 | |
| ☑ [0016] Broadcom NetXtreme Gigabit Ethernet #4 | |
| Standby Member | <not configured> |
| Team Offload Capabilities | LSO, CO, RSS |
| Team MTU | 1500 |
| VLAN Configuration | Manage VLAN(s) |
| Enable LiveLink | □ No |

**Team Name**
The team name cannot exceed 39 characters, cannot begin with spaces, and cannot contain any special characters.

□ Wizard Mode

**Note:** *The Create Team tab appears only if there are teamable adapters available.*

4. Click the Team Name field to enter a team name.
5. Click the Team Type field to select a team type.
6. Click Hyper-V Mode if you want to enable Windows virtualization services.
7. Assign any available adapter or adapters to the team by moving the adapter from the Available Adapters list to the Load Balance Members list. There must be at least one adapter in the Load Balance Members list.
8. You can assign any other available adapter to be a standby member by selecting it from the Standby Member list.

   **Note:** *There must be at least one Broadcom network adapter assigned to the team.*

   The Large Send Offload (LSO), and Checksum Offload (CO) columns indicate if the LSO, and/or the CO properties are supported for the adapter. The LSO, and CO properties are enabled for a team only when all of the members support and are configured for the feature. If this is the case, then the team offload capabilities appear on the bottom of the screen.

   Adding a network adapter to a team where its driver is disabled may negatively affect the offloading capabilities of the team. This may have an impact on the team's performance. Therefore, it is recommended that only driver-enabled network adapters be added as members to a team.

9. Type the value for Team MTU.
10. Click Create to save the team information.
11. Repeat steps 4. through 10. to define additional teams. As teams are defined, they can be selected from the team list, but they have not yet been created. Click the Preview tab to view the team structure before applying the changes.
12. Click Apply/Exit to create all the teams you have defined and exit the Manage Teams window.
13. Click Yes when the message is displayed indicating that the network connection will be temporarily interrupted.

    **Notes:**

    - The team name cannot exceed 39 characters, cannot begin with spaces, and cannot contain any of the following characters: & \ / : * ? < > |
    - Team names must be unique. If you attempt to use a team name more than once, an error message is displayed indicating that the name already exists.
    - The maximum number of team members is 8.
    - When team configuration has been correctly performed, a virtual team adapter driver is created for each configured team.
    - If you disable a virtual team and later want to re-enable it, you must first disable and re-enable all team members before you re-enable the virtual team.
    - When you create Generic Trunking and Link Aggregation teams, you cannot designate a standby member. Standby members work only with Smart Load Balancing and Failover and SLB (Auto-Fallback Disable) types of teams.
    - For an SLB (Auto-Fallback Disable) team, to restore traffic to the load balance members from the standby member, click the Fallback button on the Team Properties tab.
    - When configuring an SLB team, although connecting team members to a hub is supported for testing, it is recommended to connect team members to a switch.
    - Not all network adapters made by others are supported or fully certified for teaming.
14. Configure the team IP address.
    - From Control Panel, double-click Network Connections.
    - Right-click the name of the team to be configured, and then click Properties.
    - On the General tab, click Internet Protocol (TCP/IP), and then click

Properties.

- Configure the IP address and any other necessary TCP/IP configuration for the team, and then click OK when finished.

**Modifying a Team**

After you have created a team, you can modify the team in the following ways:

- Change the type of team
- Change the members assigned to the team
- Add a VLAN
- Modify a VLAN (using Expert Mode)
- Remove a team or a VLAN (using Expert Mode)

**To modify a team**

1. From the Team menu, click Edit Team, or right-click one of the teams in the list and select Edit Team. This option is only available if a team has already been created and is listed in the Team Management pane.
2. The wizard Welcome screen appears. Click Next to continue modifying a team using the wizard or click Expert Mode to work in Expert Mode.

   **Note:** *The Edit Team tab in Expert Mode appears only if there are teams configured on the system.*
3. Click the Edit Team tab.

