
Broadcom NetXtreme II™ Network Adapter User Guide

- Introduction
 - [Functionality and Features](#)
 - [Teaming](#)
 - [Virtual LANs \(VLANs\)](#)
 - [Manageability](#)
- [Installing the Hardware](#)
- [Installing the Driver Software](#)
 - [Broadcom Boot Agent Driver Software](#)
 - [NDIS2 Driver Software](#)
 - [ODI Driver Software](#)
 - [Linux Driver Software](#)
 - [NetWare Driver Software](#)
 - [Windows Driver Software](#)
- [Installing Management Applications](#)
- [Using iSCSI](#)
- [Configuring Teaming](#)
- [Using Broadcom Advanced Control Suite 3](#)
- [User Diagnostics](#)
- [Specifications](#)
- [Regulatory Information](#)
- [Troubleshooting](#)

Information in this document is subject to change without notice.

© 2010 Broadcom Corporation. All rights reserved.

This document is protected by copyright and is distributed under licenses restricting its use, copying, distribution, and decompilation. No part of this document may be reproduced in any form by any means without prior written authorization of Broadcom Corporation. 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. NetWare is a trademark of Novell Corporation. 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.

Initial release: December 2005

Last revised: January 2010

ENGSRV52-CDUM100-R



Functionality and Features: Broadcom NetXtreme II™ Network Adapter User Guide

- [Functional Description](#)
- [Features](#)

FUNCTIONAL DESCRIPTION

The Broadcom NetXtreme II adapter is a new class of Gigabit Ethernet (GbE) and 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 all popular protocols used in the data center, such as:

- TCP Offload Engine (TOE) for accelerating TCP over 1 GbE, 2.5 GbE, and 10 GbE
- Internet Small Computer Systems Interface (iSCSI) offload for accelerating network storage access featuring centralized boot functionality (iSCSI boot)



NOTE: Separate licences are required for all offloading technologies.

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.

The Broadcom NetXtreme II adapter includes a 10/100/1000-Mbps or 10-Gbps Ethernet MAC with both half-duplex and full-duplex capability and a 10/100/1000-Mbps or 10-Gbps PHY. The transceiver is fully compatible with the IEEE 802.3 standard for auto-negotiation of speed.

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.

FEATURES

The following is a list of the Broadcom NetXtreme II adapter features:

- TCP Offload Engine (TOE)
- Internet Small Computer Systems Interface (iSCSI) offload
- Single-chip solution
 - Integrated 10/100/1000BASE-T transceivers
 - Integrated 10GBASE-T transceivers
 - 10/100/1000 triple-speed MAC
 - Host interfaces
 - SerDes interface for optical transceiver connection
 - PCI Express 1.0a x4 (Gigabit Ethernet)
 - PCI Express Gen2 x8 (10 Gigabit Ethernet)
 - Full fast-path TCP offload
 - Zero copy capable hardware
- Other performance features
 - TCP, IP, UDP checksum
 - TCP segmentation
 - Adaptive interrupts
 - Receive Side Scaling (RSS)
- Manageability
 - Broadcom Advanced Control Suite 3 diagnostic and configuration software suite
 - Supports PXE 2.0 specification (Linux Red Hat PXE Server, SUSE Linux Enterprise Server, Windows 2000 Server, Windows Server 2003, Windows Server 2008, Windows Server 2008 R2, Intel APITEST, DOS UNDI)
 - Wake on LAN support
 - Universal Management Port (UMP) support
 - Statistics for SNMP MIB II, Ethernet-like MIB, and Ethernet MIB (IEEE Std 802.3z, Clause 30)
 - SMBus controller
 - ACPI 1.1a compliant (multiple power modes)
 - 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
 - Smart Load Balancing TOE Teaming (with the correct configuration)
 - 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 4 classes of service (CoS)
- Up to 4 send rings and receive rings

- Integrated 96 KB frame buffer memory
- GMII/MII Management Interface
- Four unique MAC unicast addresses
- Support for multicast addresses via 128 bits hashing hardware function
- Serial flash NVRAM memory
- JTAG support
- PCI Power Management Interface (v1.1)
- 64-bit BAR support
- EM64T processor support
- AMD-64 processor support
- 1.2 V core voltage, 0.13 μ m process
- iSCSI Boot support

TCP OFFLOAD ENGINE (TOE)

The TCP/IP protocol suite is used to provide transport services for a wide range of applications for the Internet, LAN, and for file transfer. Without the TCP Offload Engine, the TCP/IP protocol suite runs on the host CPU, consuming a very high percentage of its resources and leaving little resources for the applications. With the use of the Broadcom NetXtreme II adapter, the TCP/IP processing can be moved to hardware, freeing the CPU for more important tasks such as application processing.

The Broadcom NetXtreme II adapter's TOE functionality allows simultaneous operation of up to 1024 fully offloaded TCP connections for 1-Gbps network adapters and 1880 fully offloaded TCP connections for 10-Gbps network adapters. The TOE support on the adapter significantly reduces the host CPU utilization while preserving the implementation of the operating system stack.

INTERNET SMALL COMPUTER SYSTEMS INTERFACE (iSCSI)

The IETF has standardized the Internet Small Computer Systems Interface (iSCSI). SCSI is a popular protocol that enables systems to communicate with storage devices, using block-level transfer (i.e., address data stored on a storage device that is not a whole file). iSCSI maps the SCSI request/response application protocols and its standardized command set over TCP/IP networks.

As iSCSI utilizes TCP as its sole transport protocol, it greatly benefits from hardware acceleration of the TCP processing (i.e., use of a TOE). However, iSCSI as a Layer 5 protocol has additional mechanisms beyond the TCP layer. iSCSI processing can also be offloaded, thereby reducing CPU utilization even further.

The Broadcom NetXtreme II adapter targets best-system performance, maintains system flexibility to changes, and supports current and future OS convergence and integration. Therefore, the adapter's iSCSI offload architecture is unique as evident by the split between hardware and host processing.



NOTES: The iSCSI offload feature is not available for all Broadcom network adapters.

POWER MANAGEMENT

Adapter speed connection when the system is down waiting for a wake-up signal may be at 10 Mbps or 100 Mbps, but can return to 1000 Mbit/s when the system is up and running if connected to a 1000 Mbps capable switch. Systems intending to use Wake on LAN (WOL) should be connected to a switch capable of both 1000 and 10/100 Mbps speeds.



NOTES:

- For specific systems, see your system documentation for WOL support.
- WOL is supported in Broadcom NetXtreme II BCM5708 devices with silicon revisions of B2 or later. For more information, see [Limitations](#).

ADAPTIVE INTERRUPT FREQUENCY

The adapter driver intelligently adjusts host interrupt frequency based on traffic conditions to increase overall application throughput. When traffic is light, the adapter driver interrupts the host for each received packet, minimizing latency. When traffic is heavy, the adapter issues one host interrupt for multiple, back-to-back incoming packets, preserving host CPU cycles.

ASIC WITH EMBEDDED RISC PROCESSOR

The core control for Broadcom NetXtreme II adapters resides in a tightly integrated, high-performance ASIC. The ASIC includes a RISC processor. This functionality provides the flexibility to add new features to the card and adapts it to future network requirements through software downloads. This functionality also enables the adapter drivers to exploit the built-in host offload functions on the adapter as host operating systems are enhanced to take advantage of these functions.

BROADCOM ADVANCED CONTROL SUITE 3

Broadcom Advanced Control Suite 3 (BACS), a component of the Broadcom teaming software, is an integrated utility that provides useful information about each network adapter that is installed in your system. The BACS 3 utility also enables you to perform detailed tests, diagnostics, and analyses on each adapter, as well as to modify property values and view traffic statistics for each adapter.

Microsoft .NET Framework 2.0 includes the runtime and associated files needed to run BACS 3, and must be installed on your system in order for BACS 3 to operate. For optimal performance of BACS 3, Broadcom recommends .NET Framework 2.0 SP1, .NET Framework 3.0 SP1, or .NET Framework 3.5, depending on your operating system.

SUPPORTED OPERATING ENVIRONMENTS

The Broadcom NetXtreme II adapter has software support for the following operating systems:

- Microsoft® Windows® (32-bit and 64-bit extended)
- Microsoft Windows Vista™ (32-bit and 64-bit extended)
- Linux® (32-bit and 64-bit extended)
- MS-DOS®
- ESX Server (VMware)
- NetWare®
- SCO® UnixWare®
- SCO OpenServer®

NETWORK LINK AND ACTIVITY INDICATION

For copper-wire Ethernet connections, the state of the network link and activity is indicated by the LEDs on the RJ-45 connector, as described in [Table 1](#). For fiber optic Ethernet connections, the state of the network link and activity is indicated by a single LED located adjacent to the port connector, as described in [Table 2](#). Broadcom Advanced Control Suite 3 also provides information about the status of the network link and activity (see [Viewing Vital Signs](#)).

Table 1: Network Link and Activity Indicated by the RJ-45 Port LEDs

Port LED	LED Appearance	Network State
Link LED	Off	No link (cable disconnected)
	Continuously illuminated	Link
Activity LED	Off	No network activity
	Blinking	Network activity

Table 2: Network Link and Activity Indicated by the Port LED

LED Appearance	Network State
Off	No link (cable disconnected)
Continuously illuminated	Link
Blinking	Network activity



Configuring Teaming: Broadcom NetXtreme II™ Network Adapter User Guide

- [Broadcom Advanced Server Program Overview](#)
- [Load Balancing and Fault Tolerance](#)

BROADCOM ADVANCED SERVER PROGRAM OVERVIEW

Broadcom Advanced Server Program (BASP) is the Broadcom teaming software for the Windows family of operating systems. BASP runs within the Broadcom Advanced Control Suite 3 (BACS) utility.

BASP provides support for TOE teaming only for NetXtreme II adapters. BASP supports four types of teams for Layer 2 teaming:

- Smart Load Balancing and Failover
- Link Aggregation (802.3ad)
- Generic Trunking (FEC/GEC)/802.3ad-Draft Static
- SLB (Auto-Fallback Disable)

BASP supports two types of teams for TOE teaming:

- Smart Load Balancing and Failover
- SLB (Auto-Fallback Disable)



NOTE: Enabling Windows Server 2003 built-in bridging is not advisable when you are using teaming software.

For more information on network adapter teaming concepts, see [Broadcom Gigabit Ethernet Teaming Services](#).

LOAD BALANCING AND FAULT TOLERANCE

Teaming provides traffic load balancing and fault tolerance (redundant adapter operation in the event that a network connection fails). When multiple Gigabit Ethernet network adapters are installed in the same system, they can be grouped into teams, creating a virtual adapter.

A team can consist of two to eight network interfaces, and each interface can be designated as a primary interface or a standby interface (standby interfaces can be used only in a [Smart Load Balancing™ and Failover](#) type of team, and only one standby interface can be designated per SLB team). If traffic is not identified on any of the adapter team member connections due to failure of the adapter, cable, switch port, or switch (where the teamed adapters are attached to separate switches), the load distribution is reevaluated and reassigned among the remaining team members. In the event that all of the primary adapters are down, the hot standby adapter becomes active. Existing sessions are maintained and there is no impact on the user.



NOTE: Although a team can be created with one adapter, it is not recommended since this defeats the purpose of teaming. A team consisting of one adapter is automatically created when setting up VLANs on a single adapter, and this should be the only time when creating a team with one adapter.

TYPES OF TEAMS

The available types of teams for the Windows family of operating systems are:

- Smart Load Balancing and Failover
- Link Aggregation (802.3ad) (TOE is not applicable)
- Generic Trunking (FEC/GEC)/802.3ad-Draft Static (TOE is not applicable)
- SLB (Auto-Fallback Disable)

SMART LOAD BALANCING™ AND FAILOVER

Smart Load Balancing™ and Failover is the Broadcom implementation of load balancing based on IP flow. This feature supports balancing IP traffic across multiple adapters (team members) in a bidirectional manner. In this type of team, all adapters in the team have separate MAC addresses. This type of team provides automatic fault detection and dynamic failover to other team member or to a hot standby member. This is done independently of Layer 3 protocol (IP, IPX, NetBEUI); rather, it works with existing Layer 2 and 3 switches. No switch configuration (such as trunk, link aggregation) is necessary for this type of team to work.



NOTES:

- If you do not enable LiveLink™ when configuring SLB teams, disabling Spanning Tree Protocol (STP) or enabling Port Fast at the switch or port is recommended. This minimizes the downtime due to spanning tree loop determination when failing over. LiveLink mitigates such issues.
- TCP/IP is fully balanced and IPX balances only on the transmit side of the team; other protocols are limited to the primary adapter.
- If a team member is linked at a higher speed than another, most of the traffic is handled by the adapter with the higher speed rate.

LINK AGGREGATION (802.3AD)

This mode supports link aggregation and conforms to the IEEE 802.3ad (LACP) specification. Configuration software allows you to dynamically configure which adapters you want to participate in a given team. If the link partner is not correctly configured for 802.3ad link configuration, errors are detected and noted. With this mode, all adapters in the team are configured to receive packets for the same MAC address. The outbound load-balancing scheme is determined by our BASP driver. The team link partner determines the load-balancing scheme for inbound packets. In this mode, at least one of the link partners must be in active mode.



NOTE: Link Aggregation team type is not supported for TOE teaming.

GENERIC TRUNKING (FEC/GEC)/802.3AD-DRAFT STATIC

The Generic Trunking (FEC/GEC)/802.3ad-Draft Static type of team is very similar to the Link Aggregation (802.3ad) type of team in that all adapters in the team are configured to receive packets for the same MAC address. The Generic Trunking (FEC/GEC)/802.3ad-Draft Static type of team, however, does not provide LACP or marker protocol support. This type of team supports a variety of environments in which the adapter link partners are statically configured to support a proprietary trunking mechanism. For instance, this type of team could be used to support Lucent's OpenTrunk or Cisco's Fast EtherChannel (FEC). Basically, this type of team is a light version of the Link Aggregation (802.3ad) type of team. This approach is much simpler, in that there is not a formalized link aggregation control protocol (LACP). As with the other types of teams, the creation of teams and the allocation of physical adapters to various teams is done statically through user configuration software.

The Generic Trunking (FEC/GEC/802.3ad-Draft Static) type of team supports load balancing and failover for both outbound and inbound traffic.



NOTE: Generic Trunking (FEC/GEC/802.3ad-Draft Static) team type is not supported for TOE teaming.

SLB (AUTO-FALLBACK DISABLE)

The SLB (Auto-Fallback Disable) type of team is identical to the Smart Load Balancing 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.

All primary interfaces in a team participate in load-balancing operations by sending and receiving a portion of the total traffic. Standby interfaces take over in the event that all primary interfaces have lost their links.

Failover teaming provides redundant adapter operation (fault tolerance) in the event that a network connection fails. If the primary adapter in a team is disconnected because of failure of the adapter, cable, or switch port, the secondary team member becomes active, redirecting both inbound and outbound traffic originally assigned to the primary adapter. Sessions will be maintained, causing no impact to the user.

LIMITATIONS OF SMART LOAD BALANCING AND FAILOVER/SLB (AUTO-FALLBACK DISABLE) TYPES OF TEAMS

Smart Load Balancing™ (SLB) is a protocol-specific scheme. The level of support for IP, IPX, and NetBEUI protocols is listed in [Table 1](#).



NOTE: IPv6 is supported for addressing, but is not supported for load balancing.

Table 1: Smart Load Balancing

Operating System	Failover/Fallback — All Broadcom				Failover/Fallback — Multivendor			
Protocol	IP	IPv6	IPX	NetBEUI	IP	IPv6	IPX	NetBEUI
Windows 2000 Server	Y	N/S	Y	Y	Y	N/S	N	N
Windows Server 2003 SP2	Y	N/S	Y	N/S	Y	N/S	N	N/S
Windows Server 2008	Y	N/S	Y	N/S	Y	N/S	N	N/S
Windows Server 2008 R2	Y	N/S	Y	N/S	Y	N/S	N	N/S

Operating System	Load Balance — All Broadcom				Load Balance — Multivendor			
Protocol	IP	IPv6	IPX	NetBEUI	IP	IPv6	IPX	NetBEUI
Windows 2000 Server	Y	N/S	Y	N	Y	N/S	N	N
Windows Server 2003 SP2	Y	N/S	Y	N/S	Y	N/S	N	N/S
Windows Server 2008	Y	N/S	Y	N/S	Y	N/S	N	N/S
Windows Server 2008 R2	Y	N/S	Y	N/S	Y	N/S	N	N/S

Legend Y = yes
 N = no
 N/S = not supported

The Smart Load Balancing type of team works with all Ethernet switches without having to configure the switch ports to any special trunking mode. Only IP traffic is load-balanced in both inbound and outbound directions. IPX traffic is load-balanced in the outbound direction only. Other protocol packets are sent and received through one primary interface only. Failover for non-IP traffic is supported only for Broadcom network adapters. The Generic Trunking type of team requires the Ethernet switch to support some form of port trunking mode (for example, Cisco's Gigabit EtherChannel or other switch vendor's Link Aggregation mode). The Generic Trunking type of team is protocol-independent, and all traffic should be load-balanced and fault-tolerant.



NOTE: If you do not enable LiveLink™ when configuring teams, disabling Spanning Tree Protocol (STP) or enabling Port Fast at the switch is recommended. This minimizes the downtime due to the spanning tree loop determination when failing over. LiveLink mitigates such issues.

TEAMING AND LARGE SEND OFFLOAD/CHECKSUM OFFLOAD SUPPORT

Large Send Offload (LSO) and Checksum Offload are enabled for a team only when all of the members support and are configured for the feature.

Broadcom Gigabit Ethernet Teaming Services: Broadcom NetXtreme II™ Network Adapter User Guide

- [Executive Summary](#)
- [Teaming Mechanisms](#)
- [Teaming and Other Advanced Networking Properties](#)
- [General Network Considerations](#)
- [Application Considerations](#)
- [Troubleshooting Teaming Problems](#)
- [Frequently Asked Questions](#)
- [Appendix A: Event Log Messages](#)

EXECUTIVE SUMMARY

- [Glossary](#)
- [Teaming Concepts](#)
- [Software Components](#)
- [Hardware Requirements](#)
- [Teaming Support by Processor](#)
- [Configuring Teaming](#)
- [Supported Features by Team Type](#)
- [Selecting a Team Type](#)

This section describes the technology and implementation considerations when working with the network teaming services offered by the Broadcom software shipped with servers and storage products. The goal of Broadcom teaming services is to provide fault tolerance and link aggregation across a team of two or more adapters. The information in this document is provided to assist IT professionals during the deployment and troubleshooting of system applications that require network fault tolerance and load balancing.

GLOSSARY

Table 1: Glossary

<i>Item</i>	<i>Definition</i>
ARP	Address Resolution Protocol
BACS	Broadcom Advanced Control Suite
BASP	Broadcom Advanced Server Program (intermediate driver)
DNS	domain name service
G-ARP	Gratuitous Address Resolution Protocol
Generic Trunking (FEC/GEC)/802.3ad-Draft Static	Switch-dependent load balancing and failover type of team in which the intermediate driver manages outgoing traffic and the switch manages incoming traffic.
HSRP	Hot Standby Router Protocol
ICMP	Internet Control Message Protocol
IGMP	Internet Group Management Protocol
IP	Internet Protocol
IPv6	Version 6 of the IP Protocol
iSCSI	Internet Small Computer Systems Interface
L2	Layer 2. Used to describe network traffic that is not offloaded, and where hardware only performs Layer 2 operations on the traffic. Layer 3 (IP) and Layer 4 (TCP) protocols are processed in software.
L4	Layer 4. Used to describe network traffic that is heavily offloaded to the hardware, where much of the Layer 3 (IP) and Layer 4 (TCP) processing is done in the hardware to improve performance.
LACP	Link Aggregation Control Protocol
Link Aggregation (802.3ad)	Switch-dependent load balancing and failover type of team with LACP in which the intermediate driver manages outgoing traffic and the switch manages incoming traffic.
LOM	LAN on Motherboard
MAC	media access control
NDIS	Network Driver Interface Specification
NLB	Network Load Balancing (Microsoft)
PXE	Preboot Execution Environment
RAID	redundant array of inexpensive disks
Smart Load Balancing™ and Failover	Switch-independent failover type of team in which the primary team member handles all incoming and outgoing traffic while the standby team member is idle until a failover event (for example, loss of link occurs). The intermediate driver (BASP) manages incoming/outgoing traffic.
Smart Load Balancing (SLB)	Switch-independent load balancing and failover type of team, in which the intermediate driver manages outgoing/incoming traffic.
TCP	Transmission Control Protocol
TOE	TCP Offload Engine. This is the hardware that is capable of handling stateful fastpath offloading of TCP and IP processing.
UDP	User Datagram Protocol

Table 1: Glossary (Cont.)

<i>Item</i>	<i>Definition</i>
WINS	Windows name service
WLBS	Windows Load Balancing Service

TEAMING CONCEPTS

- [Network Addressing](#)
- [Teaming and Network Addresses](#)
- [Description of Teaming Types](#)
- [TOE Teaming](#)

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.

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 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.

Teaming and Network Addresses

A team of adapters function as a single virtual network interface and does 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.

Description of Teaming Types

- [Smart Load Balancing and Failover](#)
- [Generic Trunking](#)
- [Link Aggregation \(IEEE 802.3ad LACP\)](#)
- [SLB \(Auto-Fallback Disable\)](#)

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 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 team)	Link Aggregation Control Protocol Support Required on the Switch	Load Balancing	Failover
Smart Load Balancing and Failover (with two to eight load balance team members)			✓	✓
SLB (Auto-Fallback Disable)			✓	✓
Link Aggregation (802.3ad)	✓	✓	✓	✓
Generic Trunking (FEC/GEC)/802.3ad-Draft Static	✓		✓	✓

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.

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.

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.

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.

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.

TOE Teaming

All four basic teaming modes support failover of traffic from a failed adapter to other working adapters. All four teaming modes also support bidirectional load-balancing of TCP/IP traffic. A primary difference between the modes is that the SLB modes use a Broadcom proprietary algorithm to control how both inbound and outbound traffic is balanced across the network interfaces in the team. This has several advantages. First, with Generic Trunking or Link Aggregation modes, the team of network adapters must be connected to a switch that is specifically configured to support that particular mode of teaming. Since there is a dependency between the switch and the host team configuration when Generic Trunking or Link Aggregation is used, it can often lead to configuration difficulties, because both ends must be configured correctly and be synchronized. Second, with Generic Trunking or Link Aggregation modes, the switch decides how inbound traffic to the team is balanced across the adapters, while BASP only controls the balancing of outbound traffic. This is problematic for TOE environments, because in order for TOE to work, state information about a given TCP connection is stored in the hardware on a given offloaded adapter, but it is not stored in the hardware on every member of the team. So teaming and TOE cannot co-exist if the teaming software cannot steer incoming TCP/IP traffic to the adapter that contains and updates the state information for a given TCP connection.

Because Broadcom's SLB modes can control how both outbound and inbound packets are balanced across the adapters, the SLB modes are capable of ensuring that all offloaded TCP traffic for a given TCP connection goes in and out of a particular adapter. This architectural feature allows the SLB modes to also support load-balancing on adapters that have TOE enabled, since BASP is able to steer traffic on a particular TCP connection to the adapter hardware that contains offloaded state information for that TCP connection. BASP can simultaneously use TCP offload in conjunction with the SLB modes of teaming. Other teaming modes (Generic Trunking or Link Aggregation) can still be used on TOE capable devices, but if those other modes are enabled the TOE feature is disabled.

Since the TOE offloaded state is stored in only one member of a team, it might not be intuitive as to how BASP can support failover on TOE teams. When a TOE connection has been offloaded to a given adapter, and if that network interface fails in some way (that is, it loses its network link due to a cable disconnection), then BASP will detect the error and force an upload of the offloaded TCP state for each previously offloaded TCP connection on that adapter to the host. Once all of the previously offloaded state has been uploaded, BASP will rebalance the recently uploaded TCP connections and offload those connections evenly to the remaining members of the team. Basically, if there is a failure on a TOE-enabled adapter, any TCP connections that had been offloaded to that adapter are migrated to the remaining nonfailed members in the team.

For Broadcom NetXtreme II adapters, there are no specific setup requirements in order for TCP Offload Engine (TOE) to work with BASP. Once the individual adapters are configured to enable TOE, they can be added to a team and the offload is transparent to BASP. For information on configuring TOE, see [Viewing Resource Reservations](#).

Limitations of Teaming with Offloading

- TOE is enabled for a team only when all of the members support and are configured for TOE.
- TOE is only supported on SLB-type teams.
- Each virtual BASP device advertises 1024 offload connections. If the number of virtual BASP devices in a team exceeds the number of active physical members, the maximum offload connections for each virtual device may be lower.

SOFTWARE COMPONENTS

Teaming is implemented via an NDIS intermediate driver in the Windows Operating System environment. This software component works with the miniport driver, the NDIS layer, and the protocol stack to enable the teaming architecture (see [Figure 2](#)). The miniport driver controls the host LAN controller directly to enable functions such as sends, receives, and interrupt processing. The intermediate driver fits between the miniport driver and the protocol layer multiplexing several miniport driver instances, and creating a virtual adapter that looks like a single adapter to the NDIS layer. NDIS provides a

set of library functions to enable the communications between either miniport drivers or intermediate drivers and the protocol stack. The protocol stack implements IP, IPX and ARP. A protocol address such as an IP address is assigned to each miniport device instance, but when an Intermediate driver is installed, the protocol address is assigned to the virtual team adapter and not to the individual miniport devices that make up the team.

The Broadcom supplied teaming support is provided by three individual software components that work together and are supported as a package. When one component is upgraded, all the other components must be upgraded to the supported versions. [Table 3](#) describes the four software components and their associated files for supported operating systems.

Table 3: Broadcom Teaming Software Component

Software Component	Broadcom Name	Network Adapter/Operating System	System Architecture	Windows File Name
	Virtual Bus Driver (VBD)	BCM5706, BCM5708, BCM5709	32-bit	bxvbdx.sys
		BCM5706, BCM5708, BCM5709	64-bit	bxvbda.sys
		BCM57710, BCM57711	32-bit	evbdx.sys
		BCM57710, BCM57711	64-bit	evbda.sys
Miniport Driver	Broadcom Base Driver	Windows 2000 Server (NDIS 5.0)		bxnd50x.sys
		Windows Server 2003 (NDIS 5.1)	32-bit	bxnd51x.sys
		Windows Server 2003 (NDIS 5.1)	64-bit	bxnd51a.sys
		Windows Server 2003 (NDIS 5.2) Driver supports Layer 4	32-bit	bxnd52x.sys
		Windows Server 2003 (NDIS 5.2) Driver supports Layer 4	64-bit	bxnd52a.sys
		Windows Server 2008 (NDIS 6.0)	32-bit	bxnd60x.sys
		Windows Server 2008 (NDIS 6.0)	64-bit	bxnd60a.sys
		Windows Server 2008 R2 (NDIS 6.0)	64-bit	bxnd60a.sys
Intermediate Driver	Broadcom Advanced Server Program (BASP)	Windows 2000 Server	32-bit	baspw2k.sys
		Windows Server 2003	32-bit	baspxp32.sys
		Windows Server 2003	64-bit	basamd64.sys
		Windows Server 2008	32-bit, 64-bit	basp.sys
		Windows Server 2008 R2	64-bit	basp.sys
Configuration User Interface	Broadcom Advanced Control Suite 3 (BACS)	–	–	bacs.exe



HARDWARE REQUIREMENTS

- [Repeater Hub](#)
- [Switching Hub](#)
- [Router](#)

The various teaming modes described in this document place certain restrictions on the networking equipment used to connect clients to teamed systems. Each type of network interconnect technology has an effect on teaming as described in the following sections.

Repeater Hub

A Repeater Hub allows a network administrator to extend an Ethernet network beyond the limits of an individual segment. The repeater regenerates the input signal received on one port onto all other connected ports, forming a single collision domain. This means that when a station attached to a repeater sends an Ethernet frame to another station, every station within the same collision domain will also receive that message. If two stations begin transmitting at the same time, a collision occurs, and each transmitting station must retransmit its data after waiting a random amount of time.

The use of a repeater requires that each station participating within the collision domain operate in half-duplex mode. Although half-duplex mode is supported for Gigabit Ethernet adapters in the IEEE 802.3 specification, half-duplex mode is not supported by the majority of Gigabit Ethernet adapter manufacturers. Therefore, half-duplex mode is not considered here.

Teaming across hubs is supported for troubleshooting purposes (such as connecting a network analyzer) for SLB teams only.

Switching Hub

Unlike a repeater hub, a switching hub (or more simply a switch) allows an Ethernet network to be broken into multiple collision domains. The switch is responsible for forwarding Ethernet packets between hosts based solely on Ethernet MAC addresses. A physical network adapter that is attached to a switch may operate in half-duplex or full-duplex mode.

To support Generic Trunking and 802.3ad Link Aggregation, a switch must specifically support such functionality. If the switch does not support these protocols, it may still be used for Smart Load Balancing.



NOTE: All modes of network teaming are supported across switches when operating as a stackable switch.

Router

A router is designed to route network traffic based on Layer 3 or higher protocols, although it often also works as a Layer 2 device with switching capabilities. The teaming of ports connected directly to a router is not supported.

TEAMING SUPPORT BY PROCESSOR

All team types are supported by the IA-32, AMD-64, and EM64T processors.

CONFIGURING TEAMING

The Broadcom Advanced Control Suite 3 utility is used to configure teaming in the supported operating system environments.

The Broadcom Advanced Control Suite 3 (BACS) utility is designed to run on 32-bit and 64-bit Windows family of operating systems. BACS 3 is used to configure load balancing and fault tolerance teaming, and VLANs. In addition, it displays the MAC address, driver version, and status information about each network adapter. BACS 3 also includes a number of diagnostics tools such as hardware diagnostics, cable testing, and a network topology test.

SUPPORTED FEATURES BY TEAM TYPE

Table 4 provides a feature comparison across the team types. Use this table to determine the best type of team for your application. The teaming software supports up to eight ports in a single team and up to four teams in a single system. The four teams can be any combination of the supported teaming types, but each team must be on a separate network or subnet.

Table 4: Comparison of Team Types

<i>Type of Team</i>	<i>Fault Tolerance</i>	<i>Load Balancing</i>	<i>Switch-Dependent Static Trunking</i>	<i>Switch-Independent Dynamic Link Aggregation (IEEE 802.3ad)</i>
<i>Function</i>	<i>SLB with Standby^a</i>	<i>SLB</i>	<i>Generic Trunking</i>	<i>Link Aggregation</i>
Number of ports per team (same broadcast domain)	2–8	2–8	2–8	2–8
Number of teams	4	4	4	4
Adapter fault tolerance	Yes	Yes	Yes	Yes
Switch link fault tolerance (same broadcast domain)	Yes	Yes	Switch-dependent	Switch-dependent
TX load balancing	No	Yes	Yes	Yes
RX load balancing	No	Yes	Yes (performed by the switch)	Yes (performed by the switch)
Requires compatible switch	No	No	Yes	Yes
Heartbeats to check connectivity	No	No	No	No
Mixed media (adapters with different media)	Yes	Yes	Yes (switch-dependent)	Yes



Table 4: Comparison of Team Types (Cont.)

<i>Type of Team</i>	<i>Fault Tolerance</i>	<i>Load Balancing</i>	<i>Switch-Dependent Static Trunking</i>	<i>Switch-Independent Dynamic Link Aggregation (IEEE 802.3ad)</i>
<i>Function</i>	<i>SLB with Standby^a</i>	<i>SLB</i>	<i>Generic Trunking</i>	<i>Link Aggregation</i>
Mixed speeds (adapters that do not support a common speed(s), but can operate at different speeds)	Yes	Yes	No	No
Mixed speeds (adapters that support a common speed(s), but can operate at different speeds)	Yes	Yes	No (must be the same speed)	Yes
Load balances TCP/IP	No	Yes	Yes	Yes
Mixed vendor teaming	Yes ^b	Yes ^b	Yes ^b	Yes ^b
Load balances non-IP	No	Yes (IPX outbound traffic only)	Yes	Yes
Same MAC address for all team members	No	No	Yes	Yes
Same IP address for all team members	Yes	Yes	Yes	Yes
Load balancing by IP address	No	Yes	Yes	Yes
Load balancing by MAC address	No	Yes (used for no-IP/IPX)	Yes	Yes
Allows TOE functionality to co-exist when all team members support TOE ^c	Yes	Yes	No	No

^a SLB with one primary and one standby member.

^b Requires at least one Broadcom adapter in the team.


^c TOE functionality can only be achieved with SLB teams that consist of all Broadcom TOE-enabled adapters.



SELECTING A TEAM TYPE

The following flow chart provides the decision flow when planning for Layer 2 teaming. For TOE teaming, only Smart Load Balancing™ and Failover type team is supported. The primary rationale for teaming is the need for additional network bandwidth and fault tolerance. Teaming offers link aggregation and fault tolerance to meet both of these requirements. Preference teaming should be selected in the following order: Link Aggregation as the first choice, Generic Trunking as the second choice, and SLB teaming as the third choice when using unmanaged switches or switches that do not support the first two options. If switch fault tolerance is a requirement, then SLB is the only choice (see [Figure 1](#)).

Figure 1: Process for Selecting a Team Type



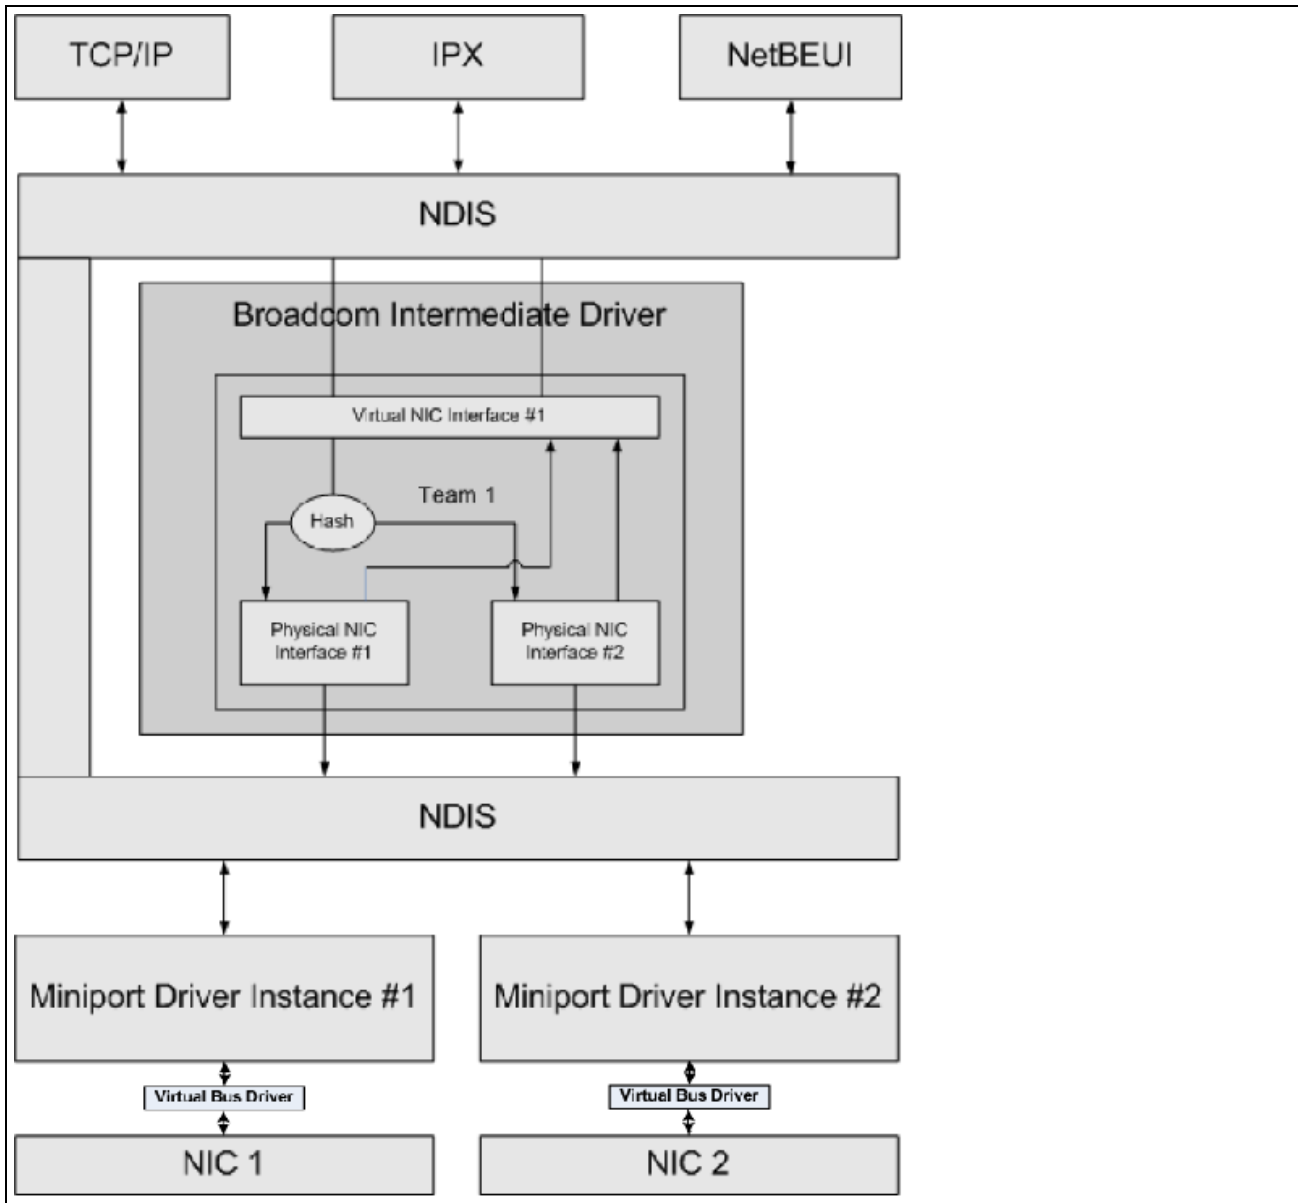
TEAMING MECHANISMS

- [Architecture](#)
- [Types of Teams](#)
- [Attributes of the Features Associated with Each Type of Team](#)
- [Speeds Supported for Each Type of Team](#)

ARCHITECTURE

The Broadcom Advanced Server Program is implemented as an NDIS intermediate driver (see [Figure 2](#)). It operates below protocol stacks such as TCP/IP and IPX and appears as a virtual adapter. This virtual adapter inherits the MAC Address of the first port initialized in the team. A Layer 3 address must also be configured for the virtual adapter. The primary function of BASP is to balance inbound (for SLB) and outbound traffic (for all teaming modes) among the physical adapters installed on the system selected for teaming. The inbound and outbound algorithms are independent and orthogonal to each other. The outbound traffic for a particular session can be assigned to a given port while its corresponding inbound traffic can be assigned to a different port.

Figure 2: Intermediate Driver



Outbound Traffic Flow

The Broadcom Intermediate Driver manages the outbound traffic flow for all teaming modes. For outbound traffic, every packet is first classified into a flow, and then distributed to the selected physical adapter for transmission. The flow classification involves an efficient hash computation over known protocol fields. The resulting hash value is used to index into an Outbound Flow Hash Table. The selected Outbound Flow Hash Entry contains the index of the selected physical adapter responsible for transmitting this flow. The source MAC address of the packets will then be modified to the MAC address of the selected physical adapter. The modified packet is then passed to the selected physical adapter for transmission.



The outbound TCP and UDP packets are classified using Layer 3 and Layer 4 header information. This scheme improves the load distributions for popular Internet protocol services using well-known ports such as HTTP and FTP. Therefore, BASP performs load balancing on a TCP session basis and not on a packet-by-packet basis.

In the Outbound Flow Hash Entries, statistics counters are also updated after classification. The load-balancing engine uses these counters to periodically distribute the flows across teamed ports. The outbound code path has been designed to achieve best possible concurrency where multiple concurrent accesses to the Outbound Flow Hash Table are allowed.

For protocols other than TCP/IP, the first physical adapter will always be selected for outbound packets. The exception is Address Resolution Protocol (ARP), which is handled differently to achieve inbound load balancing.

Inbound Traffic Flow (SLB Only)

The Broadcom intermediate driver manages the inbound traffic flow for the SLB teaming mode. Unlike outbound load balancing, inbound load balancing can only be applied to IP addresses that are located in the same subnet as the load-balancing server. Inbound load balancing exploits a unique characteristic of Address Resolution Protocol (RFC0826), in which each IP host uses its own ARP cache to encapsulate the IP Datagram into an Ethernet frame. BASP carefully manipulates the ARP response to direct each IP host to send the inbound IP packet to the desired physical adapter. Therefore, inbound load balancing is a plan-ahead scheme based on statistical history of the inbound flows. New connections from a client to the server will always occur over the primary physical adapter (because the ARP Reply generated by the operating system protocol stack will always associate the logical IP address with the MAC address of the primary physical adapter).

Like the outbound case, there is an Inbound Flow Head Hash Table. Each entry inside this table has a singly linked list and each link (Inbound Flow Entries) represents an IP host located in the same subnet.

When an inbound IP Datagram arrives, the appropriate Inbound Flow Head Entry is located by hashing the source IP address of the IP Datagram. Two statistics counters stored in the selected entry are also updated. These counters are used in the same fashion as the outbound counters by the load-balancing engine periodically to reassign the flows to the physical adapter.

On the inbound code path, the Inbound Flow Head Hash Table is also designed to allow concurrent access. The link lists of Inbound Flow Entries are only referenced in the event of processing ARP packets and the periodic load balancing. There is no per packet reference to the Inbound Flow Entries. Even though the link lists are not bounded; the overhead in processing each non-ARP packet is always a constant. The processing of ARP packets, both inbound and outbound, however, depends on the number of links inside the corresponding link list.

On the inbound processing path, filtering is also employed to prevent broadcast packets from looping back through the system from other physical adapters.

Protocol Support

ARP and IP/TCP/UDP flows are load balanced. If the packet is an IP protocol only, such as ICMP or IGMP, then all data flowing to a particular IP address will go out through the same physical adapter. If the packet uses TCP or UDP for the L4 protocol, then the port number is added to the hashing algorithm, so two separate L4 flows can go out through two separate physical adapters to the same IP address.

For example, assume the client has an IP address of 10.0.0.1. All IGMP and ICMP traffic will go out the same physical adapter because only the IP address is used for the hash. The flow would look something like this:

IGMP -----> PhysAdapter1 -----> 10.0.0.1



ICMP -----> PhysAdapter1 -----> 10.0.0.1

If the server also sends an TCP and UDP flow to the same 10.0.0.1 address, they can be on the same physical adapter as IGMP and ICMP, or on completely different physical adapters from ICMP and IGMP. The stream may look like this:

IGMP -----> PhysAdapter1 -----> 10.0.0.1

ICMP -----> PhysAdapter1 -----> 10.0.0.1

TCP-----> PhysAdapter1 -----> 10.0.0.1

UDP-----> PhysAdatper1 -----> 10.0.0.1

Or the streams may look like this:

IGMP -----> PhysAdapter1 -----> 10.0.0.1

ICMP -----> PhysAdapter1 -----> 10.0.0.1

TCP-----> PhysAdapter2 -----> 10.0.0.1

UDP-----> PhysAdatper3 -----> 10.0.0.1

The actual assignment between adapters may change over time, but any protocol that is not TCP/UDP based goes over the same physical adapter because only the IP address is used in the hash.

Performance

Modern network interface cards provide many hardware features that reduce CPU utilization by offloading certain CPU intensive operations (see [Teaming and Other Advanced Networking Properties](#)). In contrast, the BASP intermediate driver is a purely software function that must examine every packet received from the protocol stacks and react to its contents before sending it out through a particular physical interface. Though the BASP driver can process each outgoing packet in near constant time, some applications that may already be CPU bound may suffer if operated over a teamed interface. Such an application may be better suited to take advantage of the failover capabilities of the intermediate driver rather than the load balancing features, or it may operate more efficiently over a single physical adapter that provides a particular hardware feature such as Large Send Offload.

TYPES OF TEAMS

Switch-Independent

The Broadcom Smart Load Balancing type of team allows two to eight physical adapters to operate as a single virtual adapter. The greatest benefit of the SLB type of team is that it operates on any IEEE compliant switch and requires no special configuration.

Smart Load Balancing and Failover

SLB provides for switch-independent, bidirectional, fault-tolerant teaming and load balancing. Switch independence implies that there is no specific support for this function required in the switch, allowing SLB to be compatible with all switches. Under SLB, all adapters in the team have separate MAC addresses. The load-balancing algorithm operates on Layer 3 addresses of the source and destination nodes, which enables SLB to load balance both incoming and outgoing traffic.

The BASP intermediate driver continually monitors the physical ports in a team for link loss. In the event of link loss on any port, traffic is automatically diverted to other ports in the team. The SLB teaming mode supports switch fault tolerance by allowing teaming across different switches- provided the switches are on the same physical network or broadcast domain.

Network Communications

The following are the key attributes of SLB:

- Failover mechanism – Link loss detection.
- Load Balancing Algorithm – Inbound and outbound traffic are balanced through a Broadcom proprietary mechanism based on L4 flows.
- Outbound Load Balancing using MAC Address - No.
- Outbound Load Balancing using IP Address - Yes
- Multivendor Teaming – Supported (must include at least one Broadcom Ethernet adapter as a team member).

Applications

The SLB algorithm is most appropriate in home and small business environments where cost is a concern or with commodity switching equipment. SLB teaming works with unmanaged Layer 2 switches and is a cost-effective way of getting redundancy and link aggregation at the server. Smart Load Balancing also supports teaming physical adapters with differing link capabilities. In addition, SLB is recommended when switch fault tolerance with teaming is required.

Configuration Recommendations

SLB supports connecting the teamed ports to hubs and switches if they are on the same broadcast domain. It does not support connecting to a router or Layer 3 switches because the ports must be on the same subnet.

Switch-Dependent

Generic Static Trunking

This mode supports a variety of environments where the adapter link partners are statically configured to support a proprietary trunking mechanism. This mode could be used to support Lucent's *Open Trunk*, Cisco's *Fast EtherChannel* (FEC), and Cisco's *Gigabit EtherChannel* (GEC). In the static mode, as in generic link aggregation, the switch administrator needs to assign the ports to the team, and this assignment cannot be altered by the BASP, as there is no exchange of the Link Aggregation Control Protocol (LACP) frame.