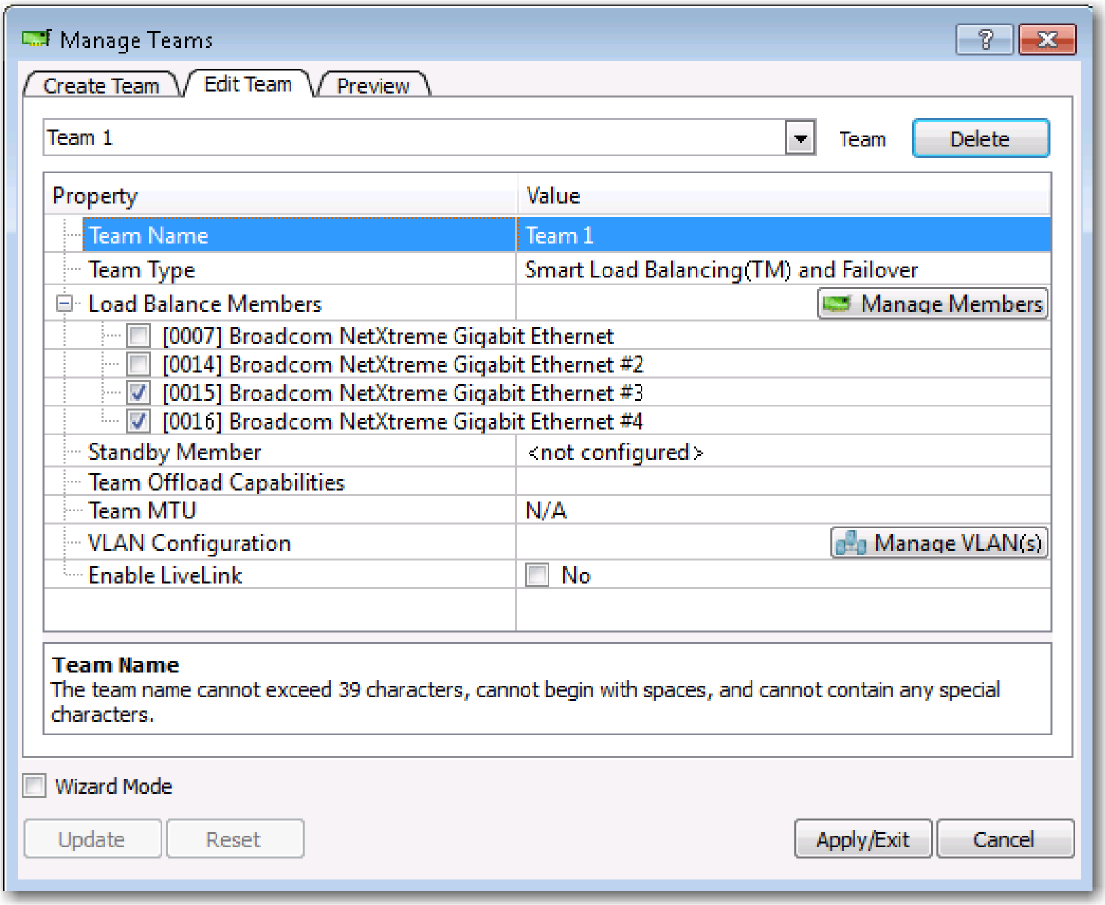

| Property | Value |
|---|---|
| Team Name | Team 1 |
| Team Type | Smart Load Balancing(TM) and Failover |
| Load Balance Members | Manage Members |
| □ [0007] Broadcom NetXtreme Gigabit Ethernet | |
| □ [0014] Broadcom NetXtreme Gigabit Ethernet #2 | |
| ☑ [0015] Broadcom NetXtreme Gigabit Ethernet #3 | |
| ☑ [0016] Broadcom NetXtreme Gigabit Ethernet #4 | |
| Standby Member | <not configured> |
| Team Offload Capabilities | |
| Team MTU | N/A |
| VLAN Configuration | Manage VLAN(s) |
| Enable LiveLink | □ No |

**Team Name**
The team name cannot exceed 39 characters, cannot begin with spaces, and cannot contain any special characters.

□ Wizard Mode

4. Make the desired changes, and then click Update. The changes have not yet been applied; click the Preview tab to view the updated team structure before applying the changes.
5. Click Apply/Exit to apply the updates and exit the Manage Teams window.
6. Click Yes when the message is displayed indicating that the network connection will be temporarily interrupted.

**Adding a VLAN**

You can add virtual LANs (VLANs) to a team. This enables you to add multiple virtual adapters that are on different subnets. The benefit of this is that your system can have one network adapter that can belong to multiple subnets. With a VLAN, you can couple the functionality of load balancing for the load balance members, and you can employ a failover adapter.

You can define up to 64 VLANs per team (63 VLANs that are tagged and 1 VLAN that is not tagged). VLANs can only be created when all teams members are DXE-820T adapters. If you try to create a VLAN with a different adapter, an error message is displayed.

**To configure a team with a VLAN**

1. From the Teams menu, select Add VLAN.
2. The Welcome screen appears.
3. Click Expert Mode.
4. On the Create Team tab of the Manage Teams window, click Manage VLAN(s).
5. Type the VLAN name, then select the type and ID.
6. Click Create to save the VLAN information. As VLANs are defined, they can be selected from the Team Name list, but they have not yet been created.
7. Continue this process until all VLANs are defined, then click OK to create them.
8. Click Yes when the message is displayed indicating that the network connection will be temporarily interrupted.

   **Note***: To maintain optimum adapter performance, your system should have 64 MB of system memory for each of the eight VLANs created per adapter.*

**Viewing VLAN Properties and Statistics and Running VLAN Tests**

To view VLAN properties and statistics and to run VLAN tests

1. Select one of the listed VLANs.
2. Click the Information tab to view the properties of the VLAN adapter.
3. Click the Statistics tab to view the statistics for the VLAN adapter.
4. Click the Diagnostics tab to run a network test on the VLAN adapter.

**Deleting a VLAN**

The procedure below applies when you are in Expert Mode.

To delete a VLAN

1. Select the VLAN to delete.
2. From the Teams menu, select Remove VLAN.
3. Click Apply.
4. Click Yes when the message is displayed indicating that the network connection will be temporarily interrupted.

   **Note:** *If you delete a team, any VLANs configured for that team will also be deleted.*

**Configuring LiveLink for a Smart Load Balancing and Failover and SLB (Auto-Fallback Disable) Team**

LiveLink is a feature of BASP that is available for the Smart Load Balancing (SLB) and SLB (Auto-Fallback Disable) type of teaming. The purpose of LiveLink is to detect link loss beyond the switch and to route traffic only through team members that have a live link.

Read the following notes before you attempt to configure LiveLink.

**Notes:**

- Before you begin configuring LiveLink™, review the description of LiveLink. Also verify that each probe target you plan to specify is available and working. If the IP address of the probe target changes for any reason, LiveLink must be reconfigured. If the MAC address of the probe target changes for any reason, you must restart the team.
- A probe target must be on the same subnet as the team, have a valid (not a broadcast, multicast, or unchaste), statically-assigned IP address, and be highly available (always on).
- To ensure network connectivity to the probe target, ping the probe target from the team.
- You can specify up to four probe targets.
- The IP address assigned to either a probe target or team member cannot have a zero as the first or last octet.

**To configure LiveLink**

1. From the Teams menu, select Edit Team.
2. Click Expert Mode (to configure LiveLink using the Teaming Wizard, see Using the Broadcom Teaming Wizard).
3. In the Manage Teams window, click the Edit Team tab.
4. Select Enable LiveLink. The LiveLink Configuration options appear below.
5. It is recommended to accept the default values for Probe interval (the number of seconds between each retransmission of a link packet to the probe target) and Probe maximum retries (the number of consecutively missed responses from a probe target before a failover is triggered). To specify different values,

click the desired probe interval in the Probe interval (seconds) list and click the desired maximum number of probe retries in the Probe maximum retries list.

6. Set the Probe VLAN ID to correspond with the VLAN where the probe target(s) resides. This will apply the appropriate VLAN tag to the link packet based on the shared configuration of the attached switch port(s).

   **Note:** *Each LiveLink enabled team can only communicate with Probe Targets on a single VLAN. Also, VLAN ID 0 is equivalent to an untagged network.*
7. Select Probe Target 1 and type the target IP address for one or all probe targets.

   **Note:** *Only the first probe target is required. You can specify up to 3 additional probe targets to serve as backups by assigning IP addresses to the other probe targets.*
8. Select one of the listed team members and type the member IP address.

   **Note:** *All of the member IP addresses must be in the same subnet as the probe targets.*
9. Click Update. Repeat these steps for each of the other listed team members.
10. Click Apply/Exit.

**Saving and Restoring a Team Configuration**

**To save a configuration**

1. From the File menu, select Team Save As.
2. Type the path and file name of the new configuration file, and then click Save.

   The configuration file is a text file that can be viewed by any text editor. The file contains information about both the adapter and the team configuration.

**To restore a configuration**

From the File menu, select Team Restore.

1. Click the name of the file to be restored, and then click Open.

   **Note:** *If necessary, go to the folder where the file is located.*
2. Click Apply.
3. Click Yes when the message is displayed indicating that the network connection will be temporarily interrupted.
4. If a configuration is already loaded, a message is displayed that asks if you want to save your current configuration. Click Yes to save the current configuration. Otherwise, the configuration data that is currently loaded is lost.

   **Note**: T*he team may take a very long time to restore if the team was configured with multiple VLANs and each VLAN was configured with one or more static IP addresses.*

### 6.11.6. Viewing BASP Statistics

The Statistics section shows performance information about the network adapters that are on a team.

To view BASP Statistics information for any team member adapter or the team as a whole, click the name of the adapter or team listed in the Team Management pane, then click the Statistics tab.

Click Refresh to get the most recent values for each statistic. Click Reset to change all values to zero.

## 6.12. Configuring With the Command Line Interface Utility

An alternate method to BACS for configuring Broadcom network adapters is with BACSCLI, which is a Broadcom utility that allows you to view information and configure network adapters using a console in either a non-interactive command line interface (CLI) mode or an interactive mode. As with BACS, BACSCLI provides information about each network adapter, and enables you to perform detailed tests, run diagnostics, view statistics, and modify property values. BACSCLI also allows you the ability to team network adapters together for load balancing and failover.

For a complete list of available commands and examples, see the BACSCLI ReadMe text file on the installation CD.

### 6.12.1. Supported Operating Systems

BACSCLI is supported on the following operating systems:

- Windows Server
- Linux Server

For information on the latest supported OS versions, see BACSCLI_Readme.txt in your software distribution.

### 6.12.2. Installation

On a system with DXE-820T network adapters, BACSCLI is installed when BACS is installed with the installer.

## 6.13. Troubleshooting BACS

Problem: When attempting to open BACS on a Linux System, the following error message displays:

"Another instance of the BACS client appears to be running on this system. Only one instance of the BACS client can be running at a time. If you are sure that no other BACS client is running, then a previous instance may have quit unexpectedly."

Solution: This message displays if you try to run a second instance of BACS. If you receive this message but are certain that no instance of BACS is currently running, a previous instance of BACS may have quit unexpectedly. To clear that instance,

remove the file:
"/dev/shm/sem.Global-BACS-{C50398EE-84A7-4bc3-9F6E-25A69603B9C0}."

# 7. User Diagnostics in DOS: D-Link DXE-820T Network Adapter User Guide

## 7.1. Introduction

Broadcom NetXtreme II User Diagnostics is a MS-DOS based application that runs a series of diagnostic tests (see Table 3) on the D-Link DXE-820T network adapters in your system. Broadcom NetXtreme II User Diagnostics also allows you to update device firmware and to view and change settings for available adapter properties. There are two versions of the Broadcom NetXtreme II User Diagnostics: uxdiag.exe (for BCM5706/BCM5708/BCM5709 network adapters) and uediag.exe (for BCM57710 network adapters).

To run Broadcom NetXtreme II User Diagnostics, create an MS-DOS 6.22 bootable disk containing the uxdiag.exe or uediag.exe file. Next, start the system with the boot disk in drive A. See Performing Diagnostics for further instructions on running diagnostic tests on Broadcom network adapters.

## 7.2. System Requirements

Operating System: MS-DOS 6.22

Software: uxdiag.exe (BCM5706/BCM5708/BCM5709), or uediag.exe (BCM57710)

## 7.3. Performing Diagnostics

At the MS-DOS prompt, type uxdiag (for BCM5706/BCM5708/BCM5709 network adapters) or uediag (for BCM577XX and BCM578XX network adapters) followed by the command options. The uxdiag command options are shown in Table 5 and the uediag command options are shown in Table 6. For example, to run all diagnostic tests on adapter #1 except Group B tests:

```
C:\>uxdiag -c 1 -t b
```

**Note:** *You must include uxdiag or uediag at the beginning of the command string each time you type a command.*

**Table 5: uxdiag Command Options**

| Command Options | Description |
|---|---|
| uxdiag | Performs all tests on all D-Link DXE-820T network adapters in your system. |
| uxdiag -c <devnum> | Specifies the adapter (devnum) to test. Use **all** in place of a specific device number to test all adapters. |
| uxdiag -cof | Allows tests to continue after detecting a failure. |

| | |
|---|---|
| uxdiag -F | Forces an upgrade of the image without checking the version. |
| uxdiag -fbc <bc_image> | Specifies the bin file to update the bootcode. |
| uxdiag -fib <ib_image> | Specifies the bin file for iSCSI boot. |
| uxdiag -fibc | Programs the iSCSI configuration block. Used only with -fib <ib_image>. |
| uxdiag -fibp | Programs the iSCSI configuration software. Used only with -fib <ib_image>. |
| uxdiag -fipmi <ipmi_image> | Specifies the bin file to update IPMI firmware. |
| uxdiag -fmba <mba_image> | Specifies the bin file to update the MBA. |
| uxdiag -fncsi <ncsi_image> | Specifies the bin file to update the NCSI firmware. |
| uxdiag -fnvm <raw_image> | Programs the raw image into NVM. |
| uxdiag -fump <ump_image> | Specifies the bin file to update UMP firmware. |
| uxdiag -help | Displays the Broadcom NetXtreme II User Diagnostics (uxdiag) command options. |
| uxdiag -I <iteration num> | Specifies the number of iterations to run on the selected tests. |
| uxdiag -idmatch | Enables matching of VID, DID, SVID, and SSID from the image file with device IDs. Used only with -fnvm <raw_image>. |
| uxdiag -log <file> | Logs the test results to a specified log file. |
| uxdiag -mba <1/0> | Enables/disables Multiple Boot Agent (MBA) protocol.<br>1 = Enable<br>0 = Disable |
| uxdiag -mbap <n> | Sets the MBA boot protocol.<br>0 = PXE<br>1 = RPL<br>2 = BOOTP<br>3 = iSCSI_Boot |
| uxdiag -mbas <n> | Sets the MBA/PXE speed.<br>0 = Auto<br>1 = 10H<br>2 = 10F<br>3 = 100H<br>4 = 100F<br>6 = 1000F |
| uxdiag -mbav <1\|0> | Enables/disables MBA VLAN.<br>1 = Enable |

| | |
|---|---|
| | 0 = Disable |
| uxdiag -mbavval <n> | Sets MBA VLAN (<65536). |
| uxdiag -mfw <1/0> | Enables/disables management firmware.<br>1 = Enable<br>0 = Disable |
| uxdiag -t <groups/tests> | Disables certain groups/tests. |
| uxdiag -T <groups/tests> | Enables certain groups/tests. |
| uxdiag -ver | Displays the version of Broadcom NetXtreme II User Diagnostics (uxdiag) and all installed adapters. |