With this mode, all adapters in the team are configured to receive packets for the same MAC address. Trunking operates on Layer 2 addresses and supports load balancing and failover for both inbound and outbound traffic. The BASP driver determines the load-balancing scheme for outbound packets, using Layer 4 protocols previously discussed, whereas the team link partner determines the load-balancing scheme for inbound packets.

The attached switch must support the appropriate trunking scheme for this mode of operation. Both the BASP and the switch continually monitor their ports for link loss. In the event of link loss on any port, traffic is automatically diverted to other ports in the team.

Network Communications

The following are the key attributes of Generic Static Trunking:

- Failover mechanism – Link loss detection
- Load Balancing Algorithm – Outbound traffic is balanced through Broadcom proprietary mechanism based L4 flows. Inbound traffic is balanced according to a switch specific mechanism.
- Outbound Load Balancing using MAC Address – No
- Outbound Load Balancing using IP Address - Yes
- Multivendor teaming – Supported (Must include at least one Broadcom Ethernet adapter as a team member)

Applications

Generic trunking works with switches that support Cisco Fast EtherChannel, Cisco Gigabit EtherChannel, Extreme Networks Load Sharing and Bay Networks or IEEE 802.3ad Link Aggregation static mode. Since load balancing is implemented on Layer 2 addresses, all higher protocols such as IP, IPX, and NetBEUI are supported. Therefore, this is the recommended teaming mode when the switch supports generic trunking modes over SLB.

Configuration Recommendations

Static trunking supports connecting the teamed ports to switches if they are on the same broadcast domain and support generic trunking. It does not support connecting to a router or Layer 3 switches since the ports must be on the same subnet.

Dynamic Trunking (IEEE 802.3ad Link Aggregation)

This mode supports link aggregation through static and dynamic configuration via the Link Aggregation Control Protocol (LACP). With this mode, all adapters in the team are configured to receive packets for the same MAC address. The MAC address of the first adapter in the team is used and cannot be substituted for a different MAC address. The BASP driver determines the load-balancing scheme for outbound packets, using Layer 4 protocols previously discussed, whereas the team's link partner determines the load-balancing scheme for inbound packets. Because the load balancing is implemented on Layer 2, all higher protocols such as IP, IPX, and NetBEUI are supported. The attached switch must support the 802.3ad Link Aggregation standard for this mode of operation. The switch manages the inbound traffic to the adapter while the BASP manages the outbound traffic. Both the BASP and the switch continually monitor their ports for link loss. In the event of link loss on any port, traffic is automatically diverted to other ports in the team.

Network Communications

The following are the key attributes of Dynamic Trunking:

- Failover mechanism – Link loss detection
- Load Balancing Algorithm – Outbound traffic is balanced through a Broadcom proprietary mechanism based on L4 flows. Inbound traffic is balanced according to a switch specific mechanism.
- Outbound Load Balancing using MAC Address - No
- Outbound Load Balancing using IP Address - Yes
- Multivendor teaming – Supported (Must include at least one Broadcom Ethernet adapter as a team member)

Applications

Dynamic trunking works with switches that support IEEE 802.3ad Link Aggregation dynamic mode using LACP. Inbound load balancing is switch dependent. In general, the switch traffic is load balanced based on L2 addresses. In this case, all network protocols such as IP, IPX, and NetBEUI are load balanced. Therefore, this is the recommended teaming mode when the switch supports LACP, except when switch fault tolerance is required. SLB is the only teaming mode that supports switch fault tolerance.

Configuration Recommendations

Dynamic trunking supports connecting the teamed ports to switches as long as they are on the same broadcast domain and supports IEEE 802.3ad LACP trunking. It does not support connecting to a router or Layer 3 switches since the ports must be on the same subnet.

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. This function is accomplished through the teaming software. The teaming software periodically probes (issues a link packet from each team member) one or more specified target network device(s). The probe target(s) responds when it receives the link packet. If a team member does not detect the response within a specified amount of time, this indicates that the link has been lost, and the teaming software discontinues passing traffic through that team member. Later, if that team member begins to detect a response from a probe target, this indicates that the link has been restored, and the teaming software automatically resumes passing traffic through that team member. LiveLink works only with TCP/IP.

LiveLink™ functionality is supported in both 32-bit and 64-bit Windows operating systems. For similar functionality in Linux operating systems, see the Channel Bonding information in your Red Hat documentation.



ATTRIBUTES OF THE FEATURES ASSOCIATED WITH EACH TYPE OF TEAM

The attributes of the features associated with each type of team are summarized in [Table 5](#).

Table 5: Attributes

Feature	Attribute
Smart Load Balancing™	
User interface	Broadcom Advanced Control Suite 3 (BACS)
Number of teams	Maximum 8
Number of adapters per team	Maximum 8
Hot replace	Yes
Hot add	Yes
Hot remove	Yes
Link speed support	Different speeds
Frame protocol	IP
Incoming packet management	BASP
Outgoing packet management	BASP
LiveLink support	Yes
Failover event	Loss of link
Failover time	<500 ms
Fallback time	1.5 s ^b (approximate)
MAC address	Different
Multivendor teaming	Yes
Generic Trunking	
User interface	Broadcom Advanced Control Suite 3 (BACS)
Number of teams	Maximum 8
Number of adapters per team	Maximum 8
Hot replace	Yes
Hot add	Yes
Hot remove	Yes
Link speed support	Different speeds ^a
Frame protocol	All
Incoming packet management	Switch
Outgoing packet management	BASP
Failover event	Loss of link only
Failover time	<500 ms
Fallback time	1.5 s ^b (approximate)
MAC address	Same for all adapters
Multivendor teaming	Yes
Dynamic Trunking	
User interface	Broadcom Advanced Control Suite 3 (BACS)
Number of teams	Maximum 8
Number of adapters per team	Maximum 8

Table 5: Attributes (Cont.)

Feature	Attribute
Hot replace	Yes
Hot add	Yes
Hot remove	Yes
Link speed support	Different speeds
Frame protocol	All
Incoming packet management	Switch
Outgoing packet management	BASP
Failover event	Loss of link only
Failover time	<500 ms
Fallback time	1.5 s ^b (approximate)
MAC address	Same for all adapters
Multivendor teaming	Yes

^a Some switches require matching link speeds to correctly negotiate between trunk connections.

^b Make sure that Port Fast or Edge Port is enabled.

SPEEDS SUPPORTED FOR EACH TYPE OF TEAM

The various link speeds that are supported for each type of team are listed in [Table 6](#). Mixed speed refers to the capability of teaming adapters that are running at different link speeds.

Table 6: Link Speeds in Teaming

Type of Team	Link Speed	Traffic Direction	Speed Support
SLB	10/100/1000/10000	Incoming/outgoing	Mixed speed
FEC	100	Incoming/outgoing	Same speed
GEC	1000	Incoming/outgoing	Same speed
IEEE 802.3ad	10/100/1000/10000	Incoming/outgoing	Mixed speed

TEAMING AND OTHER ADVANCED NETWORKING PROPERTIES

- [Checksum Offload](#)
- [IEEE 802.1p QoS Tagging](#)
- [Large Send Offload](#)
- [TCP Offload Engine \(TOE\)](#)
- [Jumbo Frames](#)
- [IEEE 802.1Q VLANs](#)
- [Wake On LAN](#)
- [Preboot Execution Environment](#)

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. Advanced adapter properties and teaming support are listed in [Table 7](#).

Table 7: Advanced Adapter Properties and Teaming Support

Adapter Properties	Supported by Teaming Virtual Adapter
Checksum Offload	Yes
IEEE 802.1p QoS Tagging	No
Large Send Offload	Yes ^a
TCP Offload Engine (TOE)	Yes ^{b, c}
Jumbo Frames	Yes ^b
IEEE 802.1Q VLANs	Yes ^c
Wake on LAN	No
Preboot Execution environment (PXE)	Yes ^d

^a All adapters on the team must support this feature. Some adapters may not support this feature if ASF/IPMI is also enabled.

^b Must be supported by all adapters in the team.

^c Only for Broadcom adapters.

^d As a PXE sever only, not as a client.

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.

CHECKSUM OFFLOAD

Checksum Offload is a property of the Broadcom network adapters that allows the TCP/IP/UDP checksums for send and receive traffic to be calculated by the adapter hardware rather than by the host CPU. In high-traffic situations, this can allow a system to handle more connections more efficiently than if the host CPU were forced to calculate the checksums. This property is inherently a hardware property and would not benefit from a software-only implementation. An adapter that supports Checksum Offload advertises this capability to the operating system so that the checksum does not need to be calculated in the protocol stack. Checksum Offload is only supported for IPv4 at this time.

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.

LARGE SEND OFFLOAD

Large Send Offload (LSO) is a feature provided by Broadcom network adapters that 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. The protocol stack need only generate a single header for a data packet as large as 64 KB, and the adapter hardware breaks the data buffer into appropriately-sized Ethernet frames with the correctly sequenced header (based on the single header originally provided).

TCP OFFLOAD ENGINE (TOE)

The TCP/IP protocol suite is used to provide transport services for a wide range of applications for the Internet, LAN, and for file transfer. Without the TCP Offload Engine, the TCP/IP protocol suite runs on the host CPU, consuming a very high percentage of its resources and leaving little resources for the applications. With the use of the Broadcom NetXtreme II adapter, the TCP/IP processing can be moved to hardware, freeing the CPU for more important tasks such as application processing.

The Broadcom NetXtreme II adapter's TOE functionality allows simultaneous operation of up to 1024 fully offloaded TCP connections for 1-Gbps network adapters and 1880 fully offloaded TCP connections for 10-Gbps network adapters. The TOE support on the adapter significantly reduces the host CPU utilization while preserving the implementation of the operating system stack.

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 Broadcom NetXtreme II 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.

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.

WAKE ON LAN

Wake on LAN (WOL) is a feature that allows a system to be awakened from a sleep state by the arrival of a specific packet over the Ethernet interface. Because a Virtual Adapter is implemented as a software only device, it lacks the hardware features to implement Wake on LAN and cannot be enabled to wake the system from a sleeping state via the Virtual Adapter. The physical adapters, however, support this property, even when the adapter is part of a team.

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.

GENERAL NETWORK CONSIDERATIONS

- [Teaming with Microsoft Virtual Server 2005](#)
- [Teaming Across Switches](#)
- [Spanning Tree Algorithm](#)
- [Layer 3 Routing/Switching](#)
- [Teaming with Hubs \(for troubleshooting purposes only\)](#)
- [Teaming with Microsoft NLB](#)

TEAMING WITH MICROSOFT VIRTUAL SERVER 2005

The only supported BASP team configuration when using Microsoft Virtual Server 2005 is with a Smart Load Balancing (TM) team-type consisting of a single primary Broadcom adapter and a standby Broadcom adapter. Make sure to unbind or deselect "Virtual Machine Network Services" from each team member prior to creating a team and prior to creating Virtual networks with Microsoft Virtual Server. Additionally, a virtual network should be created in this software and subsequently bound to the virtual adapter created by a team. Directly binding a Guest operating system to a team virtual adapter may not render the desired results.



NOTE: As of this writing, Windows Server 2008 is not a supported operating system for Microsoft Virtual Server 2005; thus, teaming may not function as expected with this combination.

TEAMING ACROSS SWITCHES

SLB teaming can be configured across switches. The switches, however, must be connected together. Generic Trunking and Link Aggregation do not work across switches because each of these implementations requires that all physical adapters in a team share the same Ethernet MAC address. It is important to note that SLB can only detect the loss of link between the ports in the team and their immediate link partner. SLB has no way of reacting to other hardware failures in the switches and cannot detect loss of link on other ports.

Switch-Link Fault Tolerance

The diagrams below describe the operation of an SLB team in a switch fault tolerant configuration. We show the mapping of the ping request and ping replies in an SLB team with two active members. All servers (Blue, Gray and Red) have a continuous ping to each other. [Figure 3](#) is a setup without the interconnect cable in place between the two switches. [Figure 4](#) has the interconnect cable in place, and [Figure 5](#) is an example of a failover event with the Interconnect cable in place. These scenarios describe the behavior of teaming across the two switches and the importance of the interconnect link.


The diagrams show the secondary team member sending the ICMP echo requests (yellow arrows) while the primary team member receives the respective ICMP echo replies (blue arrows). This illustrates a key characteristic of the teaming software. The load balancing algorithms do not synchronize how frames are load balanced when sent or received. In other words, frames for a given conversation can go out and be received on different interfaces in the team. This is true for all

types of teaming supported by Broadcom. Therefore, an interconnect link must be provided between the switches that connect to ports in the same team.

In the configuration without the interconnect, an ICMP Request from Blue to Gray goes out port 82:83 destined for Gray port 5E:CA, but the Top Switch has no way to send it there because it cannot go along the 5E:C9 port on Gray. A similar scenario occurs when Gray attempts to ping Blue. An ICMP Request goes out on 5E:C9 destined for Blue 82:82, but cannot get there. Top Switch does not have an entry for 82:82 in its CAM table because there is no interconnect between the two switches. Pings, however, flow between Red and Blue and between Red and Gray.

Furthermore, a failover event would cause additional loss of connectivity. Consider a cable disconnect on the Top Switch port 4. In this case, Gray would send the ICMP Request to Red 49:C9, but because the Bottom switch has no entry for 49:C9 in its CAM Table, the frame is flooded to all its ports but cannot find a way to get to 49:C9.

Figure 3: Teaming Across Switches Without an Interswitch Link



The addition of a link between the switches allows traffic from/to Blue and Gray to reach each other without any problems. Note the additional entries in the CAM table for both switches. The link interconnect is critical for the proper operation of the team. As a result, it is highly advisable to have a link aggregation trunk to interconnect the two switches to ensure high availability for the connection.

Figure 4: Teaming Across Switches With Interconnect

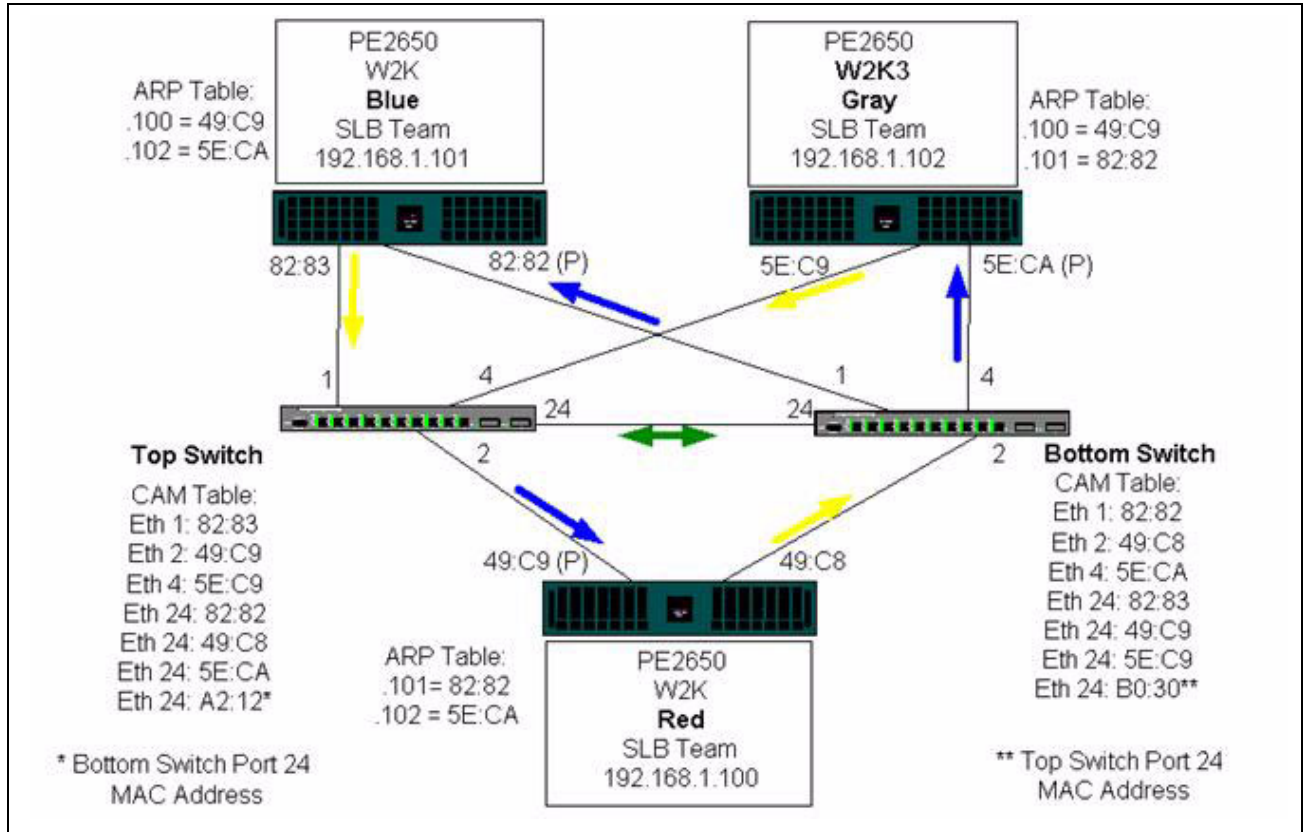
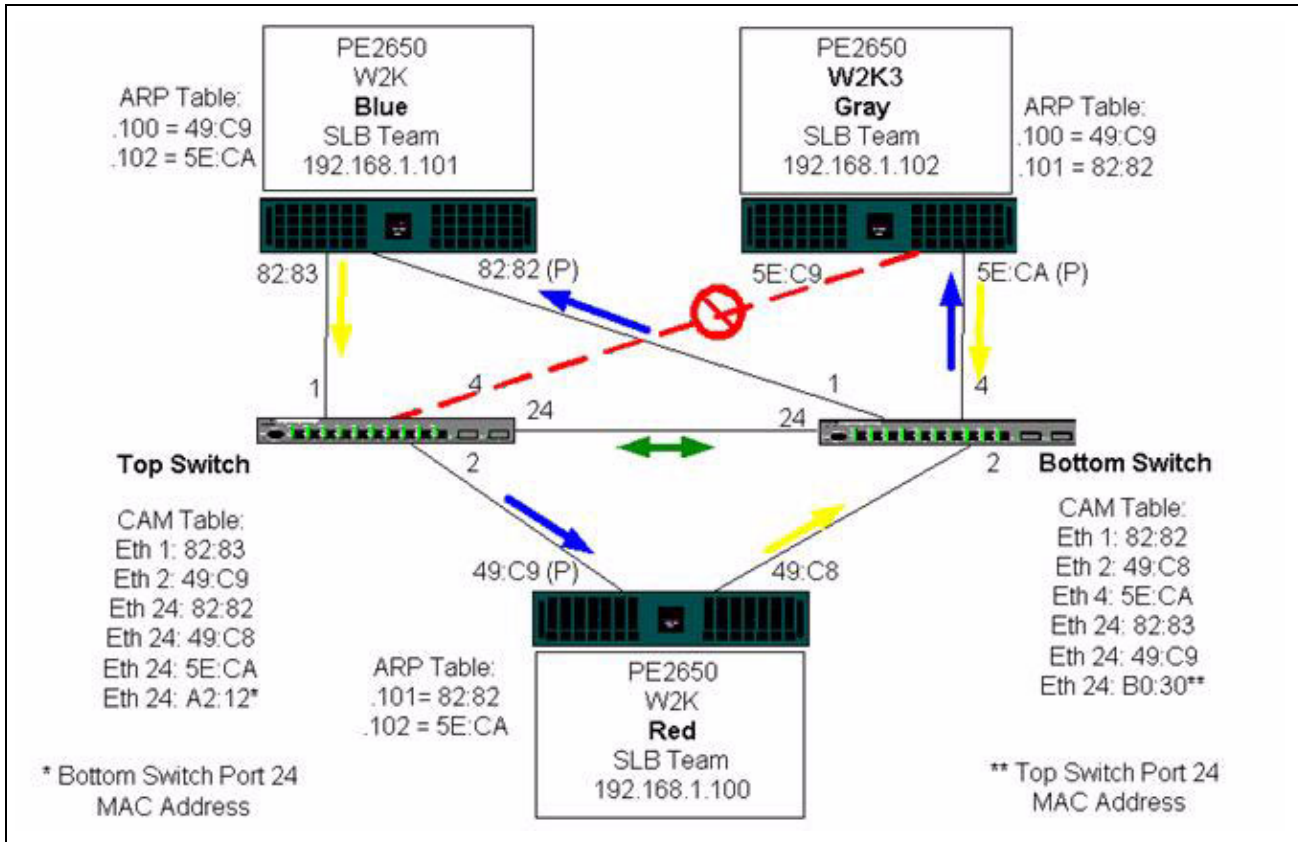


Figure 5 represents a failover event in which the cable is unplugged on the Top Switch port 4. This is a successful failover with all stations pinging each other without loss of connectivity.

Figure 5: Failover Event



SPANNING TREE ALGORITHM

- [Topology Change Notice \(TCN\)](#)
- [Port Fast/Edge Port](#)

In Ethernet networks, only one active path may exist between any two bridges or switches. Multiple active paths between switches can cause loops in the network. When loops occur, some switches recognize stations on both sides of the switch. This situation causes the forwarding algorithm to malfunction allowing duplicate frames to be forwarded. Spanning tree algorithms provide path redundancy by defining a tree that spans all of the switches in an extended network and then forces certain redundant data paths into a standby (blocked) state. At regular intervals, the switches in the network send and receive spanning tree packets that they use to identify the path. If one network segment becomes unreachable, or if spanning tree costs change, the spanning tree algorithm reconfigures the spanning tree topology and re-establishes the link by activating the standby path. Spanning tree operation is transparent to end stations, which do not detect whether they are connected to a single LAN segment or a switched LAN of multiple segments.

Spanning Tree Protocol (STP) is a Layer 2 protocol designed to run on bridges and switches. The specification for STP is defined in IEEE 802.1d. The main purpose of STP is to ensure that you do not run into a loop situation when you have redundant paths in your network. STP detects/disables network loops and provides backup links between switches or bridges. It allows the device to interact with other STP compliant devices in your network to ensure that only one path exists between any two stations on the network.

After a stable network topology has been established, all bridges listen for hello BPDUs (Bridge Protocol Data Units) transmitted from the root bridge. If a bridge does not get a hello BPDU after a predefined interval (Max Age), the bridge assumes that the link to the root bridge is down. This bridge then initiates negotiations with other bridges to reconfigure the network to re-establish a valid network topology. The process to create a new topology can take up to 50 seconds. During this time, end-to-end communications are interrupted.

The use of Spanning Tree is not recommended for ports that are connected to end stations, because by definition, an end station does not create a loop within an Ethernet segment. Additionally, when a teamed adapter is connected to a port with Spanning Tree enabled, users may experience unexpected connectivity problems. For example, consider a teamed adapter that has a lost link on one of its physical adapters. If the physical adapter were to be reconnected (also known as fallback), the intermediate driver would detect that the link has been reestablished and would begin to pass traffic through the port. Traffic would be lost if the port was temporarily blocked by the Spanning Tree Protocol.

Topology Change Notice (TCN)

A bridge/switch creates a forwarding table of MAC addresses and port numbers by learning the source MAC address that received on a particular port. The table is used to forward frames to a specific port rather than flooding the frame to all ports. The typical maximum aging time of entries in the table is 5 minutes. Only when a host has been silent for 5 minutes would its entry be removed from the table. It is sometimes beneficial to reduce the aging time. One example is when a forwarding link goes to blocking and a different link goes from blocking to forwarding. This change could take up to 50 seconds. At the end of the STP re-calculation a new path would be available for communications between end stations. However, because the forwarding table would still have entries based on the old topology, communications may not be reestablished until after 5 minutes when the affected ports entries are removed from the table. Traffic would then be flooded to all ports and re-learned. In this case it is beneficial to reduce the aging time. This is the purpose of a topology change notice (TCN) BPDU. The TCN is sent from the affected bridge/switch to the root bridge/switch. As soon as a bridge/switch detects a topology change (a link going down or a port going to forwarding) it sends a TCN to the root bridge via its root port. The root bridge then advertises a BPDU with a Topology Change to the entire network. This causes every bridge to reduce the MAC table aging time to 15 seconds for a specified amount of time. This allows the switch to re-learn the MAC addresses as soon as STP re-converges.

Topology Change Notice BPDUs are sent when a port that was forwarding changes to blocking or transitions to forwarding. A TCN BPDU does not initiate an STP recalculation. It only affects the aging time of the forwarding table entries in the switch. It will not change the topology of the network or create loops. End nodes such as servers or clients trigger a topology change when they power off and then power back on.

Port Fast/Edge Port

To reduce the effect of TCNs on the network (for example, increasing flooding on switch ports), end nodes that are powered on/off often should use the Port Fast or Edge Port setting on the switch port they are attached to. Port Fast or Edge Port is a command that is applied to specific ports and has the following effects:

- Ports coming from link down to link up will be put in the forwarding STP mode instead of going from listening to learning and then to forwarding. STP is still running on these ports.
- The switch does not generate a Topology Change Notice when the port is going up or down.

LAYER 3 ROUTING/SWITCHING

The switch that the teamed ports are connected to must not be a Layer 3 switch or router. The ports in the team must be in the same network.

TEAMING WITH HUBS (FOR TROUBLESHOOTING PURPOSES ONLY)

- [Hub Usage in Teaming Network Configurations](#)
- [SLB Teams](#)
- [SLB Team Connected to a Single Hub](#)
- [Generic and Dynamic Trunking \(FEC/GEC/IEEE 802.3ad\)](#)

SLB teaming can be used with 10/100 hubs, but it is only recommended for troubleshooting purposes, such as connecting a network analyzer in the event that switch port mirroring is not an option.

Hub Usage in Teaming Network Configurations

Although the use of hubs in network topologies is functional in some situations, it is important to consider the throughput ramifications when doing so. Network hubs have a maximum of 100 Mbps half-duplex link speed, which severely degrades performance in either a Gigabit or 100 Mbps switched-network configuration. Hub bandwidth is shared among all connected devices; as a result, when more devices are connected to the hub, the bandwidth available to any single device connected to the hub is reduced in direct proportion to the number of devices connected to the hub.

It is not recommended to connect team members to hubs; only switches should be used to connect to teamed ports. An SLB team, however, can be connected directly to a hub for troubleshooting purposes. Other team types can result in a loss of connectivity if specific failures occur and should not be used with hubs.


SLB Teams

SLB teams are the only teaming type not dependant on switch configuration. The server intermediate driver handles the load balancing and fault tolerance mechanisms with no assistance from the switch. These elements of SLB make it the only team type that maintains failover and fallback characteristics when team ports are connected directly to a hub.

SLB Team Connected to a Single Hub

SLB teams configured as shown in [Figure 6](#) maintain their fault tolerance properties. Either server connection could potentially fail, and network functionality is maintained. Clients could be connected directly to the hub, and fault tolerance would still be maintained; server performance, however, would be degraded.

Figure 6: Team Connected to a Single Hub



Generic and Dynamic Trunking (FEC/GEC/IEEE 802.3ad)

FEC/GEC and IEEE 802.3ad teams cannot be connected to any hub configuration. These team types must be connected to a switch that has also been configured for this team type.

TEAMING WITH MICROSOFT NLB

Teaming *does not* work in Microsoft's Network Load Balancing (NLB) unicast mode, only in multicast mode. Due to the mechanism used by the NLB service, the recommended teaming configuration in this environment is Failover (SLB with a standby NIC) as load balancing is managed by NLB. The TOE functionality in teaming will not operate in NLB.

APPLICATION CONSIDERATIONS

- [Teaming and Clustering](#)
- [Teaming and Network Backup](#)

TEAMING AND CLUSTERING

- [Microsoft Cluster Software](#)
- [High-Performance Computing Cluster](#)
- [Oracle](#)

Microsoft Cluster Software

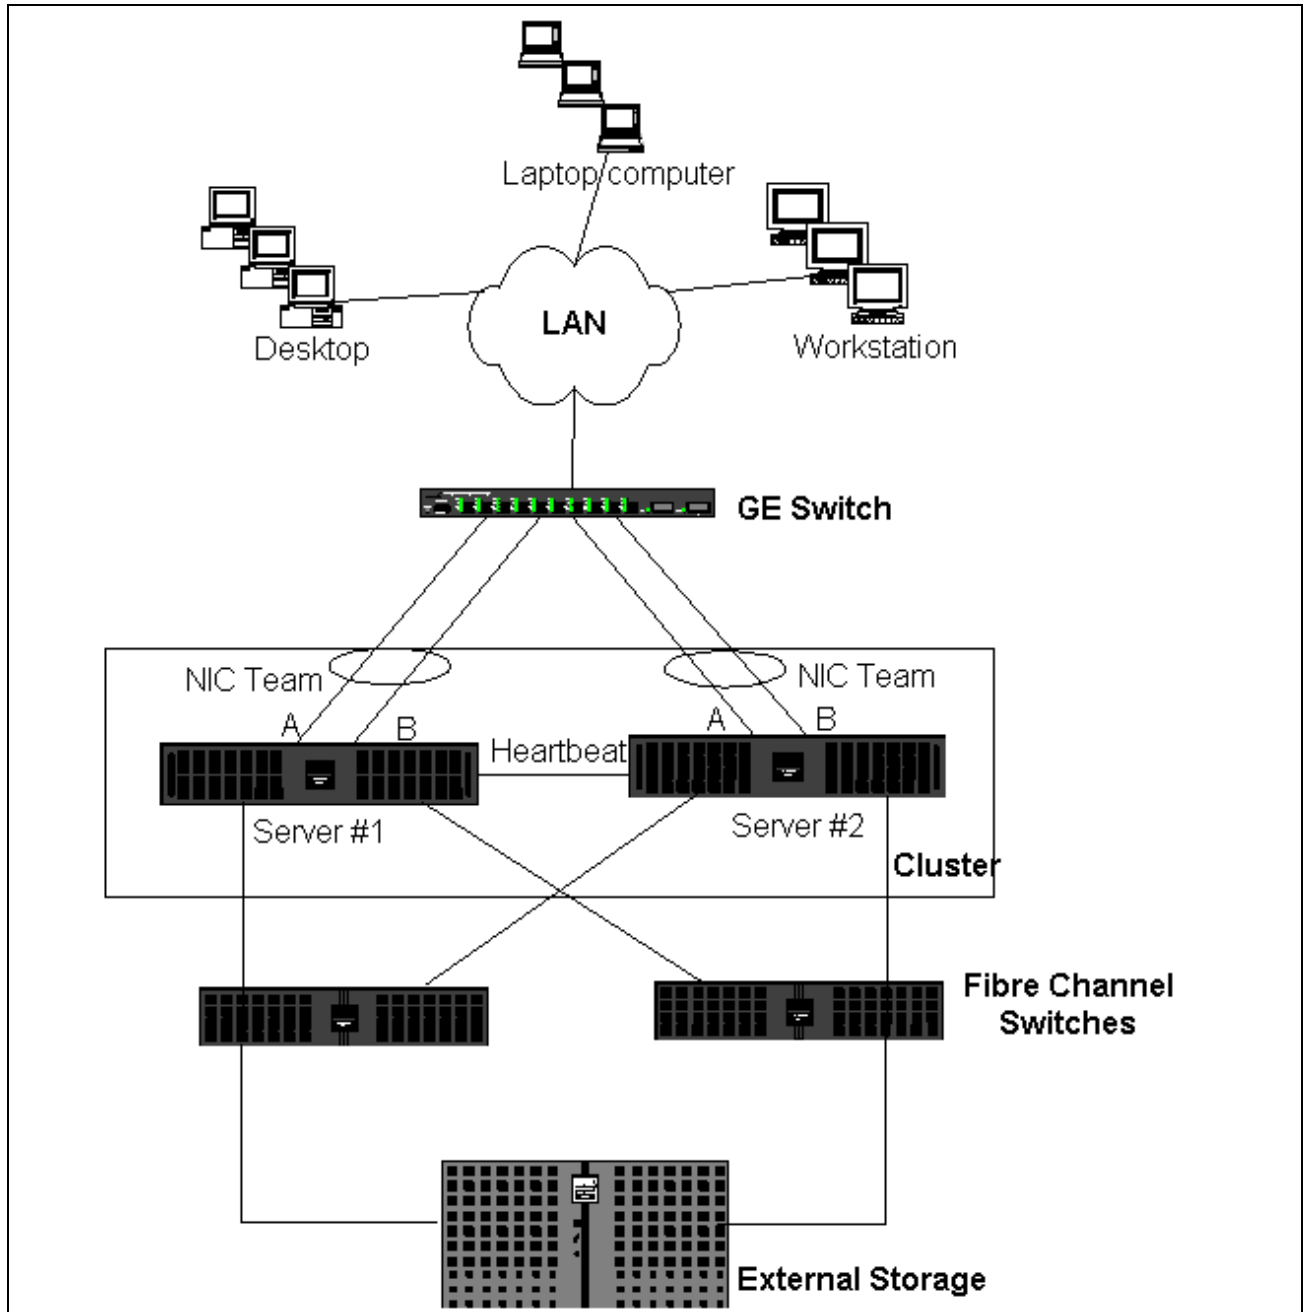
MSCS clusters support up to two nodes if you are using Windows 2000 Server. If you are using Windows Server 2003, MSCS clusters support up to eight nodes. In each cluster node, it is strongly recommended that customers install at least two network adapters (on-board adapters are acceptable). These interfaces serve two purposes. One adapter is used exclusively for intra-cluster *heartbeat* communications. This is referred to as the *private adapter* and usually resides on a separate private subnetwork. The other adapter is used for client communications and is referred to as the *public adapter*.

Multiple adapters may be used for each of these purposes: private, intracluster communications and public, external client communications. All Broadcom teaming modes are supported with Microsoft Cluster Software for the public adapter only. Private network adapter teaming is not supported. Microsoft indicates that the use of teaming on the private interconnect of a server cluster is not supported because of delays that could possibly occur in the transmission and receipt of heartbeat packets between the nodes. For best results, when you want redundancy for the private interconnect, disable teaming and use the available ports to form a second private interconnect. This achieves the same end result and provides dual, robust communication paths for the nodes to communicate over.

For teaming in a clustered environment, customers are recommended to use the same brand of adapters.

Figure 7 shows a 2-node Fibre-Channel cluster with three network interfaces per cluster node: one private and two public. On each node, the two public adapters are teamed, and the private adapter is not. Teaming is supported across the same switch or across two switches. Figure 8 shows the same 2-node Fibre-Channel cluster in this configuration.

Figure 7: Clustering With Teaming Across One Switch



NOTE: Microsoft Network Load Balancing is not supported with Microsoft Cluster Software.

High-Performance Computing Cluster

Gigabit Ethernet is typically used for the following three purposes in high-performance computing cluster (HPCC) applications:

- **Inter-Process Communications (IPC):** For applications that do not require low-latency, high-bandwidth interconnects (such as Myrinet, InfiniBand), Gigabit Ethernet can be used for communication between the compute nodes.
- **I/O:** Ethernet can be used for file sharing and serving the data to the compute nodes. This can be done simply using an NFS server or using parallel file systems such as PVFS.
- **Management & Administration:** Ethernet is used for out-of-band (ERA) and in-band (OMSA) management of the nodes in the cluster. It can also be used for job scheduling and monitoring.

In our current HPCC offerings, only one of the on-board adapters is used. If Myrinet or IB is present, this adapter serves I/O and administration purposes; otherwise, it is also responsible for IPC. In case of an adapter failure, the administrator can use the Felix package to easily configure adapter 2. Adapter teaming on the host side is neither tested nor supported in HPCC.

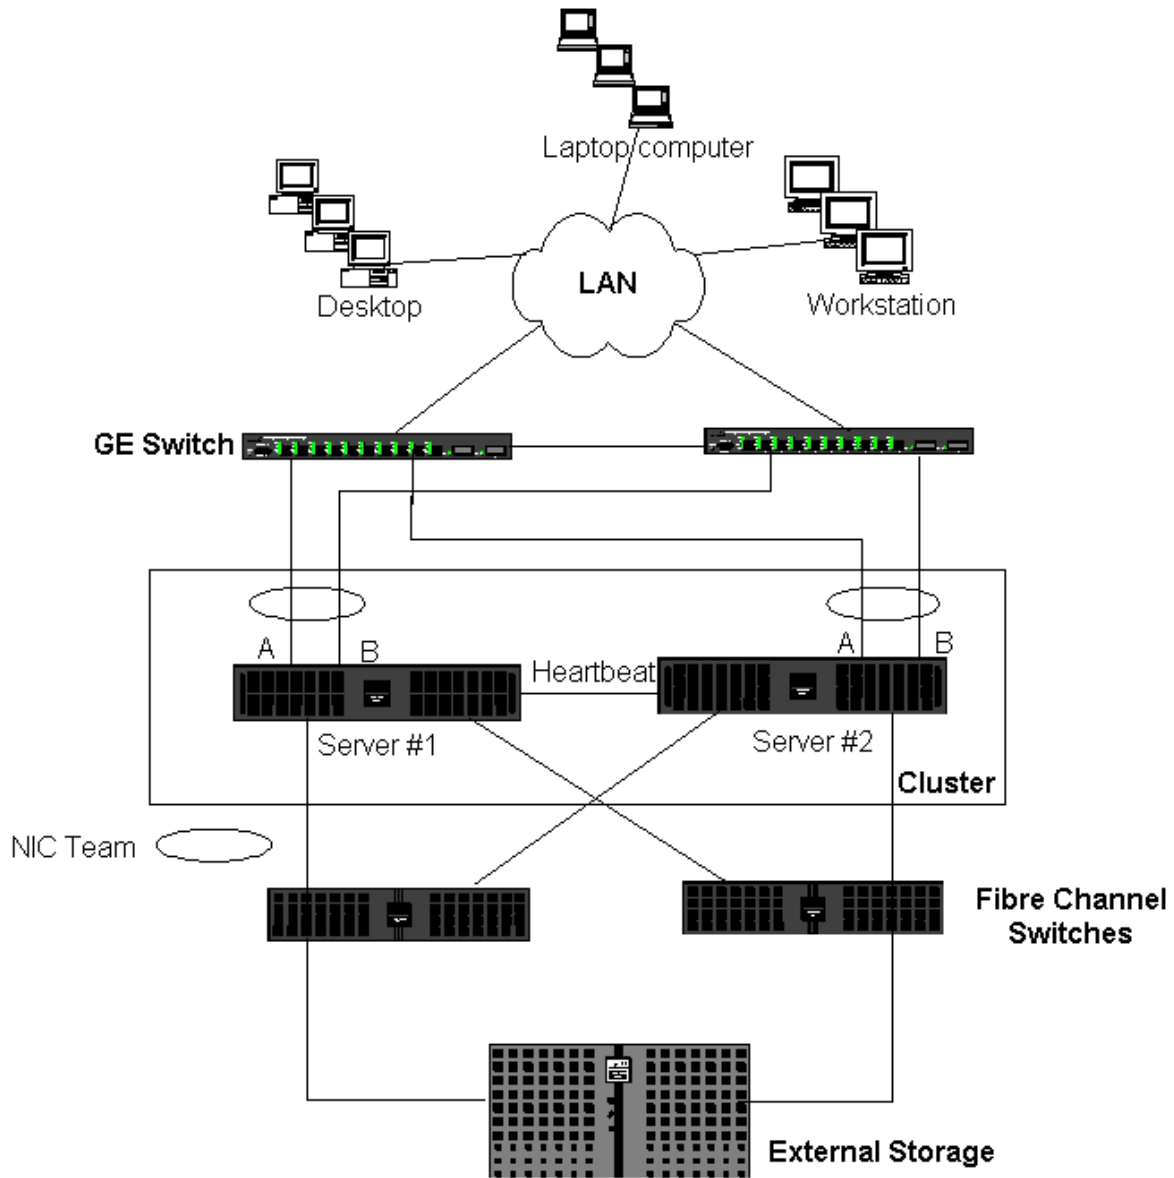
Advanced Features

PXE is used extensively for the deployment of the cluster (installation and recovery of compute nodes). Teaming is typically not used on the host side and it is not a part of our standard offering. Link aggregation is commonly used between switches, especially for large configurations. Jumbo frames, although not a part of our standard offering, may provide performance improvement for some applications due to reduced CPU overhead.

Oracle

In our Oracle Solution Stacks, we support adapter teaming in both the private network (interconnect between RAC nodes) and public network with clients or the application layer above the database layer.

Figure 8: Clustering With Teaming Across Two Switches

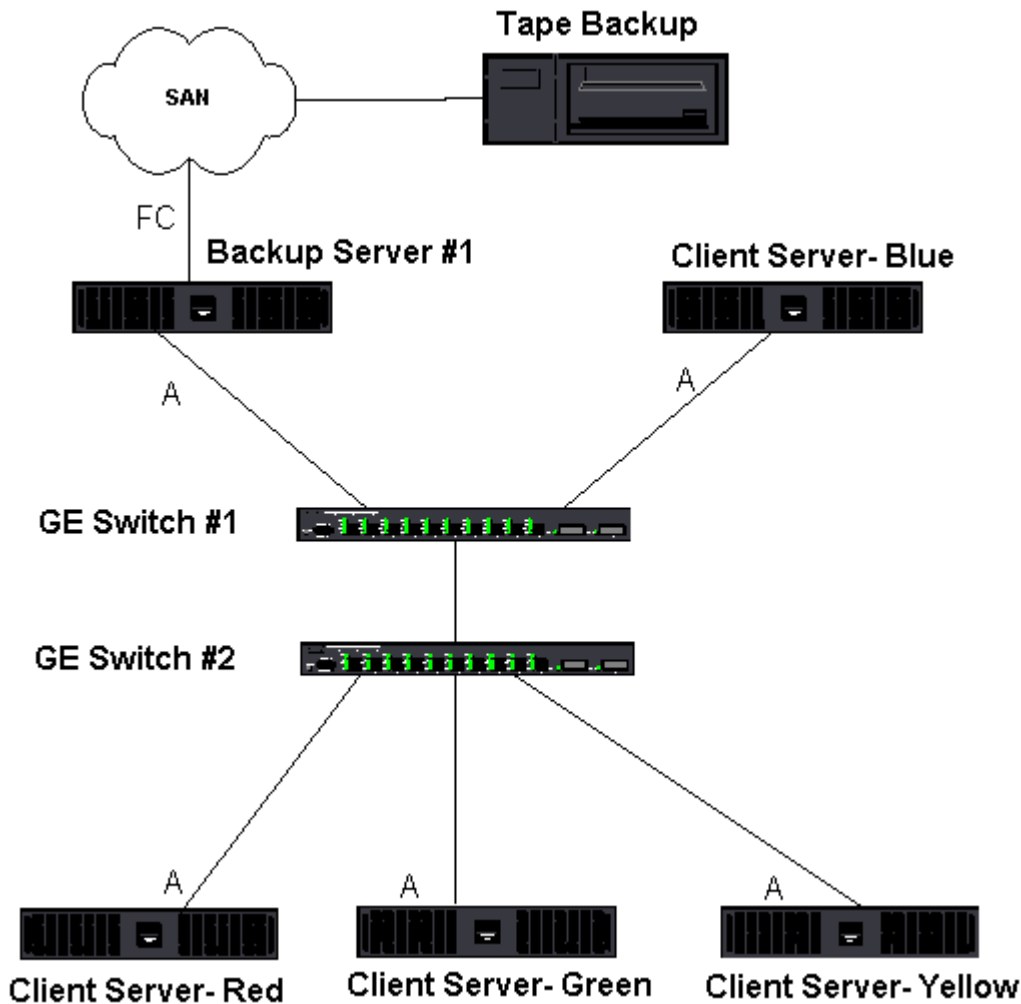


TEAMING AND NETWORK BACKUP

- [Load Balancing and Failover](#)
- [Fault Tolerance](#)

When you perform network backups in a nonteamed environment, overall throughput on a backup server adapter can be easily impacted due to excessive traffic and adapter overloading. Depending on the number of backup servers, data streams, and tape drive speed, backup traffic can easily consume a high percentage of the network link bandwidth, thus impacting production data and tape backup performance. Network backups usually consist of a dedicated backup server running with tape backup software such as NetBackup, Galaxy or Backup Exec. Attached to the backup server is either a direct SCSI tape backup unit or a tape library connected through a fiber channel storage area network (SAN). Systems that are backed up over the network are typically called clients or remote servers and usually have a tape backup software agent installed. [Figure 9](#) shows a typical 1 Gbps nonteamed network environment with tape backup implementation.

Figure 9: Network Backup without Teaming



Because there are four client servers, the backup server can simultaneously stream four backup jobs (one per client) to a multidrive autoloader. Because of the single link between the switch and the backup server; however, a 4-stream backup can easily saturate the adapter and link. If the adapter on the backup server operates at 1 Gbps (125 MB/s), and each client is able to stream data at 20 MB/s during tape backup, the throughput between the backup server and switch will be at 80 MB/s (20 MB/s x 4), which is equivalent to 64% of the network bandwidth. Although this is well within the network bandwidth range, the 64% constitutes a high percentage, especially if other applications share the same link.

Load Balancing and Failover

As the number of backup streams increases, the overall throughput increases. Each data stream, however, may not be able to maintain the same performance as a single backup stream of 25 MB/s. In other words, even though a backup server can stream data from a single client at 25 MB/s, it is not expected that four simultaneously-running backup jobs will stream at 100 MB/s (25 MB/s x 4 streams). Although overall throughput increases as the number of backup streams increases, each backup stream can be impacted by tape software or network stack limitations.

For a tape backup server to reliably use adapter performance and network bandwidth when backing up clients, a network infrastructure must implement teaming such as load balancing and fault tolerance. Data centers will incorporate redundant switches, link aggregation, and trunking as part of their fault tolerant solution. Although teaming device drivers will manipulate the way data flows through teamed interfaces and failover paths, this is transparent to tape backup applications and does not interrupt any tape backup process when backing up remote systems over the network. [Figure 10](#) shows a network topology that demonstrates tape backup in a Broadcom teamed environment and how smart load balancing can *load balance* tape backup data across teamed adapters.

There are four paths that the client-server can use to send data to the backup server, but only one of these paths will be designated during data transfer. One possible path that Client-Server Red can use to send data to the backup server is:

Example Path: Client-Server Red sends data through Adapter A, Switch 1, Backup Server Adapter A.

The designated path is determined by two factors:

- Client-Server ARP cache; which points to the backup server MAC address. This is determined by the Broadcom intermediate driver inbound load balancing algorithm.
- The physical adapter interface on Client-Server Red will be used to transmit the data. The Broadcom intermediate driver outbound load balancing algorithm determines this (see [Outbound Traffic Flow](#) and [Inbound Traffic Flow \(SLB Only\)](#)).

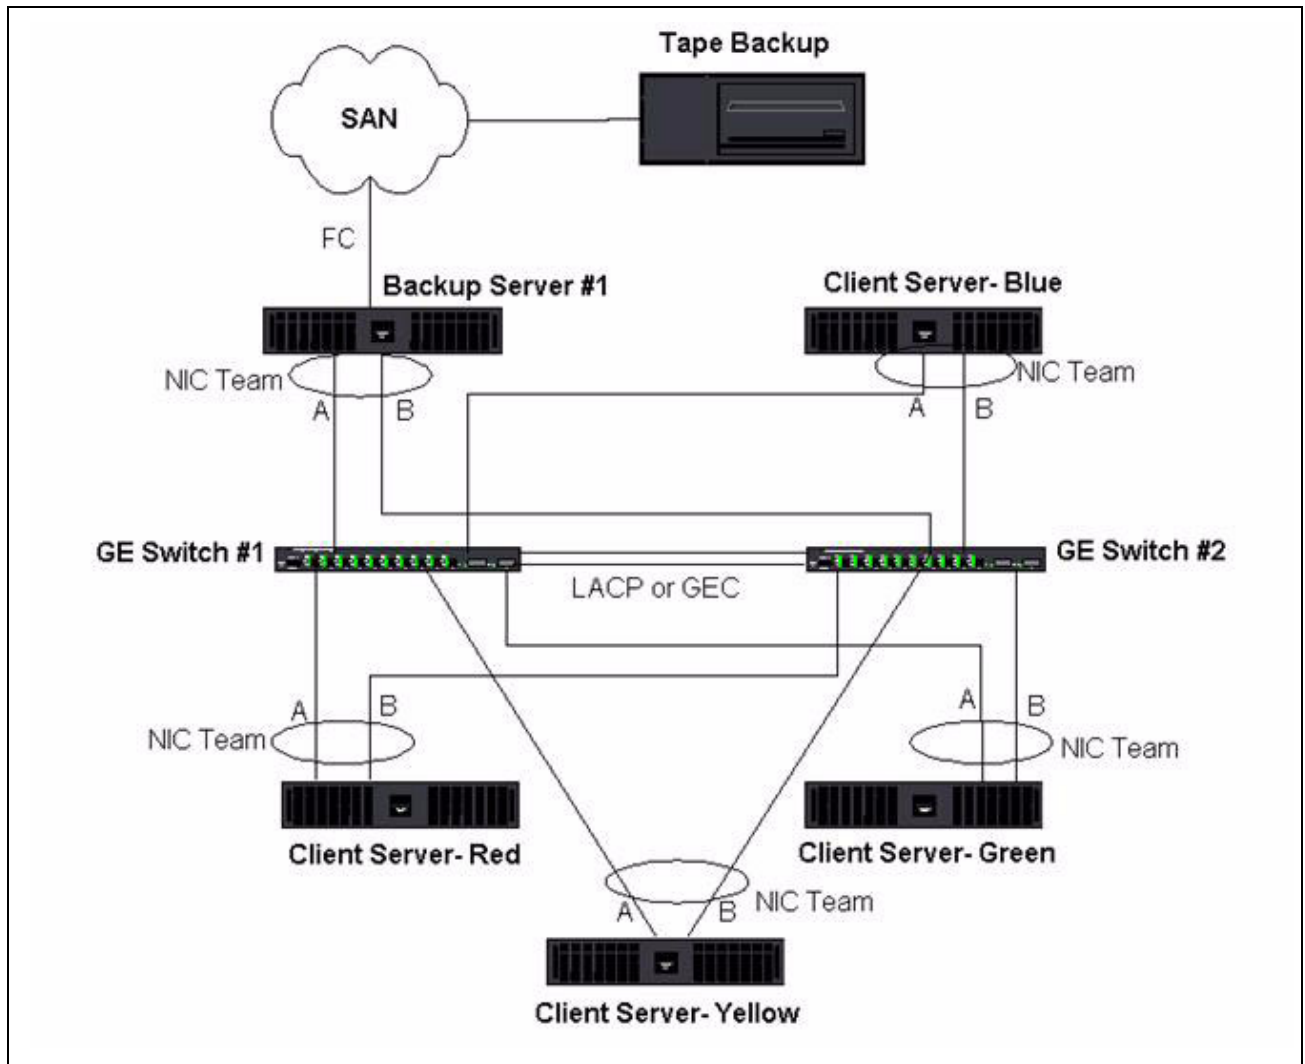
The teamed interface on the backup server transmits a gratuitous address resolution protocol (G-ARP) to Client-Server Red, which in turn, causes the client server ARP cache to get updated with the Backup Server MAC address. The load balancing mechanism within the teamed interface determines the MAC address embedded in the G-ARP. The selected MAC address is essentially the destination for data transfer from the client server. On Client-Server Red, the SLB teaming algorithm will determine which of the two adapter interfaces will be used to transmit data. In this example, data from Client Server Red is received on the backup server Adapter A interface. To demonstrate the SLB mechanisms when additional load is placed on the teamed interface, consider the scenario when the backup server initiates a second backup operation: one to Client-Server Red, and one to Client-Server Blue. The route that Client-Server Blue uses to send data to the backup server is dependant on its ARP cache, which points to the backup server MAC address. Because Adapter A of the backup server is already under load from its backup operation with Client-Sever Red, the Backup Server invokes its SLB algorithm to *inform* Client-Server Blue (through an G-ARP) to update its ARP cache to reflect the backup server Adapter B MAC address. When Client-Server Blue needs to transmit data, it uses either one of its adapter interfaces, which is determined by its own SLB algorithm. What is important is that data from Client-Server Blue is received by the Backup Server Adapter B interface, and not by its Adapter A interface. This is important because with both backup streams running simultaneously, the backup server must *load balance* data streams from different clients. With both backup streams running, each adapter interface on the backup server is processing an equal load, thus load-balancing data across both adapter interfaces.

The same algorithm applies if a third and fourth backup operation is initiated from the backup server. The teamed interface on the backup server transmits a unicast G-ARP to backup clients to inform them to update their ARP cache. Each client then transmits backup data along a route to the target MAC address on the backup server.

Fault Tolerance

If a network link fails during tape backup operations, all traffic between the backup server and client stops and backup jobs fail. If, however, the network topology was configured for both Broadcom SLB and switch fault tolerance, then this would allow tape backup operations to continue without interruption during the link failure. All failover processes within the network are transparent to tape backup software applications. To understand how backup data streams are directed during network failover process, consider the topology in [Figure 10](#). Client-Server Red is transmitting data to the backup server through Path 1, but a link failure occurs between the backup server and the switch. Because the data can no longer be sent from Switch #1 to the Adapter A interface on the backup server, the data is redirected from Switch #1 through Switch #2, to the Adapter B interface on the backup server. This occurs without the knowledge of the backup application because all fault tolerant operations are handled by the adapter team interface and trunk settings on the switches. From the client server perspective, it still operates as if it is transmitting data through the original path.

Figure 10: Network Backup With SLB Teaming Across Two Switches



TROUBLESHOOTING TEAMING PROBLEMS

- [Teaming Configuration Tips](#)
- [Troubleshooting Guidelines](#)

When running a protocol analyzer over a virtual adapter teamed interface, the MAC address shown in the transmitted frames may not be correct. The analyzer does not show the frames as constructed by BASP and shows the MAC address of the team and not the MAC address of the interface transmitting the frame. It is suggested to use the following process to monitor a team:

- Mirror all uplink ports from the team at the switch.
- If the team spans two switches, mirror the interlink trunk as well.
- Sample all mirror ports independently.
- On the analyzer, use an adapter and driver that does not filter QoS and VLAN information.

TEAMING CONFIGURATION TIPS

When troubleshooting network connectivity or teaming functionality issues, ensure that the following information is true for your configuration.

1. Although mixed-speed SLB teaming is supported, it is recommended that all adapters in a team be the same speed (either all Gigabit Ethernet or all Fast Ethernet). For speeds of 10 Gbps, it is highly recommended that all adapters in a team be the same speed.
2. If LiveLink is not enabled, disable Spanning Tree Protocol or enable an STP mode that bypasses the initial phases (for example, Port Fast, Edge Port) for the switch ports connected to a team.
3. All switches that the team is directly connected to must have the same hardware revision, firmware revision, and software revision to be supported.
4. To be teamed, adapters should be members of the same VLAN. In the event that multiple teams are configured, each team should be on a separate network.
5. Do not assign a Locally Administered Address on any physical adapter that is a member of a team.
6. Verify that power management is disabled on all physical members of any team.
7. Remove any static IP address from the individual physical team members before the team is built.
8. A team that requires maximum throughput should use LACP or GEC\FEC. In these cases, the intermediate driver is only responsible for the outbound load balancing while the switch performs the inbound load balancing.
9. Aggregated teams (802.3ad \ LACP and GEC\FEC) must be connected to only a single switch that supports IEEE 802.3a, LACP or GEC/FEC.
10. It is not recommended to connect any team to a hub, as a hub only support half duplex. Hubs should be connected to a team for troubleshooting purposes only. Disabling the device driver of a network adapter participating in an LACP or GEC/FEC team may have adverse affects with network connectivity. Broadcom recommends that the adapter first be physically disconnected from the switch before disabling the device driver in order to avoid a network connection loss.
11. Verify the base (Miniport) and team (intermediate) drivers are from the same release package. The mixing of base and teaming drivers from different releases is not supported.
12. Test the connectivity to each physical adapter prior to teaming.
13. Test the failover and fallback behavior of the team before placing into a production environment.
14. When moving from a nonproduction network to a production network, it is strongly recommended to test again for failover

and fallback.

15. Test the performance behavior of the team before placing into a production environment.
16. Network teaming is not supported when running iSCSI traffic via Microsoft iSCSI initiator or iSCSI offload. MPIO should be used instead of Broadcom network teaming for these ports.
17. For information on iSCSI boot and iSCSI offload restrictions, see [iSCSI Protocol](#).

TROUBLESHOOTING GUIDELINES

Before you call for support, make sure you have completed the following steps for troubleshooting network connectivity problems when the server is using adapter teaming.

1. Make sure the link light is ON for every adapter and all the cables are attached.
2. Check that the matching base and intermediate drivers belong to the same release and are loaded correctly.
3. Check for a valid IP Address using the Windows ipconfig command.
4. Check that STP is disabled or Edge Port/Port Fast is enabled on the switch ports connected to the team or that LiveLink is being used.
5. Check that the adapters and the switch are configured identically for link speed and duplex.
6. If possible, break the team and check for connectivity to each adapter independently to confirm that the problem is directly associated with teaming.
7. Check that all switch ports connected to the team are on the same VLAN.
8. Check that the switch ports are configured properly for Generic Trunking (FEC/GEC)/802.3ad-Draft Static type of teaming and that it matches the adapter teaming type. If the system is configured for an SLB type of team, make sure the corresponding switch ports *are not* configured for Generic Trunking (FEC/GEC)/802.3ad-Draft Static types of teams.

FREQUENTLY ASKED QUESTIONS

Question: Under what circumstances is traffic not load balanced? Why is all traffic not load balanced evenly across the team members?

Answer: The bulk of traffic does not use IP/TCP/UDP or the bulk of the clients are in a different network. The receive load balancing is not a function of traffic load, but a function of the number of clients that are connected to the server.

Question: What network protocols are load balanced when in a team?

Answer: Broadcom's teaming software only supports IP/TCP/UDP traffic. All other traffic is forwarded to the primary adapter.

Question: Which protocols are load balanced with SLB and which ones are not?

Answer: Only IP/TCP/UDP protocols are load balanced in both directions: send and receive. IPX is load balanced on the transmit traffic only.

Question: Can I team a port running at 100 Mbps with a port running at 1000 Mbps?

Answer: Mixing link speeds within a team is only supported for Smart Load Balancing™ teams and 802.3ad teams.

Question: Can I team a fiber adapter with a copper Gigabit Ethernet adapter?

Answer: Yes with SLB, and yes if the switch allows for it in FEC/GEC and 802.3ad.

Question: What is the difference between adapter load balancing and Microsoft's Network Load Balancing (NLB)?

Answer: Adapter load balancing is done at a network session level, whereas NLB is done at the server application level.



Question: Can I connect the teamed adapters to a hub?

Answer: Teamed ports can be connected to a hub for troubleshooting purposes only. However, this practice is not recommended for normal operation because the performance would be degraded due to hub limitations. Connect the teamed ports to a switch instead.

Question: Can I connect the teamed adapters to ports in a router?

Answer: No. All ports in a team must be on the same network; in a router, however, each port is a separate network by definition. All teaming modes require that the link partner be a Layer 2 switch.

Question: Can I use teaming with Microsoft Cluster Services?

Answer: Yes. Teaming is supported on the public network only, not on the private network used for the heartbeat link.

Question: Can PXE work over a virtual adapter (team)?

Answer: A PXE client operates in an environment before the operating system is loaded; as a result, virtual adapters have not been enabled yet. If the physical adapter supports PXE, then it can be used as a PXE client, whether or not it is part of a virtual adapter when the operating system loads. PXE servers may operate over a virtual adapter.

Question: Can WOL work over a virtual adapter (team)?

Answer: Wake-on-LAN functionality operates in an environment before the operating system is loaded. WOL occurs when the system is off or in standby, so no team is configured.

Question: What is the maximum number of ports that can be teamed together?

Answer: Up to eight ports can be assigned to a team.

Question: What is the maximum number of teams that can be configured on the same server?

Answer: Up to eight teams can be configured on the same server.

Question: Why does my team lose connectivity for the first 30 to 50 seconds after the Primary adapter is restored (fallback)?

Answer: Because Spanning Tree Protocol is bringing the port from blocking to forwarding. You must enable Port Fast or Edge Port on the switch ports connected to the team or use LiveLink to account for the STP delay.

Question: Can I connect a team across multiple switches?

Answer: Smart Load Balancing can be used with multiple switches because each physical adapter in the system uses a unique Ethernet MAC address. Link Aggregation and Generic Trunking cannot operate across switches because they require all physical adapters to share the same Ethernet MAC address.

Question: How do I upgrade the intermediate driver (BASP)?

Answer: The intermediate driver cannot be upgraded through the Local Area Connection Properties. It must be upgraded using the Setup installer.

Question: How can I determine the performance statistics on a virtual adapter (team)?

Answer: In Broadcom Advanced Control Suite 3, click the Statistics tab for the virtual adapter.

Question: Can I configure NLB and teaming concurrently?

Answer: Yes, but only when running NLB in a multicast mode (NLB is not supported with MS Cluster Services).

Question: Should both the backup server and client servers that are backed up be teamed?

Answer: Because the backup server is under the most data load, it should always be teamed for link aggregation and failover. A fully redundant network, however, requires that both the switches and the backup clients be teamed for fault tolerance and link aggregation.

Question: During backup operations, does the adapter teaming algorithm load balance data at a byte-level or a session-level?

Answer: When using adapter teaming, data is only load balanced at a session level and not a byte level to prevent out-of-order frames. Adapter teaming load balancing does not work the same way as other storage load balancing mechanisms such as EMC PowerPath.

Question: Is there any special configuration required in the tape backup software or hardware to work with adapter teaming?

Answer: No special configuration is required in the tape software to work with teaming. Teaming is transparent to tape backup applications.

Question: How do I know what driver I am currently using?

Answer: In all operating systems, the most accurate method for checking the driver revision is to physically locate the driver file and check the properties.

Question: Can SLB detect a switch failure in a Switch Fault Tolerance configuration?

Answer: No. SLB can only detect the loss of link between the teamed port and its immediate link partner. SLB cannot detect link failures on other ports.

Question: Why does my team lose connectivity for the first 30 to 50 seconds after the primary adapter is restored (fall-back after a failover)?

Answer: During a fall-back event, link is restored causing Spanning Tree Protocol to configure the port for blocking until it determines that it can move to the forwarding state. You must enable Port Fast or Edge Port on the switch ports connected to the team to prevent the loss of communications caused by STP.

Question: Where do I monitor real time statistics for an adapter team in a Windows server?

Answer: Use the Broadcom Advanced Control Suite 3 (BACS) to monitor general, IEEE 802.3 and custom counters.

Question: What features are not supported on a multivendor team?

Answer: TOE, VLAN tagging, and RSS are not supported on a multivendor team.

APPENDIX A: EVENT LOG MESSAGES

- [Windows System Event Log messages](#)
- [Base Driver \(Physical Adapter/Miniport\)](#)
- [Intermediate Driver \(Virtual Adapter/Team\)](#)

WINDOWS SYSTEM EVENT LOG MESSAGES

The known base and intermediate Windows System Event Log status messages for the Broadcom NetXtreme II adapters are listed in [Table 8](#) and [Table 9](#). As a Broadcom adapter driver loads, Windows places a status code in the system event viewer. There may be up to two classes of entries for these event codes depending on whether both drivers are loaded (one set for the base or miniport driver and one set for the intermediate or teaming driver).

BASE DRIVER (PHYSICAL ADAPTER/MINIPOINT)

The base driver is identified by source **L2ND**. [Table 8](#) lists the event log messages supported by the base driver, explains the cause for the message, and provides the recommended action.



Note: In [Table 8](#), message numbers 1 through 17 apply to both NDIS 5.x and NDIS 6.x drivers, message numbers 18 through 23 apply only to the NDIS 6.x driver.

Table 8: Base Driver Event Log Messages

<i>Message Number</i>	<i>Severity</i>	<i>Message</i>	<i>Cause</i>	<i>Corrective Action</i>
1	Error	Failed to allocate memory for the device block. Check system memory resource usage.	The driver cannot allocate memory from the operating system.	Close running applications to free memory.
2	Error	Failed to allocate map registers.	The driver cannot allocate map registers from the operating system.	Unload other drivers that may allocate map registers.
3	Error	Failed to access configuration information. Reinstall the network driver.	The driver cannot access PCI configuration space registers on the adapter.	For add-in adapters: reseat the adapter in the slot, move the adapter to another PCI slot, or replace the adapter.
4	Warning	The network link is down. Check to make sure the network cable is properly connected.	The adapter has lost its connection with its link partner.	Check that the network cable is connected, verify that the network cable is the right type, and verify that the link partner (for example, switch or hub) is working correctly.
5	Informational	The network link is up.	The adapter has established a link.	No action is required.
6	Informational	Network controller configured for 10Mb half-duplex link.	The adapter has been manually configured for the selected line speed and duplex settings.	No action is required.

Table 8: Base Driver Event Log Messages (Cont.)

Message Number	Severity	Message	Cause	Corrective Action
7	Informational	Network controller configured for 10Mb full-duplex link.	The adapter has been manually configured for the selected line speed and duplex settings.	No action is required.
8	Informational	Network controller configured for 100Mb half-duplex link.	The adapter has been manually configured for the selected line speed and duplex settings.	No action is required.
9	Informational	Network controller configured for 100Mb full-duplex link.	The adapter has been manually configured for the selected line speed and duplex settings.	No action is required.
10	Informational	Network controller configured for 1Gb half-duplex link.	The adapter has been manually configured for the selected line speed and duplex settings.	No action is required.
11	Informational	Network controller configured for 1Gb full-duplex link.	The adapter has been manually configured for the selected line speed and duplex settings.	No action is required.
12	Informational	Network controller configured for 2.5Gb full-duplex link.	The adapter has been manually configured for the selected line speed and duplex settings.	No action is required.
13	Error	Medium not supported.	The operating system does not support the IEEE 802.3 medium.	Reboot the operating system, run a virus check, run a disk check (chkdsk), and reinstall the operating system.
14	Error	Unable to register the interrupt service routine.	The device driver cannot install the interrupt handler.	Reboot the operating system; remove other device drivers that may be sharing the same IRQ.
15	Error	Unable to map IO space.	The device driver cannot allocate memory-mapped I/O to access driver registers.	Remove other adapters from the system, reduce the amount of physical memory installed, and replace the adapter.
16	Informational	Driver initialized successfully.	The driver has successfully loaded.	No action is required.
17	Informational	NDIS is resetting the miniport driver.	The NDIS layer has detected a problem sending/receiving packets and is resetting the driver to resolve the problem.	Run Broadcom Advanced Control Suite 3 diagnostics; check that the network cable is good.
18	Error	Unknown PHY detected. Using a default PHY initialization routine.	The driver could not read the PHY ID.	Replace the adapter.
19	Error	This driver does not support this device. Upgrade to the latest driver.	The driver does not recognize the installed adapter.	Upgrade to a driver version that supports this adapter.



Table 8: Base Driver Event Log Messages (Cont.)

Message Number	Severity	Message	Cause	Corrective Action
20	Error	Driver initialization failed.	Unspecified failure during driver initialization.	Reinstall the driver, update to a newer driver, run Broadcom Advanced Control Suite 3 diagnostics, or replace the adapter.
21	Informational	Network controller configured for 10Gb full-duplex link.	The adapter has been manually configured for the selected line speed and duplex settings.	No action is required.
22	Error	Network controller failed initialization because it cannot allocate system memory.	Insufficient system memory prevented the initialization of the driver.	Increase system memory.
23	Error	Network controller failed to exchange the interface with the bus driver.	The driver and the bus driver are not compatible.	Update to the latest driver set, ensuring the major and minor versions for both NDIS and the bus driver are the same.

INTERMEDIATE DRIVER (VIRTUAL ADAPTER/TEAM)

The intermediate driver is identified by source **BLFM**, regardless of the base driver revision. [Table 9](#) lists the event log messages supported by the intermediate driver, explains the cause for the message, and provides the recommended action.

Table 9: Intermediate Driver Event Log Messages

System Event Message Number	Severity	Message	Cause	Corrective Action
1	Informational	Event logging enabled for Broadcom Advanced Server Program driver.	–	No action is required.
2	Error	Unable to register with NDIS.	The driver cannot register with the NDIS interface.	Unload other NDIS drivers.
3	Error	Unable to instantiate the management interface.	The driver cannot create a device instance.	Reboot the operating system.
4	Error	Unable to create symbolic link for the management interface.	Another driver has created a conflicting device name.	Unload the conflicting device driver that uses the name <i>BIf</i> .
5	Informational	Broadcom Advanced Server Program Driver has started.	The driver has started.	No action is required.
6	Informational	Broadcom Advanced Server Program Driver has stopped.	The driver has stopped.	No action is required.
7	Error	Could not allocate memory for internal data structures.	The driver cannot allocate memory from the operating system.	Close running applications to free memory.
8	Warning	Could not bind to adapter.	The driver could not open one of the team physical adapters.	Unload and reload the physical adapter driver, install an updated physical adapter driver, or replace the physical adapter.
9	Informational	Successfully bind to adapter.	The driver successfully opened the physical adapter.	No action is required.
10	Warning	Network adapter is disconnected.	The physical adapter is not connected to the network (it has not established link).	Check that the network cable is connected, verify that the network cable is the right type, and verify that the link partner (switch or hub) is working correctly.
11	Informational	Network adapter is connected.	The physical adapter is connected to the network (it has established link).	No action is required.
12	Error	Broadcom Advanced Program Features Driver is not designed to run on this version of Operating System.	The driver does not support the operating system on which it is installed.	Consult the driver release notes and install the driver on a supported operating system or update the driver.
13	Informational	Hot-standby adapter is selected as the primary adapter for a team without a load balancing adapter.	A standby adapter has been activated.	Replace the failed physical adapter.



Table 9: Intermediate Driver Event Log Messages (Cont.)

System Event Message Number	Severity	Message	Cause	Corrective Action
14	Informational	Network adapter does not support Advanced Failover.	The physical adapter does not support the Broadcom NIC Extension (NICE).	Replace the adapter with one that does support NICE.
15	Informational	Network adapter is enabled via management interface.	The driver has successfully enabled a physical adapter through the management interface.	No action is required.
16	Warning	Network adapter is disabled via management interface.	The driver has successfully disabled a physical adapter through the management interface.	No action is required.
17	Informational	Network adapter is activated and is participating in network traffic.	A physical adapter has been added to or activated in a team.	No action is required.
18	Informational	Network adapter is deactivated and is no longer participating in network traffic.	The driver does not recognize the installed adapter.	No action is required.
19	Informational	The LiveLink feature in BASP connected the link for the network adapter.	The connection with the remote target(s) for the LiveLink-enabled team member has been established or restored	No action is required.
20	Informational	The LiveLink feature in BASP disconnected the link for the network adapter.	The LiveLink-enabled team member is unable to connect with the remote target(s).	No action is required.

Virtual LANs: Broadcom NetXtreme II™ Network Adapter User Guide

- [VLAN Overview](#)
- [Adding VLANs to Teams](#)

VLAN OVERVIEW

Virtual LANs (VLANs) allow you to split your physical LAN into logical parts, to create logical segmentation of workgroups, and to enforce security policies for each logical segment. Each defined VLAN behaves as its own separate network with its traffic and broadcasts isolated from the others, increasing bandwidth efficiency within each logical group. Up to 64 VLANs (63 tagged and 1 untagged) can be defined for each Broadcom adapter on your server, depending on the amount of memory available in your system.

VLANs can be added to a team to allow multiple VLANs with different VLAN IDs. A virtual adapter is created for each VLAN added.

Although VLANs are commonly used to create individual broadcast domains and/or separate IP subnets, it is sometimes useful for a server to have a presence on more than one VLAN simultaneously. Broadcom adapters support multiple VLANs on a per-port or per-team basis, allowing very flexible network configurations.

Figure 1: Example of Servers Supporting Multiple VLANs with Tagging




Figure 1 shows an example network that uses VLANs. In this example network, the physical LAN consists of a switch, two servers, and five clients. The LAN is logically organized into three different VLANs, each representing a different IP subnet. The features of this network are described in Table 1.

Table 1: Example VLAN Network Topology

Component	Description
VLAN #1	An IP subnet consisting of the Main Server, PC #3, and PC #5. This subnet represents an engineering group.
VLAN #2	Includes the Main Server, PCs #1 and #2 via shared media segment, and PC #5. This VLAN is a software development group.
VLAN #3	Includes the Main Server, the Accounting Server and PC #4. This VLAN is an accounting group.
Main Server	A high-use server that needs to be accessed from all VLANs and IP subnets. The Main Server has a Broadcom adapter installed. All three IP subnets are accessed via the single physical adapter interface. The server is attached to one of the switch ports, which is configured for VLANs #1, #2, and #3. Both the adapter and the connected switch port have tagging turned on. Because of the tagging VLAN capabilities of both devices, the server is able to communicate on all three IP subnets in this network, but continues to maintain broadcast separation between all of them.
Accounting Server	Available to VLAN #3 only. The Accounting Server is isolated from all traffic on VLANs #1 and #2. The switch port connected to the server has tagging turned off.
PCs #1 and #2	Attached to a shared media hub that is then connected to the switch. PCs #1 and #2 belong to VLAN #2 only, and are logically in the same IP subnet as the Main Server and PC #5. The switch port connected to this segment has tagging turned off.
PC #3	A member of VLAN #1, PC #3 can communicate only with the Main Server and PC #5. Tagging is not enabled on PC #3 switch port.
PC #4	A member of VLAN #3, PC #4 can only communicate with the servers. Tagging is not enabled on PC #4 switch port.
PC #5	A member of both VLANs #1 and #2, PC #5 has a Broadcom adapter installed. It is connected to switch port #10. Both the adapter and the switch port are configured for VLANs #1 and #2 and have tagging enabled.



NOTE: VLAN tagging is only required to be enabled on switch ports that create trunk links to other switches, or on ports connected to tag-capable end-stations, such as servers or workstations with Broadcom adapters.

ADDING VLANs TO TEAMS

Each team supports up to 64 VLANs (63 tagged and 1 untagged). Note that only Broadcom adapters and Alteon® AceNIC adapters can be part of a team with VLANs. With multiple VLANs on an adapter, a server with a single adapter can have a logical presence on multiple IP subnets. With multiple VLANs in a team, a server can have a logical presence on multiple IP subnets and benefit from load balancing and failover. For instructions on adding a VLAN to a team, see [Adding a VLAN](#) for Windows operating systems.



NOTE: Adapters that are members of a failover team can also be configured to support VLANs. Because VLANs are not supported for an Intel LOM, if an Intel LOM is a member of a failover team, VLANs cannot be configured for that team.

Manageability: Broadcom NetXtreme II™ Network Adapter User Guide

- CIM
- SNMP

CIM

The Common Information Model (CIM) is an industry standard defined by the Distributed Management Task Force (DMTF). Microsoft implements CIM on Windows server platforms. Broadcom will support CIM on Windows server platforms.

Broadcom's implementation of CIM will provide various classes to provide information to users through CIM client applications. Note that Broadcom CIM data provider will provide data only, and users can choose their preferred CIM client software to browse the information exposed by Broadcom CIM provider.

Broadcom CIM provider provides information through `BRCM_NetworkAdapter` and `BRCM_ExtraCapacityGroup` classes. `BRCM_NetworkAdapter` class provides network adapter information pertaining to a group of adapters including Broadcom and other vendors' controllers. `BRCM_ExtraCapacityGroup` class provides team configuration for the Broadcom Advanced Server Program. Current implementation will provide team information and information of physical network adapters in the team.

Broadcom Advanced Server Program provides events through event logs. Users can use the "Event Viewer" provided by Windows server platforms, or use CIM to inspect or monitor these events. Broadcom CIM provider will also provide event information through the CIM generic event model. These events are `__InstanceCreationEvent`, `__InstanceDeletionEvent` and `__InstanceModificationEvent`, and are defined by CIM. CIM requires the client application to register the events from the client application, using queries as examples shown below in order to receive events properly.

```
SELECT * FROM __InstanceModificationEvent
where TargetInstance ISA "BRCM_NetworkAdapter"
SELECT * FROM __InstanceModificationEvent
where TargetInstance ISA "BRCM_ExtraCapacityGroup"
SELECT * FROM __InstanceCreationEvent
where TargetInstance ISA "BRCM_NetworkAdapter"
SELECT * FROM __InstanceDeletionEvent
where TargetInstance ISA "BRCM_NetworkAdapter"
SELECT * FROM __InstanceCreationEvent
where TargetInstance ISA "BRCM_ActsAsSpare"
SELECT * FROM __InstanceDeletionEvent
where TargetInstance ISA "BRCM_ActsAsSpare"
```

For detailed information about these events, see the CIM documentation at http://www.dmtf.org/standards/published_documents/DSP0004V2.3_final.pdf.

SNMP

BASP SUBAGENT

The BASP subagent, `baspmgmt.dll`, is designed for the Windows 2000 Server and Windows Server 2003 SNMP service. It is required to install the SNMP service before installing the BASP subagent.

The BASP subagent allows an SNMP manager software to actively monitor the configurations and performance of the Broadcom Advanced Server features. The subagent also provides an alarm trap to an SNMP manager to inform the manager of any changes to the conditions of the BASP component.

The BASP subagent allows monitoring of the configurations and statistics for the BASP teams, the physical NIC adapters participating in a team, and the virtual NIC adapters created as the result of teaming. Non-teamed NIC adapters are not monitored at this time. The BASP configuration data includes information such as team IDs, physical/virtual/VLAN/team adapter IDs, physical/virtual/VLAN/team/ adapter descriptions, and MAC addresses of the adapters.

The statistics include detailed information such as data packets transmitted and received for the physical/virtual/VLAN/team adapters.

The alarm trap forwards information about the changes in configuration of the physical adapters participating in a team, such as physical adapter link up/down, and adapter installed/removed events.

To monitor this information, an SNMP manager must load the Broadcom BASP MIB database files to allow monitoring of the information described above. These files, which are shown below, are included with the driver source media.

- `baspcfg.mib`
- `baspcfg.mib`
- `baspcfg.mib`

BASP EXTENSIBLE-AGENT

The Broadcom NetXtreme II Gigabit Ethernet Controller Extended Information SNMP extensible-agent (bcmif.dll) is designed for Windows 2000 Server, Windows Server 2003, and Windows Server 2008 SNMP service. It is required that Windows 2000 Server SNMP service is installed before installing the extensible-agent.

The extensible-agent allows the SNMP manager software to actively monitor the configurations of the Broadcom NetXtreme II adapter. It is intended to supplement the information already provided by the standard SNMP Management Network Interface information.

The extensible-agent provides in-depth information about a Broadcom NetXtreme II adapter such as:

- MAC address
- Bound IP address
- IP subnet mask
- Physical link status
- Adapter state
- Line speed
- Duplex mode
- Memory range
- Interrupt setting
- Bus number
- Device number
- Function number

To monitor this information, a SNMP manager needs to load the Broadcom Extended information MIB file to allow monitoring of the information described above. This file, bcmif.mib, is included on the installation CD.

The monitored workstation requires the installation of the Broadcom Extended Information SNMP extensible-agent, bcmif.dll, and requires the Microsoft Windows 2000 Server, Windows Server 2003, and Windows Server 2008 SNMP service to be installed and loaded.

Installing the Hardware: Broadcom NetXtreme II™ Network Adapter User Guide

- [Overview](#)
- [System Requirements](#)
- [Safety Precautions](#)
- [Preinstallation Checklist](#)
- [Installation of the Add-In NIC](#)

OVERVIEW

This section applies to Broadcom NetXtreme II add-in network interface cards.

SYSTEM REQUIREMENTS

Before you install a Broadcom NetXtreme II adapter, verify that your system meets the following hardware and operating system requirements:

HARDWARE REQUIREMENTS

- IA32- or EMT64-based computer that meets operating system requirements
- One open slot: PCI Express 1.0a x4 or PCI Express Gen2 x8
- 128-MB RAM (minimum)

OPERATING SYSTEM REQUIREMENTS

General

- PCI Express v1.0a, x4 (or greater) Host Interface

Microsoft Windows

One of the following versions of Microsoft Windows:

- Windows 2000 Server family
- Windows Server 2003 family
- Windows Server 2008 family
- Windows Server 2008 R2 family

Novell NetWare

Novell NetWare 6.5 with the latest support pack.

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.

VMware ESX

VMware ESX 3.5

SAFETY PRECAUTIONS



CAUTION! The adapter is being installed in a system that operates with voltages that can be lethal. Before you open the case of your system, observe the following precautions to protect yourself and to prevent damage to the system components.

- Remove any metallic objects or jewelry from your hands and wrists.
- Make sure to use only insulated or nonconducting tools.
- Verify that the system is powered OFF and is unplugged before you touch internal components.
- Install or remove adapters in a static-free environment. The use of a properly grounded wrist strap or other personal antistatic devices and an antistatic mat is strongly recommended.

PREINSTALLATION CHECKLIST

1. Verify that your system meets the hardware and software requirements listed under [System Requirements](#).
2. Verify that your system is using the latest BIOS.



NOTE: If you acquired the adapter software on a disk, verify the path to the adapter driver files.

1. If your system is active, shut it down.
2. When system shutdown is complete, turn off the power and unplug the power cord.
3. Remove the adapter from its shipping package and place it on an antistatic surface.
4. Check the adapter for visible signs of damage, particularly on the edge connector. Never attempt to install a damaged adapter.

INSTALLATION OF THE ADD-IN NIC

The following instructions apply to installing the Broadcom NetXtreme II adapter (add-in NIC) in most systems. Refer to the manuals that were supplied with your system for details about performing these tasks on your particular system.

INSTALLING THE ADD-IN NIC

1. Review [Safety Precautions](#) and [Preinstallation Checklist](#). Before you install the adapter, ensure that the system power is OFF, the power cord is unplugged from the power outlet, and that you are following proper electrical grounding procedures.
2. Open the system case and select the slot based on the adapter: PCI Express 1.0a x4, PCI Express Gen2 x8, or other appropriate slot. A lesser width adapter can be seated into a greater width slot (x1 in a x4), but a greater width adapter cannot be seated into a lesser width slot (x4 in a x1). If you do not know how to identify a PCI Express slot, refer to your system documentation.
3. Remove the blank cover-plate from the slot that you selected.
4. Align the adapter connector edge with the PCI Express connector slot in the system.
5. Applying even pressure at both corners of the card, push the adapter card into the slot until it is firmly seated. When the adapter is properly seated, the adapter port connectors are aligned with the slot opening, and the adapter faceplate is flush against the system chassis.



CAUTION! Do not use excessive force when seating the card, as this may damage the system or the adapter. If you have difficulty seating the adapter, remove it, realign it, and try again.

6. Secure the adapter with the adapter clip or screw.
7. Close the system case and disconnect any personal antistatic devices.

CONNECTING THE NETWORK CABLES

The Broadcom NetXtreme II adapter has either an RJ-45 connector used for attaching the system to an Ethernet copper-wire segment or a fiber optic connector for attaching the system to an Ethernet fiber optic segment.



NOTE: This section does not apply to blade servers.

Copper Wire



NOTE: The Broadcom NetXtreme II adapter 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 1: 10/100/1000BASE-T and 10GBASE-T Cable Specifications

Port Type	Connector	Media	Maximum Distance
10BASE-T	RJ-45	Category 3, 4, or 5 unshielded twisted pairs (UTP)	100m (328 ft)
100/1000BASE-T ¹	RJ-45	Category 5 ² UTP	100m (328 ft)
10GBASE-T	RJ-45	Category 6 ³ UTP	50m (164 ft)
		Category 6A ³ UTP	100m (328 ft)

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.

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.

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.

Fiber Optic

1. Select an appropriate cable. [Table 2](#) lists the fiber optic cable requirements for connecting to 1000/2500BASE-X ports:

Table 2: 1000/2500BASE-X Fiber Optic Specifications

Port Type	Connector	Media	Maximum Distance
1000BASE-X	Small form factor (SFF) transceiver with LC™ connection system (Infineon p/n V23818-K305-L57)	Multimode fiber (MMF) System optimized for 62.5/50 µm graded index fiber	550m (1804 ft)
2500BASE-X ¹	Small form factor (SFF) transceiver with LC™ connection system (Finisar p/n FTLF8542E2KNV)	Multimode fiber (MMF) System optimized for 62.5/50 µm graded index fiber	550m (1804 ft)

¹ Electricals leveraged from IEEE 802.3ae-2002 (XAUI). 2500BASE-X is term used by Broadcom to describe 2.5 Gbit/s (3.125GBd) operation. LC is a trademark of Lucent Technologies.

2. Connect one end of the cable to the fiber optic connector on the adapter.
3. Connect the other end of the cable to an fiber optic Ethernet network port.



Broadcom Boot Agent Driver Software: Broadcom NetXtreme II™ Network Adapter User Guide

- [Overview](#)
- [Setting Up MBA in a Client Environment](#)
- [Setting Up MBA in a Server Environment](#)

OVERVIEW

Broadcom NetXtreme II adapters support Preboot Execution Environment (PXE), Remote Program Load (RPL), iSCSI, and Bootstrap Protocol (BootP). Multi-Boot Agent (MBA) is a software module that allows your network computer to boot with the images provided by remote servers across the network. The Broadcom MBA driver complies with the PXE 2.1 specification and is released with both monolithic and split binary images. This provides flexibility to users in different environments where the motherboard may or may not have built-in base code.

The MBA module operates in a client/server environment. A network consists of one or more boot servers that provide boot images to multiple computers through the network. The Broadcom implementation of the MBA module has been tested successfully in the following environments:

- **Linux Red Hat PXE Server.** Broadcom PXE clients are able to remotely boot and use network resources (NFS mount, and so forth) and to perform Linux installations. In the case of a remote boot, the Linux universal driver binds seamlessly with the Broadcom Universal Network Driver Interface (UNDI) and provides a network interface in the Linux remotely-booted client environment.
- **Intel APITEST.** The Broadcom PXE driver passes all API compliance test suites.
- **MS-DOS UNDI.** The MS-DOS Universal Network Driver Interface (UNDI) seamlessly binds with the Broadcom UNDI to provide a network adapter driver interface specification (NDIS2) interface to the upper layer protocol stack. This allows computers to connect to network resources in an MS-DOS environment.
- **Remote Installation Service (RIS).** The Broadcom PXE clients are able to remotely boot to a Windows 2000 Server or a Windows Server 2003 (SP1 and older) system running RIS to initialize and install Windows Server 2003 and prior operating systems. To extend functionalities beyond basic network connectivity when loading an operating system through RIS, see [Using the NetXtreme II Monolithic Driver](#).
- **Windows Deployment Service (WDS).** For Windows Server 2003 SP2, RIS was replaced by WDS, which offers a Broadcom PXE client to install Windows operating systems, including Windows Vista and Windows Server 2008 and Windows Server 2008 R2. To extend functionalities beyond basic network connectivity when loading an operating system through WDS, see [Using the NetXtreme II Monolithic Driver](#).
- **Automated Deployment Service (ADS).** The Broadcom PXE client can connect to a Windows Server 2003 system and run a deployment agent that allows one to perform some administrative functions, including, but not limited to, deploying a Windows Server 2003 image. To extend functionalities beyond basic network connectivity when loading an operating system through ADS, see [Using the NetXtreme II Monolithic Driver](#).

SETTING UP MBA IN A CLIENT ENVIRONMENT

Setting up MBA in a client environment involves the following steps:

1. Enabling the MBA driver.
2. Configuring the MBA driver.
3. Setting up the BIOS for the boot order.

ENABLING THE MBA DRIVER

To enable or disable the MBA driver:

1. Insert the installation CD in the CD-ROM drive and boot up in DOS mode.



NOTE: The `uxdiag.exe` file is on the installation CD.

2. Type:

```
uxdiag -mba [ 0-disable | 1-enable ] -c devnum
```

where

devnum is the specific device(s) number (0,1,2, ...) to be programmed.

CONFIGURING THE MBA DRIVER

This section pertains to configuring the MBA driver on add-in NIC models of the Broadcom network adapter. For configuring the MBA driver on LOM models of the Broadcom network adapter, check your system documentation.



NOTE: You can use the MBA Configuration Menu to configure the MBA driver one adapter at a time as described below, or you can use the [Broadcom NetXtreme II User Diagnostics](#) MS-DOS based application to simultaneously configure the MBA driver for multiple adapters.

1. Restart your system.
2. Press CTRL+S within 4 seconds after you are prompted to do so.



NOTE: The message prompting you to press CTRL+S is displayed once for each Broadcom NetXtreme II adapter you have in your system that has MBA enabled. The messages are displayed in the same order as the assigned adapter device number.

3. Use the UP ARROW and DOWN ARROW keys to move to the Boot Protocol menu item. Then use the RIGHT ARROW or LEFT ARROW key to select the boot protocol of choice if other boot protocols besides Preboot Execution Environment (PXE) are available. If available, other boot protocols include Remote Program Load (RPL), and Bootstrap Protocol (BOOTP).



NOTE: If you have multiple adapters in your system and you are unsure which adapter you are configuring, press CTRL+F6, which causes the port LEDs on the adapter to start blinking.

4. Use the UP ARROW, DOWN ARROW, LEFT ARROW, and RIGHT ARROW keys to move to and change the values for other menu items, as desired.
5. Press F4 to save your settings.
6. Press ESC when you are finished.

SETTING UP THE BIOS

To boot from the network with the MBA, make the MBA enabled adapter the first bootable device under the BIOS. This procedure depends on the system BIOS implementation. Refer to the user manual for the system for instructions.

SETTING UP MBA IN A SERVER ENVIRONMENT

RED HAT LINUX PXE SERVER

The Red Hat Enterprise Linux distribution has PXE Server support. It allows users to remotely perform a complete Linux installation over the network. The distribution comes with the boot images *boot kernel* (*vmlinuz*) and *initial ram disk* (*initrd*), which are located on the Red Hat disk#1:

```
/images/pxeboot/vmlinuz  
/images/pxeboot/initrd.img
```

Refer to the Red Hat documentation for instructions on how to install PXE Server on Linux.

The *initrd.img* file distributed with Red Hat Enterprise Linux, however, does not have a Linux network driver for the Broadcom NetXtreme II adapters. This version requires a driver disk for drivers that are not part of the standard distribution. You can create a driver disk for the Broadcom NetXtreme II adapter from the image distributed with the installation CD. Refer to the *Linux Readme.txt* file for more information.

MS-DOS UNDI/INTEL APITEST

To boot in MS-DOS mode and connect to a network for the MS-DOS environment, download the Intel PXE PDK from the Intel website. This PXE PDK comes with a TFTP/ProxyDHCP/Boot server. The PXE PDK can be downloaded from Intel at <http://downloadcenter.intel.com/SearchResult.aspx?lang=eng&ProductFamily=Network+Connectivity&ProductLine=Boot+Agent+Software&ProductProduct=Intel%20ae+Boot+Agent>.



iSCSI Protocol: Broadcom NetXtreme II™ Network Adapter User Guide

- [iSCSI Boot](#)
- [iSCSI Crash Dump](#)
- [iSCSI Offload](#)

iSCSI BOOT

Broadcom NetXtreme II Gigabit Ethernet adapters support iSCSI boot to enable network boot of operating systems to diskless systems. The iSCSI boot allows a Windows or Linux operating system boot from an iSCSI target machine located remotely over a standard IP network.

For Windows operating systems, iSCSI boot can be configured to boot with two distinctive paths: non-offload (also known as Microsoft initiator) and offload (Broadcom's offload iSCSI driver or HBA). Configuration of the path is set with the **Windows HBA Boot Mode** option located on the **General Parameters** screen of the iSCSI Configuration utility. See [Table 1](#) for more information on all **General Parameters** screen configuration options.

SUPPORTED OPERATING SYSTEMS FOR iSCSI BOOT

The Broadcom NetXtreme II Gigabit Ethernet adapters support iSCSI boot on the following operating systems:

- Windows Server 2003 32-bit and 64-bit, SP1 and SP2
- Windows Server 2008 32-bit and 64-bit
- Windows Server 2008 R2 64-bit
- Linux RHEL 5.x, SLES 10.x, and SLES 11 (limited distributions with open-iscsi)

For Linux iSCSI boot, see [Linux iSCSI Boot Setup](#).

WINDOWS iSCSI BOOT SETUP

The Windows iSCSI boot setup consists of:

- [Configuring the iSCSI Target](#)
- [Configuring the iSCSI Boot Parameters](#)
- [Preparing the Image on the Local Hard Drive](#)
- [Transferring the OS Image to the iSCSI Target](#)
- [Booting](#)

Configuring the iSCSI Target

Configuring the iSCSI target varies by target vendors. For information on configuring the iSCSI target, refer to the documentation provided by the vendor. The general steps include:



1. Create an iSCSI target.
2. Create a virtual disk.
3. Map the virtual disk to the iSCSI target created in step 1.
4. Associate an iSCSI initiator with the iSCSI target.
5. Record the iSCSI target name, TCP port number, iSCSI Logical Unit Number (LUN), initiator Internet Qualified Name (IQN), and CHAP authentication details.