## Table 6: uediag Command Options

| Command Options | Description |
|---|---|
| uediag | Performs all tests on all D-link DXE-820T network adapters in your system. |
| uediag -c <device#> | Specifies the adapter (device#) to test. Similar to -dev (for backward compatibility). |
| uediag -cof | Allows tests to continue after detecting a failure. |
| uediag -dev <device#> | Specifies the adapter (device#) to test. |
| uediag -F | Forces an upgrade of the image without checking the version. |
| uediag -fbc <bc_image> | Specifies the bin file to update the bootcode. |
| uediag -fbc1 <bc1_image> | Specifies the bin file to update bootcode 1. |
| uediag -fbc2 <bc2_image> | Specifies the bin file to update bootcode 2. |
| uediag -fl2b <l2b_image> | Specifies the bin file for L2B firmware. |
| uediag -fib <ib_image> | Specifies the bin file for iSCSI boot. |
| uediag -fibc | Programs iSCSI configuration block 0. Used only with -fib <ib_image>. |
| uediag -fibc2 | Programs iSCSI configuration block 1. Used only with -fib <ib_image>. |
| uediag -fibp | Programs iSCSI configuration software. Used only with -fib <ib_image>. |
| uediag -fipmi <ipmi_image> | Specifies the bin file to update IPMI firmware. |
| uediag -fmba <mba_image> | Specifies the bin file to update the MBA. |
| uediag -fnvm <raw_image> | Programs the raw image into NVM. |
| uediag -fump <ump_image> | Specifies the bin file to update UMP firmware. |

| | |
|---|---|
| uediag -help | Displays the Broadcom NetXtreme II User Diagnostics (uediag) command options. |
| uediag -I <iteration#> | Specifies the number of iterations to run on the selected tests. |
| uediag -idmatch | Enables matching of VID, DID, SVID, and SSID from the image file with device IDs: Used only with -fnvm <raw_image>. |
| uediag -log <logfile> | Logs the tests results to a specified log file. |
| uediag -mba <1/0> | Enables/disables Multiple Boot Agent (MBA) protocol.<br>1 = Enable<br>0 = Disable |
| uediag -mbap <n> | Sets the MBA boot protocol.<br>0 = PXE<br>1 = RPL<br>2 = BOOTP<br>3 = iSCSI_Boot |
| uediag -mbav <1/0> | Enables/disables MBA VLAN.<br>1 = Enable<br>0 = Disable |
| uediag -mbavval <n> | Sets MBA VLAN (<65536). |
| uediag -mfw <1/0> | Enables/disables management firmware.<br>1 = Enable<br>0 = Disable |
| uediag -t <groups/tests> | Disables certain groups/tests. |
| uediag -T <groups/tests> | Enables certain groups/tests. |
| uediag -ver | Displays the version of Broadcom NetXtreme II User Diagnostics (uediag) and all installed adapters. |

## 7.4. Diagnostic Test Descriptions

The diagnostic tests are divided into four groups: Basic Functional Tests (Group A), Memory Tests (Group B), Block Tests (Group C), and Ethernet Traffic Tests (Group D). The diagnostic tests are listed and described in Table 7.

**Table 7: Diagnostic Tests**

| **Test** | | **Description** |
|---|---|---|
| **Number** | **Name** | |
| **Group A: Basic Functional Tests** | | |
| A1 | Register | Verifies that registers accessible through the PCI/PCIe interface implement the expected read-only or read/write attributes by attempting to modify those |

| | | |
|---|---|---|
| | | registers. |
| A2 | PCI Configuration | Checks the functionality of the PCI Base Address Register (BAR) by varying the amount of memory requested by the BAR and verifying that the BAR actually requests the correct amount of memory (without actually mapping the BAR into system memory). Refer to PCI or PCI-E specifications for details on the BAR and its addressing space. |
| A3 | Interrupt | Generates a PCI interrupt and verifies that the system receives the interrupt and invokes the correct ISR. A negative test is also performed to verify that a masked interrupt does not invoke the ISR. |
| A5 | MSI | Verifies that a Message Signaled Interrupt (MSI) causes an MSI message to be DMA'd to host memory. A negative test is also performed to verify that when an MSI is masked, it does not write an MSI message to host memory. |
| A6 | Memory BIST | Invokes the internal chip Built-In Self Test (BIST) command to test internal memory. |
| **Group B: Memory Tests** | | |
| B1 | TXP Scratchpad | The Group B tests verify all memory blocks of the D-Link DXE-820T network adapter by writing various data patterns (0x55aa55aa, 0xaa55aa55, walking zeroes, walking ones, address, etc.) to each memory location, reading back the data, and then comparing it to the value written. The fixed data patterns are used to ensure that no memory bit is stuck high or low, while the walking zeroes/ones and address tests are used to ensure that memory writes do not corrupt adjacent memory locations. |
| B2 | TPAT Scratchpad | |
| B3 | RXP Scratchpad | |
| B4 | COM Scratchpad | |
| B5 | CP Scratchpad | |
| B6 | MCP Scratchpad | |
| B7 | TAS Header Buffer | |
| B8 | TAS Payload Buffer | |
| B9 | RBUF via GRC | |
| B10 | RBUF via Indirect Access | |
| B11 | RBUF Cluster List | |
| B12 | TSCH List | |
| B13 | CSCH List | |
| B14 | RV2P Scratchpads | |
| B15 | TBDC Memory | |
| B16 | RBDC Memory | |
| B17 | CTX Page Table | |
| B18 | CTX Memory | |
| **Group C: Block Tests** | | |
| C1 | CPU Logic and DMA Interface | Verifies the basic logic functionality of all the on-chip CPUs. It also exercises the DMA interface exposed to those CPUs. The internal CPU tries to initiate DMA activities (both read and write) to system memory and then compares the |

| | | |
|---|---|---|
| | | values to confirm that the DMA operation completed successfully. |
| C2 | RBUF Allocation | Verifies the RX buffer (RBUF) allocation interface by allocating and releasing buffers and checking that the RBUF block maintains an accurate count of the allocated and free buffers. |
| C3 | CAM Access | Verifies the content-addressable memory (CAM) block by performing read, write, add, modify, and cache hit tests on the CAM associative memory. |
| C4 | TPAT Cracker | Verifies the packet cracking logic block (i.e., the ability to parse TCP, IP, and UDP headers within an Ethernet frame) as well as the checksum/CRC offload logic. In this test, packets are submitted to the chip as if they were received over Ethernet and the TPAT block cracks the frame (identifying the TCP, IP, and UDP header data structures) and calculates the checksum/CRC. The TPAT block results are compared with the values expected by Broadcom NetXtreme II User Diagnostics and any errors are displayed. |
| C5 | FIO Register | The Fast IO (FIO) verifies the register interface that is exposed to the internal CPUs. |
| C6 | NVM Access and Reset-Corruption | Verifies non-volatile memory (NVM) accesses (both read and write) initiated by one of the internal CPUs. It tests for appropriate access arbitration among multiple entities (CPUs). It also checks for possible NVM corruption by issuing a chip reset while the NVM block is servicing data. |
| C7 | Core-Reset Integrity | Verifies that the chip performs its reset operation correctly by resetting the chip multiple times, checking that the bootcode and the internal uxdiag driver loads/unloads correctly. |
| C8 | DMA Engine | Verifies the functionality of the DMA engine block by performing numerous DMA read and write operations to various system and internal memory locations (and byte boundaries) with varying lengths (from 1 byte to over 4 KB, crossing the physical page boundary) and different data patterns (incremental, fixed, and random). CRC checks are performed to ensure data integrity. The DMA write test also verifies that DMA writes do not corrupt the neighboring host memory. |
| C9 | VPD | Exercises the Vital Product Data (VPD) interface using PCI configuration cycles and requires a proper bootcode to be programmed into the non-volatile memory. If no VPD data is present (i.e., the VPD NVM area is all 0s), the test first initializes the VPD data area with non-zero data before starting the test and restores the original data after the test completes. |
| C11 | FIO Events | Verifies that the event bits in the CPU's Fast IO (FIO) interface are triggering correctly when a particular chip event occurs, such as a VPD request initiated by the host, an expansion ROM request initiated by the host, a timer event generated internally, toggling any GPIO bits, or accessing NVM. |