After configuring the iSCSI target, obtain the following:

- Target IQN
- Target IP address
- Target TCP port number
- Target LUN
- Initiator IQN
- CHAP ID and secret

Configuring the iSCSI Boot Parameters

Configure the Broadcom iSCSI boot software for either static or dynamic configuration. Refer to [Table 1](#) for configuration options available from the General Parameters screen.

[Table 1](#) lists parameters for both IPv4 and IPv6. Parameters specific to either IPv4 or IPv6 are noted.



NOTE: Availability of IPv6 iSCSI boot is platform/device dependent.

Table 1: Configuration Options

Option	Description
TCP/IP parameters via DHCP	This option is specific to IPv4. Controls whether the iSCSI boot host software acquires the IP address information using DHCP (Enabled) or use a static IP configuration (Disabled).
IP Autoconfiguration	This option is specific to IPv6. Controls whether the iSCSI boot host software will configure a stateless link-local address and/or stateful address if DHCPv6 is present and used (Enabled). Router Solicit packets are sent out up to three times with 4 second intervals in between each retry. Or use a static IP configuration (Disabled).
iSCSI parameters via DHCP	Controls whether the iSCSI boot host software acquires its iSCSI target parameters using DHCP (Enabled) or through a static configuration (Disabled). The static information is entered through the iSCSI Initiator Parameters Configuration screen.
CHAP Authentication	Controls whether the iSCSI boot host software uses CHAP authentication when connecting to the iSCSI target. If CHAP Authentication is enabled, the CHAP ID and CHAP Secret are entered through the iSCSI Initiator Parameters Configuration screen.
Boot to iSCSI target	Controls whether the iSCSI boot host software attempts to boot from the iSCSI target after successfully connecting to it. When the option is enabled, the iSCSI boot host software immediately attempts to boot from the iSCSI target. If set to disabled, the iSCSI boot host software does not attempt to boot from the iSCSI target and control returns to the system BIOS so that the next boot device may be used. The One Time Disabled option is used when you want to do a remote install of the OS to an iSCSI target. As the option is named, it is set to disable on the first boot, then changes to enabled on subsequent reboots to indicate that iSCSI boots from the iSCSI target.

Table 1: Configuration Options

Option	Description
DHCP Vendor ID	Controls how the iSCSI boot host software interprets the Vendor Class ID field used during DHCP. If the Vendor Class ID field in the DHCP Offer packet matches the value in the field, the iSCSI boot host software looks into the DHCP Option 43 fields for the required iSCSI boot extensions. If DHCP is disabled, this value does not need to be set.
Link Up Delay Time	Controls how long the iSCSI boot host software waits, in seconds, after an Ethernet link is established before sending any data over the network. The valid values are 0 to 255. As an example, a user may need to set a value for this option if a network protocol, such as Spanning Tree, is enabled on the switch interface to the client system.
Use TCP Timestamp	Controls if the TCP Timestamp option is enabled or disabled.
Target as First HDD	Allows specifying that the iSCSI target drive will appear as the first hard drive in the system.
LUN Busy Retry Count	Controls the number of connection retries the iSCSI Boot initiator will attempt if the iSCSI target LUN is busy.
IP Version	This option specific to IPv6. Toggles between the IPv4 or IPv6 protocol. All IP settings will be lost when switching from one protocol version to another.
Windows HBA Boot Mode	Set to disable when the host OS is configured for software initiator mode and to enable for HBA mode. This option is available on NetXtreme II adapters.

MBA Boot Protocol Configuration

To configure the boot protocol

1. Restart your system.
2. From the PXE banner, select **CTRL+S**. The MBA Configuration Menu appears (see [Broadcom Boot Agent](#)).
3. From the MBA Configuration Menu, use the **UP ARROW** or **DOWN ARROW** to move to the **Boot Protocol** option. Use the **LEFT ARROW** or **RIGHT ARROW** to change the **Boot Protocol** option to **iSCSI**.
4. Select **CTRL+K** to access the iSCSI Configuration Utility.



NOTE: If iSCSI boot firmware is not programmed in the NetXtreme II network adapter, selecting **CTRL+K** will not have any effect.

Static iSCSI Boot Configuration

In a static configuration, you must enter data for the system's IP address, the system's initiator IQN, and the target parameters obtained in [Configuring the iSCSI Target](#). For information on configuration options, see [Table 1](#).

To configure the iSCSI boot parameters using static configuration

1. From the **MBA Configuration Menu**, select **CTRL+K**.
2. From the **Main** menu, select **General Parameters**.
3. From the **General Parameters** screen, set the following:
 - **TCP/IP parameters via DHCP:** Disabled. (For IPv4.)
 - **IP Autoconfiguration:** Disabled. (For IPv6, non-offload.)
 - **iSCSI parameters via DHCP:** Disabled
 - **CHAP Authentication:** Disabled

- **Boot to iSCSI target:** Disabled
 - **DHCP Vendor ID:** BCM ISAN
 - **Link Up Delay Time:** 0
 - **Use TCP Timestamp:** Enabled
 - **Target as First HDD:** Disabled
 - **LUN Busy Retry Count:** 0
 - **IP Version:** IPv6. (For IPv6, non-offload.)
 - **Windows HBA Boot Mode:** Disabled
4. Select **ESC** to return to the **Main** menu.
 5. From the **Main** menu, select **Initiator Parameters**.
 6. From the **Initiator Parameters** screen, type values for the following:
 - IP Address
 - Subnet Mask Prefix
 - Default Gateway
 - Primary DNS
 - Secondary DNS
 - iSCSI Name (corresponds to the iSCSI initiator name to be used by the client system)



NOTE: Carefully enter the IP address. There is no error-checking performed against the IP address to check for duplicates or incorrect segment/network assignment.

7. Select **ESC** to return to the **Main** menu.
8. From the **Main** menu, select **1st Target Parameters**.
9. From the **1st Target Parameters** screen, enable **Connect** to connect to the iSCSI target. Type values for the following using the values used when configuring the iSCSI target:
 - IP Address
 - TCP Port
 - Boot LUN
 - iSCSI Name
10. Select **ESC** to return to the **Main** menu.
11. A second iSCSI boot adapter can be configured for redundancy in the event the primary adapter fails to boot. To configure the secondary device parameters, select **Secondary Device Parameters** from the **Main** menu (see [Configure Parameters for a Secondary Adapter](#)). Otherwise, go to step 12.
12. Select **ESC** and select **Exit and Save Configuration**.
13. Select **F4** to save your MBA configuration.
14. If necessary, return to the iSCSI Boot Configuration Utility to configure a second iSCSI target.

Dynamic iSCSI Boot Configuration

In a dynamic configuration, you only need to specify that the system's IP address and target/initiator information are provided by a DHCP server (see IPv4 and IPv6 configurations in [Configuring the DHCP Server to Support iSCSI Boot](#)). For IPv4, with the exception of the initiator iSCSI name, any settings on the Initiator Parameters, 1st Target Parameters, or 2nd Target Parameters screens are ignored and do not need to be cleared. For IPv6, with the exception of the CHAP ID and Secret, any settings on the Initiator Parameters, 1st Target Parameters, or 2nd Target Parameters screens are ignored and do not need to be cleared. For information on configuration options, see [Table 1](#).



NOTE: When using a DHCP server, the DNS server entries are overwritten by the values provided by the DHCP server. This occurs even if the locally provided values are valid and the DHCP server provides no DNS server information. When the DHCP server provides no DNS server information, both the primary and secondary DNS server values are set to 0.0.0.0. When the Windows OS takes over, the Microsoft iSCSI initiator retrieves the iSCSI Initiator parameters and configures the appropriate registries statically. It will overwrite whatever is configured. Since the DHCP daemon runs in the Windows environment as a user process, all TCP/IP parameters have to be statically configured before the stack comes up in the iSCSI Boot environment.

If DHCP Option 17 is used, the target information is provided by the DHCP server, and the initiator iSCSI name is retrieved from the value programmed from the Initiator Parameters screen. If no value was selected, then the controller defaults to the name:

```
iqn.1995-05.com.broadcom.<11.22.33.44.55.66>.iscsiboot
```

where the string 11.22.33.44.55.66 corresponds to the controller's MAC address.

If DHCP option 43 (IPv4 only) is used, then any settings on the Initiator Parameters, 1st Target Parameters, or 2nd Target Parameters screens are ignored and do not need to be cleared.

To configure the iSCSI boot parameters using dynamic configuration

1. From the **MBA Configuration Menu**, select **CTRL+K**.
2. From the **Main** menu, select **General Parameters**.
3. From the **General Parameters** screen, set the following:
 - **TCP/IP parameters via DHCP:** Enabled. (For IPv4.)
 - **IP Autoconfiguration:** Enabled. (For IPv6, non-offload.)
 - **iSCSI parameters via DHCP:** Enabled
 - **CHAP Authentication:** Disabled
 - **Boot to iSCSI target:** Disabled
 - **DHCP Vendor ID:** BCM ISAN
 - **Link Up Delay Time:** 0
 - **Use TCP Timestamp:** Enabled
 - **Target as First HDD:** Disabled
 - **LUN Busy Retry Count:** 0
 - **IP Version:** IPv6. (For IPv6, non-offload.)
 - **Windows HBA Boot Mode:** Disabled
4. Select **ESC** to return to the **Main** menu.



NOTE: Information on the **Initiator Parameters**, **1st Target Parameters**, and **2nd Target Parameters** screens are ignored and do not need to be cleared.

5. A second iSCSI boot adapter can be configured for redundancy in the event the primary adapter fails to boot. To configure the secondary device parameters, select **Secondary Device Parameters** from the **Main** menu (see [Configure Parameters for a Secondary Adapter](#)). Otherwise, go to step 12.
6. Select **Exit and Save Configurations**.

Configure Parameters for a Secondary Adapter

A second iSCSI boot adapter can be configured for redundancy in the event the primary adapter fails to boot.

To configure the iSCSI boot parameters for a secondary adapter

1. From the **MBA Configuration Menu**, select **CTRL+K**.
2. From the **Main** menu, select **Secondary Device Parameters**.
3. From the **Secondary Device Parameters** screen, select **Secondary Device**.
4. From the **Device List**, select the adapter that will be used as the secondary adapter.
5. From the **Secondary Device Parameters** screen, set **Use Independent Target Portal** to Enabled (or Disabled if MPIO mode is not required) and set **Use Independent Target Name** to Enabled.
6. Select **Invoke** to configure the secondary adapter.
7. Configure the secondary adapter parameters.



NOTE: The IP addresses for the primary and secondary adapters must be in two different subnets.

8. Select **ESC** and select **Exit and Save Configuration**.
9. Select **F4** to save your MBA configuration.

Preparing the Image on the Local Hard Drive



NOTE: Prior to creating the disk image when using KMDf VBD driver version 3.1 and above, it is necessary to change the load order of the wdf01000.sys driver. In the registry at HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Wdf0100, modify the registry entry by changing **Group** to "base" and **Start** to "0".

When the software initiator mode operating system installation is used, the Windows operating system install must be performed in two stages. In the first stage, the OS is installed to a local hard drive on the system. In the second stage after the OS has been completely installed, an image of the local drive must be transferred to the iSCSI target for use in subsequent boots.

Initial Windows Install

Before beginning, verify that the desired version of Windows supports iSCSI boot. Proceed to install Windows to a local hard drive. When partitioning the local hard drive, make sure that the Windows boot drive (normally C:) is partitioned to a size less than or equal to the size of the iSCSI target to be used. Depending on the method used to copy an image of the local hard drive to the iSCSI target, this may or may not be an actual requirement. Proceed to install Windows with the desired options.

1. Install Windows 2003 32-bit or Windows 2003 64-bit OS on the local hard drive.
2. Install the Broadcom drivers using the Setup installer.



NOTE: Do not install the drivers through Windows Plug-and-Play (PnP). Failure to install the drivers through the Setup installer might blue screen your system.

3. Install Microsoft iSCSI Software Initiator with integrated software boot support (version 2.06 or later). To download from Microsoft, go to <http://www.microsoft.com/downloads/details.aspx?familyid=12cb3c1a-15d6-4585-b385-befd1319f825&displaylang=en> and locate the direct link for your system on the page's "Overview" section.

4. Select support for Multipath I/O (MPIO), if needed. Refer to Microsoft's Initiator documentation for more information on MPIO.
5. Select the iSCSI boot option.
6. Select the Broadcom adapter as the iSCSI boot device.

**NOTES:**

- Do not manually create connections to the iSCSI target for iSCSI boot adapters.
- If the image is used on other hardware, Sysprep is required.
- It is recommended to always run `iscsibcg.exe /verify /fix` at a command prompt after updating the driver and before restarting the system. For information on configuring the `iscsibcg` utility to run automatically when a system shuts down, see <http://support.microsoft.com/kb/934235>.

Transferring the OS Image to the iSCSI Target

1. Create a new FAT32 partition on the local hard drive using the disk management console.
2. Boot to a bootable device such as a diskette drive, media, or USB key and run a disk imaging software, such as Symantec Ghost.
3. Clone the OS partition to the FAT32 partition (partition to image).
4. Place the iSCSI boot adapter before the hard drive in the boot menu.
5. Reboot the host and boot into the OS in the local hard drive.
6. Launch Windows compatible cloning software, such as Ghost32, and write the image in the FAT32 partition to the remote LUN.

Booting

After that the system has been prepared for an iSCSI boot and the operating system is present on the iSCSI target, the last step is to perform the actual boot. The system will boot to Windows over the network and operate as if it were a local disk drive.

1. Reboot the server.
2. Select **CTRL+S** and **CTRL+K**.
3. From the **Main** menu, select **General Parameters** and configure the **Boot to iSCSI target** option to **Enabled**.

If CHAP authentication is needed, enable CHAP authentication after determining that booting is successful (see [Enabling CHAP Authentication](#)).

Enabling CHAP Authentication

Ensure that CHAP authentication is enabled on the target.

To enable CHAP authentication

1. From the **MBA Configuration Menu**, select **CTRL+K**.
2. From the **Main** menu, select **General Parameters**.
3. From the **General Parameters** screen, set **CHAP Authentication** to Enabled.
4. From the **Initiator Parameters** screen, type values for the following:
 - CHAP ID (up to 128 bytes)
 - CHAP Secret (if authentication is required, and must be 12 characters in length or longer)

5. Select **ESC** to return to the **Main** menu.
6. From the **Main** menu, select **1st Target Parameters**.
7. From the **1st Target Parameters** screen, type values for the following using the values used when configuring the iSCSI target:
 - CHAP ID (optional if two-way CHAP)
 - CHAP Secret (optional if two-way CHAP, and must be 12 characters in length or longer)
8. Select **ESC** to return to the **Main** menu.
9. Select **ESC** and select **Exit and Save Configuration**.

Configuring the DHCP Server to Support iSCSI Boot

The DHCP server is an optional component and it is only necessary if you will be doing a dynamic iSCSI Boot configuration setup (see [Dynamic iSCSI Boot Configuration](#)).

Configuring the DHCP server to support iSCSI boot is different for IPv4 and IPv6.

- [DHCP iSCSI Boot Configurations for IPv4](#)
- [DHCP iSCSI Boot Configuration for IPv6](#)

DHCP iSCSI Boot Configurations for IPv4

The DHCP protocol includes a number of options that provide configuration information to the DHCP client. For iSCSI boot, Broadcom adapters support the following DHCP configurations:

- [DHCP Option 17, Root Path](#)
- [DHCP Option 43, Vendor-Specific Information](#)

DHCP Option 17, Root Path

Option 17 is used to pass the iSCSI target information to the iSCSI client.

The format of the root path as defined in IETF RFC 4173 is:

```
"iscsi:"<servername>": "<protocol>": "<port>": "<LUN>": "<targetname>"
```

The parameters are defined below.

Table 2: DHCP Option 17 Parameter Definition

Parameter	Definition
"iscsi:"	A literal string
<servername>	The IP address or FQDN of the iSCSI target
": "	Separator
<protocol>	The IP protocol used to access the iSCSI target. Currently, only TCP is supported so the protocol is 6.
<port>	The port number associated with the protocol. The standard port number for iSCSI is 3260.
<LUN>	The Logical Unit Number to use on the iSCSI target. The value of the LUN must be represented in hexadecimal format. A LUN with an ID OF 64 would have to be configured as 40 within the option 17 parameter on the DHCP server.
<targetname>	The target name in either IQN or EUI format (refer to RFC 3720 for details on both IQN and EUI formats). An example IQN name would be "iqn.1995-05.com.broadcom:iscsi-target".



DHCP Option 43, Vendor-Specific Information

DHCP option 43 (vendor-specific information) provides more configuration options to the iSCSI client than DHCP option 17. In this configuration, three additional suboptions are provided that assign the initiator IQN to the iSCSI boot client along with two iSCSI target IQNs that can be used for booting. The format for the iSCSI target IQN is the same as that of DHCP option 17, while the iSCSI initiator IQN is simply the initiator's IQN.



NOTE: DHCP Option 43 is supported on IPv4 only.

The suboptions are listed below.

Table 3: DHCP Option 43 Suboption Definition

Suboption	Definition
201	First iSCSI target information in the standard root path format "iscsi:<servername>:<protocol>:<port>:<LUN>:<targetname>"
202	Second iSCSI target information in the standard root path format "iscsi:<servername>:<protocol>:<port>:<LUN>:<targetname>"
203	iSCSI initiator IQN

Using DHCP option 43 requires more configuration than DHCP option 17, but it provides a richer environment and provides more configuration options. Broadcom recommends that customers use DHCP option 43 when performing dynamic iSCSI boot configuration.

Configuring the DHCP Server

Configure the DHCP server to support option 17 or option 43.



NOTE: If using Option 43, you also need to configure Option 60. The value of Option 60 should match the **DHCP Vendor ID** value. The **DHCP Vendor ID** value is BRCM ISAN, as shown in **General Parameters** of the iSCSI Boot Configuration menu.

DHCP iSCSI Boot Configuration for IPv6

The DHCPv6 server can provide a number of options, including stateless or stateful IP configuration, as well s information to the DHCPv6 client. For iSCSI boot, Broadcom adapters support the following DHCP configurations:

- [DHCPv6 Option 16, Vendor Class Option](#)
- [DHCPv6 Option 17, Vendor-Specific Information](#)



NOTE: The DHCPv6 standard Root Path option is not yet available. Broadcom suggests using Option 16 or Option 17 for dynamic iSCSI Boot IPv6 support.

DHCPv6 Option 16, Vendor Class Option

DHCPv6 Option 16 (vendor class option) must be present and must contain a string that matches your configured **DHCP Vendor ID** parameter. The **DHCP Vendor ID** value is BRCM ISAN, as shown in **General Parameters** of the iSCSI Boot Configuration menu.

The content of Option 16 should be <2-byte length> <DHCP Vendor ID>.

DHCPv6 Option 17, Vendor-Specific Information

DHCPv6 Option 17 (vendor-specific information) provides more configuration options to the iSCSI client. In this configuration, three additional suboptions are provided that assign the initiator IQN to the iSCSI boot client along with two iSCSI target IQNs that can be used for booting.

The suboptions are listed below.

Table 4: DHCP Option 17 Suboption Definition

Suboption	Definition
201	First iSCSI target information in the standard root path format "iscsi:" [<servername>] ":" <protocol> ":" <port> ":" <LUN> ":" <targetname> "
202	Second iSCSI target information in the standard root path format "iscsi:" [<servername>] ":" <protocol> ":" <port> ":" <LUN> ":" <targetname> "
203	iSCSI initiator IQN



NOTE: In [Table 4](#), the brackets [] are required for the IPv6 addresses.

The content of option 17 should be <2-byte Option Number 201|202|203> <2-byte length> <data>.

Configuring the DHCP Server

Configure the DHCP server to support Option 16 and Option 17.

**NOTE:**

- The format of DHCPv6 Option 16 and Option 17 are fully defined in RFC 3315.

LINUX ISCSI BOOT SETUP

Linux iSCSI boot is supported on Red Hat Enterprise Linux 5 update 1 and later and SUSE Linux Enterprise Server 10 SP1 and later.

The Linux iSCSI boot setup consists of:

- [Configuring the iSCSI Target](#)
- [Configuring the iSCSI Boot Parameters](#)
- [Changing the BIOS Boot Order](#)
- [Installing to Target](#)
- [Bootting](#)

Configuring the iSCSI Target

See [Configuring the iSCSI Target](#).

Configuring the iSCSI Boot Parameters

See [Configuring the iSCSI Boot Parameters](#).

Changing the BIOS Boot Order

Change the BIOS boot order to the following:

- iSCSI Option ROM
- CDROM/DVD

Installing to Target

1. Update the iSCSI boot Option ROM with the target IQN and target IP address.
2. Insert the first Red Hat 5, SUSE 10, or SUSE 11 CD into the CDROM drive.
3. Ensure that the Option ROM is able to log into the target disk.
4. For SUSE 10 and SUSE 11, select **installation** at the first screen and enter `withiscsi=1 netsetup=1`.
5. For Red Hat 5, select **Installation** at the first screen and **Enter**.
 - a. At the **Installation requires partitioning of your hard drive** screen, click **Advanced storage configuration**.
 - b. Select **Add iSCSI target** and click **Add drive**.
 - c. Select the boot adapter and configure your network information.
 - d. Enter your iSCSI Target IP address at the iSCSI parameters screen.
6. Follow the typical installation procedure.

Booting

See [Booting](#).

OTHER iSCSI BOOT CONSIDERATIONS

There are several other factors that should be considered when configuring a system for iSCSI boot.

Virtual LANs

Virtual LAN (VLAN) tagging is not supported for iSCSI boot with the Microsoft iSCSI Software Initiator.

Teaming

The use of any form of teaming (Smart Load Balancing, Generic Trunking, or Link Aggregation) with the iSCSI boot device is not supported; however, teaming can still be configured on other devices in the system. For more information, see [Configuring Teaming](#).

iSCSI BOOT REMOTE INSTALLATION

This section discusses the procedures for creating and installing a Microsoft OS directly to the target through Broadcom iSCSI solution.

- [Windows Server 2003 \(OIS\)](#)
- [Windows Server 2008 \(Non-OIS\)](#)
- [Windows Server 2008 \(OIS\)](#)

Windows Server 2003 (OIS)

Installation of the OIS driver, along with a clean Windows Server 2003 SP2 installation requires that the Storport hotfix KB957910 (or later) be "slipstreamed" into the standard Windows Sever 2003 SP2 CD. The Microsoft Storport hotfix KB957910 can be found at <http://support.microsoft.com/kb/957910>. The Microsoft knowledge base article that details the slipstreaming procedure can be found at <http://support.microsoft.com/kb/814847>.

The OIS F6 installation is composed of two floppy disks per architecture. The contents of these disks are:

x86 Architecture

- Disk1: "WDF Installation Disk"
 - bxvbd.cat
 - bxvbd.inf
 - bxvbdx.sys
 - txtsetup.oem
 - wdf01000.sys
 - wdfldr.sys
- Disk2: "Broadcom iSCSI Driver"
 - bxois.cat
 - bxois.inf
 - bxois.sys
 - txtsetup.oem

x64 Architecture

- Disk1: "WDF Installation Disk"
 - bxvbd.cat
 - bxvbd.inf
 - bxvbda.sys
 - txtsetup.oem
 - wdf01000.sys
 - wdfldr.sys
- Disk2: "Broadcom iSCSI Driver"
 - bxois.cat
 - bxois.inf
 - bxois.sys
 - txtsetup.oem

To perform an iSCSI remote installation

1. During system POST, enter the MBA configuration menu **CTRL+S** when prompted.
2. Set **Boot protocol** to **iSCSI**
3. Enter the iSCSI configuration menu by selecting **CTRL+K**.
4. Fill in General Parameters as required. Set **Boot to iSCSI target** to **Disable**.
5. Fill in **Initiator** and **Target** parameters as required.
6. Exit and save the iSCSI configuration.
7. Exit and save the MBA configuration.
8. Reorder the BIOS boot order so that the Broadcom iSCSI adapter precedes the CDROM.
9. Insert the Windows Server 2003 SP2 Slipstreamed CD.
10. Verify the offload iSCSI BIOS boot code successfully opens a session with the remote iSCSI target.
11. Continue to boot with the Windows Server installation CD.
12. Press **F6** when prompted to install additional drivers.
13. Press **S** to install additional drivers. Insert Disk 1 "WDF Installation Disk" when prompted. Select **Load 1st: wdfldr**.
14. Press **S** to install additional drivers. Insert Disk 1 "WDF Installation Disk" when prompted. Select **Load 2nd: wdf01000**.
15. Press **S** to install additional drivers. Insert Disk 1 "WDF Installation Disk" when prompted. Select **Load 3rd: Broadcom Virtual Bus Driver**.
16. Press **S** to install additional drivers. Insert Disk 1 "WDF Installation Disk" when prompted. Select **Load 4th: Broadcom iSCSI Driver**.
17. Continue with Windows text mode installation as usual.
18. The system will reboot after text mode setup. During POST, reenter the MBA configuration menu by selecting **CTRL+S** when prompted.
19. Enter the iSCSI configuration menu by selecting **CTRL+K**.
20. Enter General Parameters and set **Boot to iSCSI target** to **Enable**.
21. Exit and save the iSCSI configuration.
22. Exit and save the MBA configuration.
23. Continue with Windows GUI mode installation as usual.

Windows Server 2008 (Non-OIS)

To perform a remote install for non-offload

1. Configure iSCSI OpROM with all proper settings.
 - **Boot to iSCSI Target:** Disable
 - **Target as First HDD:** Enable
 - **HBA:** Disable
2. Order the boot order so that iSCSI precedes the DVD.
3. Copy the VBD and NDIS driver to a USB flash or a diskette.
4. Boot the system and connect to the target. The Windows Server 2008 DVD installation begins.
5. Continue with the installation. When you receive the **Where do you want to install Windows?** window, click on **load driver**.
6. Browse to the location of the VBD and install.
7. At the **Where do you want to install Windows?**, again, click **load driver**. This time, browse to the location of the NDIS driver and install.
8. At the **Where do you want to install Windows?** window, click **Refresh** if you do not see the target HDD.
9. Select the HDD/Partition to continue.
10. The system will reboot after text mode setup. During POST, reenter the MBA configuration menu by selecting **CTRL+S** when prompted.
11. Enter the iSCSI configuration menu by selecting **CTRL+K**.
12. Enter General Parameters and set **Boot to iSCSI target** to **Enable**.
13. Exit and save the iSCSI configuration.
14. Exit and save the MBA configuration.
15. Continue with Windows GUI mode installation as usual.

Windows Server 2008 (OIS)

To perform a remote install for offload

1. Configure iSCSI OpROM with all proper settings.
 - **Boot to iSCSI Target:** Disable
 - **Target as First HDD:** Enable
 - **HBA:** Enable
2. Order the boot order so that iSCSI precedes the DVD.
3. Copy the VBD and OIS driver to a USB flash or a diskette.
4. Boot the system and connect to the target. The Windows Server 2008 DVD installation begins.
5. Continue with the installation. When you receive the **Where do you want to install Windows?** window, click on **load driver**.
6. Browse to the location of the VBD and install.
7. At the **Where do you want to install Windows?**, again, click **load driver**. This time, browse to the location of the OIS driver and install.
8. At the **Where do you want to install Windows?** window, click **Refresh** if you do not see the target HDD.
9. Select the HDD/Partition to continue.
10. The system will reboot after text mode setup. During POST, reenter the MBA configuration menu by selecting **CTRL+S**

when prompted.

11. Enter the iSCSI configuration menu by selecting **CTRL+K**.
12. Enter General Parameters and set **Boot to iSCSI target** to **Enable**.
13. Exit and save the iSCSI configuration.
14. Exit and save the MBA configuration.
15. Continue with Windows GUI mode installation as usual.

TROUBLESHOOTING iSCSI BOOT

The following troubleshooting tips are useful for iSCSI boot.

Problem: When switching iSCSI boot from the Microsoft standard path to Broadcom iSCSI offload, the booting fails to complete.

Solution: Install or upgrade the Broadcom Virtual Bus Device (VBD) driver to 5.0.x, along with the OIS driver, prior to switching the iSCSI boot path.

Problem: If a Windows 2003-based system is booted with an MPIO configuration in the first boot, the interface booted without a cable attached will not be functional since the IP address of the interface will be statically configured to a value of 0. This problem only occurs if the TCP/IP parameters are configured as DHCP. In other words, it will not occur if the IP addresses are statically configured in the iSCSI configuration program.

Solution: Perform the first boot with both cables attached and ensure that both interfaces are able to acquire IP addresses.

Problem: The iSCSI configuration utility will not run.

Solution: Ensure that the iSCSI Boot firmware is installed in the NVRAM.

Problem: A system blue screen occurs when installing the Broadcom drivers through Windows Plug-and-Play (PnP).

Solution: Install the drivers through the Setup installer.

Problem: For static IP configuration when switching from Layer 2 iSCSI boot to Broadcom iSCSI HBA, then you will receive an IP address conflict.

Solution: Change the IP address of the network property in the OS.

Problem: After configuring the iSCSI boot LUN to 255, a system blue screen appears when performing iSCSI boot.

Solution: Although Broadcom's iSCSI solution supports a LUN range from 0 to 255, the Microsoft iSCSI software initiator does not support a LUN of 255. Configure a LUN value from 0 to 254.

iSCSI CRASH DUMP

If you will use the Broadcom iSCSI Crash Dump utility, it is important to follow the installation procedure to install the iSCSI Crash Dump driver. See [Using the Installer](#) for more information.

ISCSI OFFLOAD

iSCSI offload is a technology that offloads iSCSI protocol processing overhead from host processors to the iSCSI host bus adapter to increase network performance and throughput while helping to optimize server processor utilization.

This section covers Windows iSCSI offload for the NetXtreme II family of network adapters. For Linux iSCSI offload, see [Linux iSCSI Offload](#).

CONFIGURING ISCSI OFFLOAD

With the proper iSCSI offload licensing, you can configure your iSCSI-capable NetXtreme II network adapter to offload iSCSI processing from the host processor. The following process enables your system to take advantage of Broadcom's iSCSI offload feature.

- [Installing Broadcom Drivers and Management Applications](#)
- [Installing the Microsoft iSCSI Initiator](#)
- [Configuring Broadcom iSCSI Using BACS 3](#)
- [Configure Microsoft Initiator to Use Broadcom's iSCSI Offload](#)

Installing Broadcom Drivers and Management Applications

1. Install the Windows drivers. See [Windows Driver Software](#).
2. Install the management applications. See [Installing Management Applications](#).

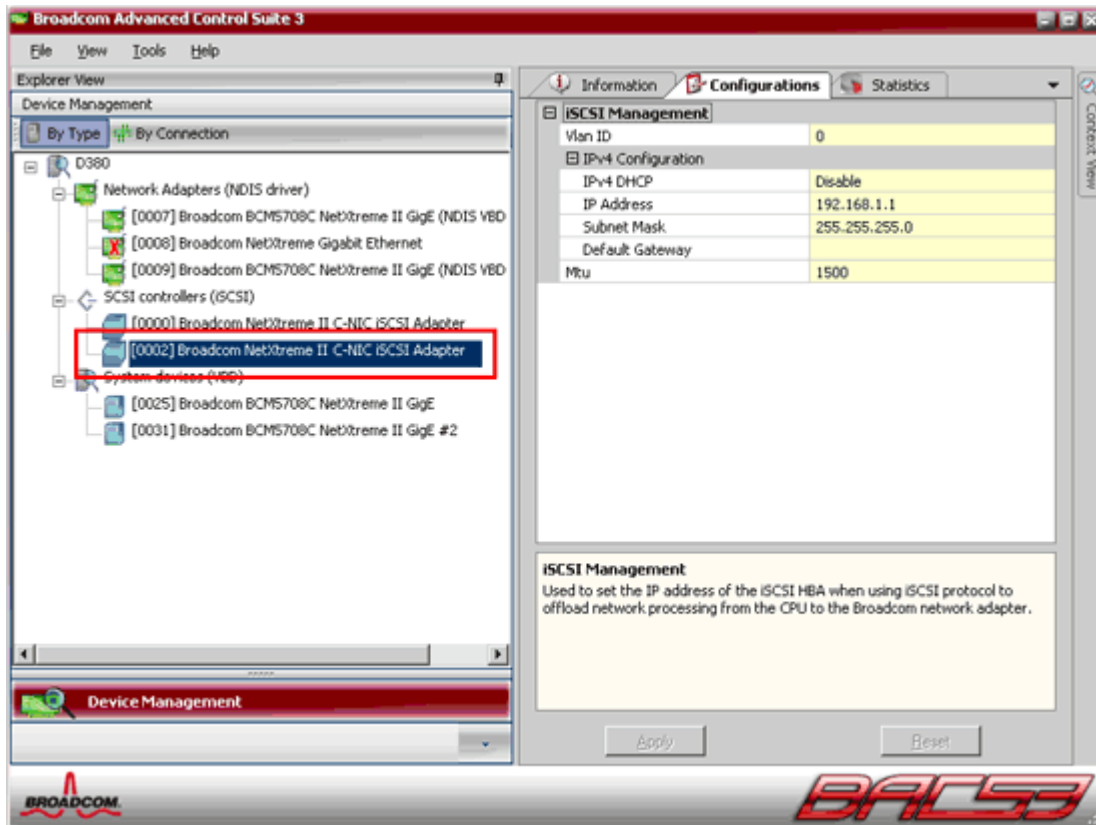
Installing the Microsoft iSCSI Initiator

Install the Microsoft iSCSI Software Initiator, version 2.08 or later, on Windows Server 2003. For Windows Server 2008 and later, the iSCSI initiator is included in-box. To download the iSCSI initiator from Microsoft, go to <http://www.microsoft.com/downloads/details.aspx?familyid=12cb3c1a-15d6-4585-b385-befd1319f825&displaylang=en> and locate the direct link for your system.


Configuring Broadcom iSCSI Using BACS 3

The Broadcom Advanced Control Suite (BACS 3) is used to manage all of Broadcom’s network adapters and advanced features. For more information, see [Using Broadcom Advanced Control Suite 3](#).

1. Open BACS 3.
2. Select the Broadcom NetXtreme II C-NIC iSCSI adapter. If the C-NIC iSCSI adapter is not present, then select the VBD device and enable iSCSI offload by selecting **iSCSI Offload Engine** from the **Resource Reservations** area of the Configuration tab.



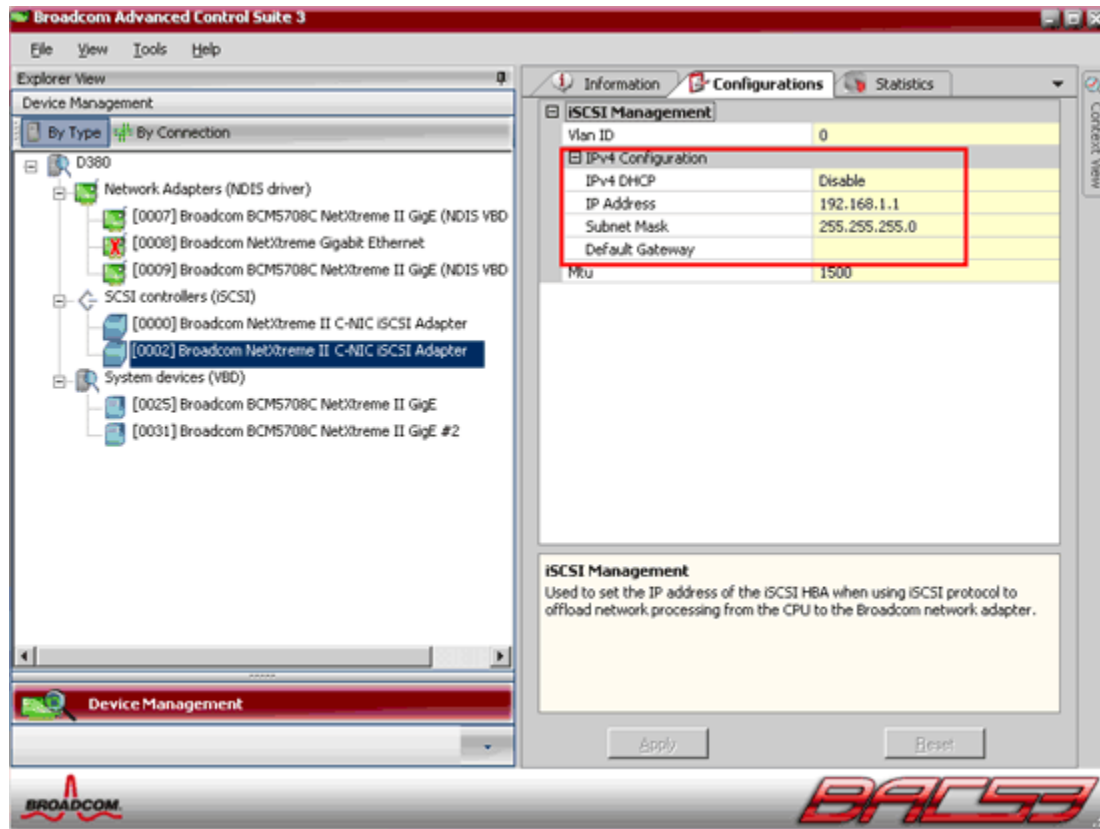
- 3. Select the Configuration tab.



- 4. DHCP is the default for IP address assignment, but this can be changed to static IP address assignment, if this is the preferred method of IP address assignment.



NOTE: The IP address assignment method cannot be changed if the adapter was used for boot.

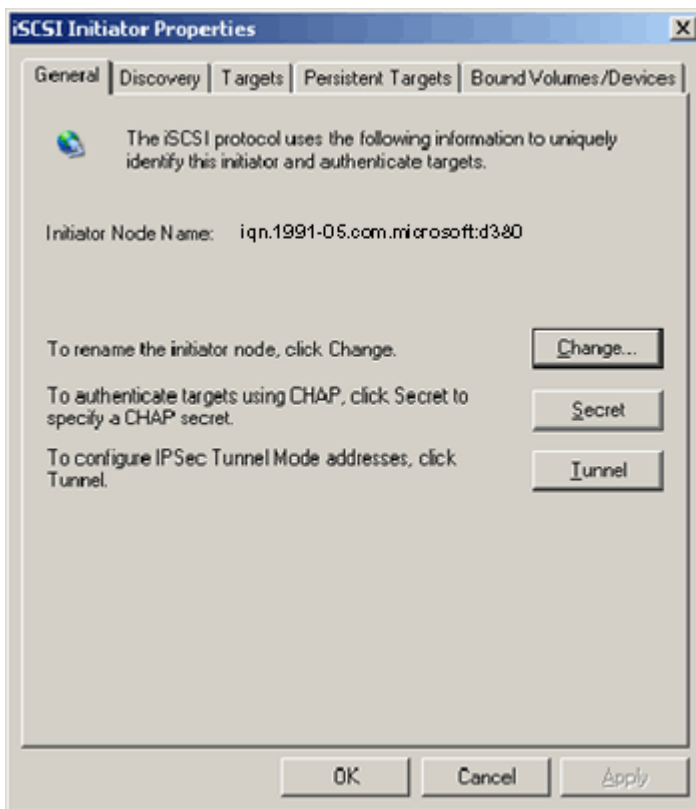


5. Select **Apply** and close BACS.

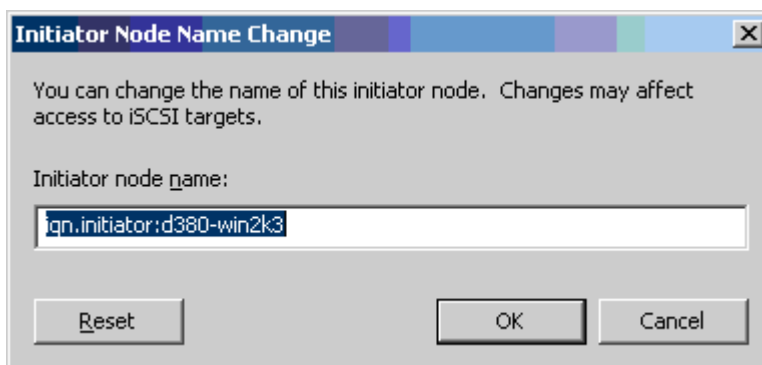
Configure Microsoft Initiator to Use Broadcom's iSCSI Offload

Now that the IP address has been configured for the iSCSI adapter, you need to use Microsoft Initiator to configure and add a connection to the iSCSI target using Broadcom iSCSI adapter. See Microsoft's user guide for more details on Microsoft Initiator.


1. Open Microsoft Initiator.
2. Configure the initiator IQN name according to your setup. To change, click on **Change**.



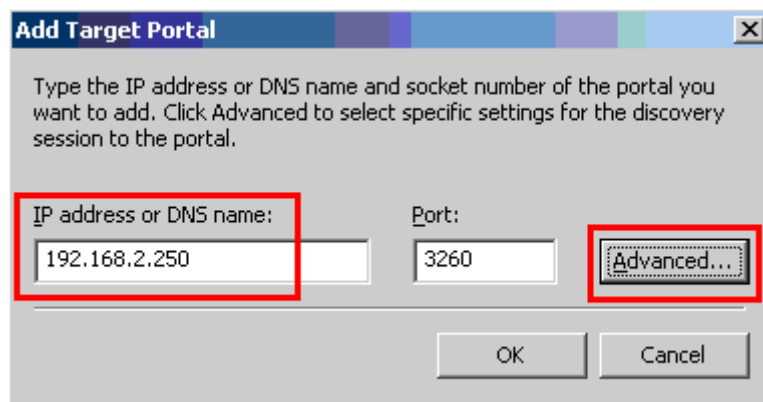
3. Enter the initiator IQN name.




4. Select the Discovery tab and click **Add** to add a target portal.




5. Enter the IP address of the target and click **Advanced**.



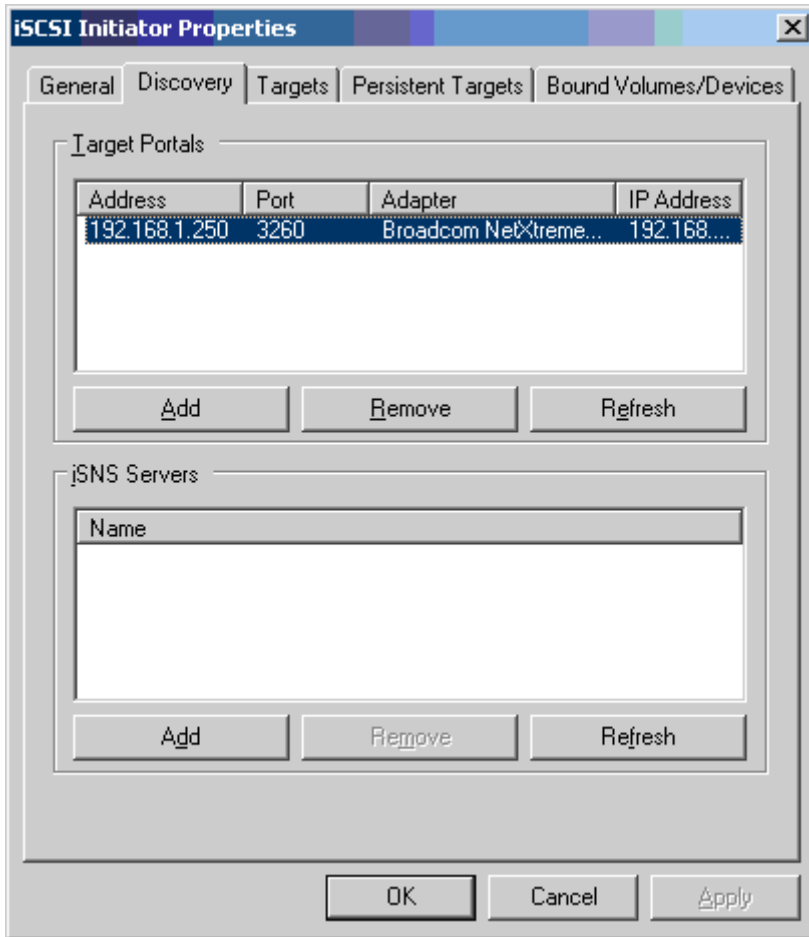
- From the General tab, select Broadcom NetXtreme II C-NIC iSCSI Adapter from **Local adapter**.



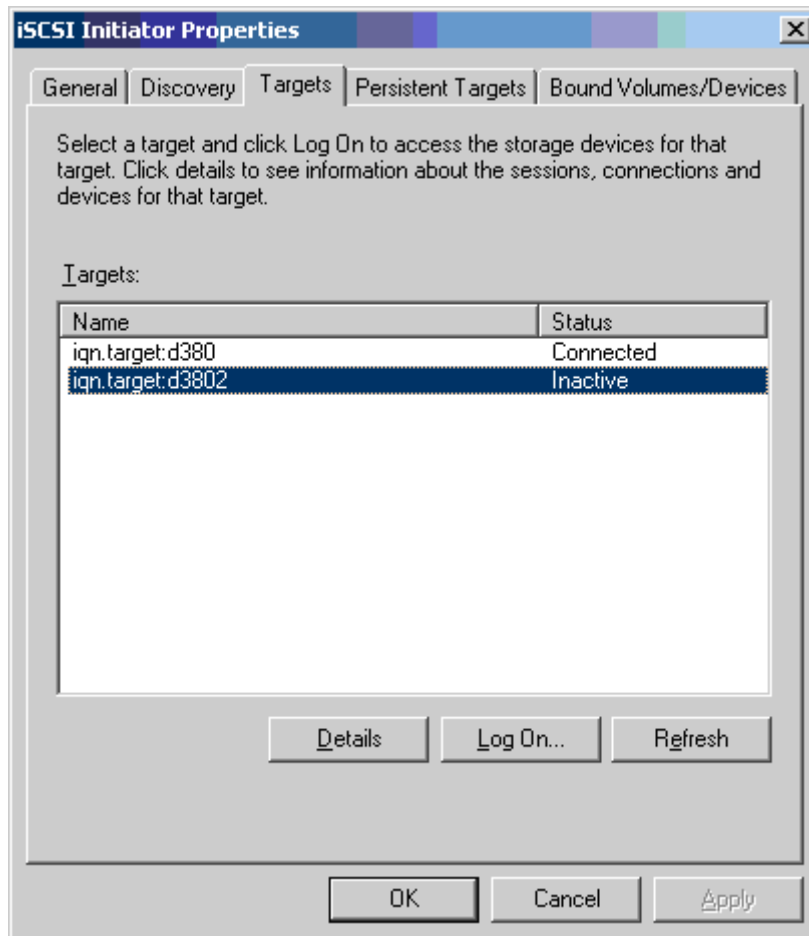
7. Select the IP address for the adapter from **Source IP**.




8. Click **OK** to close Advanced setting and then **OK** to add the target portal.



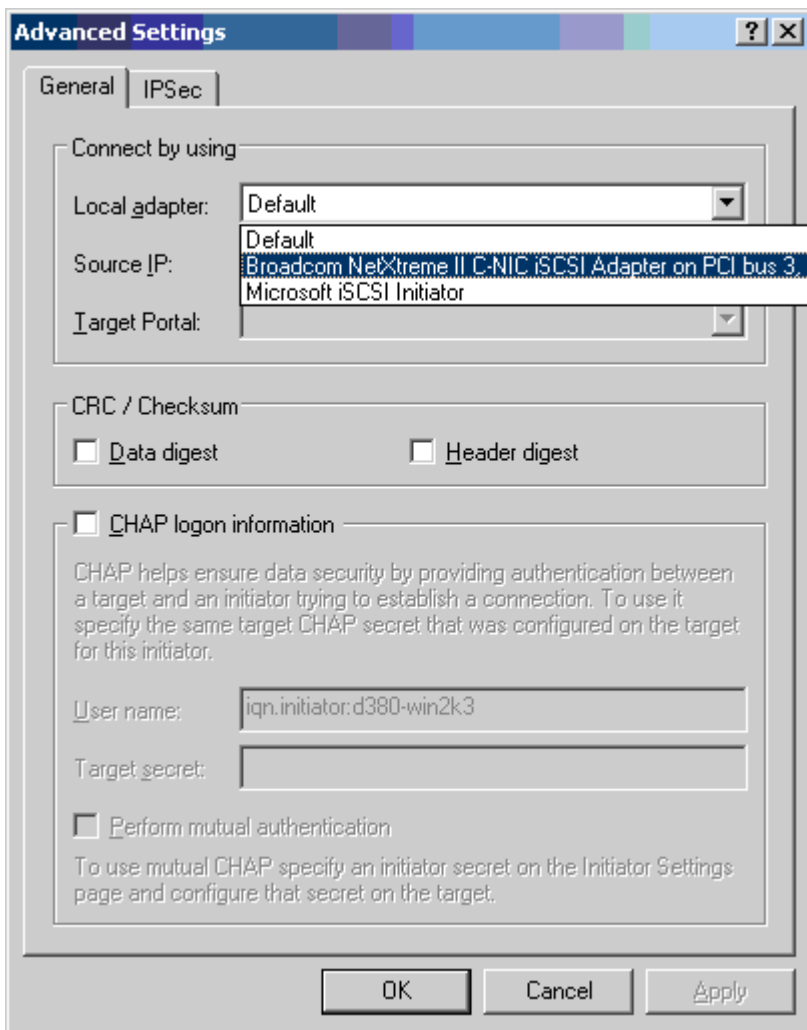
- From the Targets tab, select the target and click **Log On** to log into your iSCSI target using the Broadcom iSCSI adapter.



- Click on **Advanced**.



11. On the General tab, select the Broadcom NetXtreme II C-NIC iSCSI adapter from **Local adapter**.
12. Click **OK** to close Advanced settings.



13. Click **OK** to close the Microsoft Initiator.
14. To format your iSCSI partition, use Disk Manager.

**NOTES:**

- No teaming support for iSCSI-enabled adapters is available at this point. MPIO handles all failover and link aggregation for iSCSI traffic.
- Network teaming and MPIO are supported concurrently on different ports. Ports running iSCSI use MPIO and ports running LAN traffic use network teaming.
- If a Broadcom network adapter has iSCSI offload enabled, the teaming software, BASP, will not allow the iSCSI offload enabled network port interface to be a part of a GEC or LACP (802.3ad) team, but SLB teams are allowed.

iSCSI OFFLOAD FAQs

Q: How do I assign an IP address for iSCSI offload?

A: Use the Configurations tab in Broadcom Advanced Control Suite 3 (BACS).

Q: What tools should be used to create the connection to the target?

A: Use Microsoft iSCSI Software Initiator (version 2.08 or later).

Q: How do I know that the connection is offloaded?

A: Use Microsoft iSCSI Software Initiator. From a command line, type `iscsicli sessionlist`. From **Initiator Name**, an iSCSI offloaded connection will display an entry beginning with "B06BDRV...". A non-offloaded connection will display an entry beginning with "Root...".

Q: What configurations should be avoided?

A: The IP address should not be the same as the LAN.

Installing Management Applications: Broadcom NetXtreme II™ Network Adapter User Guide

- [Installing Broadcom Advanced Control Suite 3 and Related Management Applications](#)
- [Modifying Management Applications](#)
- [Repairing Management Applications](#)
- [Removing Management Applications](#)

INSTALLING BROADCOM ADVANCED CONTROL SUITE 3 AND RELATED MANAGEMENT APPLICATIONS

The Broadcom Advanced Control Suite 3 (BACS 3) 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 3 (BACS 3).
- **BASP.** Installs Broadcom Advanced Server Program.
- **SNMP.** Installs the Simple Network Management Protocol subagent.
- **CIM Provider.** Installs the Common Information Model provider.



NOTES:

- Ensure that the Broadcom network adapter(s) is physically installed in the system before installing BACS 3.
- Before installing Broadcom Advance Control Suite 3, verify that .NET Framework 2.0, or above, is installed.
- Before you begin the installation, close all applications, windows, or dialog boxes.
- To use the TCP/IP Offload Engine (TOE), you must have Windows Server 2003 with Service Pack 2 (SP2) or later, Windows Server 2008, or Windows Server 2008 R2. You must also have a license key preprogrammed in the hardware. If supported, for iSCSI, you only need a license key.
- BASP is not available on Windows Small Business Server (SBS) 2008.

.NET FRAMEWORK REQUIREMENT

Microsoft .NET Framework 2.0 includes the runtime and associated files needed to run BACS 3, and must be installed on your system in order for BACS 3 to operate. For information on the minimum and recommended .NET Framework versions for your operating system, see [Table 1](#).



NOTE: For optimal performance of BACS 3, Broadcom recommends .NET Framework 2.0 SP1, .NET Framework 3.0 SP1, or .NET Framework 3.5, depending on your operating system.

Table 1: Minimum and Recommended .NET Framework Versions for BACS 3

Operating System	Minimum Required		Recommended	
	Version	How to Obtain	Version	How to Obtain
Windows 2000 Server	.NET 2.0	Download from Microsoft	.NET 2.0 SP1	Download from Microsoft
Windows Server 2003	.NET 2.0	Download from Microsoft	.NET 2.0 SP1	Download from Microsoft
Windows Server 2003 R2	.NET 2.0	Included in OS package, but by default, not included during installation	.NET 2.0 SP1	Download from Microsoft
Windows Server 2008	.NET 2.0	Included in .NET 2.0 SP1, which is included during OS installation	.NET 2.0 SP1	Included in OS package



NOTES:

- Starting BACS 3 without .NET Framework (version 2.0 or above) installed on your system will result in an error.
- Long startup times for BACS 3 may be caused by a limitation of .NET Framework 2.0, as described in Microsoft knowledge base article KB936707. To improve the startup time of BACS 3, apply .NET Framework 2.0 SP1 (or appropriate version for your OS as shown in the Recommended column of [Table 1](#)), and reboot your system.

USING THE INSTALLER

To install the management applications



NOTE: Before starting the installation, verify that .NET Framework 2.0 or above is installed. This ensures optimal performance of BACS 3.

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 **I accept the terms in the license agreement** and then click **Next** to continue.
5. Select the features you want installed. If you receive a warning regarding .NET Framework 2.0, then press OK to continue installing BACS 3 and manually install .NET Framework 2.0 when installation completes, or Cancel to quit.
6. Click **Next**.



7. Click **Install**.
8. Click **Finish** to close the wizard.

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 3 (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 and wait for the install to complete before continuing with the next command line, type the following:

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

MODIFYING MANAGEMENT APPLICATIONS

To modify 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 **Modify** to change program features.
5. Click **Next** to continue.
6. Click on an icon to change how a feature is installed.
7. Click **Next**.
8. Click **Install**.
9. Click **Finish** to close the wizard.
10. Reboot your system to complete the modification of the management applications.

REPAIRING MANAGEMENT APPLICATIONS

To repair or reinstall 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.

NDIS2 Driver Software: Broadcom NetXtreme II™ Network Adapter User Guide

- [Overview](#)
- [Preinstallation Requirements](#)
- [Installing the NDIS2 Driver Software for Use on MS-DOS Platforms](#)
- [Using Keywords for the Drivers](#)

OVERVIEW

Two drivers are discussed in this section:

- BXND20X: Broadcom NetXtreme II Gigabit Ethernet driver
- BNX2EV: Broadcom NetXtreme II 10 Gigabit Ethernet driver

The examples used in this section refer to the BXND20X driver, but also apply to the BNX2EV driver.

PREINSTALLATION REQUIREMENTS

Before you can successfully install the NDIS2 driver software, the Broadcom network adapter must be physically installed in the server. Networking software that is appropriate to the operating system (such as Microsoft LAN Manager 2.2 for MS-DOS) must already be running on your server.

INSTALLING THE NDIS2 DRIVER SOFTWARE FOR USE ON MS-DOS PLATFORMS

The NDIS2 driver software can be run from an MS-DOS startup disk using Microsoft Network Client 3.0 or from the hard disk using Microsoft LAN Manager 2.2.

CREATING A STARTUP DISK TO RUN MICROSOFT NETWORK CLIENT

To perform this installation you must have the following items

- Windows NT Server 4.0 CD-ROM
- A blank MS-DOS system disk (3.5" high-density floppy disk)
- Access to the Broadcom NDIS2 driver file (BXND20X.dos). This file is located on the driver source media.



NOTES:

- Windows NT Server 4.0 users. When running Setup for Microsoft Network Client v3.0 for MS-DOS, click any network card from the list (NE2000 Compatible, for example) to create the startup disk.
- After creating the startup disk, follow the instructions in [Modifying the Startup Disk](#).

To create a startup disk

1. Create a folder called NCADMIN in the root of the C drive.
2. Copy the NCADMIN.CN_, NCADMIN.EX_, and NCADMIN.HL_ files from the I386 folder on the Windows NT Server 4.0 CD-ROM.
3. Open a command prompt window and change the directory to C:\NCADMIN.
4. Type `expand -r ncadmin.*` and press **ENTER**.
5. Close the command prompt window by typing `exit` and then pressing **ENTER**.
6. Start Windows Explorer.
7. Open the NCADMIN folder and double-click **ncadmin.exe**.
8. Follow the on-screen instructions to make the network startup disk (choose **NE2000 Compatible** from the list of adapters).

Modifying the Startup Disk

To modify the startup disk

1. Edit A:\Net\Protocol.ini with Notepad or a similar text editor.
 - a. Change `DriverName=$` to `DriverName=BXND20X$`.
 - b. Remove all other parameter entries under the `[MS$NE2CLONE]` or equivalent section such as `IOBASE=0x300` or `INTERRUPT=3`, and so on.

Example Protocol.ini file for IP

```
[network.setup]
version=0x3110
netcard=ms$ne2clone,1,MS$NE2CLONE,1
transport=tcPIP,TCPIP
lana0=ms$ne2clone,1,tcPIP
[MS$NE2CLONE]
DriverName=BXND20X$
[protman]
DriverName=PROTMAN$
PRIORITY=MS$NDISHLP
[tcPIP]
NBSessions=6
DefaultGateway=0
SubNetMask=255 0 0 0
IPAddress=192 168 0 1
DisableDHCP=0
DriverName=TCPIP$
BINDINGS=MS$NE2CLONE
LANABASE=0
```

Example Protocol.ini file for IPX

```
[network.setup]
version=0x3110
netcard=ms$ne2clone,1,MS$NE2CLONE,1
transport=ms$ndishlp,MS$NDISHLP
transport=ms$nwlink,MS$NWLINK
lana0=ms$ne2clone,1,ms$nwlink
lana1=ms$ne2clone,1,ms$ndishlp
[MS$NE2CLONE]
DriverName=BXND20X$
[protman]
DriverName=PROTMAN$
PRIORITY=MS$NDISHLP
[MS$NDISHLP]
DriverName=ndishlp$
BINDINGS=ms$ne2clone
[ms$nwlink]
DriverName=nwlink$
FRAME=Ethernet_802.2
BINDINGS=MS$NE2CLONE
LANABASE=0
```

Example Protocol.ini file for NetBEUI

```
[network.setup]
version=0x3110
netcard=ms$ne2clone,1,MS$NE2CLONE,1
transport=ms$ndishlp,MS$NDISHLP
transport=ms$netbeui,MS$NETBEUI
lana0=ms$ne2clone,1,ms$ndishlp
lana1=ms$ne2clone,1,ms$netbeui
[MS$NE2CLONE]
DriverName=BXND20X$
[protman]
DriverName=PROTMAN$
PRIORITY=MS$NDISHLP
[MS$NDISHLP]
DriverName=ndishlp$
BINDINGS=MS$NE2CLONE
[MS$NETBEUI]
DriverName=netbeui$
SESSIONS=10
NCBS=12
BINDINGS=MS$NE2CLONE
LANABASE=0
```

2. Edit A:\Net\System.ini.
 - a. Change netcard= to netcard=BXND20X.dos.
 - b. Check for references to C:\NET and change C:\NET to A:\NET if necessary.

Example System.ini file

```
[network]
sizeworkbuf=1498
filesharing=no
printsharing=no
autologon=yes
computername=MYPC
lanroot=A:\NET
username=USER1
workgroup=WORKGROUP
reconnect=yes
dospophotkey=N
lmlogon=0
logondomain=
preferredredir=basic
autostart=basic
maxconnections=8
[network drivers]
netcard=BXND20X.dos
transport=ndishlp.sys,*netbeui
devdir=A:\NET
LoadRMDrivers=yes
```

3. Copy **BXND20X.dos** to A:\Net.
4. Create the appropriate Autoexec.bat file in drive A for the chosen protocol as shown below.

For TCP/IP

```
path=a:\net
a:\net\net initialize
a:\net\netbind.com
a:\net\umb.com
a:\net\tcptsr.exe
a:\net\tinyrfc.exe
a:\net\nmtsr.exe
a:\net\emsbfr.exe
a:\net\net start basic
net use z: \\SERVERNAME\SHARENAME
```

For IPX

```
SET PATH=A:\NET
A:\NET\net initialize
A:\NET\nwlink
A:\NET\NET START BASIC
net use z: \\SERVERNAME\SHARENAME
```

For NetBEUI

```
SET PATH=A:\NET
A:\NET\NET START BASIC
net use z: \\SERVERNAME\SHARENAME
```

5. Create a Config.sys file on the startup disk in drive A as shown below.


```
files=30
device=a:\net\ifshlp.sys
```

```
lastdrive=z
```

INSTALLING THE DOS NDIS2 DRIVER SOFTWARE ON THE HARD DISK

To install the DOS NDIS2 Driver Software on the hard disk

1. Verify that the system has Microsoft LAN Manager 2.2 installed, with a protocol such as NetBEUI configured.
2. Create a folder on your hard disk to store the NDIS 2.01 driver.

Example: C:\LANMAN

3. Copy the **BXND20X.dos** file to this folder.

Edit the **Config.sys** file by adding the following lines:

```
DEVICE = C:\LANMAN\PROTMAN.DOS
```

```
DEVICE = C:\LANMAN\BXND20X.DOS
```

```
DEVICE = C:\LANMAN\NETBEUI.DOS
```

4. Edit the **Autoexec.bat** file by adding the following lines:

```
C:\ LANMAN\NETBIND.EXE
```

```
C:\LANMAN\NET START WORKSTATION
```

```
C:\LANMAN\NET USE drive letter: \\server name\resource name
```

5. Edit the **Protocol.ini** file (located in C:\LANMAN) to configure the driver to bind with NetBEUI or any other protocols.

Example:

```
[PROTOCOL MANAGER]
```

```
DriverName = PROTMAN$
```

```
[NETBEUI_XIF]
```

```
DriverName = netbeui$
```

```
BINDINGS = BXND20X
```

```
[BXND20X]
```

```
DriverName = "BXND20X$"
```

6. Restart the computer to complete the installation.



NOTE: The driver loads during system configuration and displays the Broadcom banner, controller name, MAC address, IRQ number, detected line speed, and the controller BusNum and DevNum. If the driver fails to load, an *initialization fail* message is displayed.

USING KEYWORDS FOR THE DRIVERS

The Protocol.ini file contains certain keywords that are used by the BXND20X.dos AND BXND20X.dos drivers. These keywords are listed below:

BusNum. Specifies the number of the PCI bus on which the network adapter is located. Requires a decimal number having a value ranging from 0 to 255.

DevNum. Specifies the device number assigned to the network adapter when it is configured by the PCI BIOS. Requires a decimal number having a value ranging from 0 to 255.

FuncNum or PortNum. Specifies the PCI function or port number assigned to the network controller. Requires a decimal number having a value ranging from 0 to 7.



NOTE: These keywords, **BusNum**, **DevNum**, and **FuncNum** (or **PortNum**) are needed when multiple adapters are installed in the server and when a specific controller must be loaded in a certain order. These keywords are used concurrently and are included for manufacturing purposes. Do not use them unless you are familiar with how to configure PCI devices. A PCI device scan utility is needed to find this information.

LineSpeed. Specifies the speed of the network connection in Mbit/s. Requires the decimal number 10, 100, or 1000. Technically, a line speed of 1000 Mbit/s cannot be forced and is achievable only through auto-negotiation. For the sake of simplicity, the driver performs auto-negotiation when the line speed is set to a value of 1000.



NOTE: LineSpeed is not available with the Broadcom NetXtreme II 10 Gigabit Ethernet driver.

Duplex. Specifies the duplex mode of the network adapter. Requires a setting of either **Half** or **Full**. When this keyword is used, the **LineSpeed** keyword must also be used. If neither keyword is used, the network adapter defaults to auto-negotiation mode.



NOTE: LineSpeed is not available with the Broadcom NetXtreme II 10 Gigabit Ethernet driver.

NodeAddress. Specifies the network address used by the network adapter. If a multicast address or a broadcast address is specified, the adapter uses the default MAC address.

Example:

```
[BXND20X]
DriverName = "BXND20X$"
BusNum = 3
DevNum = 14
PortNum = 2
LineSpeed = 1000
Duplex = Full
NodeAddress = 001020304050
```

FixChecksumOff. Turns off the driver's workaround for the TCP/IP stack to recognize the 1s complemented version of the checksum.

AcceptAllMC. Informs the driver to deliver all multicast packets to the upper protocol.



ODI Driver Software: Broadcom NetXtreme II™ Network Adapter User Guide

- [Overview](#)
- [ODI Client32](#)
- [DOS ODI](#)

OVERVIEW

Three drivers are discussed in this section:

- BX2.LAN: Broadcom NetXtreme II Gigabit Ethernet 32-bit driver
- BX2.COM: Broadcom NetXtreme II Gigabit Ethernet 16-bit driver
- BNX2EV.COM: Broadcom NetXtreme II 10 Gigabit Ethernet 16-bit driver

The examples used in the [ODI Client32](#) section only refer to the BX2.LAN driver. The examples used in the [DOS ODI](#) section refer to the BX2.COM driver, but also apply to the BNX2EV.COM driver.

ODI CLIENT32

The Open Data-Link Interface (ODI) driver software allows LAN drivers to work together irrespective of differences between the hardware and/or the protocol stacks of communicating systems. An intervening layer, called the link support layer (LSL), is responsible for routing different packet types between local area network (LAN) drivers and protocol stacks.

INSTALLING THE DRIVER SOFTWARE

The ODI Driver software must be installed on a system having an MS-DOS operating system. MS-DOS 6.22 is recommended.

Preinstallation Requirements

1. Copy the Broadcom ODI driver files to a floppy disk. Copy the BX2.LAN and BX2.LDI files for the Client32 driver.
2. Go to the Novell web site <http://www.novell.com/cool solutions/tools/13555.html> and download the NetWare DOS Client files to the hard disk: download the dw271e.zip file for the Client32 installation.

Installing the Novell NetWare Client32 Driver for MS-DOS

1. Extract the dw271.zip file to a temporary directory on the hard disk.
2. Double-click the dw271e.exe file.
3. Type *y* and then press **ENTER** when you are prompted for a response.
4. Restart your system in MS-DOS mode.
5. Change to the directory where the extracted files are located.

6. Type `install`, and then press **ENTER**.
7. Read the agreement, and then press **ENTER** to accept the agreement.
8. Using the UP ARROW or DOWN ARROW key as necessary, highlight **Novell Client for DOS (required)**, and then press **F10** to save and continue.
9. Highlight **32-bit LAN Drivers**, and then press **F10** to save and continue.
10. Insert the floppy disk containing the Broadcom 32-bit LAN driver files into the floppy disk drive.
11. Highlight **USER SPECIFIED 32-BIT DRIVER**, and then press **ENTER**.
12. Press **ENTER**.
13. Highlight the NetXtreme II Gigabit Ethernet Driver, and then press **ENTER**.
14. Press **ENTER**.
15. Select **Receive Buffers**, type `32` for the value, and then press **F10** to save and continue.



NOTE: Do not change the default settings for any of the other items in the Installation Configuration Summary.

16. Press **F10** to save and continue. The `Bx2.LAN` and `Bx2.LDI` files are copied to the directory.
17. Restart the system to complete the install.
18. As the files are being loaded, the system attempts to attach to a NetWare server using the frame types specified in `Net.cfg`. If the connection is successful, a dialog box is displayed.
19. Change to drive F and log on to the NetWare server.



NOTE: The driver loads during system configuration, and the Broadcom banner, the adapter name, the MAC address, the IRQ number, the detected line speed, and the bus number (`BusNum`) and device number (`DevNum`) of the adapter are displayed. If the driver fails to load, an *initialization fail* message appears.

CONFIGURING THE DRIVER SOFTWARE

The ODI driver software can be configured by adding specific optional keywords to the `Net.cfg` file.

If multiple (or multiport) Broadcom NetXtreme II adapters are installed in a system, the ODI driver software loads by default on the adapter/port that has a good link. If two or more adapters have a good link, the ODI driver software loads on the adapter having the latest Device ID. If two or more adapters that have a good link have the same Device ID, the ODI driver software loads on the adapter that is located in the slot having the lowest bus number.

If it is necessary to have the ODI driver load on certain adapters in a certain order, the **BusNum**, **DevNum**, and **FuncNum** keywords can be used. Do not use these keywords unless you know how to configure PCI devices.

The **BusNum** keyword value, which represents the PCI bus number in which the adapter is located, is a hexadecimal number ranging from 0 to FF.

The **FuncNum** keyword value, which represents the function (port) number of a multiport adapter, is a hexadecimal number with 0 representing the first port and 1 representing the second port.

The **DevNum** keyword value, which represents the assigned device number, is a hexadecimal number ranging from 0 to 1F.



**NOTES:**

- At the end of the ODI driver software installation process, note the **BusNum** and **DevNum** values that are displayed. Alternatively, use Broadcom Advanced Control Suite 3 (see [Viewing Resource Information](#)) to view the bus number, function (port) number, and device number assigned to each adapter (Windows users only).
- The values for the **BusNum**, **DevNum**, and **FuncNum** keywords are displayed in hexadecimal notation when the driver loads on the adapter.

Example **BusNum**, **DevNum**, and **FuncNum** keyword entries for loading the ODI driver on multiple adapters in a certain order are shown below:

```
PROTOCOL <protocol name>
  BIND BX2
  BIND BX2
LINK DRIVER BX2
  BUSNUM 1
  DEVNUM 9
  FRAME <frame #1 name>
LINK DRIVER BX2
  BUSNUM 1
  DEVNUM A
  FRAME <frame #2 name>
```

The **LineSpeed** keyword is used to force the speed of the network connection. The **LineSpeed** keyword requires a decimal number and of either 10, 100, or 1000. Technically, a line speed of 1000 Mbit/s cannot be forced and can be achieved only through auto-negotiation. For the sake of simplicity, the driver performs auto-negotiation when the line speed is set to a value of 1000.

The **Duplex** keyword is used to force the duplex mode of the adapter. The **Duplex** keyword requires a text string of either HALF or FULL. When the **Duplex** keyword is used, the **LineSpeed** keyword must also be used. If neither keyword is used, the network adapter defaults to auto-negotiation mode.

Example keyword entries for the **LineSpeed**, and **Duplex** keywords are shown below:

```
LINK DRIVER BX2
LINESPEED 100
DUPLEX FULL
```

DOS ODI

The Open Data-Link Interface (ODI) driver software allows LAN drivers to work together irrespective of differences between the hardware and/or the protocol stacks of communicating computers. An intervening layer, called the link support layer (LSL), is responsible for routing different packet types between LAN drivers and protocol stacks.



NOTE: The examples used in this section refer to the BX2.COM driver, but also apply to the BNX2EV.COM driver.

INSTALLING THE DRIVER SOFTWARE

Installing the DOS/ODI driver software involves the following steps:

1. Verifying that the DOS 6.22 operating system files are installed on your computer hard disk or are available on a removable storage device.
2. Installing the DOS/ODI driver software in the DOS environment.
3. If necessary, reconfiguring the driver software after installation.

Preinstallation Requirements

1. Verify that the Broadcom Ethernet adapter driver software has been installed on your computer and install it if it has not (see [Installing the Driver Software](#)). Otherwise, you cannot use a Broadcom Ethernet adapter in the DOS environment.



NOTE: To successfully install the adapter driver for DOS, the adapter must be installed in the computer before you install the driver.

2. Install the DOS 6.22 files on your computer hard disk or have them available on a removable storage device.

Installing the DOS ODI Driver on DOS 6.22

1. Copy the DOS ODI driver file to the root directory of the hard disk.
2. Copy Lsl.com, lpxodi.com, and Vlm.com to the same directory (these files are provided by Novell).
3. Edit the Autoexec.bat file by adding the following lines:
C:\LSL
C:\BX2
C:\IPXODI
C:\VLM
4. Edit the Net.cfg file and place the file in the same location as the Autoexec.bat file (located in drive C) to configure the driver to bind with different frame types.

Example:

```
LINK DRIVER BX2
FRAME ETHERNET_802.2
FRAME ETHERNET_802.3
FRAME ETHERNET_II
FRAME ETHERNET_SNAP
```

```
NetWare DOS Requester
FIRST NETWORK DRIVE=F
NETWARE PROTOCOL=NDS BIND
```

PROTOCOL IPXODI

5. Restart the computer to complete the installation.
6. As the files are being loaded, the computer attempts to attach to a NetWare server using the frame types specified in Net.cfg. If the connection is successful, a dialog box appears.
7. Change to drive F and log onto the NetWare server.
 F: <enter>
 F:\login admin <enter>



NOTE: The driver loads during system configuration, and the Broadcom banner, the adapter name, the MAC address, the IRQ number, the detected line speed, and the BUSNUM and DEVNUM of the adapter are displayed. If the driver fails to load, an *initialization fail* message appears.

CONFIGURING PARAMETERS FOR THE BX2.COM/BX2EV.COM DRIVER (IN THE NET.CFG FILE)

The following are optional custom keywords used in the Net.cfg file:

BUSNUM
DEVNUM
FUNNUM
SUBVID
MODEL

The first two keywords (**BUSNUM** and **DEVNUM**) are used together and have been included for manufacturing purposes. Do not use them unless you are familiar with PCI device configuration. These two keywords are needed if multiple Broadcom NetXtreme II adapters are installed on a system, and each adapter must be loaded in a specific order.

BUSNUM specifies the PCI bus number and **DEVNUM** specifies the device number that is assigned to the board when it is configured by the PCI BIOS.

The values for these adapter parameters are displayed in hexadecimal notation when the driver loads (the Net.cfg file expects hexadecimal values). A PCI device scan utility can also be used to find this information. If the driver loads on a single adapter, the displayed values can be used in the Net.cfg file.

If a PCI scan utility displays card 1 at bus 4 and device 11 (hexadecimal), and card 3 at bus 3 and device 13 (hexadecimal), the following load instructions are written to the Net.cfg file (load card 3 first, and optionally load card 1 second):

```
LINK DRIVER BX2
BUSNUM 3
DEVNUM 13 (13 hex = 19 decimal)
LINK DRIVER BX2
BUSNUM 4
DEVNUM 11 (11 hex = 17 decimal)
```

The **FUNNUM** keyword is used to determine which function/port the driver is to load on when using a multifunction (multiport) adapter. Function number values are limited to 0 and 1, for the first and second function of the adapter.

The **SUBVID** keyword is used to for load the driver only on the first network interface card found that has a subvendor ID (SVID) that matches the one set in Net.cfg. The driver ignores all other network interface cards that do not have a matching SVID. **SUBVID** expects a hexadecimal number.

Example:

```
LINK DRIVER BX2
SUBVID 14E4
```

SUBVID cannot be used with **BUSNUM**, **DEVNUM** or **FUNNUM**.

The **MODEL** keyword is for loading the driver only on the first network interface card found that has a subvendor ID (SVID) that matches the one set in **MODEL**. The driver ignores all other NICs that do not have a matching SVID. **MODEL** expects a decimal number.

Example:

```
LINK DRIVER B44
MODEL 5348
```


Linux Driver Software: Broadcom NetXtreme II™ Network Adapter User Guide

- [Introduction](#)
- [Limitations](#)
- [Packaging](#)
- [Installing Linux Driver Software](#)
- [Unloading/Removing the Linux Driver](#)
- [Patching PCI Files \(Optional\)](#)
- [Network Installations](#)
- [Setting Values for Optional Properties](#)
- [Driver Defaults](#)
- [Driver Messages](#)
- [Teaming with Channel Bonding](#)
- [Statistics](#)
- [Linux iSCSI Offload](#)

INTRODUCTION

This section discusses the Linux drivers for the Broadcom NetXtreme II network adapters.

Table 1: Broadcom NetXtreme II Linux Drivers

Linux Driver	Description
bnx2 cnic	Linux drivers for the BCM5706/BCM5708/BCM5709 10/100/1000/2500 Mbps network adapters. The bnx2 driver is the networking driver and the cnic driver supports additional features required by the bnx2i iSCSI offload driver.
bnx2x	Linux driver for the BCM57710/BCM57711 10 Gbps network adapters
bnx2i	Linux driver to enable iSCSI offload on the BCM5706/BCM5708/BCM5709/BCM57710/BCM57711 10/100/1000/2500/10000 Mbps network adapters

LIMITATIONS

- [bnx2 Driver](#)
- [bnx2x Driver](#)
- [bnx2i Driver](#)

BNX2 DRIVER

The current version of the driver has been tested on 2.4.x kernels (starting from 2.4.24) and all 2.6.x kernels. The driver may not compile on kernels older than 2.4.24.



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

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. Additionally, the Makefile will not compile the cnic driver on kernels older than 2.6.16. iSCSI offload is only supported on 2.6.16 and newer kernels.



NOTE: For Broadcom NetXtreme II BCM5708 devices with a silicon revision prior to B2, the open source bnx2 driver does not support the reporting and configuration of NetXtreme II WOL settings via ethtool. For silicon revisions of B2 or later, the bnx2 driver reports support for Magic Packet WOL via ethtool. Enabling support via ethtool is mandatory to successfully wake the system. To determine the silicon revision of your Broadcom NetXtreme II device, use the `lspci` command, where “10” = revision B0, “11” = revision B1, and “12” = revision B2.

BNX2X DRIVER

The current version of the driver has been tested on 2.6.x kernels starting from 2.6.9. The driver may not compile on kernels older than 2.6.9. Testing is concentrated on i386 and x86_64 architectures. Only limited testing has been done on some other architectures. Minor changes to some source files and Makefile may be needed on some kernels.

BNX2I DRIVER

The current version of the driver has been tested on 2.6.x kernels, starting from 2.6.18 kernel. The driver may not compile on older kernels with the exception of SLES 10 SP1, which runs 2.6.16.46 kernel. SUSE upgraded the `iscsi_transport` kernel module in SLES 10 SP1, and Broadcom iSCSI offload initiators is supported on SLES 10 SP1. Testing is concentrated on i386 and x86_64 architectures, Red Hat EL5, and SUSE 10 SP1 distributions. Support added for SLES11U0 running 2.6.27 kernel.

PACKAGING

The Linux driver is released in the packaging formats shown in [Table 2](#). The netxtreme2 package contains the bnx2 (1 Gb network adapter) driver, the bnx2x (10 Gb network adapter) driver, and the bnx2i (iSCSI) driver for source RPM and compressed tar.

Table 2: Linux Driver Packaging

Format	Drivers
Source RPM	netxtreme2- <i>version</i> .src.rpm
Compressed tar	netxtreme2- <i>version</i> .tar.gz
Supplemental tar	netxtreme2_sup- <i>version</i> .tar.gz

Identical source files to build the driver are included in both RPM and TAR source packages. The supplemental tar file contains additional utilities such as patches and driver diskette images for network installation.

Updated open-iSCSI components are released in source RPM format.

- iscsi-initiator-utils-6.2.0.868-0.18e.src.rpm (Red Hat EL 5.3)
- iscsi-initiator-utils-6.2.0.871-0.10d.src.rpm (Red Hat EL 5.4)
- open-iscsi-2.0.707-0.25c.src.rpm
- open-iscsi-2.0.707-0.44f.src.rpm (SLES 10 SP2)
- open-iscsi-2.0.868-0.5.1b.src.rpm (SLES 10 SP3)

INSTALLING LINUX DRIVER SOFTWARE

- [Installing the Source RPM Package](#)
- [Building the Driver from the Source TAR File](#)



NOTE: If a bnx2/bnx2x/bnx2i driver is loaded and the Linux kernel is updated, the driver module must be recompiled if the driver module was installed using the source RPM or the TAR package.

INSTALLING THE SOURCE RPM PACKAGE



NOTE: The examples in this procedure refer to the bnx2 driver, but also apply to the bnx2x and bnx2i drivers.

1. Install the source RPM package:

```
rpm -ivh netxtreme2-version.src.rpm
```

2. Change the directory to the RPM path and build the binary driver for your kernel (the RPM path is different for different Linux distributions):

```
cd /usr/src/redhat,OpenLinux,turbo,packages,rpm ...
```

```
rpm -bb SPECS/bnx2.spec
```

or

```
rpmbuild -bb SPECS/netxtreme2.spec (for RPM version 4.x.x)
```



NOTE: During your attempt to install a source RPM package, the following message may be displayed:

```
error: cannot create %sourcedir /usr/src/redhat/SOURCES
```

The most likely cause of the error is that the rpm-build package has not been installed. Locate the rpm-build package on the Linux installation media and install it using the following command:

```
rpm -ivh rpm-build-version.arch.rpm
```

Complete the installation of the source RPM.

3. Install the newly built package (driver and man page):

```
rpm -ivh RPMS/i386/netxtreme2-version.arch.rpm
```

The --force option is needed if installing over an existing distribution that may already contain an older version of the driver.

Depending on the kernel, the driver is installed to one of the following paths:

2.4.x kernels

```
/lib/modules/kernel_version/kernel/drivers/net/bnx2.o
```

2.6.x kernels

```
/lib/modules/kernel_version/kernel/drivers/net/bnx2.ko
```

For the bnx2i driver:

2.6.16 kernels and newer (bnx2 driver)

```
/lib/modules/kernel_version/kernel/drivers/net/bnx2.ko
```

```
/lib/modules/kernel_version/kernel/drivers/net/cnic.ko
```



2.6.16 kernels and newer (bnx2i driver)

```
/lib/modules/kernel_version/kernel/drivers/iscsi/bnx2i.ko
```

4. Load the driver:

```
insmod bnx2
```

or

```
modprobe bnx2
```

5. Load the cnic driver (if applicable):

```
insmod cnic.ko
```

or

```
modprobe cnic
```

To configure the network protocol and address, refer to the documentation provided with your operating system.

BUILDING THE DRIVER FROM THE SOURCE TAR FILE



NOTE: The examples used in this procedure refer to the bnx2 driver, but also apply to the bnx2x driver.

1. Create a directory and extract the TAR files to the directory:

```
tar xvzf netxtreme2-version.tar.gz
```

2. Build the driver bnx2.ko (or bnx2.o) as a loadable module for the running kernel:

```
cd netxtreme2-version  
make
```

3. Test the driver by loading it (first unload the existing driver, if necessary):

```
rmmod bnx2  
insmod bnx2.o  
modprobe crc32 && insmod bnx2.o  
or, for Linux 2.6 kernels:  
rmmod bnx2  
insmod bnx2.ko
```

No message should be returned if this command runs properly

4. Load the cnic driver (if applicable):

```
insmod cnic.ko
```

5. Install the driver and man page:

```
make install
```



NOTE: See the RPM instructions above for the location of the installed driver.

To configure the network protocol and address after building the driver, refer to the manuals supplied with your operating system.

LOAD AND RUN NECESSARY iSCSI SOFTWARE COMPONENTS

Broadcom iSCSI Offload software suite consists of 3 kernel modules and a user daemon. Required software components can be loaded either manually or through system services.

1. Unload existing driver and kill user daemon, if necessary:

Manual:

```
rmmod bnx2i  
pkill -9 bnx2id
```

Using service:

```
service bnx2id stop
```

2. Load the iSCSI driver and the user daemon:

Manual:

```
bnx2id  
insmod bnx2i.ko
```

or

```
modprobe bnx2i
```

Using service:

```
service bnx2id start
```

UNLOADING/REMOVING THE LINUX DRIVER

- [Unloading/Removing the Driver from an RPM Installation](#)
- [Removing the Driver from a TAR Installation](#)

UNLOADING/REMOVING THE DRIVER FROM AN RPM INSTALLATION

**NOTES:**

- The examples used in this procedure refer to the bnx2 driver, but also apply to the bnx2x driver.
- On 2.6 kernels, it is not necessary to bring down the eth# interfaces before unloading the driver module.
- If the cnic driver is loaded, unload the cnic driver before unloading the bnx2 driver.
- Prior to unloading the bnx2i driver, disconnect all active iSCSI sessions to targets.

To unload the driver, use ifconfig to bring down all eth# interfaces opened by the driver, and then type the following:

```
rmmod bnx2
```

If the driver was installed using rpm, do the following to remove it:

```
rpm -e netxtreme2
```

REMOVING THE DRIVER FROM A TAR INSTALLATION



NOTE: The examples used in this procedure refer to the bnx2 driver, but also apply to the bnx2x and bnx2i drivers.

If the driver was installed using make install from the tar file, the bnx2.o or bnx2.ko driver file has to be manually deleted from the operating system. See [Installing the Source RPM Package](#) for the location of the installed driver.

PATCHING PCI FILES (OPTIONAL)



NOTE: The examples used in this procedure refer to the bnx2 driver, but also apply to the bnx2x and bnx2i drivers.

For hardware detection utilities such as Red Hat kudzu to properly identify bnx2 supported devices, a number of files containing PCI vendor and device information may need to be updated.

Apply the updates by running the scripts provided in the supplemental tar file. For example, on Red Hat Enterprise Linux, apply the updates by doing the following:

```
./patch_pcitbl.sh /usr/share/hwdata/pcitable pci.updates  
/usr/share/hwdata/pcitable.new bnx2  
./patch_pciids.sh /usr/share/hwdata/pci.ids pci.updates  
/usr/share/hwdata/pci.ids.new
```

Next, the old files can be backed up and the new files can be renamed for use.

```
cp /usr/share/hwdata/pci.ids /usr/share/hwdata/old.pci.ids
cp /usr/share/hwdata/pci.ids.new /usr/share/hwdata/pci.ids
cp /usr/share/hwdata/pcitable /usr/share/hwdata/old.pcitable
cp /usr/share/hwdata/pcitable.new /usr/share/hwdata/pcitable
```

NETWORK INSTALLATIONS

For network installations through NFS, FTP, or HTTP (using a network boot disk or PXE), a driver disk that contains the bnx2/bnx2x driver may be needed. The driver disk images for the most recent Red Hat and SuSE versions are included. Boot drivers for other Linux versions can be compiled by modifying the Makefile and the make environment. Further information is available from the Red Hat website, <http://www.redhat.com>.

SETTING VALUES FOR OPTIONAL PROPERTIES

Optional properties exist for the different drivers:

- [bnx2 Driver](#)
- [bnx2x Driver](#)
- [bnx2i Driver](#)

BNX2 DRIVER

disable_msi

The **disable_msi** optional property can be supplied as a command line argument to the insmod or modprobe command. The property can also be set in modprobe.conf. See the man page for more information. All other driver settings can be queried and changed using the ethtool utility. See the ethtool man page for more information. The ethtool settings do not persist across a reboot or module reload. The ethtool commands can be put in a startup script such as /etc/rc.local to preserve the settings across a reboot.



NOTE: Some combinations of property values may conflict and result in failures. The driver cannot detect all such conflicting combinations.

This property is used to disable Message Signal Interrupts (MSI), and the property is valid only on 2.6 kernels that support MSI. On 2.4 kernels, this property cannot be used. By default, the driver enables MSI if it is supported by the kernel. It runs an interrupt test during initialization to determine if MSI is working. If the test passes, the driver enables MSI. Otherwise, it uses legacy INTx mode.

```
insmod bnx2.ko disable_msi=1
```

or

```
modprobe bnx2 disable_msi=1
```


BNX2X DRIVER

disable_tpa

The **disable_tpa** parameter can be supplied as a command line argument to disable the Transparent Packet Aggregation (TPA) feature. By default, the driver will aggregate TCP packets. Use `disable_tpa` to disable the advanced TPA feature.

Set the **disable_tpa** parameter to 1 as shown below to disable the TPA feature on all NetXtreme II network adapters in the system. The parameter can also be set in `modprobe.conf`. See the man page for more information.

```
insmod bnx2x.ko disable_tpa=1
```

or

```
modprobe bnx2x disable_tpa=1
```

int_mode

The **int_mode** parameter is used to force using an interrupt mode.

Set the **int_mode** parameter to 1 to force using the legacy INTx mode on all NetXtreme II adapters in the system.

```
insmod bnx2x.ko int_mode=1
```

or

```
modprobe bnx2x int_mode=1
```

Set the **int_mode** parameter to 2 to force using MSI mode on all NetXtreme II adapters in the system.

```
insmod bnx2x.ko int_mode=2
```

or

```
modprobe bnx2x int_mode=2
```

Set the **int_mode** parameter to 3 to force using MSI-X mode on all NetXtreme II adapters in the system.

dropless_fc

The **dropless_fc** parameter can be used to enable a complementary flow control mechanism on BCM57711 adapters. The default flow control mechanism is to send pause frames when the on-chip buffer (BRB) is reaching a certain level of occupancy. This is a performance targeted flow control mechanism. On BCM57711 adapters, one can enable another flow control mechanism to send pause frames, where one of the host buffers (when in RSS mode) are exhausted.

This is a "zero packet drop" targeted flow control mechanism.

Set the **dropless_fc** parameter to 1 to enable the dropless flow control mechanism feature on all BCM57711 NetXtreme II adapters in the system.

```
insmod bnx2x.ko dropless_fc=1
```

or

```
modprobe bnx2x dropless_fc=1
```



BNX2I DRIVER

Optional parameters **en_tcp_dack**, **error_mask1**, and **error_mask2** can be supplied as command line arguments to the `insmod` or `modprobe` command for `bnx2i`.

error_mask1 and **error_mask2**

"Config FW iSCSI Error Mask #", use to configure certain iSCSI protocol violation to be treated either as a warning or a fatal error. All fatal iSCSI protocol violations will result in session recovery (ERL 0). These are bit masks.

Defaults: All violations will be treated as errors.



CAUTION! Do not use **error_mask** if you are not sure about the consequences. These values are to be discussed with Broadcom development team on a case-by-case basis. This is just a mechanism to work around iSCSI implementation issues on the target side and without proper knowledge of iSCSI protocol details, users are advised not to experiment with these parameters.

en_tcp_dack

"Enable TCP Delayed ACK", enables/disables TCP delayed ACK feature on offloaded iSCSI connections.

Defaults: TCP delayed ACK is ENABLED. For example:

```
insmod bnx2i.ko en_tcp_dack=0
```

or

```
modprobe bnx2i en_tcp_dack=0
```

sq_size

"Configure SQ size", used to choose send queue size for offloaded connections and SQ size determines the maximum SCSI commands that can be queued. SQ size also has a bearing on the number of connections that can be offloaded; as QP size increases, the number of connections supported will decrease. With the default values, the BCM5706/BCM5708 adapter can offload 28 connections.

Defaults: 128

Range: 32 to 128

Note that Broadcom validation is limited to a power of 2; for example, 32, 64, 128.

rq_size

"Configure RQ size", used to choose the size of asynchronous buffer queue size per offloaded connections. RQ size is not required greater than 16 as it is used to place iSCSI ASYNC/NOP/REJECT messages and SCSI sense data.

Defaults: 16

Range: 16 to 32



Note that Broadcom validation is limited to a power of 2; for example, 16, 32.

event_coal_div

"Event Coalescing Divide Factor", performance tuning parameter used to moderate the rate of interrupt generation by the iscsi firmware.

Defaults: 1

Valid values: 1, 2, 4, 8

bnx2i_nopout_when_commands_active

"iSCSI NOOP even when connection is not idle", this parameter enables the offload initiator to send iSCSI NOP-OUT on wire even when the link is not truly idle. This was introduced to avoid unnecessary session recoveries induced by some older targets when iSCSI NOP-OUT and iSCSI CMD pdus are intermixed. Newer iSCSI target systems are immune to this condition and this parameter is turned ON for quite some time.

Defaults: 1

Values: Binary parameter, 0/1

The parameters can also be set in modprobe.conf. See the man page for more information.