| **Group D: Ethernet Traffic Tests** | | |
|---|---|---|
| D1 | MAC Loopback | Enables MAC loopback mode in the adapter and transmits 5000 Layer 2 packets of various sizes. As the packets are received back by Broadcom NetXtreme II User Diagnostics, they are checked for errors. Packets are returned through the MAC receive path and never reach the PHY. The adapter should not be connected to a network. |
| D2 | PHY Loopback | Enables PHY loopback mode in the adapter and transmits 5000 Layer 2 packets of various sizes. As the packets are received back by Broadcom NetXtreme II User Diagnostics, they are checked for errors. Packets are returned through the PHY receive path and never reach the wire. The adapter should not be connected to a network. |
| D4 | LSO | Verifies the functionality of the adapter's Large Send Offload (LSO) support by enabling MAC loopback mode and transmitting large TCP packets. As the packets are received back by Broadcom NetXtreme II User Diagnostics, they are checked for proper segmentation (according to the selected MSS size) and any other errors. The adapter should not be connected to a network. |
| D5 | EMAC Statistics | Verifies that the basic statistics information maintained by the chip is correct by enabling MAC loopback mode and sending Layer 2 packets of various sizes. The adapter should not be connected to a network. |
| D6 | RPC | Verifies the Receive Path Catch-up (RPC) block by sending packets to different transmit chains. The packets traverse the RPC logic (though not the entire MAC block) and return to the receive buffers as received packets. This is another loopback path that is used by Layer 4 and Layer 5 traffic within the MAC block. As packets are received back by Broadcom NetXtreme II User Diagnostics, they are checked for errors. The adapter should not be connected to a network. |

# 8. Specifications: D-Link DXE-820T Network Adapter User Guide

## 8.1. 10/100/1000BASE-T and 10GBASE-T Cable Specifications

**Table 8: 10/100/1000BASE-T Cable Specifications**

| Port Type | Connector | Media | Maximum Distance |
|---|---|---|---|
| 10BASE-T | RJ-45 | Category 3, 4, or 5 unshielded twisted pairs (UTP) | 100m (328 feet) |
| 100/1000BASE-T[1] | RJ-45 | Category 5[2] UTP | 100m (328 feet) |

[1] 1000BASE-T signaling requires four twisted pairs of Category 5 balanced cabling, as specified in ISO/IEC 11801:2002 and ANSI/EIA/TIA-568-B.

[2] Category 5 is the minimum requirement. Category 5e and Category 6 are fully supported.

**Table 9: 10GBASE-T Cable Specifications**

| Port Type | Connector | Media | Maximum Distance |
|---|---|---|---|
| 10GBASE-T | RJ-45 | Category 6[1] UTP<br>Category 6A[1] UTP | 50m (164 feet)<br>100m (328 feet) |

[1] 10GBASE-T signaling requires four twisted pairs of Category 6 or Category 6A (augmented Category 6) balanced cabling, as specified in ISO/IEC 11801:2002 and ANSI/TIA/EIA-568-B.

## 8.2. Interface Specifications

**Table 10: 10GBASE-T Performance Specifications**

| Feature | Specification |
|---|---|
| PCI Express Interface | x8 link width |
| 10GBASE-T | 10 Gbps |

# 9. Troubleshooting: D-Link DXE-820T Network Adapter User Guide

## 9.1. Hardware Diagnostics

Loopback diagnostic tests are available for testing the adapter hardware. These tests provide access to the adapter internal/external diagnostics, where packet information is transmitted across the physical link (for instructions and information on running tests in an MS-DOS environment, see User Diagnostics; for Windows environments, see Running Diagnostic Tests in Windows).

## 9.2. Checking Port LEDs

See Network Link and Activity Indication to check the state of the network link and activity.

## 9.3. Troubleshooting Checklist

CAUTION! Before you open the cabinet of your server to add or remove the adapter, review Safety Precautions.

The following checklist provides recommended actions to take to resolve problems installing the D-Link DXE-820T network adapter or running it in your system.

- Inspect all cables and connections. Verify that the cable connections at the network adapter and the switch are attached properly. Verify that the cable length and rating comply with the requirements listed in Connecting the Network Cables.
- Check the adapter installation by reviewing Installation of the Add-In NIC. Verify that the adapter is properly seated in the slot. Check for specific hardware problems, such as obvious damage to board components or the PCI edge connector.
- Check the configuration settings and change them if they are in conflict with another device.
- Verify that your server is using the latest BIOS.
- Try inserting the adapter in another slot. If the new position works, the original slot in your system may be defective.
- Replace the failed adapter with one that is known to work properly. If the second adapter works in the slot where the first one failed, the original adapter is probably defective.
- Install the adapter in another functioning system and run the tests again. If the adapter passed the tests in the new system, the original system may be defective.
- Remove all other adapters from the system and run the tests again. If the adapter

passes the tests, the other adapters may be causing contention.

## 9.4. Checking if Current Drivers are Loaded

### 9.4.1. Windows

See Viewing Vital Signs to view vital information about the adapter, link status, and network connectivity.

### 9.4.2. Linux

To verify that the bnx2.o driver is loaded properly, run:

```
lsmod | grep -i <module name>
```

If the driver is loaded, the output of this command shows the size of the driver in bytes and the number of adapters configured and their names. The following example shows the drivers loaded for the bnx2 module:

```
[root@test1]# lsmod | grep -i bnx2
bnx2                    199238    0
bnx2fc                  133775    0
libfcoe                  39764    2 bnx2fc,fcoe
libfc                   108727    3 bnx2fc,fcoe,libfcoe
scsi_transport_fc        55235    3 bnx2fc,fcoe,libfc
bnx2i                    53488    11
cnic                     86401    6 bnx2fc,bnx2i
libiscsi                 47617    8
be2iscsi,bnx2i,cxgb4i,cxgb3i,libcxgbi,ib_iser,iscsi_tcp,libiscsi_tcp
scsi_transport_iscsi     53047    8 be2iscsi,bnx2i,libcxgbi,ib_iser,iscsi_tcp,libiscsi
bnx2x                  1417947    0
libcrc32c                 1246    1 bnx2x
mdio                      4732    2 cxgb3,bnx2x
```

If you reboot after loading a new driver, you can use the following command to verify that the currently loaded driver is the correct version.

```
modinfo bnx2
[root@test1]# lsmod | grep -i bnx2
bnx2 199238 0
```

Or, you can use the following command:

```
[root@test1]# ethtool -i eth2
driver: bnx2x
version: 1.78.07
firmware-version: bc 7.8.6
```

bus-info: 0000:04:00.2

if you loaded a new driver but have not yet booted, the modinfo command will not show the updated driver information. Instead, you can view the logs to verify that the proper driver is loaded and will be active upon reboot:

dmesg | grep -i "Broadcom" | grep -i "bnx2"

## 9.5. Running a Cable Length Test

For Windows operating systems, see Analyzing Cables in Windows for information on running a cable length test. Cable analysis is not available for DXE-820T network adapters.

## 9.6. Testing Network Connectivity

**Note:** *When using forced link speeds, verify that both the adapter and the switch are forced to the same speed.*

### 9.6.1. Windows

Network connectivity can be tested using the Testing the Network feature in Broadcom Advanced Control Suite.

An alternate method is to use the ping command to determine if the network connection is working.

1. Click Start, and then click Run.
2. Type cmd in the Open box, and then click OK.
3. Type ipconfig /all to view the network connection to be tested.
4. Type ping IP address, and then press ENTER.

The ping statistics that are displayed indicate whether the network connection is working or not.

### 9.6.2. Linux

To verify that the Ethernet interface is up and running, run ifconfig to check the status of the Ethernet interface. It is possible to use netstat -i to check the statistics on the Ethernet interface. See Linux Driver Software for information on ifconfig and netstat.

Ping an IP host on the network to verify connection has been established.

From the command line, type ping IP address, and then press ENTER.

The ping statistics that are displayed indicate whether or not the network connection is working.

## 9.7. Microsoft Virtualization with Hyper-V

Microsoft Virtualization is a hypervisor virtualization system for Windows Server 2008 and Windows Server 2008 R2. This section is intended for those who are familiar with Hyper-V, and it addresses issues that affect the configuration of DXE-820T network adapters and teamed network adapters when Hyper-V is used. For more information on Hyper-V, see http://www.microsoft.com/windowsserver2008/en/us/hyperv.aspx.

Table 11 identifies Hyper-V supported features that are configurable for DXE-820T network adapters. This table is not an all-inclusive list of Hyper-V features.

**Table 11: Configurable Network Adapter Hyper-V Features**

| Feature | Supported in Windows Server | | | Comments/Limitation |
|---|---|---|---|---|
| | 2008 | 2008 R2 | 2012 | |
| IPv4 | Yes | Yes | Yes | – |
| IPv6 | Yes | Yes | Yes | – |
| IPv4 Large Send Offload (LSO) (parent and child partition) | Yes | Yes | Yes | – |
| IPv4 Checksum Offload (CO) (parent and child partition) | Yes | Yes | Yes | – |
| IPv6 LSO (parent and child partition) | No* | Yes | Yes | *When bound to a virtual network, OS limitation. |
| IPv6 CO (parent and child partition) | No* | Yes | Yes | *When bound to a virtual network, OS limitation. |
| Jumbo frames | No* | Yes | Yes | *OS limitation. |
| RSS | No* | No* | Yes | *OS limitation. |
| RSC | No* | No* | Yes | *OS limitation. |
| SRIOV | No* | No* | Yes | *OS limitation. |

**Note**: *Ensure that Integrated Services, which is a component of Hyper-V, is installed in the guest operating system (child partition) for full functionality.*

### 9.7.1. Single Network Adapter

**Windows Server 2008**

When configuring a DXE-820T network adapter on a Hyper-V system, be aware of the following:

- An adapter that is to be bound to a virtual network should not be configured for VLAN tagging through the driver's advanced properties. Instead, Hyper-V should manage VLAN tagging exclusively.

- Since Hyper-V does not support Jumbo Frames, it is recommended that this feature not be used or connectivity issues may occur with the child partition.
- The Locally Administered Address (LAA) set by Hyper-V takes precedence over the address set in the adapter's advanced properties.
- In an IPv6 network, a team that supports CO and/or LSO and is bound to a Hyper-V virtual network will report CO and LSO as an offload capability in BACS; however, CO and LSO will not work. This is a limitation of Hyper-V. Hyper-V does not support CO and LSO in an IPv6 network.

**Windows Server 2008 R2 and 2012**

When configuring a DXE-820T network adapter on a Hyper-V system, be aware of the following:

- An adapter that is to be bound to a virtual network should not be configured for VLAN tagging through the driver's advanced properties. Instead, Hyper-V should manage VLAN tagging exclusively.
- The Locally Administered Address (LAA) set by Hyper-V takes precedence over the address set in the adapter's advanced properties.
- The LSO and CO features in the guest OS are independent of the network adapter properties.
- To allow jumbo frame functionality from the guest OS, both the network adapter and the virtual adapter must have jumbo frames enabled. The Jumbo MTU property for the network adapter must be set to allow traffic of large MTU from within the guest OS. The jumbo packet of the virtual adapter must be set in order to segment the sent and received packets.

### 9.7.2. Teamed Network Adapters

Table 12 identifies Hyper-V supported features that are configurable for DXE-820T teamed network adapters. This table is not an all-inclusive list of Hyper-V features.