DRIVER DEFAULTS

- [bnx2 Driver](#)
- [bnx2x Driver](#)

BNX2 DRIVER

Speed: Autonegotiation with all speeds advertised

Flow Control: Autonegotiation with RX and TX advertised

MTU: 1500 (range is 46–9000)

RX Ring Size: 255 (range is 0–4080)

RX Jumbo Ring Size: 0 (range 0–16320) adjusted by the driver based on MTU and RX Ring Size

TX Ring Size: 255 (range is (MAX_SKB_FRAGS+1)–255). MAX_SKB_FRAGS varies on different kernels and different architectures. On a 2.6 kernel for x86, MAX_SKB_FRAGS is 18.

Coalesce RX Microseconds: 18 (range is 0–1023)

Coalesce RX Microseconds IRQ: 18 (range is 0–1023)

Coalesce RX Frames: 6 (range is 0–255)



Coalesce RX Frames IRQ: 6 (range is 0–255)

Coalesce TX Microseconds: 80 (range is 0–1023)

Coalesce TX Microseconds IRQ: 80 (range is 0–1023)

Coalesce TX Frames: 20 (range is 0–255)

Coalesce TX Frames IRQ: 20 (range is 0–255)

Coalesce Statistics Microseconds: 999936 (approximately 1 second) (range is 0–16776960 in increments of 256)

MSI: Enabled (if supported by the 2.6 kernel and the interrupt test passes)

TSO: Enabled (on 2.6 kernels)

WoL: Initial setting based on NVRAM's setting

BNX2X DRIVER

Speed: Autonegotiation with all speeds advertised

Flow control: Autonegotiation with RX and TX advertised

MTU: 1500 (range is 46–9000)

RX Ring Size: 4078 (range is 0–4078)

TX Ring Size: 4078 (range is (MAX_SKB_FRAGS+4)–4078). MAX_SKB_FRAGS varies on different kernels and different architectures. On a 2.6 kernel for x86, MAX_SKB_FRAGS is 18.

Coalesce RX Microseconds: 25 (range is 0–3000)

Coalesce TX Microseconds: 50 (range is 0–12288)

Coalesce Statistics Microseconds: 999936 (approximately 1 second) (range is 0–16776960 in increments of 256)

MSI-X: Enabled (if supported by the 2.6 kernel and the interrupt test passes)

TSO: Enabled

WoL: Disabled

DRIVER MESSAGES

The following are the most common sample messages that may be logged in the `/var/log/messages` file. Use `dmesg -n <level>` to control the level at which messages appear on the console. Most systems are set to level 6 by default. To see all messages, set the level higher.

- [bnx2 and bnx2x Driver](#)
- [bnx2i Driver](#)

BNX2 AND BNX2X DRIVER



NOTE: The examples used in this procedure refer to the bnx2 driver, but also apply to the bnx2x driver.

Driver Sign on

```
Broadcom NetXtreme II Gigabit Ethernet Driver
bnx2 v1.6.3c (July 23, 2007)
```

CNIC Driver Sign on (bnx2 only)

```
Broadcom NetXtreme II cnic v1.1.19 (Sep 25, 2007)
```

NIC Detected

```
eth#: Broadcom NetXtreme II BCM5708 1000Base-T (B0)
PCI-X 64-bit 133MHz found at mem f6000000, IRQ 16, node addr 0010180476ae
```

```
cnic: Added CNIC device: eth0
```

Link Up and Speed Indication

```
bnx2: eth# NIC Link is Up, 1000 Mbps full duplex
```

Link Down Indication

```
bnx2: eth# NIC Link is Down
```

MSI enabled successfully (bnx2 only)

```
bnx2: eth0: using MSI
```

MSI-X enabled successfully (bnx2x only)

```
bnx2x: eth0: using MSI-X
```

BNX2I DRIVER

BNX2I Driver signon

```
Broadcom NetXtreme II iSCSI Driver bnx2i v1.0.30 (Sep 29, 2007)
```

Network port to iSCSI transport name binding

```
bnx2i: netif=eth2, iscsi=bcm570x-050000
bnx2i: netif=eth1, iscsi=bcm570x-030c00
```

Driver completes handshake with iSCSI offload-enabled CNIC device

```
bnx2i [05:00.00]: ISCSI_INIT passed
```





NOTE: This message is displayed only when the user attempts to make an iSCSI connection.

Driver detects iSCSI offload is not enabled on the CNIC device

```
bnx2i: iSCSI not supported, dev=eth3
bnx2i: bnx2i: LOM is not enabled to offload iSCSI connections, dev=eth0
```

Driver unable to allocate TCP port for iSCSI connection

```
bnx2i: run 'bnx2id' daemon to alloc TCP ports
```

Exceeds maximum allowed iSCSI connection offload limit

```
bnx2i: unable to allocate iSCSI context resources
```

Network route to target node and transport name binding are two different devices

```
bnx2i: conn bind, ep=0x... ($ROUTE_HBA) does not belong to hba $USER_CHOSEN_HBA
```

where ROUTE_HBA --> net device on which connection was offloaded based on route information USER_CHOSEN_HBA --> HBA to which target node is bound (using iscsi transport name)

Target cannot be reached on any of the CNIC devices

```
bnx2i: check route, cannot connect using cnic
```

Network route is assigned to network interface, which is down

```
bnx2i: check route, hba not found
```

Attempting to offload iSCSI connection onto a Jumbo Frame-enabled device

```
bnx2i: eth# network i/f mtu is set to #mtu
bnx2i: iSCSI HBA can support mtu of 1500
```



NOTE: Change **mtu** to **1500** using ifconfig and restart the interface in order to offload iSCSI connections.

SCSI-ML initiated host reset (session recovery)

```
bnx2i: attempting to reset host, #3
```

CNIC detects iSCSI protocol violation - Fatal errors

```
bnx2i: iscsi_error - wrong StatsN rcvd
bnx2i: iscsi_error - hdr digest err
bnx2i: iscsi_error - data digest err
bnx2i: iscsi_error - wrong opcode rcvd
bnx2i: iscsi_error - AHS len > 0 rcvd
bnx2i: iscsi_error - invalid ITT rcvd
bnx2i: iscsi_error - wrong StatsN rcvd
bnx2i: iscsi_error - wrong DataSN rcvd
bnx2i: iscsi_error - pend R2T violation
bnx2i: iscsi_error - ERL0, U0
bnx2i: iscsi_error - ERL0, U1
bnx2i: iscsi_error - ERL0, U2
```

```

bnx2i: iscsi_error - ERL0, U3
bnx2i: iscsi_error - ERL0, U4
bnx2i: iscsi_error - ERL0, U5
bnx2i: iscsi_error - ERL0, U
bnx2i: iscsi_error - invalid resi len
bnx2i: iscsi_error - MRDSL violation
bnx2i: iscsi_error - F-bit not set
bnx2i: iscsi_error - invalid TTT
bnx2i: iscsi_error - invalid DataSN
bnx2i: iscsi_error - burst len violation
bnx2i: iscsi_error - buf offset violation
bnx2i: iscsi_error - invalid LUN field
bnx2i: iscsi_error - invalid R2TSN field
bnx2i: iscsi_error - invalid cmd len1
bnx2i: iscsi_error - invalid cmd len2
bnx2i: iscsi_error - pend r2t exceeds MaxOutstandingR2T value
bnx2i: iscsi_error - TTT is rsvd
bnx2i: iscsi_error - MBL violation
bnx2i: iscsi_error - data seg len != 0
bnx2i: iscsi_error - reject pdu len error
bnx2i: iscsi_error - async pdu len error
bnx2i: iscsi_error - nopin pdu len error
bnx2i: iscsi_error - pend r2t in cleanup
bnx2i: iscsi_error - IP fragments rcvd
bnx2i: iscsi_error - IP options error
bnx2i: iscsi_error - urgent flag error

```

CNIC detects iSCSI protocol violation - non-FATAL, warning

```

bnx2i: iscsi_warning - invalid TTT
bnx2i: iscsi_warning - invalid DataSN
bnx2i: iscsi_warning - invalid LUN field

```



NOTE: The driver needs to be configured to consider certain violation to treat as warning and not as a critical error.

Driver puts a session through recovery

```
conn_err - hostno 3 conn 03fbcd00, iscsi_cid 2 cid a1800
```

Reject iSCSI PDU received from the target

```

bnx2i - printing rejected PDU contents
[0]: 1 ffffffffal 0 0 0 20 0
[8]: 0 7 0 0 0 0 0 0
[10]: 0 0 40 24 0 0 ffffffff80 0
[18]: 0 0 3 ffffffff88 0 0 3 4b
[20]: 2a 0 0 2 ffffffff8c 14 0 0
[28]: 40 0 0 0 0 0 0 0

```

Open-iSCSI daemon handing over session to driver

```
bnx2i: conn update - MBL 0x800 FBL 0x800MRDSL_I 0x800 MRDSL_T 0x2000
```

TEAMING WITH CHANNEL BONDING

With the Linux drivers, you can team adapters together using the bonding kernel module and a channel bonding interface. For more information, see the Channel Bonding information in your operating system documentation.

STATISTICS

Detailed statistics and configuration information can be viewed using the ethtool utility. See the ethtool man page for more information.

LINUX ISCSI OFFLOAD

- [User Application - bnx2id](#)
- [Open iSCSI User Applications](#)
- [Bind iSCSI Target to Broadcom NX2 iSCSI Transport Name](#)
- [Making Connections to iSCSI Targets](#)
- [Maximize Offload iSCSI Connections](#)
- [Linux iSCSI Offload FAQ](#)

USER APPLICATION - BNx2ID

Run the **bnx2id** daemon before attempting to create iSCSI connections. The driver will not be able to establish connections to the iSCSI target without the daemon's assistance.

```
bnx2id
```

The bnx2id daemon requires **mknod** and **sh** shell, which are standard on any regular server. For iSCSI boot using NetXtreme II offload support, binaries for **mknod** and **sh** need to be bundled into initrd image.

OPEN ISCSI USER APPLICATIONS

Install and run the open-iscsi programs, **iscsid** and **iscsiadm**, from the Broadcom distributed open-iscsi packages. Refer to [Packaging](#) for details. All pre-installed open-iscsi packages need to be removed before the Broadcom iSCSI-supported packages can be installed.

1. Remove all existing open-iscsi packages.

RHEL5

```
rpm -e iscsi-initiator-utils
```

SLES10 SP1

```
rpm -e open-iscsi
```

2. Install the source RPM package.

```
rpm -ivh <open-iscsi-package-name>.src.rpm
```

3. CD to the RPM path and build the binary driver for your kernel.

```
cd /usr/src/{redhat,OpenLinux,turbo,packages,rpm ..}
```

```
rpm -bb SPECS/<open-iscsi-package-name>.spec
```

or

```
rpmbuild -bb SPECS/<open-iscsi-package-name>.spec (for RPM version 4.x.x)
```

Note that the RPM path is different for different Linux distributions.

4. Install the newly built package.

```
rpm -ivh RPMS/<arch>/<open-iscsi-package-name>.<arch>.rpm
```

where <arch> is the machine architecture such as i386.

```
rpm -ivh RPMS/i386/<open-iscsi-package-name>.i386.rpm
```

5. Start the daemon.

```
iscsid
```



BIND ISCSI TARGET TO BROADCOM NX2 ISCSI TRANSPORT NAME

By default, the open-iscsi daemon connects to discovered targets using software initiator (transport name = 'tcp'). Users who wish to offload iSCSI connection onto CNIC device should explicitly change transport binding of the iSCSI node. This can be done using **iscsiadm** cli utility as follows,

```
iscsiadm --mode node --targetname iqn.2004-06.com.broadcom:tg1 \
  --portal 192.168.1.100 --op=update \
  --name=node.transport_name --value=${XPORT_NAME}
```

where XPORT_NAME=bcm570x-xyyyzz

```
xx - pci bus number of the NX2 device
yy - pci device number of the NX2 device
zz - pci function number of the NX2 device
```

Network interface to iscsi transport name binding can be obtained by executing the following.

```
dmesg | grep "bnx2i: netif"
```

Sample output in a system with two NetXtreme II devices:

```
bnx2i: netif=eth1, iscsi=bcm570x-050000
bnx2i: netif=eth0, iscsi=bcm570x-030000
```

If you wish to switch back to use the software initiator, use the following:

```
iscsiadm --mode node --targetname iqn.2004-06.com.broadcom:tg1 \
  --portal 192.168.1.100 --op=update \
  --name=node.transport_name --value=tcp
```

MAKING CONNECTIONS TO ISCSI TARGETS

Refer to open-iscsi documentation for a comprehensive list of **iscsiadm** commands. This is a sample list of commands to discovery targets and to create iscsi connections to a target.

Add static entry

```
iscsiadm -m node -p <ipaddr[:port]> -T iqn.2007-05.com.broadcom:target1 -o new
```

iSCSI target discovery using 'SendTargets'

```
iscsiadm -m discovery --type sendtargets -p <ipaddr[:port]>
```

Login to target using 'iscsiadm' command

```
iscsiadm --mode node --targetname <iqn.targetname> --portal <ipaddr[:port]> --login
```

List all drives active in the system

```
'fdisk -l'
```



MAXIMIZE OFFLOAD ISCSI CONNECTIONS

With default driver parameters set, which includes 128 outstanding commands, bnx2i can offload a maximum of 28 iSCSI connections. This is not a hard limit, but just a simple on-chip resource allocation math. bnx2i will be able to offload > 28 connections by reducing the shared queue size which in turn limits the maximum outstanding tasks on a connection. See [Setting Values for Optional Properties](#) for information on `sq_size` and `rq_size`. Driver logs the following message to syslog when the maximum allowed connection offload limit is reached - "bnx2i: unable to allocate iSCSI context resources"

LINUX ISCSI OFFLOAD FAQ

- Not all Broadcom NetXtreme II adapters support iSCSI offload.
- The iSCSI session will not recover after a hot remove and hot plug.
- For MPIO to work properly, iSCSI nopout should be enabled on each iSCSI session. Refer to open-iscsi documentation for procedures on setting up `noop_out_interval` and `noop_out_timeout` values.
- In the scenario where multiple CNIC devices are in the system and the system is booted via Broadcom's iSCSI boot solution, ensure that the iscsi node under `/etc/iscsi/nodes` for the boot target is bound to the NIC that is used for booting.

NetWare Driver Software: Broadcom NetXtreme II™ Network Adapter User Guide

- [Overview](#)
- [Preinstallation Requirements](#)
- [Installing NetWare Server 6.5 Driver Software](#)
- [Viewing or Modifying Adapter Properties](#)
- [Removing Drivers from Autoexec.ncf](#)

OVERVIEW

This section discusses three NetWare drivers:

- BX2.LAN: Broadcom NetXtreme II Gigabit Ethernet 32-bit driver
- BX2.COM: Broadcom NetXtreme II Gigabit Ethernet 16-bit driver
- BNX2EV.COM: Broadcom NetXtreme II 10 Gigabit Ethernet 16-bit driver

PREINSTALLATION REQUIREMENTS

A network device driver must be installed before the Broadcom NetXtreme II Gigabit Ethernet adapter can be used with your Novell NetWare system.

Before you can successfully install the adapter driver for Novell NetWare, the adapter card must be physically installed in the system and, typically, NetWare operating system software must already be running on the system. Make sure that your system meets the hardware and operating system software requirements described in [Installing the Hardware](#).

To enable the Broadcom NetXtreme II adapter to function correctly, you must install the latest Novell NetWare support pack files. The NetWare support pack or patch file(s) needed for the operating system running on your server are indicated below:

Table 1: NetWare Support Files

NetWare Version	Support Pack	Files to be Installed
NetWare 6.5	Latest support pack	The latest support pack can be found at: http://support.novell.com/tools/csp/index.html

INSTALLING NETWARE SERVER 6.5 DRIVER SOFTWARE

You may want to create an archive disk by copying all the files from the CDROM\NetWare\Driver directory onto a floppy disk. If you elect to use the CD directly, ensure that the CDROM.NLM file is loaded and that you know the NetWare volume name for the CD.

1. From the NetWare Server console, type `HDETECT`, and then press `ENTER`.
2. From Options, press `ENTER` to accept Platform Support Module and Hot-Plug Support Module and continue detecting devices.
3. From Options, scroll down to select `Modify`, and then press `ENTER` to change focus to detected devices.
4. Scroll down to Network boards and press `ENTER`.
5. From Additional Driver Options, scroll down to `Modify`, and then press `ENTER`.
6. Insert the CD or archive floppy disk you created, and then select the **Add** option by pressing `INSERT`.
7. Select the **Install an unlisted driver** option, by pressing `INSERT`.
8. If you inserted the archive floppy disk, press `ENTER`. If you have the CD mounted as a NetWare volume, press `F3` and type *Volume Name: NetWare\Driver* as the source path, where *Volume Name* is the name of the NetWare volume for the CD, and *NetWare\Driver* is the directory of the specific files on the CD needed for this installation.
9. Press `DOWN ARROW` or `TAB` to switch back to **Additional Driver Options**.
10. Choose `Return to driver summary`, and then press `ENTER`.
11. Choose `Continue in Options`.
12. Choose `Continue from Protocol Options` unless you want to configure and bind protocols before continuing. Installation is now complete.

VIEWING OR MODIFYING ADAPTER PROPERTIES

When an adapter configuration is saved, the NetWare installation program adds load and bind statements to the `Autoexec.ncf` file. By accessing this file, you can view and modify properties or add properties for each network adapter.



NOTE: The Novell monitor program and the `config` command are also useful for viewing driver configuration properties. For information on how to use these programs, see the Utilities Reference in your Novell NetWare online documentation.

The properties that can be defined in the load statements are described below:

CONFIGURATION PROPERTIES FOR BX2.LAN DRIVER

TxDescriptors=

Initializes Descriptor resources on the adapter for transmits.

Min = 100

Max = 512

Default = 120



RxBuffers=

This is to pre-allocate receive ECBs & Receive adapter resources. This setting may be affected by the NetWare Server maximum/minimum packet receive buffer settings.

Min = 32

Max = 1000

Default = 200

Speed=

This keyword forces the line speed of the adapter. When this keyword is used, the switch to which the adapter cable is connected must be forced. Otherwise, the adapter and switch may run at different duplex modes. It is important to have the switch and adapter set properly. The best way to ensure good connectivity is to not set the speed keyword and allow for automatic link setup (auto-negotiation).

AUTO

10FD

10HD

100FD

100HD

The default setting is AUTO. Note that 1000 Mbit/s speed is auto-detected and cannot be forced.

Link=

This keyword is set to AUTO to allow the adapter to auto-negotiate with a switch using the Speed keyword as an advertised speed upper limit. It is best to allow for auto-negotiation of the adapter and switch by not setting this keyword or the Speed keyword.

AUTO

FORCE

The default is Auto. When set to FORCE, auto-negotiation is disabled and the link speed is forced to a specific value.

RxTicks=

Enables the use of batching receives within a specific time period.

Min = 0 (Disabled)

Max = 5000000 (5 seconds)

Units are in microseconds, The default value is 300.

TxTicks=

Enables the use of a transmit tick threshold interrupt within a specific time period.

Min = 0 (Disabled)

Max = 5000000 (5 seconds)

Units are in microseconds. The default value is 200.

TxPacketsPer=

Enables allowing an interrupt to occur after a specific amount of packets is transmitted.

Min = 0 (Disabled)

Max = 100

The default value is 20.

RxPacketsPer=

Enables allowing an interrupt to occur after a specific amount of packets is received.

Min = 0 (Disabled)

Max = 100

The default value is 75.

Checksum=

Enables or disables the transmit and receive checksum offload feature. Checksum offload support is only for TCP/IP packets; therefore, the default setting for checksum offload is OFF.

OFF

ON

TX

RX

FlowCtrl=

This keyword allows enabling/disabling of transmit and receive flow control.

BOTH

RX

TX

OFF

The default setting is OFF.

PDriver=

Enables the driver to operate in persistent driver mode. Use only if the adapter is placed in a hot-plug slot and only if you must swap with an adapter that is exactly the same.

OFF

ON

The default setting is OFF.

NODE=

A Novell NetWare keyword that allows an input Ethernet node address to replace the adapter factory-programmed Ethernet node address until a subsequent reboot.

NODE=*nnnnnnnnnnnn*

FRAME=

A Novell NetWare keyword for specifying the frame type.

ETHERNET_II
ETHERNET_802.3
ETHERNET_802.2
ETHERNET_SNAP

The default value is ETHERNET_802.2.

SLOT=

A Novell NetWare keyword used to set a system-wide unique hardware instance number (HIN) that may be the physical slot number on a slot-based bus such as PCI.

SLOT=*n*

Jumbo=

Enables jumbo frame support. When enabled, jumbo packets of up to 9000 bytes are supported.

Jumbo=*1536–9000*

The default setting is no jumbo packets.

**NOTES:**

- Jumbo frames are supported only on NetWare 6.0 and later. Also, the first frame loaded must be ETHERNET_II.
- Jumbo frame support must have the following text in the Startup.ncf file: "Set maximum physical receive packet size = 18000."

P3=

This keyword is used when running Perform3.exe tests and may be used to increase performance on networks running with many IPX clients.

Set P3=1 when running the Perform3 test back to back to a single client.

Set P3=2 when running the Perform3 test with many clients.

The default setting is P3=0.

Spuriousfix=

When this keyword is set to 1 (Spurious Fix is ON), the spurious interrupts count that is sometimes displayed on the NetWare console may be reduced. Setting this keyword to 0 (Spurious Fix is OFF) may enhance performance.

Spuriousfix = 0 (Spurious Fix is off)

Spuriousfix = 1 (Spurious Fix is on)

The default value is 1 (Spurious Fix is on).

Poll=

Disables the interrupt driven mode. When set to Poll=1, the driver does not use interrupts, but is polled by the NetWare operating system. This is a common feature supported in NetWare. The poll mode may increase driver performance in some environments.

Poll = 1 (on)

Poll = 0 (off)

The default value is 0 (Poll mode is off).

Wirespeed=

This feature provides adapter link and data integrity even if the adapter is attached to a questionable cable and/or switch. For example, an adapter that is trying to connect at 1000 Mbit/s on a Category 3 cable ordinarily would not link. With Wirespeed set to 1, the link is made at 100 Mbit/s.

Wirespeed = 1 (on)

Wirespeed = 0 (off)

The default value is 1 (Wirespeed is on).

Model=

This keyword allows the addition of a subsystem ID of a specific NIC so that the driver loads only on the first NIC found with a matching subsystem ID. For example, MODEL= 0x14e4. The default value is 0

MagicP=

When MagicP=1, the driver enables the adapter to wake up the system when a Magic Packet™ frame is received after the system is shut down. MagicP=0 is the default setting with the adapter having no wake-up ability.

MagicP = 1 (on)

MagicP = 0 (off)

The default value is 0 (No wake-up capability).

Fiber=

The driver has support for the 1000FD fiber adapter. The fiber auto-negotiates link with a fiber switch even though it only supports 1000FD. In some cases, the user may want to force the adapter to 1000FD.

AUTO (default)
FORCE



NOTE: If you modify any of the adapter properties, you must reboot the system before the changes take effect. If you make changes and do not reboot, you may experience problems.

A valid Autoexec.ncf file is shown below. One set of load and bind commands (in **bold**) is added for each frame type the adapter is configured to support.

```

SET Daylight Savings Time Offset = 1:00:00
SET Start Of Daylight Savings Time = (MARCH SUNDAY LAST 2:00:00 AM)
SET End Of Daylight Savings Time = (OCTOBER SUNDAY LAST 2:00:00 AM)
SET Time Zone = TPST8TPDT
# Note: The Time zone information mentioned above
# should always precede the SERVER name.
# WARNING!!
FILE SERVER NAME DL760
# WARNING!!
# If you change the name of this server, you must update
# the server name in all the licenses that are assigned
# to it using iManager.
SERVERID 487B10D
LOAD ODINEB.NLM
LOAD IPXRTR
LOAD IPXRTRNM
LOAD TCPIP
LOAD BX2 SLOT=2 FRAME=Ethernet_802.2 NAME=BX2_1_E82
BIND IPX BX2_1_E82 NET=FAFD3D25
LOAD BX2 SLOT=2 FRAME=Ethernet_802.3 NAME=BX2_1_E83
BIND IPX BX2_1_E83 NET=5A2D8D6D
LOAD BX2 SLOT=2 FRAME=Ethernet_SNAP NAME=BX2_1_ESP
BIND IPX BX2_1_ESP NET=477A35BD
LOAD BX2 SLOT=2 FRAME=Ethernet_II NAME=BX2_1_EII
BIND IPX BX2_1_EII NET=C3C8F2E4
BIND IP BX2_1_EII ADDR=172.16.1.1 MASK=ff.ff.ff.0
mount all
SEARCH ADD SYS:\JAVA\BIN
SEARCH ADD SYS:\JAVA\NWGFX

```

REMOVING DRIVERS FROM AUTOEXEC.NCF

To remove the drivers from the Autoexec.ncf file, locate the load and bind command lines associated with the Broadcom driver and remark them out by inserting the # symbol at the beginning of each command line, or by deleting the statement.

Example:

```
# LOAD BX2 SLOT=2 FRAME=Ethernet_802.2 NAME=BX2_1_E82
# BIND IPX BX2_1_E82 NET=FAFD3D25
# LOAD BX2 SLOT=2 FRAME=Ethernet_802.3 NAME=BX2_1_E83
# BIND IPX BX2_1_E83 NET=5A2D8D6D
# LOAD BX2 SLOT=2 FRAME=Ethernet_SNAP NAME=BX2_1_ESP
# BIND IPX BX2_1_ESP NET=477A35BD
# LOAD BX2 SLOT=2 FRAME=Ethernet_II NAME=BX2_1_EII
# BIND IPX BX2_1_EII NET=C3C8F2E4
# BIND IP BX2_1_EII ADDR=172.16.1.1 MASK=ff.ff.ff.0
```



NOTE: If you modify the Autoexec.ncf file, you must reboot the server before the changes take effect.

Windows Driver Software: Broadcom NetXtreme II™ Network Adapter User Guide

- [Installing the Driver Software](#)
- [Removing the Device Drivers](#)
- [Using the NetXtreme II Monolithic Driver](#)
- [Inserting the NetXtreme II Monolithic Driver in a WinPE 2.0 Image](#)
- [Viewing or Changing the Properties of the Adapter](#)
- [Setting Power Management Options](#)

INSTALLING THE DRIVER SOFTWARE



NOTE: These instructions are based on the assumption that your Broadcom NetXtreme II adapter was not factory installed. If your controller was installed at the factory, the driver software has been installed for you.

When Windows first starts after a hardware device (such as a Broadcom NetXtreme II Adapter) has been installed, or after the existing device driver has been removed, the operating system automatically detects the hardware and prompts you to install the driver software for that device.

Both a graphical interactive installation mode (see [Using the Installer](#)) and a command-line silent mode for unattended installation (see [Using Silent Installation](#)) are available.



NOTES:

- Before installing the driver software, verify that the Windows operating system has been upgraded to the latest version with the latest service pack applied.
- A network device driver must be physically installed before the Broadcom NetXtreme II Controller can be used with your Windows operating system. Drivers are located on the installation CD.
- To use the TCP/IP Offload Engine (TOE), you must have Windows Server 2003 with Service Pack 2 (SP2) or later, Windows Server 2008, or Windows Server 2008 R2. You must also have a license key preprogrammed in the hardware. If supported, for iSCSI, you only need a license key.

USING THE INSTALLER

To install the Broadcom NetXtreme II drivers

1. When the **Found New Hardware Wizard** appears, click **Cancel**.
2. Insert the installation CD into the CD or DVD drive.
3. On the installation CD, open the folder for your operating system, open the DrvInst folder, and then double-click **Setup.exe** to open the InstallShield Wizard.



4. Click **Next** to continue.
5. After you review the license agreement, click **I accept the terms in the license agreement** and then click **Next** to continue.
6. Select how you want to install the NetXtreme II drivers and then click **Next**.
7. Click **Install**.
8. Click **Finish** to close the wizard.
9. The installer will determine if a system restart is necessary. Follow the on-screen instructions.

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 DrvInst folder.

To perform a silent install from within the installer source folder

Type the following:

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

or

```
msiexec /i "BDrv5706.msi" /qn
```

To perform a silent upgrade from within the installer source folder

Type the following:

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

To perform a silent uninstall from within the installer source folder

Type the following:

```
msiexec /x "BDrv5706.msi" /qn
```

To perform a silent uninstall from any folder

```
msiexec /x "{F0DA8A3F-1457-419E-96F4-235DD3EF41E1}" /qn
```



NOTE: The hexadecimal number above may differ from your current installer. Check the Key name corresponding with the Broadcom Advanced Control Suite 3 (BACS) application in HKLM\Software\Microsoft\Windows\CurrentVersion\Uninstall for the correct hexadecimal number.

To perform a silent reinstall of the same installer

Type the following:

```
setup /s /v"/qn REINSTALL=ALL"
```



NOTE: The REINSTALL switch should only be used if the same installer is already installed on the system. If upgrading an earlier version of the installer, use `setup /s /v/qn` as listed above.

To perform a GUI reinstall of the same installer

Type the following:

```
setup /V"REINSTALL=ALL"
```

If performing a silent upgrade or uninstall, ensure to do a manual reboot afterwards to avoid leaving the system in an inconsistent state.

During silent upgrade or uninstall, your system may reboot automatically. If you wish to suppress the reboot, please append `REBOOT=ReallySuppress` to the end of the corresponding upgrade or uninstall command listed above.

In some circumstances, reboot is required before uninstallation can continue. If you used `REBOOT=ReallySuppress` to suppress the reboot, the uninstallation may be suspended. You will need to reboot manually for the uninstallation to continue.

To perform a silent install and create a log file at (f:\1testlog.txt)

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

REMOVING THE DEVICE DRIVERS



NOTE: Windows Server 2003, Windows Server 2008, and Windows Server 2008 R2 provide the Device Driver Rollback feature to replace a device driver with one that was previously installed. However, the complex software architecture of the NetXtreme II device may present problems if the rollback feature is used on one of the individual components. Therefore, we recommend that changes to driver versions be made only through the use of a driver installer.

To remove the device drivers

1. In Control Panel, double-click **Add or Remove Programs**.
2. Click **Broadcom NetXtreme II GigE Driver Installer**, and then click **Remove**. Follow the onscreen prompts.
3. Click **Yes** to restart your system.
- or -
4. Click **No** to restart your system at a later time. Click **OK** to acknowledge that the installation has been suspended. The uninstallation of the driver is postponed until the next restart of your system.

USING THE NETXTREME II MONOLITHIC DRIVER

The NetXtreme II, based on its advanced functionalities, uses a software architecture that includes a Virtual Bus Device (VBD) to extend functionalities beyond basic network connectivity. Microsoft, however, does not currently support this architecture when loading an operating system through its Windows Deployment Services (WDS), which was previously

known as Remote Installation Services (RIS), or for the deployment agent used in the Automated Deployment Services (ADS). Therefore, a separate driver was created to accommodate these Microsoft deficiencies. This driver is known as the NetXtreme II monolithic driver, but it is sometimes referred to as the "RIS" driver.

The NetXtreme II monolithic driver was developed to work only for the text mode portion of a WDS legacy installation and to establish connectivity with a deployment agent for ADS. It is not intended to be used as a driver loaded in the running state of an operating system. The exception to this would be when used for the Windows Preinstallation Environment (WinPE).

For WDS, this driver is used similarly to any other network adapter driver for supporting network connectivity after the PXE boot to the WDS server. When placed in the I386 or AMD64 directory (depending on the version of the operating system being deployed), the monolithic driver is called to establish that there is driver support for the NetXtreme II adapter included in the WDS legacy image.

For ADS, the driver is placed in the PreSystem directory on the server running ADS to establish connectivity with the deployment agent on remote systems with NetXtreme II adapters when booting from PXE.

While Windows PE 2005 natively supports the VBD architecture, it was found that using the "minint" switch in the startnet.cmd file does not. The minint switch performs a limited scan of the system bus to identify network devices only and, therefore, does not support the VBD architecture. Since only network connectivity is required in Windows PE, the only supported driver is the monolithic driver for the NetXtreme II adapter in this environment as well. Place the b06nd.inf file in the INF directory within the Windows PE image, and place the appropriate driver file (b06nd51a.sys for x64-based builds or b06nd51.sys for x86-based builds) in the driver's directory. If Windows PE is deployed as a flat image from a RIS or WDS server, you must also place both the b06nd.inf and the appropriate driver file in the I386 or AMD64 directory containing the image. If the RIS or WDS server is running Windows 2000 Server and deploying an x86 WinPE image, you may need to include the Windows 2000 monolithic driver file (b06nd50x.sys) in the I386 directory. In cases where adding the Windows 2000 monolithic driver still does not work, apply the following modification to the b06nd.inf file located in the I386 directory as follows:

1. Locate [Manufacturer] header within the file.
2. Review the line below it which reads: %brcm% = broadcom, ntx86, ntamd64, ntia64 or equivalent.
3. Modify that line to read: %brcm% = broadcom.ntx86, ntamd64, ntia64. The change made replaces the comma and space after "broadcom" with a period.
4. Save the file.
5. Restart the RIS service (binlsvc) or WDS services (wdsserver).

INSERTING THE NETXTREME II MONOLITHIC DRIVER IN A WINPE 2.0 IMAGE

The Microsoft Windows Server 2008 method of inserting the NetXtreme II monolithic driver in a WinPe 2.0 image is different from the Windows Server 2008 R2 method, as discussed below.

By default, the monolithic driver is not included in the boot.wim and install.wim files that come with the Microsoft Windows Server 2008/Vista CD. Microsoft's Windows Automated Installation Kit (AIK) allows you to modify the default boot.wim and install.wim files, and create WinPE 2.0 images to include the NetXtreme II monolithic driver in the Windows Server 2008/Vista installation.

To insert the monolithic driver into a WinPE 2.0 boot image (Windows Server 2008)



To insert Broadcom's NetXtreme II monolithic driver in a WinPE 2.0 image, download AIK from <http://www.microsoft.com/downloads/en/default.aspx> and install.

After installing AIK, copy the latest monolithic driver to a directory on the local hard drive of the system you installed the AIK. Follow the procedure below to insert the monolithic driver into a WinPE 2.0 boot image.

1. From All Programs, open Windows AIK and select **Windows PE Tools Command prompt**.
2. At the command prompt, run the `copype.cmd` script. The script requires two arguments: hardware architecture and destination location.

```
copype.cmd <arch> <destination>
```

For example: `copype x86 c:\VistaPEx86`



NOTE: The directory structure `c:\VistaPEx86` is used throughout this procedure.

3. Mount the base image to a local directory so that you can add or remove packages by typing:
`imagex /mountrw c:\VistaPEx86\winpe.wim 1 c:\VistaPEx86\mount`
4. Place the monolithic driver and inf file in `c:\drivers\x32\` by typing:
`peimg /inf=c:\Drivers\x32\b06nd.inf c:\VistaPEx86\mount\windows`
AIK inserts the driver into the WinPE 2.0 image.
5. To complete the customization of the image, prepare the image for deployment, type:
`peimg /prep c:\VistaPEx86\mount\windows`
6. When asked to continue and have the program prepare the image for deployment, type:
`yes`
7. To commit the changes to the original image file (`Winpe.wim`), type:
`imagex /unmount c:\VistaPEx86\mount /commit`
8. To replace the default `Boot.wim` file in the `\ISO` directory with your new custom image, type:
`copy c:\VistaPEx86\winpe.wim c:\VistaPEx86\ISO\sources\boot.wim`

To add a device driver to an offline Windows PE image (Windows Server 2008 R2)

This procedure demonstrates how to use the Deployment Image Servicing and Management (DISM) tool to add a device driver (.inf) to an offline Windows PE image. Before running a DISM command, first mount the Windows PE image.

1. Mount the base image by using the DISM tool to a local Windows PE directory. For example:

```
Dism /Mount-WIM /WimFile:c:\winpe_x86\winpe.wim /index:1 /MountDir:c:\winpe_x86\mount
```



NOTE: The directory structure c:\winpe_x86 is used throughout this procedure.

2. Add the .inf file to the base image by using the **dism** command with the **/Add-Driver** option. For example Driver.inf is the Broadcom driver, evnd.inf is the driver for the 10 Gbps devices, and b06nd.inf is the driver for the 1 Gbps devices.

```
Dism /image:<path_to_image> /Add-Driver /Driver:c:\winpe_x86\mount\Windows\driver.inf
```

3. Repeat steps 1 and 2 for each additional device driver.

4. Unmount the image after modifying it.

```
dism /unmount-wim /Mountdir:c:\winpe_x86\mount /commit
```

5. After unmounting the image, burn it to the designated media.

To create a bootable CD-ROM

1. On your technician computer, at the command prompt, create an .iso file by typing:

```
oscdimg -n -bc:\VistaPEx86\etfsboot.com c:\VistaPEx86\ISO C:\VistaPEx86\VistaPEx86.iso
```

2. Burn the iso image to a CD.

CONFIGURING THE SPEED/DUPLEX SETTING FOR THE NETXTREME II MONOLITHIC DRIVER

Since the typical environment where the NetXtreme II monolithic driver is used does not provide the means to configure advanced network adapter properties, the driver file (b06nd.inf) was modified to include a section that allows it to be configured for a specific speed and/or duplex. This provides a more robust connection to the network as it allows the adapter to match the settings of its link partner (e.g., a switch, router, etc.).

To manually configure the speed and duplex

1. Open the b06nd.inf file with a text editor like Microsoft Notepad or WordPad.
2. Perform a search on the file for "Registry parameters" to locate the section that will allow you to configure the adapter speed/duplex.
3. Once located, notice the following information shown.

```
[params_utp]
hkr, , req_medium,          2, "0"
[params_fiber]
hkr, , req_medium,          2, "65283"
```

These make up two separate sections that can be configured: one for standard RJ-45 copper interfaces (params_utp) and one for fiber devices (params_fiber).

4. As described in the file, replace the value above in quotation marks under the correct section, depending upon the network adapter in your system. The available values are shown below.

Options for copper interfaces:

- Auto (1 Gbps is enabled when that speed is supported) = "0"
- 10 Mbps Half Duplex = "65794"
- 10 Mbps Full Duplex = "258"
- 100 Mbps Half Duplex = "66050"
- 100 Mbps Full Duplex = "514"

Options for fiber interfaces:

- Auto (1 Gbps is enabled when that speed is supported) = "0"
- 1 Gbps Full Duplex = "771"
- Auto with 1 Gbps Fallback = "33539"
- Hardware default = "65283"

An example is provided in the file showing how to configure a copper interface for a 10 Mbps Full Duplex connection. The example is shown below.

```
hkr, , req_medium,          2, "258"
```

VIEWING OR CHANGING THE PROPERTIES OF THE ADAPTER

To view or change the properties of the Broadcom network adapter


1. In Control Panel, click **Broadcom Control Suite 3**.
2. Click the Advanced section of the **Configurations** tab.
3. See [Setting Adapter Properties](#) for a detailed description of the available properties as well as for instructions for viewing and changing the value of a particular property.

SETTING POWER MANAGEMENT OPTIONS

You can set power management options to allow the operating system to turn off the controller to save power or to allow the controller to wake up the computer. If the device is busy doing something (servicing a call, for example) however, the operating system will not shut down the device. The operating system attempts to shut down every possible device only when the computer attempts to go into hibernation. To have the controller stay on at all times, do not click the **Allow the computer to turn off the device to save power** check box.



NOTE: Power management options are not available on blade servers.

**NOTES:**

- The Power Management tab is available only for servers that support power management.
- To enable Wake on LAN (WOL) when the computer is on standby, click **Allow the device to bring the computer out of standby** box.
- If you select **Only allow management stations to bring the computer out of standby**, the computer can be brought out of standby *only by Magic Packet*.



CAUTION! Do not select **Allow the computer to turn off the device to save power** for any adapter that is a member of a team.

Using Broadcom Advanced Control Suite 3: Broadcom NetXtreme II™ Network Adapter User Guide

- [Broadcom Advanced Control Suite 3 Overview](#)
- [Starting Broadcom Advanced Control Suite 3](#)
- [Using Broadcom Advanced Control Suite 3](#)
- [Configuring Teaming](#)
- [Configuring With the Command Line Interface Utility](#)

BROADCOM ADVANCED CONTROL SUITE 3 OVERVIEW

Broadcom Advanced Control Suite 3 (BACS 3) is an integrated utility that provides useful information about each network adapter that is installed in your system. BACS 3 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 each adapter.

Broadcom Advanced Server Program (BASP), which runs within Broadcom Advanced Control Suite 3, 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.

STARTING BROADCOM ADVANCED CONTROL SUITE 3

In Control Panel, click **Broadcom Control Suite 3**.

Microsoft .NET Framework 2.0 includes the runtime and associated files needed to run BACS 3, and must be installed on your system in order for BACS 3 to operate. For information on the minimum and recommended .NET Framework versions for your operating system, see [Table 1](#).



NOTES:

- Starting BACS 3 without .NET Framework (version 2.0 or above) installed on your system will result in an error.
- Long startup times for BACS 3 may be caused by a limitation of .NET Framework 2.0, as described in Microsoft knowledge base article KB936707. To improve the startup time of BACS 3, apply .NET Framework 2.0 SP1 (or appropriate version for your OS as shown in the Recommended column of [Table 1](#)), and reboot your system.
- For optimal performance of BACS 3, Broadcom recommends .NET Framework 2.0 SP1, .NET Framework 3.0 SP1, or .NET Framework 3.5, depending on your operating system.

Table 1: Minimum and Recommended .NET Framework Versions for BACS 3

Operating System	Minimum Required		Recommended	
	Version	How to Obtain	Version	How to Obtain
Windows 2000 Server	.NET 2.0	Download from Microsoft	.NET 2.0 SP1	Download from Microsoft
Windows Server 2003	.NET 2.0	Download from Microsoft	.NET 2.0 SP1	Download from Microsoft
Windows Server 2003 R2	.NET 2.0	Included in OS package, but by default, not included during installation	.NET 2.0 SP1	Download from Microsoft
Windows Vista	.NET 2.0	Included in .NET 3.0, which is included during OS installation	.NET 2.0 SP1	Download from Microsoft
Windows Server 2008	.NET 2.0	Included in .NET 2.0 SP1, which is included during OS installation	.NET 2.0 SP1	Included in OS package
Windows Server 2008 R2	.NET 2.0	Included in .NET 2.0 SP1, which is included during OS installation	.NET 2.0 SP1	Included in OS package

USING BROADCOM ADVANCED CONTROL SUITE 3

Start BACS 3. Click the tab that provides the information of interest or from which to perform a desired test, diagnostic, analysis, or set adapter properties. To create a team, see [Configuring Teaming](#).

USING THE INTERFACE

BACS 3 features a context-sensitive interface. When an item is selected in the Explorer View, such as an adapter, driver, team, or VLAN, the interface changes to display the information and options available for the selected item. Certain tabs,

menu options, and device information are only available based on the types of devices installed or the item selected in the interface.

Two buttons at the bottom of the interface control what you want to manage: click the [Device Management](#) button to manage the connected devices, or click the [Team Management](#) button to create and manage teams.

Device Management

The Device Management pane lists the devices that can be viewed, analyzed, tested, or adjusted by BACS 3. When an item is selected in the Device Management pane, the tabs showing the information and options that are available for the item appear. Because the BACS 3 interface is context-sensitive, only the information that applies to the selected item can be viewed or adjusted.


Items in Device Management can be listed by type or by connection.

To list items in Device Management by type

- At the top of the Device Management pane, click **By Type**.

A single hardware device may be using several different drivers, so Device Management breaks each item out into separate categories.

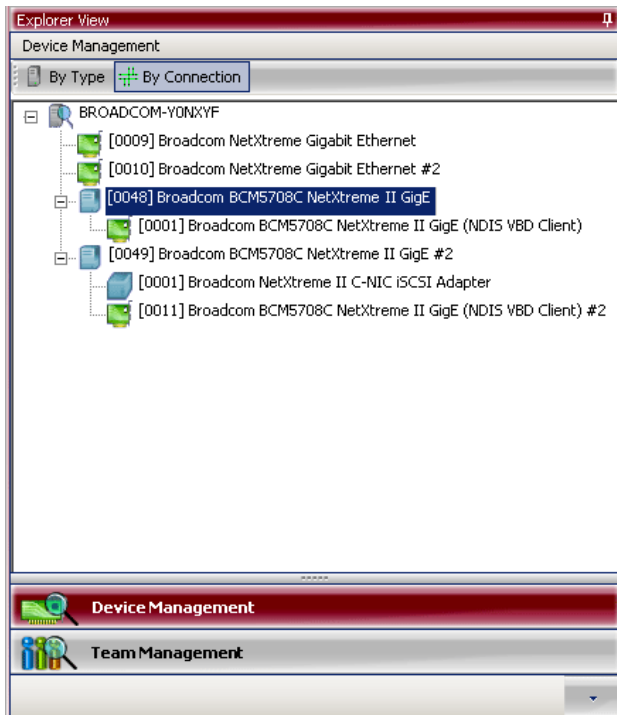
- **Network Adapters.** Lists the available NDIS client adapters.
- **iSCSI Controllers.** Lists client items using iSCSI network protocols.
- **System Devices.** Lists VBD server drivers.



To list items in Device Management by connection

- At the top of the Device Management pane, click **By Connection**.

The Connection View lists the server and client devices in a parent/child hierarchy, which makes it easier to see the relationships between the adapters. Just as in the Device Management pane, selecting items in the Connection View shows the options available for that item.



Device Icons

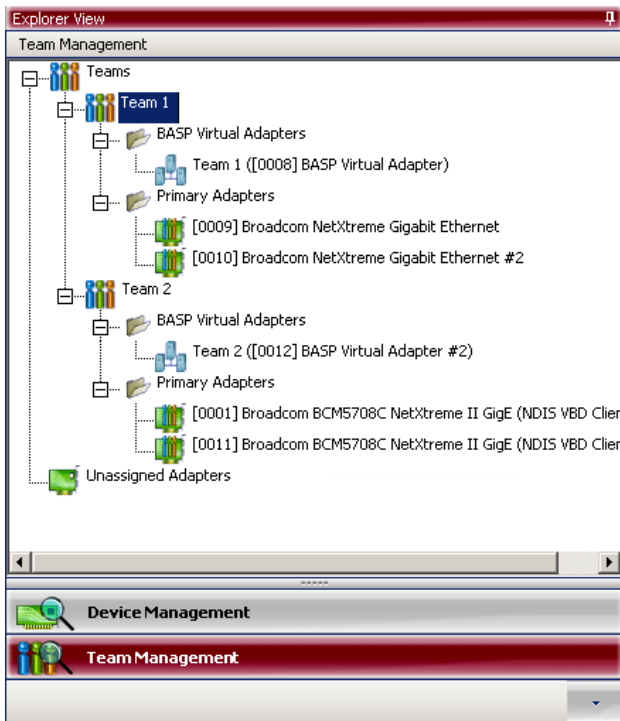
The icon next to each device in the Device Management pane shows its status. An icon next to a device name that appears normal means the device is connected and working.