**Table 12: Configurable Teamed Network Adapter Hyper-V Features**

| Feature | Supported in Windows Server Version | | | Comments/Limitation |
|---|---|---|---|---|
| | 2008 | 2008 R2 | 2012 | |
| Smart Load Balancing and Failover (SLB) team type | Yes | Yes | Yes | Multi-member SLB team allowed with latest BASP6 version.<br>**Note**: *VM MAC is not presented to external switches.* |
| Link Aggregation (IEEE 802.3ad LACP) team type | Yes | Yes | Yes | – |

| | | | | |
|---|---|---|---|---|
| Generic Trunking (FEC/GEC) 802.3ad Draft Static team type | Yes | Yes | Yes | – |
| Failover | Yes | Yes | Yes | – |
| LiveLink | Yes | Yes | Yes | – |
| Large Send Offload (LSO) | Limited* | Yes | Yes | *Conforms to miniport limitations outlines in Table 1. |
| Checksum Offload (CO) | Limited* | Yes | Yes | *Conforms to miniport limitations outlines in Table 1. |
| Hyper-V VLAN over an adapter | Yes | Yes | Yes | – |
| Hyper-V VLAN over a teamed adapter | Yes | Yes | Yes | – |
| Hyper-V VLAN over a VLAN | Limited* | Limited* | Limited* | Only an untagged VLAN. |
| Hyper-V virtual switch over an adapter | Yes | Yes | Yes | – |
| Hyper-V virtual switch over a teamed adapter | Yes | Yes | Yes | – |
| Hyper-V virtual switch over a VLAN | Yes | Yes | Yes | – |
| iSCSI boot | No | No* | No* | *Remote boot to SAN is supported. |
| Virtual Machine Queue (VMQ) | No | Yes | Yes | See Configuring VMQ with SLB Teaming. |
| RSC | No | No | Yes | |

### Windows Server 2008

When configuring a team of DXE-820T network adapters on a Hyper-V system, be aware of the following:

- Create the team prior to binding the team to the Hyper-V virtual network.
- Create a team only with an adapter that is not already assigned to a Hyper-V virtual network.
- In an IPv6 network, a team that supports CO and/or LSO and is bound to a Hyper-V virtual network will report CO and LSO as an offload capability in BACS; however, CO and LSO will not work. This is a limitation of Hyper-V. Hyper-V does not support CO and LSO in an IPv6 network.
- To successfully perform VLAN tagging for both the host (parent partition) and the guest (child partition) with the BASP teaming software, you must configure the team for tagging. Unlike VLAN tagging with a single adapter, tagging cannot be managed by Hyper-V when using BASP software.
- When making changes to a team or removing a team, remove the team's binding

from all guest OSs that use any of the VNICs in the team, change the configuration, and then rebind the team's VNICs to the guest OS. This can be done in the Hyper-V Manager.

**Windows Server 2008 R2**

When configuring a team of DXE-820T network adapters on a Hyper-V system, be aware of the following:

- Create the team prior to binding the team to the Hyper-V virtual network.
- Create a team only with an adapter that is not already assigned to a Hyper-V virtual network.
- A BASP virtual adapter configured for VLAN tagging can be bound to a Hyper-V virtual network, and is a supported configuration. However, the VLAN tagging capability of BASP cannot be combined with the VLAN capability of Hyper-V. In order to use the VLAN capability of Hyper-V, the BASP team must be untagged.
- When making changes to a team or removing a team, remove the team's binding from all guest OSs that use any of the VNICs in the team, change the configuration, and then rebind the team's VNICs to the guest OS. This can be done in the Hyper-V Manager.

**Configuring VMQ with SLB Teaming**

When Hyper-V server is installed on a system configured to use Smart Load Balance and Failover (SLB) type teaming, you can enable Virtual Machine Queuing (VMQ) to improve overall network performance. VMQ enables delivering packets from an external virtual network directly to virtual machines defined in the SLB team, eliminating the need to route these packets and, thereby, reducing overhead.

To create a VMQ-capable SLB team:

1. Create an SLB team. If using the Teaming Wizard, when you select the SLB team type, also select Enable HyperV Mode. If using Expert mode, enable the property in the Create Team or Edit Team tabs. See Configuring Teaming for additional instructions on creating a team.
2. Follow these instructions to add the required registry entries in Windows: http://technet.microsoft.com/en-us/library/gg162696%28v=ws.10%29.aspx
3. For each team member on which you want to enable VMQ, modify the following registry entry and configure a unique instance number (in the following example, it is set to 0026):

```
[HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Class\
        {4D36E972-E325-11CE-BFC1-08002BE10318}\0026]
"*RssOrVmqPreference"="1"
```

## 9.8. Removing the Device Drivers

Uninstall the device drivers from your system only through the InstallShield wizard. Uninstalling the device drivers with Device Manager or any other means may not provide a clean uninstall and may cause the system to become unstable. For information on uninstalling device drivers, see Removing the Device Drivers.

## 9.9. Upgrading Windows Operating Systems

This section covers Windows upgrades for the following:

- From Windows Server 2003 to Windows Server 2008
- From Windows Server 2008 to Windows Server 2008 R2
- From Windows Server 2008 R2 to Windows Server 2012

Prior to performing an OS upgrade when a D-Link DXE-820T network adapter is installed on your system, Broadcom recommends the procedure below.

1. Save all team and adapter IP information.
2. Uninstall all Broadcom drivers using the installer.
3. Perform the Windows upgrade.
4. Reinstall the latest drivers and the BACS application.

## 9.10. Broadcom Boot Agent

- Problem: Unable to obtain network settings through DHCP using PXE.

  Solution: For proper operation make sure that the Spanning Tree Protocol (STP) is disabled or that portfast mode (for Cisco) is enabled on the port to which the PXE client is connected. For instance, set spantree portfast 4/12 enable.

## 9.11. Broadcom Advanced Server Program (BASP)

- Problem: After physically removing a NIC that was part of a team and then rebooting, the team did not perform as expected.

  Solution: To physically remove a teamed NIC from a system, you must first delete the NIC from the team. Not doing this before shutting down could result in breaking the team on a subsequent reboot, which may result in unexpected team behavior.
- Problem: After deleting a team that uses IPv6 addresses and then re-creating the team, the IPv6 addresses from the old team are used for the re-created team.

  Solution: This is a third-party issue. To remove the old team's IPv6 addresses, locate the General tab for the team's TCP/IP properties from your system's Network Connections. Either delete the old addresses and type in new IPv6 addresses or select the option to automatically obtain IP addresses.
- Problem: Adding an NLB-enabled DXE-820T adapter to a team may cause unpredictable results.

Solution: Prior to creating the team, unbind NLB from the DXE-820T adapter, create the team, and then bind NLB to the team.

- Problem: A system containing an 802.3ad team causes a Net logon service failure in the system event log and prevents it from communicating with the domain controller during boot up.

  Solution: Microsoft Knowledge Base Article 326152 (http://support.microsoft.com/kb/326152/en-us) indicates that Gigabit Ethernet adapters may experience problems with connectivity to a domain controller due to link fluctuation while the driver initializes and negotiates link with the network infrastructure. The link negotiation is further affected when the Gigabit adapters are participating in an 802.3ad team due to the additional negotiation with a switch required for this team type. As suggested in the Knowledge Base Article above, disabling media sense as described in a separate Knowledge Base Article 938449 (http://support.microsoft.com/kb/938449) has shown to be a valid workaround when this problem occurs.
- Problem: The 802.3ad team member links disconnect and reconnect continuously (applies to all operating systems).