- **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.

Team Management

The Team Management pane displays how network adapters have been grouped together to function as a team, or virtual LAN. The pane separates the adapters that have been grouped into a team from those that remain unassigned.

Team Management only becomes available when more than one Broadcom network adapters are installed and available to be grouped into a team.



To create or edit teams

- See [Configuring Teaming](#).

CONFIGURING BACS USER INTERFACE OPTIONS

Enabling/Disabling the BACS Tray Icon

BACS 3 places an icon in the Windows taskbar when the program is installed. Use the Options window to turn this icon on or off.

To enable or disable the BACS tray icon

1. From the **Tools** menu, select **Options**.
2. In the Options window, select **General**.
3. Select or clear **Enable Tray Icon** (the option is enabled by default).
4. Click **OK**.

CONFIGURING BACS USER INTERFACE APPEARANCE

Changing the Look and Feel

The appearance of the BACS 3 interface can be customized.

To change the interface appearance

1. From the **Tools** menu, select **Options**.



2. In the Options window, select **Appearance**.
3. Choose the appearance settings. Click **Apply** to set changes. Click **OK** to set changes and close the Options window.

To reset the interface to the default appearance

1. From the **Tools** menu, select **Options**.
2. In the Options window, select **General**.
3. Click **Restore Default**.
4. Close and then restart BACS 3.

VIEWING VITAL SIGNS

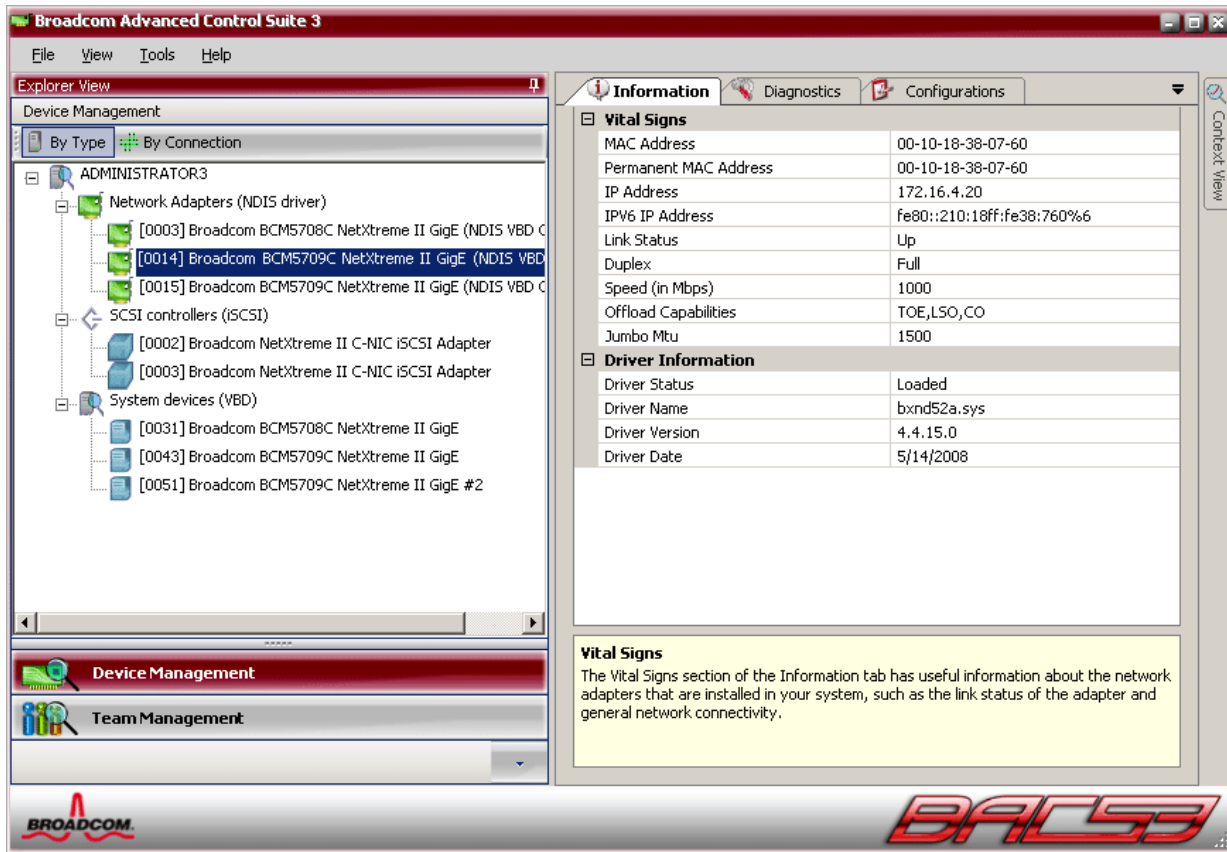
The **Vital Signs** section of the **Information** tab has useful information about the network adapters that are installed in your system, such as the link status of the adapter and general network connectivity.

To view Vital Signs information for any installed network adapter, click the name of the adapter listed in the Device Management pane, then click the **Information** tab.

- If the **Information** tab is not visible, then from the **View** menu, select **Navigate**, then **Device Management**, and then **Information**.
- If the Vital Signs section within the **Information** tab is not visible, then from the **Context View** tab on the right side of the window, select **Information** and then select **Vital Signs**.

**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 IP Address. The IPv6 network address associated with the adapter.

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 Broadcom NetXtreme II adapters.

- **TOE.** TCP Offload Engine (TOE) allows simultaneous operation of up to 1024 fully offloaded TCP connections for 1-Gbps network adapters and 1880 fully offloaded TCP connections for 10-Gbps network adapters to the hardware.
- **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

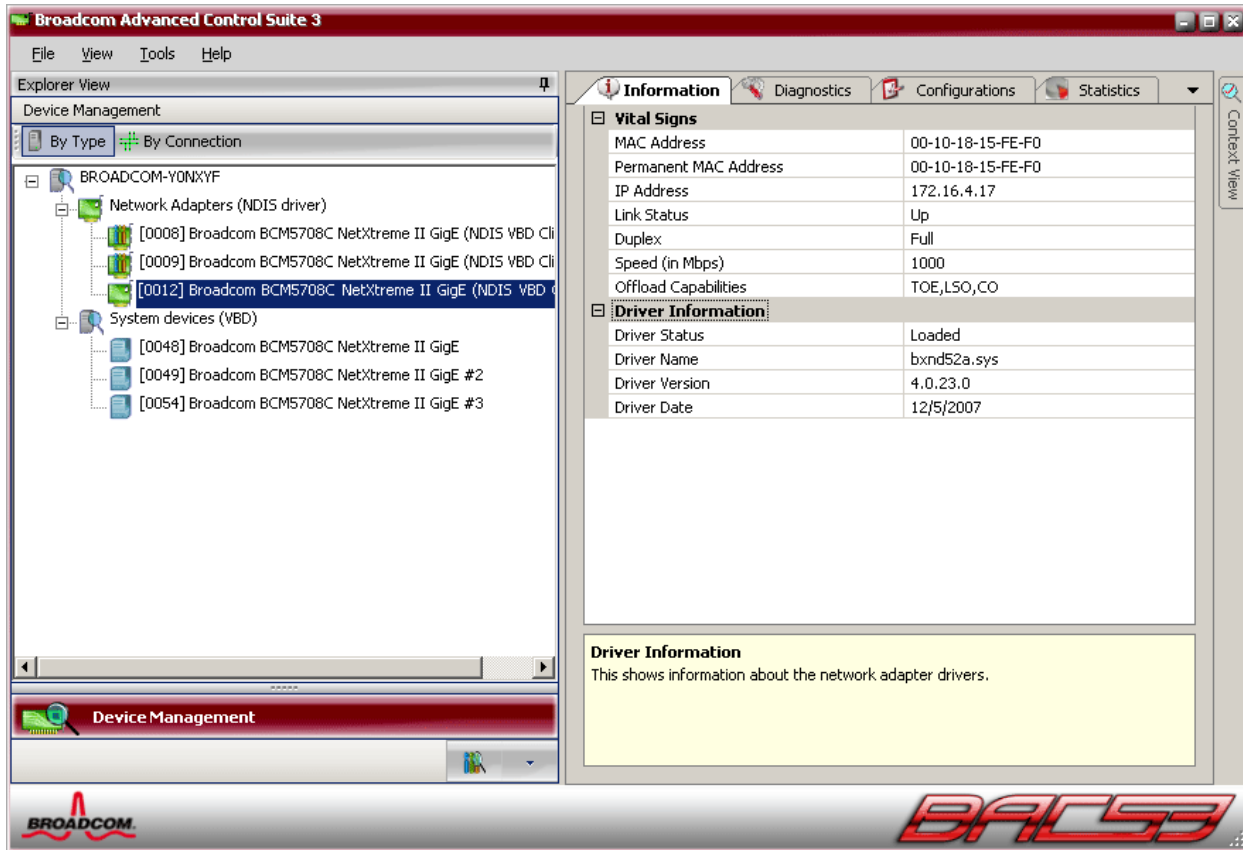
BASP State. Information about the status of the BASP application. This information is displayed only when there is a team (see [Configuring Teaming](#)).

VIEWING DRIVER INFORMATION

The **Driver Information** section of the **Information** tab displays data about the driver for the selected network adapter.

To view Driver Information for any installed network adapter, click the name of the adapter listed in the Device Management pane, then click the **Information** tab.

- If the Information tab is not visible, then from the **View** menu, select **Navigate**, then **Device Management**, and then **Information**.
- If the Driver Information section within the Information tab is not visible, then from the **Context View** tab on the right side of the window, select **Information** and then select **Driver**.



Driver Status. The status of the adapter driver.

- **Loaded.** Normal operating mode. The adapter driver has been loaded by Windows and is functioning.
- **Not Loaded.** The driver associated with the adapter has not been loaded by Windows.
- **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.

VIEWING RESOURCE INFORMATION

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

To view Resources for any installed network adapter, click the name of the adapter listed in the Device Management pane, then click the **Information** tab.


- If the Information tab is not visible, then from the **View** menu, select **Navigate**, then **Device Management**, and then **Information**.
- If the Resources section within the Information tab is not visible, then from the **Context View** tab on the right side of the



window, select **Information** and then select **Resources**.



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



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 Broadcom NetXtreme II adapters.

Bridge Lanes. The number of PCI-E lanes connected to the bridge. This information is only available for Broadcom NetXtreme II adapters.

Bridge Speed. The clock speed on PCI-E bus. This information is only available for Broadcom NetXtreme II adapters.

Slot No. The slot number on the system board occupied by the adapter. This item is not available for PCI Express type adapters.

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

Bus Width (bit). 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 No. Indicates the number of the bus where the adapter is installed.

Device No. The number assigned to the adapter by the operating system.

Function No. 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.

VIEWING HARDWARE INFORMATION

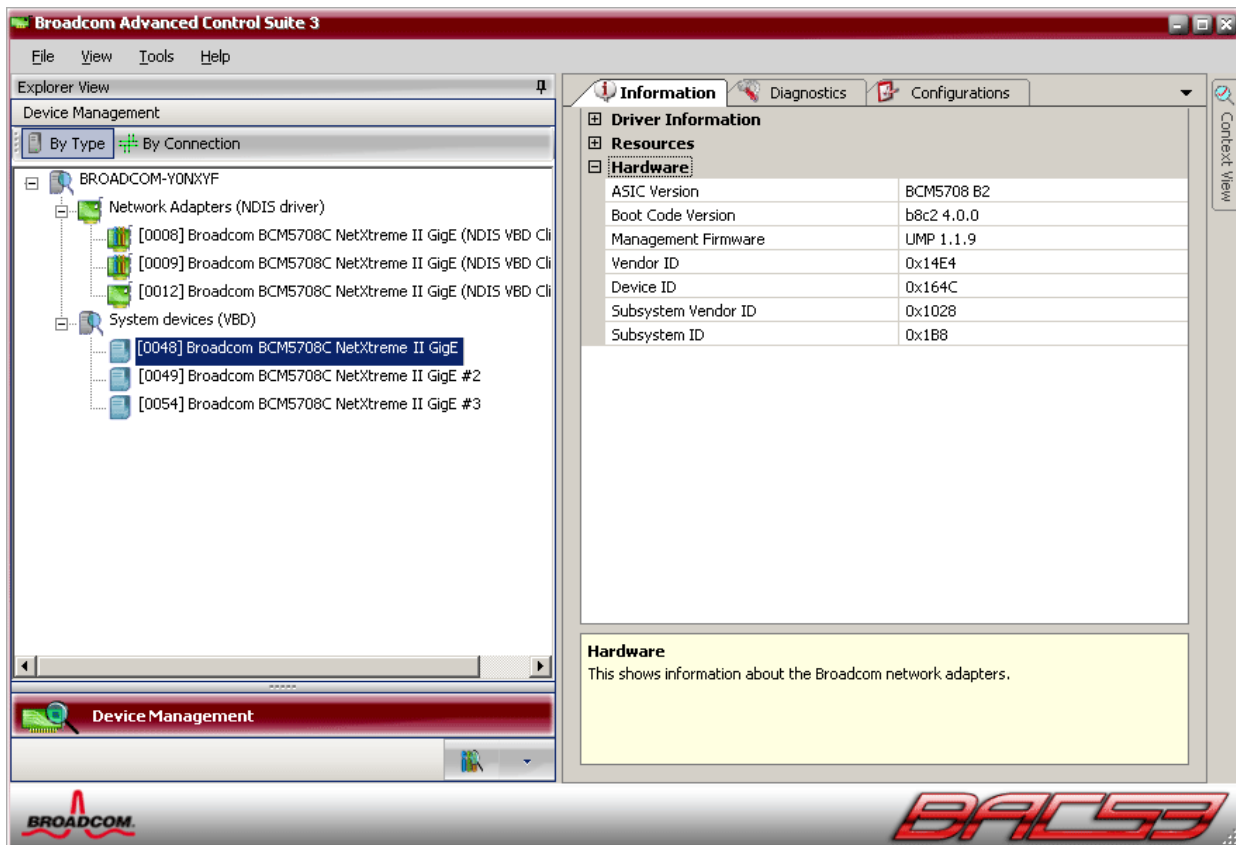
The Hardware section of the **Information tab** displays information about the hardware settings for the selected network adapter.

To view Hardware for any installed network adapter, click the name of the adapter listed in the Device Management pane, then click the Information tab.

- If the Information tab is not visible, then from the **View** menu, select **Navigate**, then **Device Management**, and then **Information**.
- If the Hardware section within the Information tab is not visible, then from the **Context View** tab on the right side of the window, select **Information** and then select **Hardware**.



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



ASIC Version. The chip version of the Broadcom adapter (this information is not available for adapters made by others).

Boot Code Version. The version of the boot code. This information is only available for Broadcom NetXtreme II adapters.

Firmware Family 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.

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.

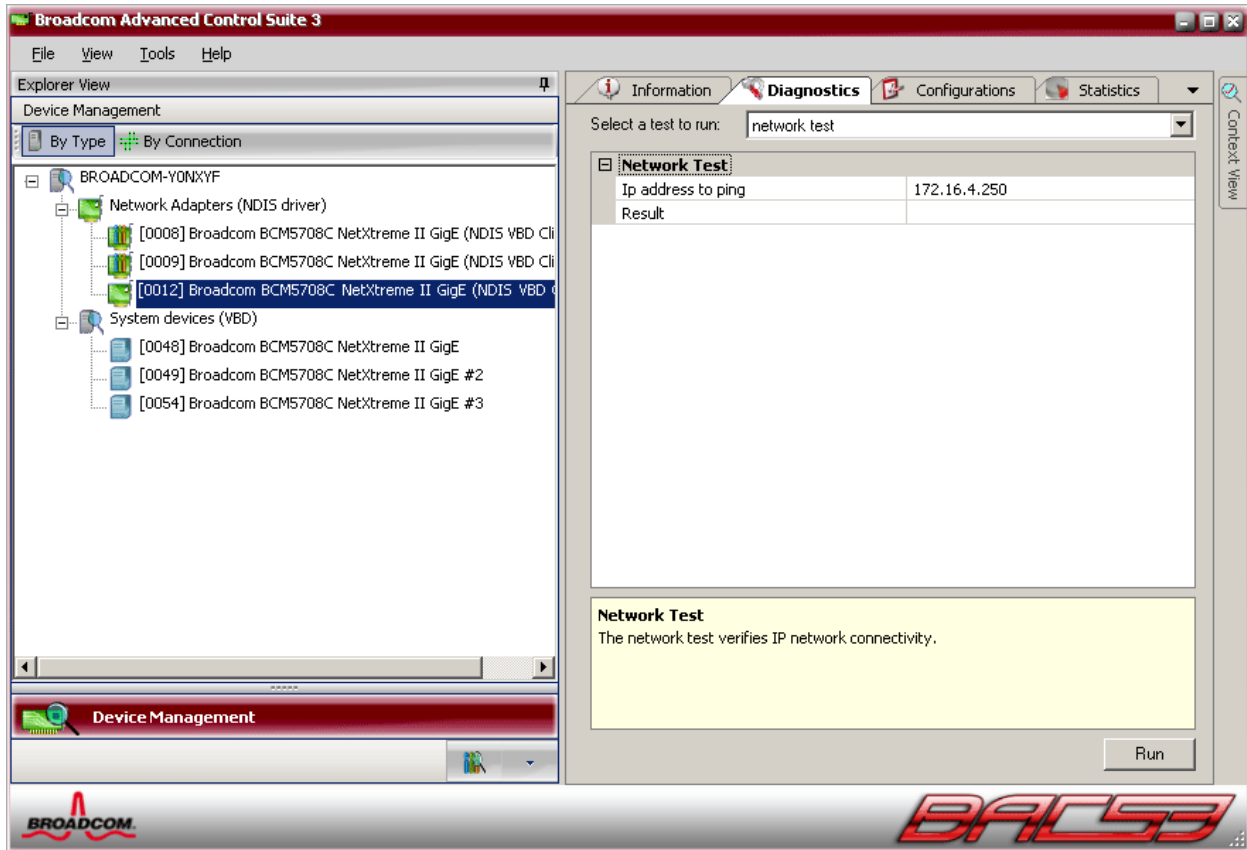


NOTE: The network test option is not available on adapters that are grouped into a team (see [Configuring Teaming](#)).

To run the network test

1. Click the name of the adapter to test in the Device Management pane.
2. Click the **Diagnostics** tab. If the Diagnostics tab is not visible, then from the **View** menu, select **Navigate**, then **Device Management**, and then **Diagnostics**.
3. From the **Select a test to run** list, select **Network Test**. If the **Network Test** option is not available, then from the **Context View** tab on the right side of the window, select **Diagnostics** and then select **Network Test**.
4. 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**.
5. Click **Run**.

The results of the network test are displayed in the **Result** field.



RUNNING DIAGNOSTIC TESTS

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 3 continuously perform them. If the test are performed continuously, then the number of passes and fails in the **Status** 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 **Status** field for that test is 4/0. However, if there were 3 passes and 1 fail, the value in the **Status** field is 3/1.



NOTES:

- You must have administrator privileges to run diagnostic tests.
- The network connection is temporarily lost while these tests are running.
- This option is available for Broadcom NetXtreme II VBD drivers.
- Some tests are not supported on all Broadcom adapters.

To run the diagnostic tests once

1. Click the name of the adapter to test in the Device Management pane.



NOTE: For Broadcom NetXtreme II adapters, select a VBD driver; for other adapters, select an NDIS driver.

2. Click the **Diagnostics** tab. If the **Diagnostics** tab is not visible, then from the **View** menu, select **Navigate**, then **Device Management**, and then **Diagnostics**.
3. From the **Select a test to run** list, select **Diagnostic Tests**. If the **Diagnostic Tests** option is not available, then from the **Context View** tab on the right side of the window, select **Diagnostics** and then select **Diagnostic Tests**.
4. Select the diagnostic tests you want to run. Click **Select All** to select all tests or **Clear All** to clear all test selections.
5. Click **Run**.
6. In the error message window that warns of the network connection being temporarily interrupted, click **Yes**. The results are displayed in the **Status** field for each test.


To run the diagnostic tests continuously

1. Click the name of the adapter to test in the Device Management pane.



NOTE: For Broadcom NetXtreme II adapters, select a VBD driver; for other adapters, select an NDIS driver.

2. Click the **Diagnostics** tab. If the **Diagnostics** tab is not visible, then from the **View** menu, select **Navigate**, then **Device Management**, and then **Diagnostics**.
3. From the **Select a test to run** list, select **Diagnostic Tests**. If the **Diagnostic Tests** option is not available, then from the **Context View** tab on the right side of the window, select **Diagnostics** and then select **Diagnostic Tests**.
4. Select the diagnostic tests you want to run. Click **Select All** to select all tests or **Clear All** to clear all test selections.
5. Select **Run Continuously**.
6. Click **Run**.
7. In the error message window that warns of the network connection being temporarily interrupted, click **Yes**. The diagnostic tests will run continuously, and the **Status** field for each test displays the accumulated total of passes and failures.
8. To stop the continuous testing, click **Abort**.



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.

ANALYZING CABLES

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:

- 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 NetXtreme II 10 GbE network adapters.
- This option is not available for all Broadcom network adapters.
- This option is available for Broadcom NetXtreme II VBD drivers.

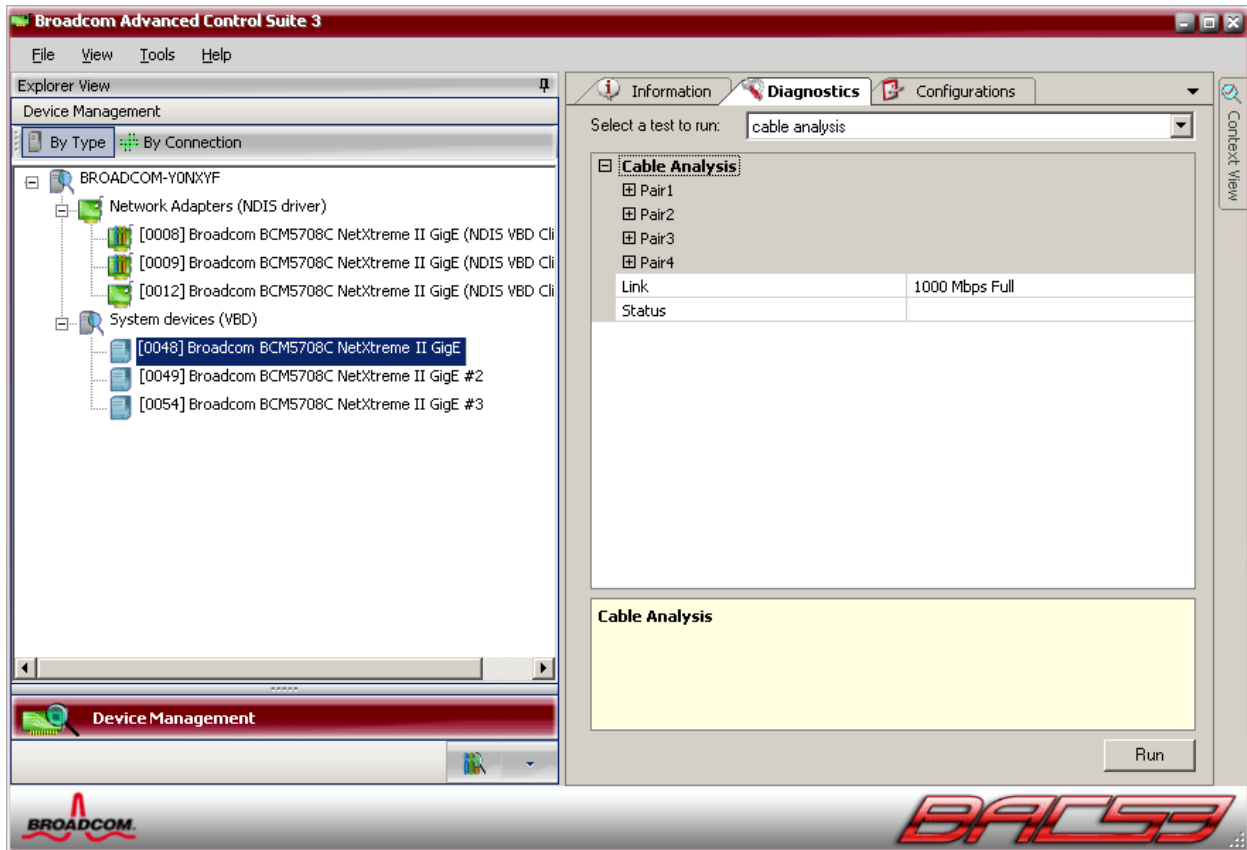
To run a cable analysis

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 Device Management pane.



NOTE: For Broadcom NetXtreme II adapters, select a VBD driver; for other adapters, select an NDIS driver.

3. Click the **Diagnostics** tab. If the Diagnostics tab is not visible, then from the **View** menu, select **Navigate**, then **Device Management**, and then **Diagnostics**.
4. From the **Select a test to run** list, select **Cable Analysis**. If the **Cable Analysis** option is not available, then from the **Context View** tab on the right side of the window, select **Diagnostics** and then select **Cable Analysis**.
5. Click **Run**.
6. 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.

SETTING ADAPTER PROPERTIES

Advanced and **iSCSI Boot Configuration** on the **Configurations** tab allow you to view and change the values of the available properties of the selected adapter. The potentially available properties and their respective settings are described below.



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.


To set adapter properties

1. Click the name of the adapter in the Device Management pane.
2. Click the **Configurations** tab. If the Configurations tab is not visible, then from the **View** menu, select **Navigate**, then **Device Management**, and then **Configurations**.
3. From the **Advanced** or **iSCSI Boot Configuration** section, select the property you want to set. If the **Advanced** or **iSCSI Boot Configuration** section is not available, then from the **Context View** tab on the right side of the window, select **Configurations** and then select **Advanced** or **iSCSI Boot Configuration**.
4. To change the value of a property, select an item from the property's list or type a new value, as appropriate (selection options are different for different properties).
5. 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**.



802.1p QOS. Enables *quality of service*, which is an Institute of Electrical and Electronics Engineering (IEEE) specification that treats different types of network traffic diversely to ensure required levels or reliability and latency according to the type of traffic. This property is disabled by default. Unless the network infrastructure supports QoS, do not enable this property. Otherwise, problems may occur.

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. By enabling TOE, network performance improves, but with the increased performance, TOE performance is more susceptible to packet loss when flow control is disabled. Enable flow control to reduce the number of packets lost.



NOTE: If **Jumbo MTU** 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 PAUSE.** Receive PAUSE frame is enabled.
- **Rx/Tx PAUSE.** Receive and transmit PAUSE frame are enabled.
- **Tx PAUSE.** 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 Broadcom NetXtreme II adapters.

Jumbo MTU. 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 Broadcom NetXtreme II 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 MTU** 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.

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.



NOTE: For all network adapters, IPv6 Receive Side Scaling is not supported on Windows Server 2003 due to a limitation in the operating system. However, BCM5709, BCM57710, and BCM57711 network adapters support IPv6 Receive Side Scaling on Windows Server 2008 and later.

Switch Configuration. Allows configuring of the connected switch for BCM57710 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 Auto.** Sets the speed to Auto-Negotiate at 1000 Mb Full-Duplex mode only.
- **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.
- **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, the Wake on LAN feature may not work. Additionally, 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).

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**.

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 1, 2, 4, and 8, with 8 RSS queues as the default option. For 10 Gbps network adapters, the RSS queue options are Auto, 1, 2, 4, 8, and 16, with Auto being the default option.

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 2000 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 1 with 0 receive buffers as the default value.

Transmit Buffers. 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 2500 in increments of 1 with 1500 transmit buffers as the default value.

TCP Connection Offload (IPv4). Enables and disables TOE offload when using the IPv4 protocol. The default is Enabled.

TCP Connection Offload (IPv6). Enables and disables TOE offload when using the IPv6 protocol. The default is Enabled.

Pause on Exhausted Host Ring. For BCM57711 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.

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.

Chimney Offload State. Enables TCP Offload Engine (TOE) for the entire system. On Windows Server 2003 and Windows Server 2008 operating systems, the options are Enable (default) and Disable. For Windows Server 2008 R2, the options are Enable, Disable, and Auto (default). If Chimney Offload State is configured for Auto, then a 10 Gbps network adapter will have TOE enabled, but not for a 1 Gbps network adapter.

To enable TOE for individual network adapters, configure Chimney Offload State to Enable and also enable TCP Connection Offload (IPv4) or TCP Connection Offload (IPv6) from the Advanced area of the Configuration tab.



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 Broadcom adapters.

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

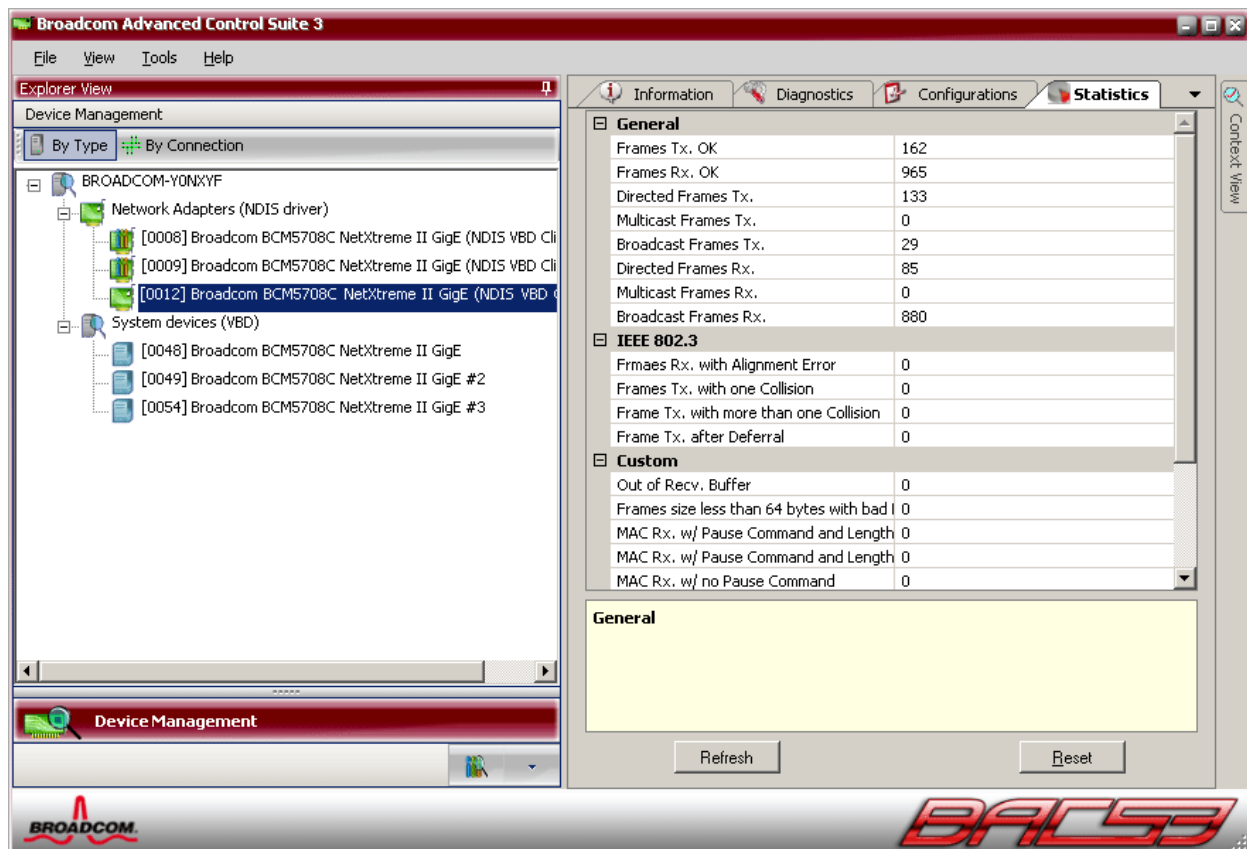
- If the Statistics tab is not visible, then from the **View** menu, select **Navigate**, then **Device Management**, and then **Statistics**.
- If any of the sections described below is not visible, then from the **Context View** tab on the right side of the window, select **Statistics** and then select 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 BACS3 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 Broadcom NetXtreme II 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.

VIEWING RESOURCE RESERVATIONS



NOTES:

- Resource Reservation information is only available for Broadcom NetXtreme II adapters and VBD drivers.
- Not all offload technologies are available with all adapters.

The Resource Reservations section shows the number of connections allocated to an offload technology: TOE and iSCSI.

- TCP Offload Engine (TOE) for accelerating TCP over 1 GbE, 2.5 GbE, and 10 GbE.
- Internet Small Computer Systems Interface (iSCSI) offload for accelerating network storage access featuring centralized boot functionality (iSCSI boot).


You can also view the number of unlicensed resources and unallocated resources.

TOE and iSCSI can only be configured on certain adapters and require a license key. License keys are preprogrammed in the hardware.

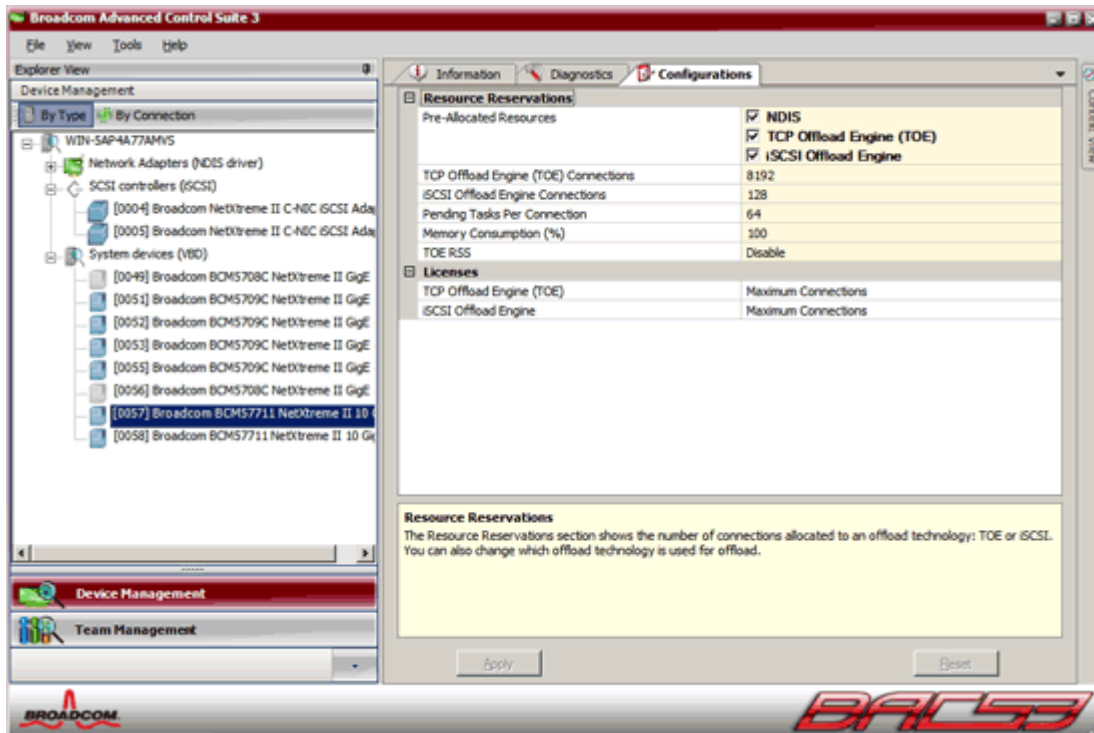


To view resource reservations

1. Click the name of the Broadcom NetXtreme II system device in the Device Management pane.
 2. Click the Configurations tab. If the Configurations tab is not visible, then from the **View** menu, select **Navigate**, then **Device Management**, and then **Configurations**.
 3. From the **Resource Reservations** section, select the property you want to set. If the **Resource Reservations** section is not available, then from the **Context View** tab on the right side of the window, select **Configurations** and then select **Resource Reservations**.
 4. Click **Apply** to confirm the changes to all properties. Click **Reset** to return the properties to their original values.
- The Configurations tab for 1 Gbps network adapters is shown below.



The Configurations tab for 10 Gbps network adapters is shown below.



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 Broadcom NetXtreme II iSCSI device in the SCSI controller section of the Device Management 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 0 and 4094.
4. After the configurations are complete, click **Apply** to save the settings or click **Reset** to revert back to the previous settings.

VIEWING LICENSES



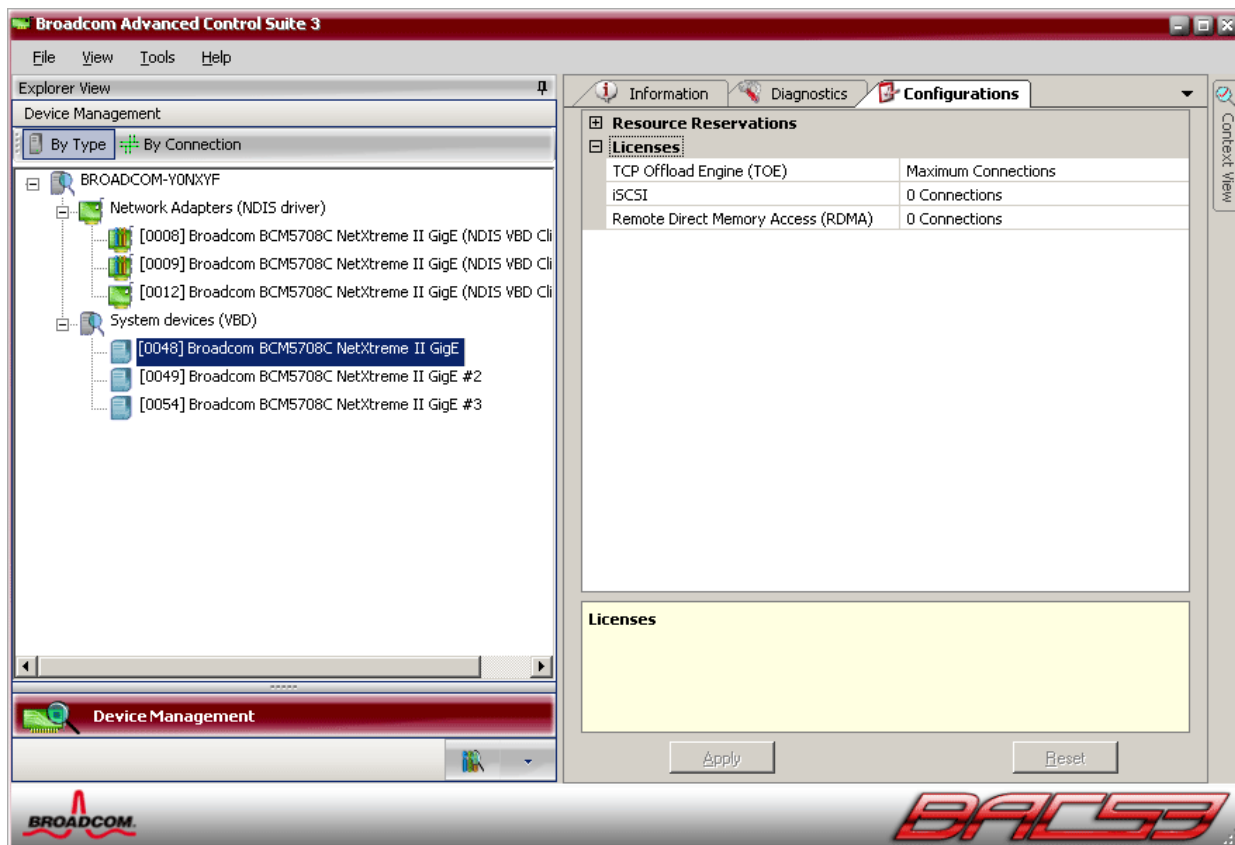
NOTES:

- The **Licenses** section of the **Configurations** tab is only available for Broadcom NetXtreme II adapters and VBD drivers.
- Not all offload technologies are available with all adapters.

The **Licenses** section shows the number of connections available for TOE and iSCSI offload technologies.

To view licenses

1. Click the name of the Broadcom NetXtreme II system device in the Device Management pane.
2. Click the **Configurations** tab. If the **Configurations** tab is not visible, then from the **View** menu, select **Navigate**, then **Device Management**, and then **Configurations**.



CONFIGURING TEAMING

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.
- The TCP Offload Engine (TOE), 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.
- 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.
- Prior to creating a team on a Windows Server 2003 or Windows Server 2008 system, ensure that RSS is disabled in all members of the team. See [Setting Adapter Properties](#) for more information. If you include an adapter in a team that has RSS enabled, then disable RSS and rebuild the team to ensure that the team's advanced properties are properly set.

TEAM TYPES

You can create four types of load balance teams:

- Smart Load Balance and Failover
- Link Aggregation (802.3ad) (TOE is not applicable)
- Generic Trunking (FEC/GEC)/802.3ad-Draft Static (TOE is not applicable)
- SLB (Auto-Fallback Disable) – The Auto-Fallback Disable feature is configured for Smart Load Balance and Failover type teams in the Teaming Wizard.



NOTE: NetXtreme II 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.



NOTE: TOE is not applicable for Link Aggregation team type. NetXtreme II network adapters with iSCSI enabled is not supported for Link Aggregation team type.

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.



NOTE: TOE is not applicable for Generic Trunking (FEC/GEC)/802.3ad-Draft Static team type. NetXtreme II network adapters with iSCSI enabled is not supported for Generic Trunking (FEC/GEC)/802.3ad-Draft Static team type.

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.

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.

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.

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. Click the **Team Management** button at the bottom of the BACS 3 window, or from the **View** menu, select **Navigate**, then **Team Management**.
2. Click **Teams** at the top of the Team Management pane, or select one of the listed adapters.
3. 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, select **Edit Team** from the **Team** menu, 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.





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](#).

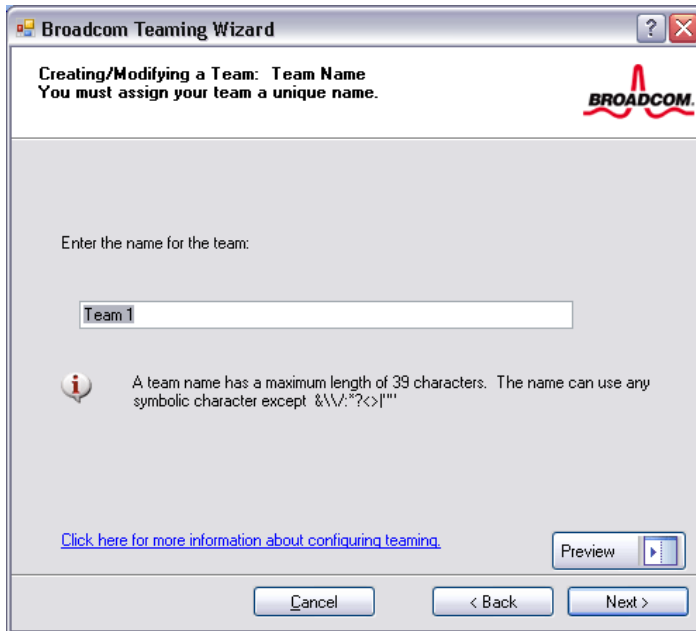
4. To continue using the wizard, click **Next**.



5. 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.



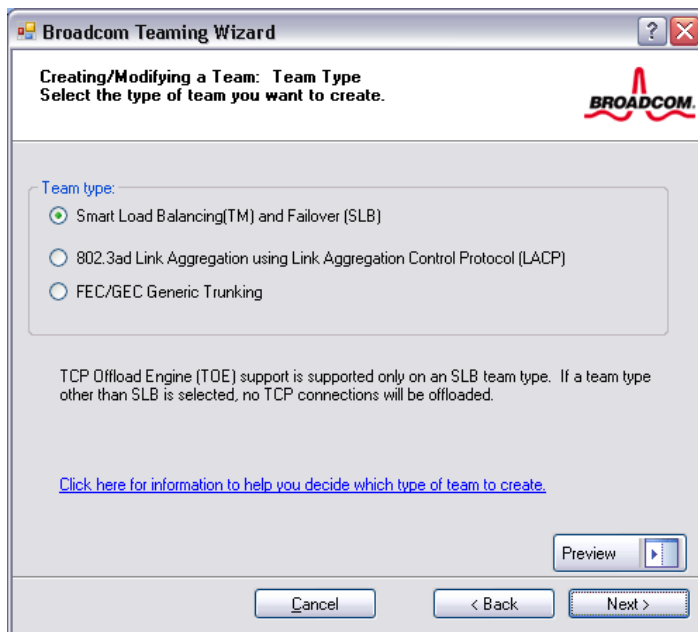
NOTE: The team name cannot exceed 39 characters, cannot begin with spaces, and cannot contain any of the following characters: & \ / : * ? < > |



6. Select the type of team you want to create. 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: NetXtreme II network adapters with iSCSI enabled is 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.



- 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**.



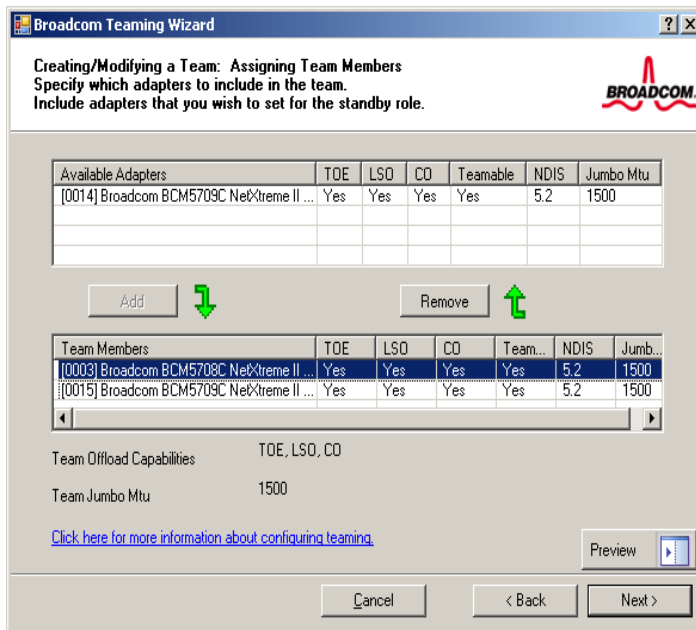
NOTE: There must be at least one Broadcom network adapter assigned to the team.