  Solution: This is a third-party issue. It is seen only when configuring an 802.3ad team with greater than two members on the server and connecting an HP2524 switch, with LACP enabled as passive or active. The HP switch shows an LACP channel being brought up successfully with only two team members. All other team member links disconnect and reconnect. This does not occur with a Cisco Catalyst 6500.
- Problem: A Generic Trunking (GEC/FEC) 802.3ad-Draft Static type of team may lose some network connectivity if the driver to a team member is disabled.

  Solution: If a team member supports underlying management software (ASF/IPMI/UMP) or Wake-On-LAN, the link may be maintained on the switch for the adapter despite its driver being disabled. This may result in the switch continuing to pass traffic to the attached port rather than route the traffic to an active team member port. Disconnecting the disabled adapter from the switch will allow traffic to resume to the other active team members.
- Problem: Large Send Offload (LSO) and Checksum Offload are not working on my team.

  Solution: If one of the adapters on a team does not support LSO, LSO does not function for the team. Remove the adapter that does not support LSO from the team, or replace it with one that does. The same applies to Checksum Offload.
- Problem: The advanced properties of a team do not change after changing the advanced properties of an adapter that is a member of the team.

Solution: If an adapter is included as a member of a team and you change any advanced property, then you must rebuild the team to ensure that the team's advanced properties are properly set.

## 9.12. Linux

- Problem: BCM5771x devices with SFP+ Flow Control default to Off rather than Rx/Tx Enable.

  Solution: The Flow Control default setting for revision 1.6.x and newer has been changed to Rx Off and Tx Off because SFP+ devices do not support Auto negotiation for Flow Control.
- Problem: On kernels older than 2.6.16 when 16 partitions are created on a server containing two BCM57711 network adapters, not all partitions would come up and an error indicating a shortage of space would display.

  Solution: On architectures where the default vmalloc size is relatively small and not sufficient to load many interfaces, use vmalloc=<size> during boot to increase the size.
- Problem: Routing does not work for DXE-820T network adapters installed in Linux systems.

  Solution: For DXE-820T network adapters installed in systems with Linux kernels older than 2.6.26, disable TPA with either ethtool (if available) or with the driver parameter (see disable_tpa). Use ethtool to disable TPA (LRO) for a specific DXE-820T network adapter.
- Problem: On a DXE-820T network adapter in a CNIC environment, flow control does not work.

  Solution: Flow control is working, but in a CNIC environment, it has the appearance that it is not. The network adapter is capable of sending pause frames when the on-chip buffers are depleted, but the adapter also prevents the head-of-line blocking of other receive queues. Since the head-of-line blocking causes the on-chip firmware to discard packets inside the on-chip receive buffers, in the case a particular host queue is depleted, the on-chip receive buffers will rarely be depleted, therefore, it may appear that flow control is not functioning.
- Problem: Errors appear when compiling driver source code.

  Solution: Some installations of Linux distributions do not install the development tools by default. Ensure the development tools for the Linux distribution you are using are installed before compiling driver source code.

## 9.13. NPAR

- Problem: The following error message displays if the storage configurations are not consistent for all four ports of the device in NPAR mode:

PXE-M1234: NPAR block contains invalid configuration during boot.

A software defect can cause the system to be unable to BFS boot to an iSCSI or FCoE target if an iSCSI personality is enabled on the first partition of one port, whereas an FCoE personality is enabled on the first partition of another port. The MBA driver performs a check for this configuration and prompts the user when it is found.

Solution: If using the 7.6.x firmware and driver, to work around this error, configure the NPAR block such that if iSCSI or FCoE is enabled on the first partition, the same must be enabled on all partitions of all four ports of that device.

## 9.14. Miscellaneous

- Problem: When setting the Jumbo MTU property to 5000 bytes or greater and forcing Flow Control on network adapters that support a link speed of 10 Gbps, the system performance performs at less than optimal levels.

  Solution: If Jumbo MTU is set to 5000 bytes or greater, ensure that Flow Control is set to Auto.
- Problem: iSCSI Crash Dump is not working in Windows.

  Solution: After upgrading the device drivers using the installer, the iSCSI crash dump driver is also upgraded, and iSCSI Crash Dump must be re-enabled from the Advanced section of the BACS Configuration tab.
- Problem: In Windows Server 2008 R2, if the OS is running as an iSCSI boot OS, the VolMgr error, "The system could not successfully load the crash dump driver," appears in the event log.

  Solution: Enable iSCSI Crash Dump from the Advanced section of the BACS Configuration tab.
- Problem: The D-Link DXE-820T network adapter may not perform at optimal levels on some systems if it is added after the system has booted.

  Solution: The system BIOS in some systems does not set the cache line size and the latency timer if the adapter is added after the system has booted. Reboot the system after the adapter has been added.
- Problem: Intelligent Platform Management Interface (IPMI) is not functioning properly.

  Solution: IPMI works only when LiveLink™ is disabled. See Configuring LiveLink for a Smart Load Balancing and Failover and SLB (Auto-Fallback Disable) Team.
- Problem: Cannot configure Resource Reservations in BACS after SNP is uninstalled.

  Solution: Reinstall SNP. Prior to uninstalling SNP from the system, ensure that NDIS is enabled via the checkbox on the Resource Configuration screen, available from the Resource Reservations section of the Configurations tab (see Viewing Resource Reservations). If NDIS is disabled and SNP is removed, there is no access to re-enable the device.

- Problem: A DCOM error message (event ID 10016) appears in the System Even Log during the installation of the drivers.

  Solution: This is a Microsoft issue. For more information, see Microsoft knowledge base KB913119 at http://support.microsoft.com/kb/913119.
- Problem: Performance is degraded when multiple BCM57710 network adapters are used in a system.

  Solution: Ensure that the system has at least 2 GB of main memory when using up to four network adapters and 4 GB of main memory when using four or more network adapters.
- Problem: Remote installation of Windows Server 2008 to an iSCSI target via iSCSI offload fails to complete, and the computer restarts, repeatedly.

  Solution: This is a Microsoft issue. For more information on applying the Microsoft hotfix, see Microsoft knowledge base article KB952942 at http://support.microsoft.com/kb/952942.
- Problem: Performance drops when a BCM5709C network adapter is connected back-to-back to a switch, MTU = 9000, and Tx and Rx Flow Control are enabled.

  Solution: When enable_cu_rate_limiter is enabled, the device performs flow control in the catchup path to prevent catchup frames from dropping. The catchup path is used in processing iSCSI out-of-order PDUs. When enable_cu_rate_limiter is disabled, there is a potential for some drops of iSCSI out-of-order PDUs, which reduces performance. This feature does not work well when jumbo frame is enabled on any of the client devices. Enable_cu_rate_limiter should be set to disabled when jumbo frame is enabled.
- Problem: When using a BCM57840 4-port adapter in a blade server, ports 3 and 4 show no link.

  Solution: The I/O (switch) module must support 32 internal ports. If it does not, ports 3 and 4 cannot establish a link.