The TCP Offload Engine (TOE), Large Send Offload (LSO) and Checksum Offload (CO) columns indicate if the TOE, LSO, Jumbo MTU, and/or the CO properties are supported for the adapter. The TOE, 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.

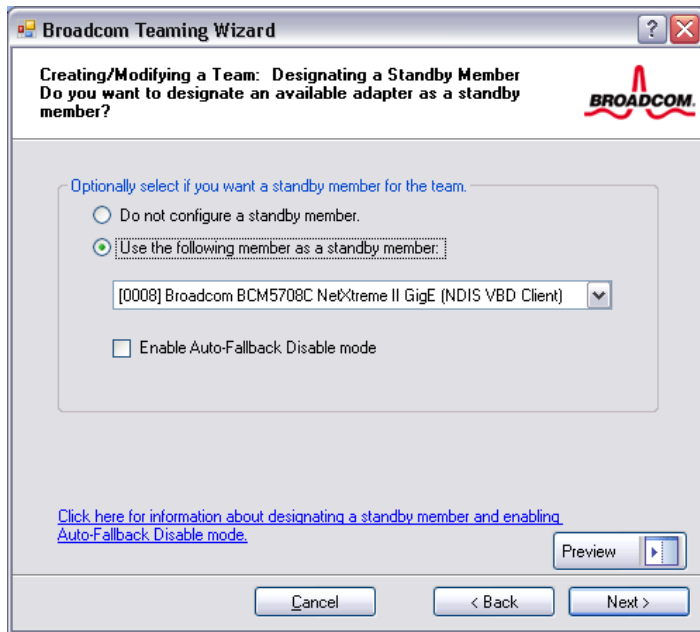


NOTES:

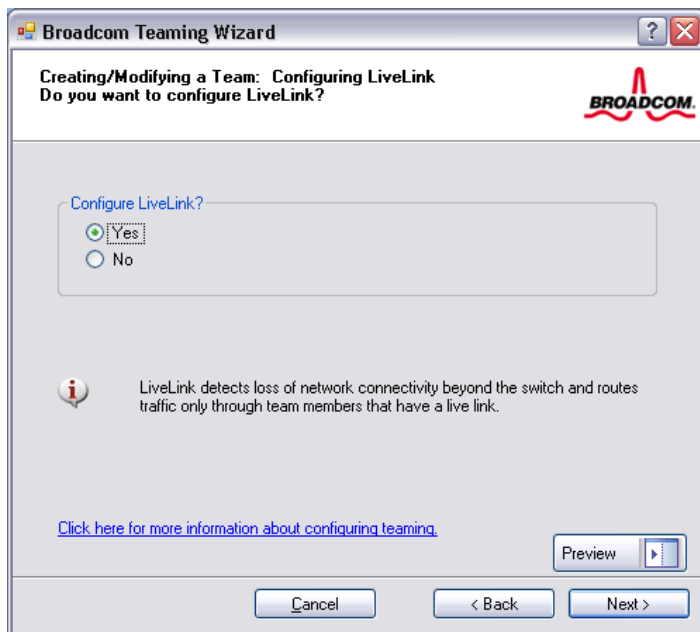
- Adding a network adapter to a team where its driver is disabled may negatively affect the offloading capabilities of the team. This may have an impact on the team's performance. Therefore, it is recommended that only driver-enabled network adapters be added as members to a team.



- 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.
- 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**.



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



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.



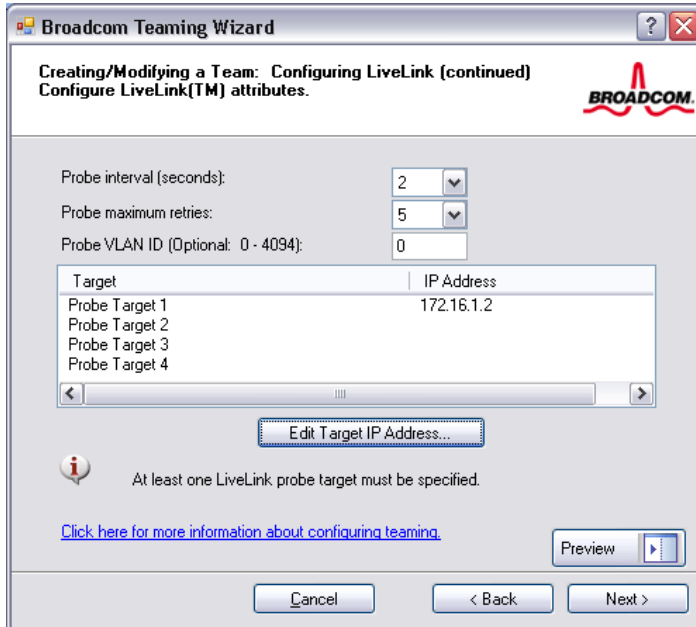
NOTE: Each LiveLink enabled team can only communicate with Probe Targets on a single VLAN. Also, VLAN ID 0 is equivalent to an untagged network. If the Probe VLAN ID is set to a value other than 0, then a VLAN must be

created with an identical VLAN tag value (see [Step 18](#)).

- 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**.



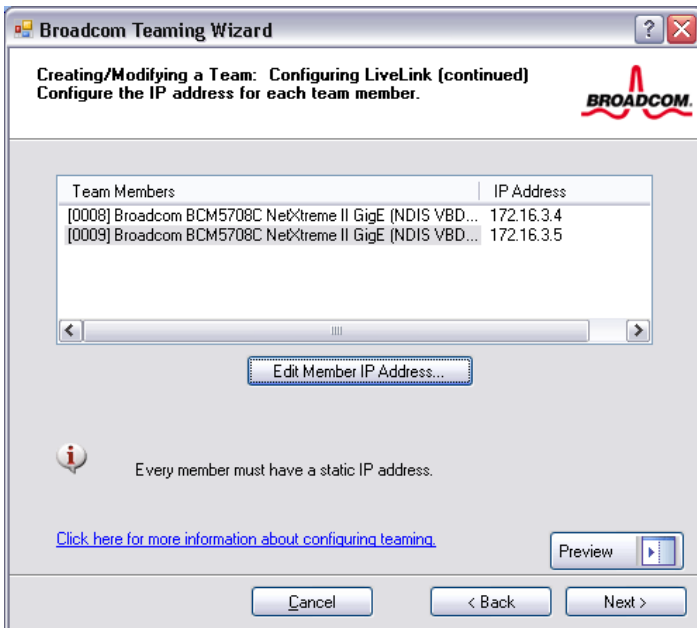
NOTE: Only the first probe target is required. You can specify up to three additional probe targets to serve as backups by assigning IP addresses to the other probe targets.



- 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.

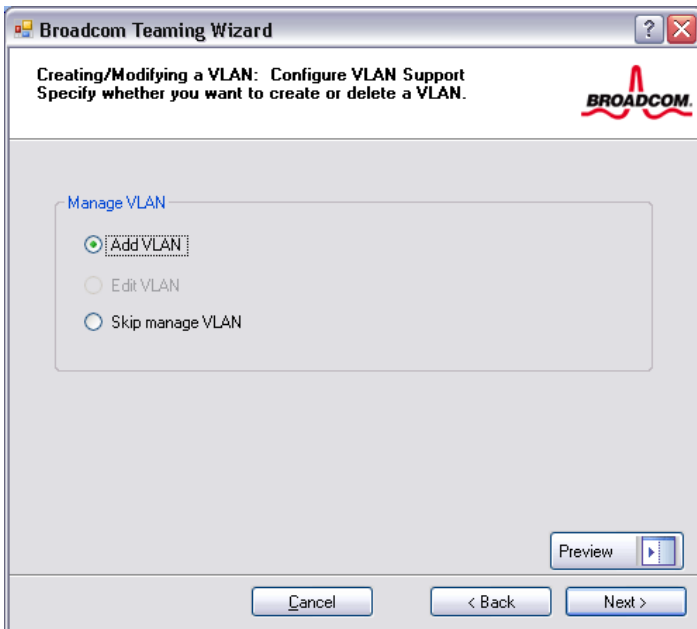


15. If you want to create a VLAN on the team, select **Add VLAN**, or if you want to change the settings of an existing VLAN, select **Edit VLAN**, then click **Next**. If you do not want to create or edit a VLAN, select **Skip Manage VLAN**, then click **Next**, and continue with the wizard from the 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.



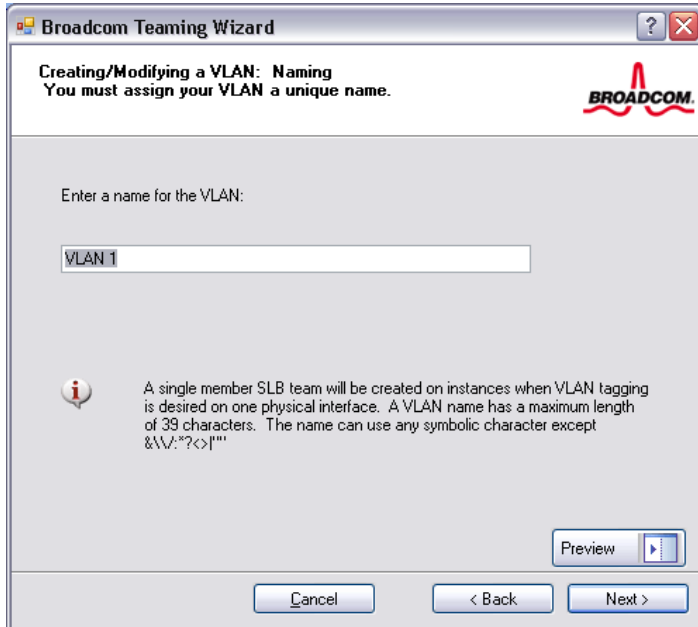
NOTE: VLANs can only be created when all team members are Broadcom adapters.



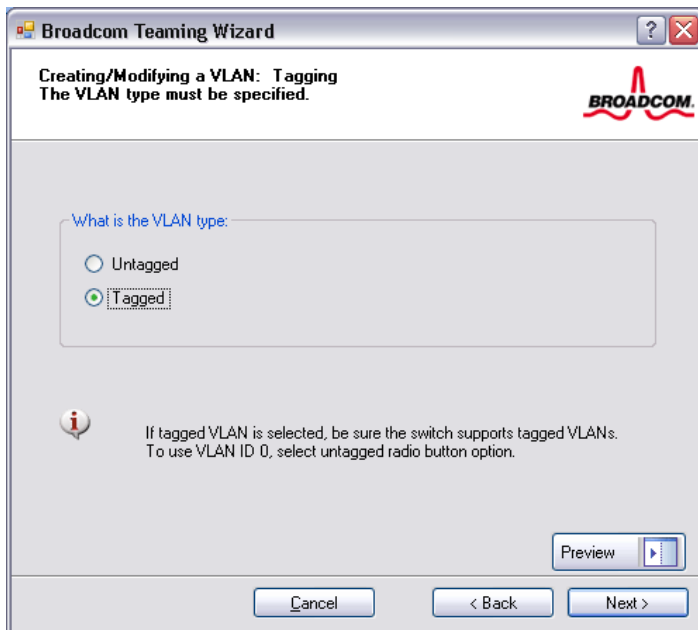
16. Type the VLAN name and then click **Next**.




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



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 (see [Step 19](#) of this procedure).




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



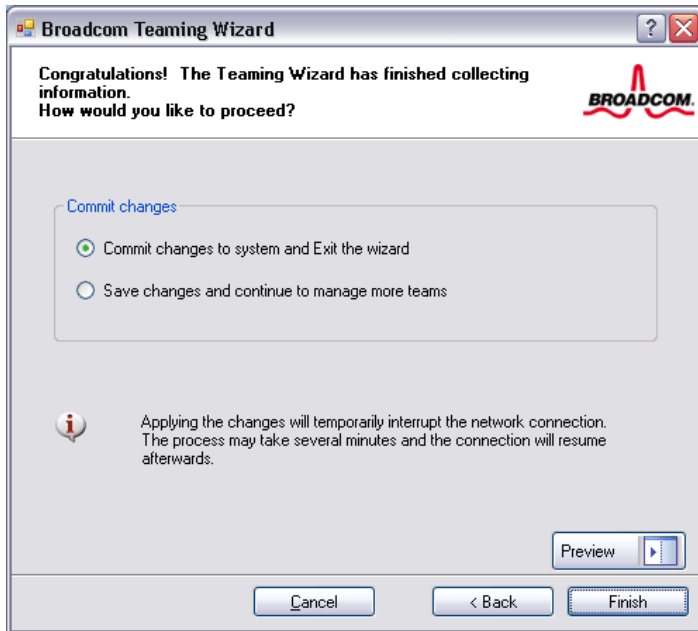
19. Select **Yes** to add or manage another VLAN and then click **Next**. Repeat until you do not want to add or manage any additional VLANs.




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.



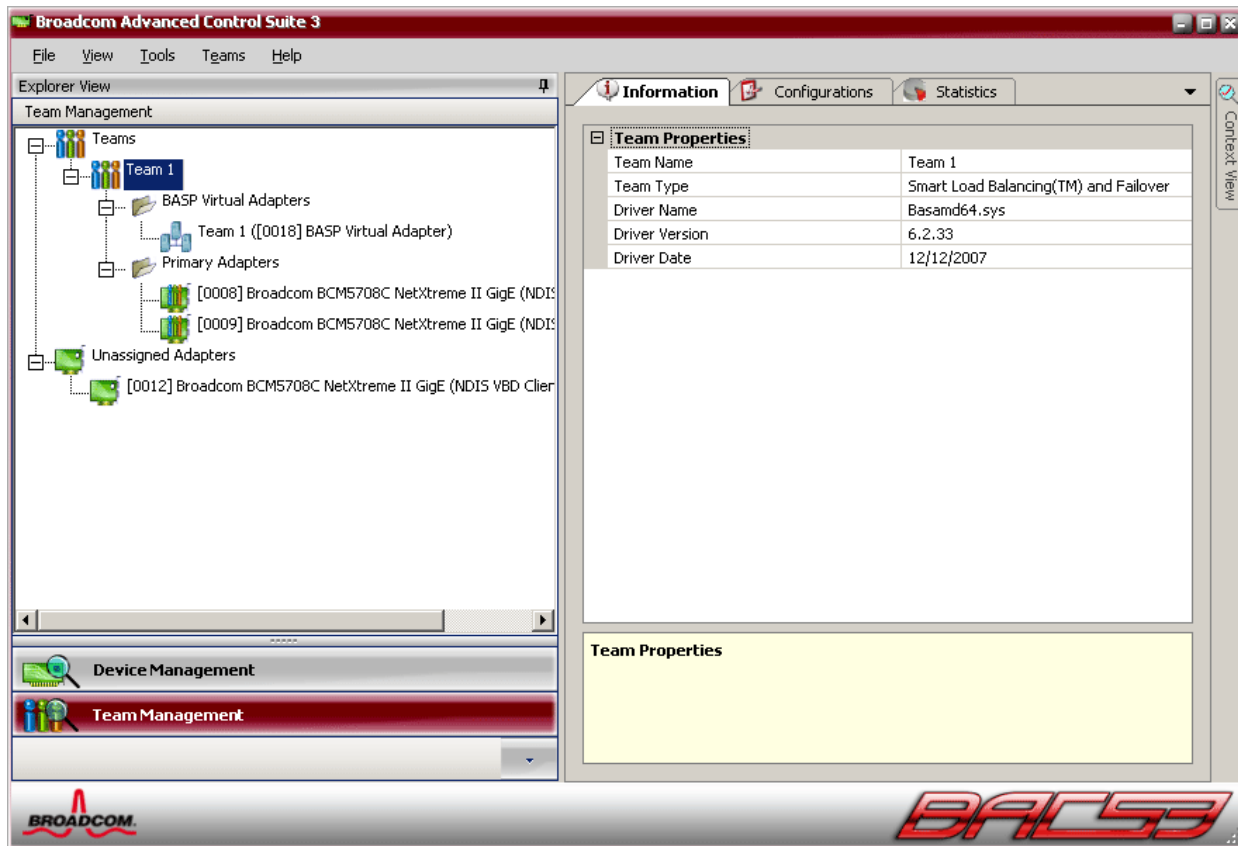
20. To apply and commit the changes to the team, select **Commit changes to system and Exit the wizard**. To apply your changes but continue using the wizard, select **Save changes and continue to manage more teams**. Click Finish.



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.



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.



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.

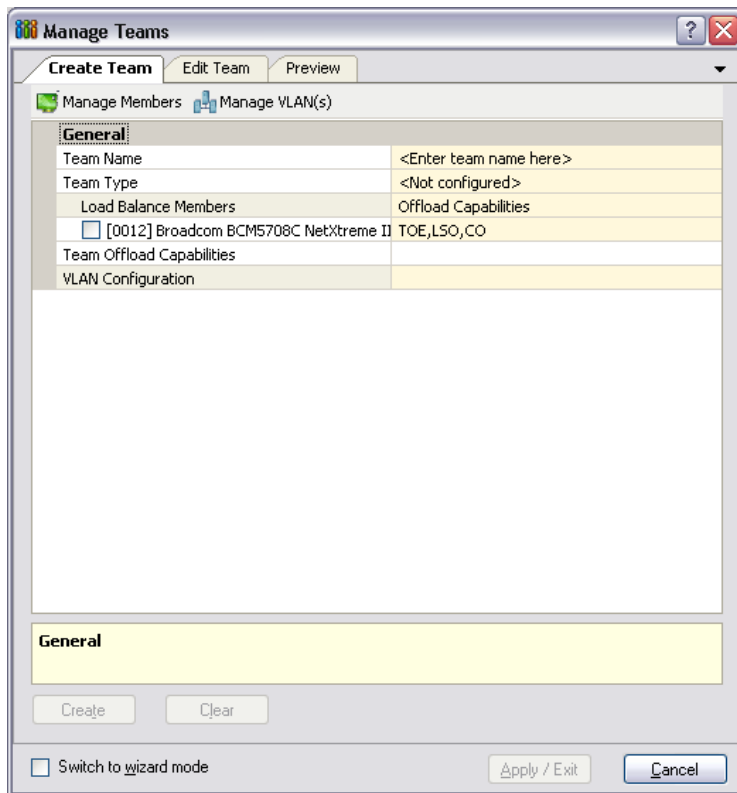
1. Click the **Team Management** button at the bottom of the BACS 3 window, or from the **View** menu, select **Navigate**, then **Team Management**.
2. Click **Teams** at the top of the Team Management pane, or select one of the listed devices in the "Unassigned Adapters" section.
3. From the **Team** 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.



4. Click **Expert Mode**.



NOTE: If you want to always use Expert Mode to create a team, click **Default to Expert Mode on next start**.

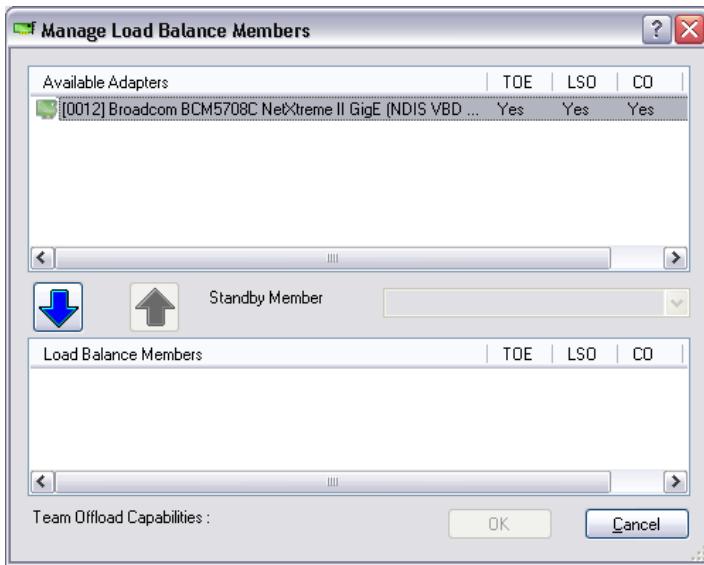


5. Click the **Create Team** tab.



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

6. Click the **Team Name** field to enter a team name.
7. Click the **Team Type** field to select a team type.
8. Click **Manage Members** at the top of the window.



9. 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.
10. 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 TCP Offload Engine (TOE), Large Send Offload (LSO), and Checksum Offload (CO) columns indicate if the TOE, LSO, and/or the CO properties are supported for the adapter. The TOE, 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.



NOTES:

- 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.

11. Click **OK** to accept your changes to the team members.
12. Click **Create** to save the team information.
13. Repeat steps 6. through 12. 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.
14. Click **Apply/Exit** to create all the teams you have defined and exit the Manage Teams window.
15. 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 reenoble it, you must first disable and reenoble all team members before you reenoble 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.

16. 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

- Click the **Team Management** button at the bottom of the BACS 3 window, or from the **View** menu, select **Navigate**, then **Team Management**.
- Select one of the listed teams.
- 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.

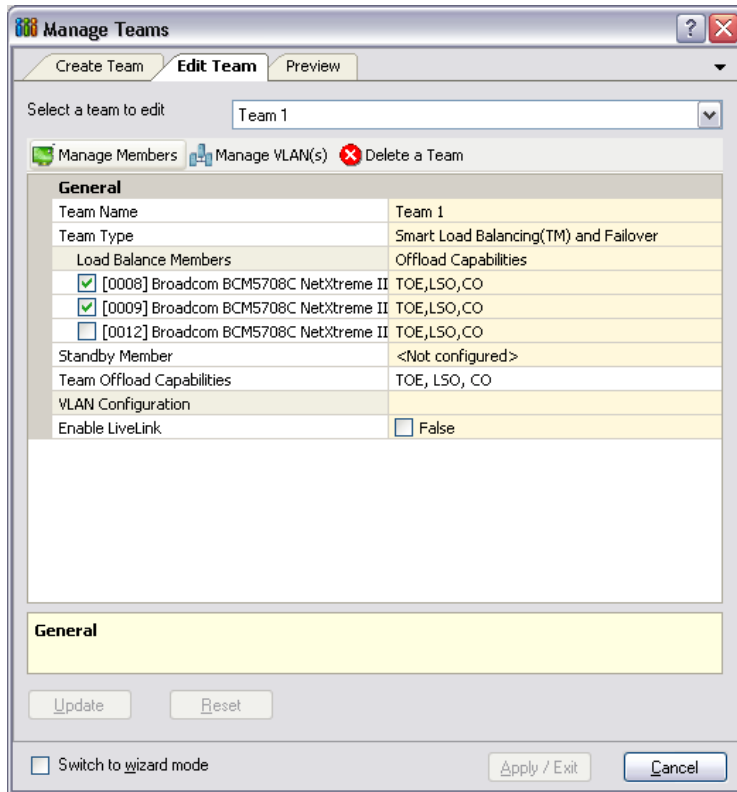


- The wizard Welcome screen appears. Click **Next** to continue modifying a team using the wizard or click **Expert Mode** to work in Expert Mode.



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

- Click the **Edit Team** tab.



6. 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.
7. Click **Apply/Exit** to apply the updates and exit the Manage Teams window.
8. 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.

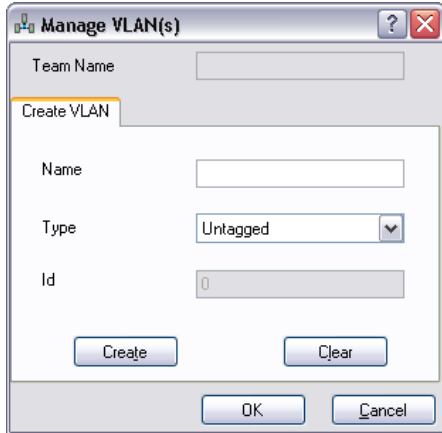
You can define up to 64 VLANs per team (63 VLANs that are tagged and 1 VLAN that is not tagged). VLANs can only be created when all teams members are Broadcom adapters. If you try to create a VLAN with a non-Broadcom adapter, an error message is displayed.

To configure a team with a VLAN

1. Click the **Team Management** button at the bottom of the BACS 3 window, or from the **View** menu, select **Navigate**, then **Team Management**.
2. Select one of the listed teams.
3. From the **Teams** menu, select **Add VLAN**.
4. The Welcome screen appears.
5. Click **Expert Mode**.
6. On the **Create Team** tab of the **Manage Teams** window, click **Manage VLAN(s)**.



7. Type the VLAN name, then select the type and ID.
8. 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.
9. Continue this process until all VLANs are defined, then click **OK** to create them.



10. 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. Click the **Team Management** button at the bottom of the BACS 3 window, or from the **View** menu, select **Navigate**, then **Team Management**.
2. Select one of the listed VLANs.
3. Click the **Information** tab to view the properties of the VLAN adapter.
4. Click the **Statistics** tab to view the statistics for the VLAN adapter.
5. 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. Click the **Team Management** button at the bottom of the BACS 3 window, or from the **View** menu, select **Navigate**, then **Team Management**.
2. Select the VLAN to delete.
3. From the **Teams** menu, select **Remove VLAN**.
4. Click **Apply**.
5. 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 are also 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 (see “Troubleshooting”).
- A probe target must be on the same subnet as the team, have a valid (not a broadcast, multicast, or unicast), 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. Click the **Team Management** button at the bottom of the BACS 3 window, or from the **View** menu, select **Navigate**, then **Team Management**.
2. Select the Smart Load Balance and Failover or SLB (Auto-Fallback Disable) team.
3. From the **Teams** menu, select **Edit Team**.
4. Click Expert Mode (to configure LiveLink using the Teaming Wizard, see [Using the Broadcom Teaming Wizard](#)).
5. In the Manage Teams window, click the **Edit Team** tab.
6. Select **Enable LiveLink**. The LiveLink Configuration options appear below.
7. 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.
8. 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.

9. Select **Probe Target 1** and type the target IP address for one or all probe targets.



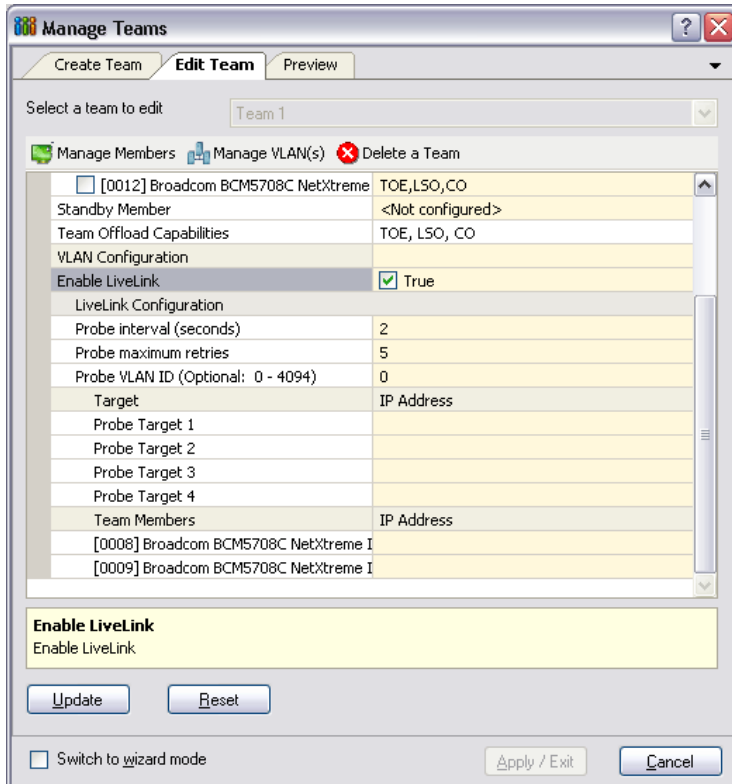
NOTE: Only the first probe target is required. You can specify up to 3 additional probe targets to serve as backups by assigning IP addresses to the other probe targets.

10. 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.

11. Click **Update**. Repeat these steps for each of the other listed team members.
12. Click **Apply/Exit**.



Saving and Restoring a 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** (a .bcg extension is added).

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

1. From the **File** menu, select **Team Restore**.
2. 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.

3. Click **Apply**.
4. Click **Yes** when the message is displayed indicating that the network connection will be temporarily interrupted.

- 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.

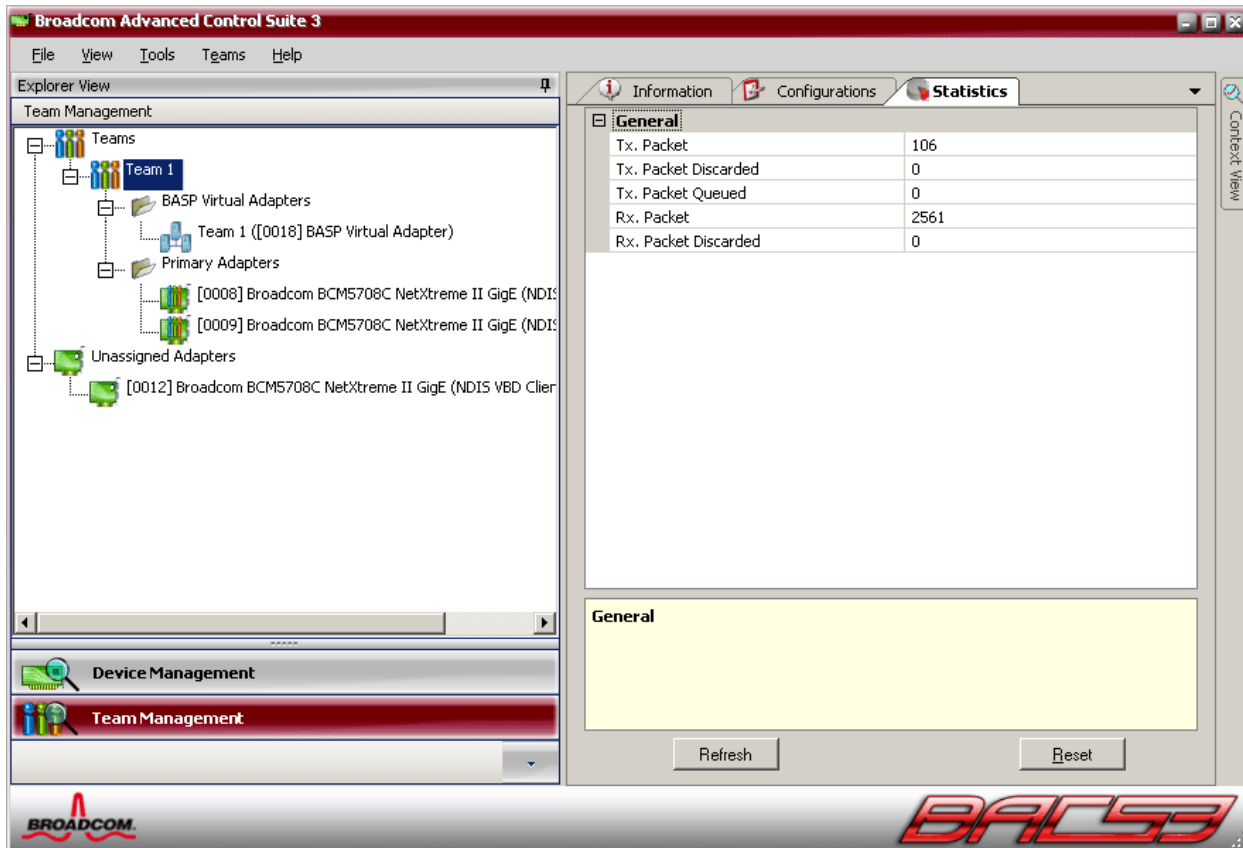
VIEWING BASP STATISTICS

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.

- If the **Statistics** tab is not visible, then from the **View** menu, select **Navigate**, then **Team Management**, and then **Statistics**.
- If the **General** section is not visible, then from the **Context View** tab on the right side of the window, select **Statistics** and then select **General**.

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



CONFIGURING WITH THE COMMAND LINE INTERFACE UTILITY

An alternate method to BACS 3 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 3, 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.

SUPPORTED OPERATING SYSTEMS

BACSCLI is supported on the following operating systems:

- Windows 2000
- Windows XP
- Windows Server 2003
- Windows Server 2008 (including Server Core)
- Windows Server 2008 R2 (including Server Core)

INSTALLATION

On a system with Broadcom NetXtreme I and NetXtreme II network adapters, BACSCLI is installed when BACS 3 is installed with the installer.

Specifications: Broadcom NetXtreme II™ Network Adapter User Guide

- [10/100/1000BASE-T and 10GBASE-T Cable Specifications](#)
- [1000/2500BASE-X Fiber Optic Specifications](#)
- [Interface Specifications](#)
- [NIC Physical Characteristics](#)
- [NIC Power Requirements](#)
- [Wake On LAN Power Requirements](#)
- [Environmental Specifications](#)

10/100/1000BASE-T AND 10GBASE-T CABLE SPECIFICATIONS

Table 1: 10/100/1000BASE-T Cable Specifications

<i>Port Type</i>	<i>Connector</i>	<i>Media</i>	<i>Maximum Distance</i>
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 2: 10GBASE-T Cable Specifications

<i>Port Type</i>	<i>Connector</i>	<i>Media</i>	<i>Maximum Distance</i>
10GBASE-T	RJ-45	Category 6 ¹ UTP	50m (164 feet)
		Category 6A ¹ UTP	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.



1000/2500BASE-X FIBER OPTIC SPECIFICATIONS

Table 3: 1000/2500BASE-X Fiber Optic Specifications

Port Type	Connector	Media	Maximum Distance
1000BASE-X	Small form-factor pluggable (SFP) transceiver with LC™ connection system (Infineon p/n V23818-K305-L57)	Multimode fiber (MMF) System optimized for 62.5/ 50 µm graded index fiber	550m (1804 feet)
2500BASE-X ¹	Small form-factor pluggable (SFP) transceiver with LC™ connection system (Finisar p/n FTLF8542E2KNV)	Multimode fiber (MMF) System optimized for 62.5/ 50 µm graded index fiber	550m (1804 feet)

¹ Electricals leveraged from IEEE 802.3ae-2002 (XAUI). 2500BASE-X is term used by Broadcom to describe 2.5 Gbit/s (3.125GBd) operation. LC is a trademark of Lucent Technologies.

INTERFACE SPECIFICATIONS

Table 4: 10/100/1000BASE-T Performance Specifications

Feature	Specification
PCI Express Interface	x4 link width
10/100/1000BASE-T	10/100/1000 Mbps

Table 5: 10GBASE-T Performance Specifications

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

NIC PHYSICAL CHARACTERISTICS

Table 6: NIC Physical Characteristics

NIC Type	NIC Length	NIC Width
BCM5708 PCI Express	14.7 cm (5.79 inches)	6.4 cm (2.52 inches)
BCM5709 PCI Express x4 low-profile	11.9 cm (4.7 inches)	6.9 cm (2.7 inches)
BCM57710/BCM57711 PCI Express x8 low profile	16.8 cm (6.6 inches)	5.1 cm (2.0 inches)

NIC POWER REQUIREMENTS

Table 7 shows the power requirements for the BCM5708C add-in NIC.

Table 7: BCM5708C NIC Power Requirements

Link	NIC 3.3V Current Draw (A)	NIC Power (W)
Idle (no link)	1.44	4.75
1 Gbit	1.97	6.50
100 Mbit	1.60	5.28
10 Mbit	1.62	5.35

Table 8 shows the power requirements for the BCM5709 add-in NIC.

Table 8: BCM5709C NIC Power Requirements

Link	NIC 3.3V Current Draw (A)	NIC Power (W)
Idle (no link)	1.01	3.32
1 Gbit	1.43	4.71
100 Mbit	1.16	3.81
10 Mbit	1.12	3.71

Table 9 shows the power requirements for the BCM57710/BCM57711 add-in NIC.

Table 9: BCM57710/BCM57711 NIC Power Requirements

Link	NIC 12V Current Draw (A)	NIC 3.3V Current Draw (A)	NIC Power (W)^a
Idle (no link)	0.60	0.28	8.12
Low power mode	0.50	0.35	7.16
10GBASE-T link	1.23	1.79	20.67
10GBASE-T traffic	1.24	1.95	21.32

a. Power, measured in watts (W), is a direct calculation of total current draw (A) multiplied by voltage (V). The maximum power consumption for the adapter will not exceed 30W.

WAKE ON LAN POWER REQUIREMENTS

Table 10: BCM5708C Wake On LAN Power Requirements (Nominal Conditions)

100 Mbit Link		10 Mbit Link	
NIC 3.3V Current (mA)	NIC Power (W)	NIC 3.3V Current (mA)	NIC Power (W)
236	0.78	150	0.5

Table 11: BCM5709C Wake On LAN Power Requirements (Nominal Conditions)

100 Mbit Link		10 Mbit Link	
NIC 3.3V Current (mA)	NIC Power (W)	NIC 3.3V Current (mA)	NIC Power (W)
0	0.87	0	0.85

ENVIRONMENTAL SPECIFICATIONS

Table 12: BCM5708 Environmental Specifications

Condition	Operating Specification	Storage Specification
Temperature	0°C to 55°C (+32°F to +131°F)	-40°C to +85°C (-40°F to +185°F)
Relative humidity	5% to 85% (noncondensing) 40°C, 16 hour dwells at extremes	5% to 95% (noncondensing) 10°C/hour
Altitude	Up to 10,000 ft.	Up to 35,000 ft.
Shock	10g, 1/2 sine wave, 11 ms	60g, 1/2 sine wave, 11 ms
Vibration, peak to peak displacement	0.005 in. max (5 Hz to 32 Hz)	0.1 in. max (5 Hz to 17 Hz)
Vibration, peak acceleration	0.25g (5 Hz to 500 Hz) (Sweep Rate = 1 octave/min.)	0.25g (5 Hz to 500 Hz) (Sweep Rate = 1 octave/min.)

Table 13: BCM5709 Environmental Specifications

Parameter	Condition
Operating Temperature	0°C to 55°C
Air Flow Requirement (LFM)	0
Storage Temperature	-40°C to +65°C
Storage Humidity	5% to 95% condensing
Vibration and Shock	IEC 68, FCC Part 68.302, NSTA, 1A
Electrostatic/Electromagnetic Susceptibility	EN 61000-4-2, EN 55024



Regulatory Information: Broadcom NetXtreme II™ Network Adapter User Guide

- [FCC Notice](#)
- [VCCI Notice](#)
- [CE Notice](#)
- [Canadian Regulatory Information \(Canada Only\)](#)
- [Korea Communications Commission \(KCC\) Notice \(Republic of Korea Only\)](#)

FCC NOTICE

FCC, CLASS B

Broadcom NetXtreme II Gigabit Ethernet Controller

Broadcom Corporation
190 Mathilda Place
Sunnyvale, California 94086 USA

The equipment complies with Part 15 of the FCC Rules. Operation is subject to the following two conditions: 1) The device may not cause harmful interference, and 2) This equipment must accept any interference received, including interference that may cause undesired operation.

The equipment has been tested and found to comply with the limits for a Class B digital device, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference in a residential installation. The equipment generates, uses and can radiate radio-frequency energy and, if not installed and used in accordance with the instructions, may cause harmful interference to radio communications. However, there is no guarantee that interference will not occur in a particular installation. If the equipment does cause harmful interference to radio or television reception, which can be determined by turning the equipment off and on, the user is encouraged to try to correct the interference by one or more of the following measures:

- Reorient or relocate the receiving antenna.
- Increase the separation between the equipment and the receiver.
- Connect the equipment into an outlet on a circuit different from that to which the receiver is connected.
- Consult the dealer or an experienced radio/TV technician for assistance.

Do not make mechanical or electrical modifications to the equipment.



NOTE: If the device is changed or modified without permission of Broadcom, the user may void his or her authority to operate the equipment.

FCC, CLASS A

Broadcom NetXtreme II Gigabit Ethernet Controller

Broadcom NetXtreme II 10 Gigabit Ethernet Controller

Broadcom Corporation
190 Mathilda Place
Sunnyvale, California 94086 USA

This device complies with Part 15 of the FCC Rules. Operations is subject to the following two conditions: 1) This device may not cause harmful interference, and 2) This device must accept any interference received, including interference that may cause undesired operation.

This product has been tested and found to comply with the limits for a Class A digital device pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This product generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the manufacturer's instruction manual, may cause harmful interference with radio communications. Operation of this product in a residential area is likely to cause harmful interference, in which case you will be required to correct the interference at your own expense.

These limits are designed to provide reasonable protection against harmful interference in a non-residential installation. However, there is no guarantee that interference will not occur in a particular installation. If this equipment does cause harmful interference with radio or television reception, which can be determined by turning the equipment off and on, you are encouraged to try to correct the interference by one or more of the following measures:

- Reorient the receiving antenna.
- Relocate the system with respect to the receiver.
- Move the system away from the receiver.
- Plug the system into a different outlet so that the system and receiver are on different branch circuits.

Do not make mechanical or electrical modifications to the equipment.



NOTE: If the device is changed or modified without permission of Broadcom, the user may void his or her authority to operate the equipment.

VCCI NOTICE**CLASS B**

Broadcom NetXtreme II Gigabit Ethernet Controller

Broadcom Corporation
190 Mathilda Place
Sunnyvale, California 94086 USA



The equipment is a Class B product based on the standard of the Voluntary Control Council for Interference from Information Technology Equipment (VCCI). If used near a radio or television receiver in a domestic environment, it may cause radio interference. Install and use the equipment according to the instruction manual.



CAUTION! The potential exists for this equipment to become impaired in the presence of conducted radio frequency energy between the frequency range of 59–66 MHz. Normal operation will return upon removal of the RF energy source.

VCCI Class B Statement (Japan)

この装置は、情報処理装置等電波障害自主規制協議会（VCCI）の基準に基づくクラス B 情報技術装置です。この装置は、家庭環境で使用することを目的としていますが、この装置がラジオやテレビジョン受信機に近接して使用されると、電波障害を引き起こすことがあります。取扱説明書に従って正しい取り扱いをして下さい。

CLASS A

Broadcom NetXtreme II Gigabit Ethernet Controller

Broadcom NetXtreme II 10 Gigabit Ethernet Controller

Broadcom Corporation
190 Mathilda Place
Sunnyvale, California 94086 USA

This equipment is a Class A product based on the standard of the Voluntary Control Council for interference by Information Technology Equipment (VCCI). If used in a domestic environment, radio disturbance may arise. Install and use the equipment according to the instruction manual.

VCCI Class A Statement (Japan)

この装置は、情報処理装置等電波障害自主規制協議会（VCCI）の基準に基づくクラス A 情報技術装置です。この装置を家庭環境で使用すると電波障害を引き起こす可能性があります。この場合には使用者が適切な対策を講ずるよう要求されることがあります。

CE NOTICE**CLASS B**

Broadcom NetXtreme II Gigabit Ethernet Controller

CLASS A

Broadcom NetXtreme II Gigabit Ethernet Controller

Broadcom NetXtreme II 10 Gigabit Ethernet Controller

БЪЛГАРСКИ Bulgarian	<p>Този продукт отговаря на 2006/95/ЕС (Нисковолтова директива), 2004/108/ЕС (Директива за електромагнитна съвместимост) и измененията на Европейския съюз.</p> <p>Изготвена е "Декларация за съответствие" според горепосочените директиви и стандарти, която се съхранява в Broadcom Corporation, 190 Mathilda Place, Sunnyvale, California 94086, USA.</p> <p>Европейски съюз, Клас В Това устройство на Broadcom е класифицирано за използване в типичната за Клас В жилищна среда.</p> <p>Европейски съюз, Клас А ВНИМАНИЕ: Това е продукт от Клас А. В жилищна среда този продукт може да създаде радиочестотни смущения, в който случай потребителят ще трябва да вземе съответните мерки.</p>
ČESKÝ Czech	<p>Bylo ustanoveno, že tento produkt splňuje směrnici 2006/95/EC (niskonapěťová směrnice), směrnici 2004/108/EC (směrnice EMC) a dodatky Evropské unie.</p> <p>„Prohlášení o shodě“ v souladu s výše uvedenými směrnici a normami bylo zpracováno a je uloženo v archivu společnosti Broadcom Corporation, 190 Mathilda Place, Sunnyvale, California 94086, USA.</p> <p>Evropská unie, třída B Toto zařízení společnosti Broadcom je klasifikováno pro použití v obvyklém prostředí domácnosti (třída B).</p> <p>Evropská unie, třída A VAROVÁNÍ: Toto je produkt třídy A. V domácím prostředí může tento produkt způsobovat rušení rádiových frekvencí, a v takovém případě se od uživatele vyžaduje, aby učinil odpovídající opatření.</p>
Danish	<p>Denne produkt er fundet i overensstemmelse med 2006/95/EC (Lavvoltage-direktivet), 2004/108/EC (EMC-direktivet) og den Europæiske Unions ændringer.</p> <p>En "Overensstemmelseserklæring", som er i henhold til foregående direktiver og standarder, er udført og arkiveret hos Broadcom Corporation, 190 Mathilda Place, Sunnyvale, California 94086, USA.</p> <p>Den Europæiske Union, Klasse B Denne Broadcom-enhed er klassificeret til anvendelse i et typisk Klasse B-hjemligt miljø.</p> <p>Europæiske Union, Klasse A ADVARSEL: Dette er et Klasse A-produkt. I et hjemligt miljø kan dette produkt medføre forstyrrelse af radiofrekvens, og i det tilfælde må brugeren fortage passende foranstaltninger.</p>
NEDERLANDS Dutch	<p>Dit product is in overeenstemming bevonden met 2006/95/EC (Laagspanningsrichtlijn), 2004/108/EC (EMC-richtlijn) en amendementen van de Europese Unie.</p> <p>Een "Verklaring van conformiteit" in overeenstemming met de voorgenoemde richtlijnen en standaarden is beschikbaar bij Broadcom Corporation, 190 Mathilda Place, Sunnyvale, California 94086, USA.</p> <p>Europese Unie/Klasse B Dit Broadcom-apparaat is geclassificeerd voor gebruik in een typische klasse B woonomgeving.</p> <p>Europese Unie/Klasse A VOORZICHTIG: Dit is een Klasse A-product. Binnen een woonomgeving kan dit product radiofrequente storingen veroorzaken, in welk geval de gebruiker passende maatregelen dient te nemen.</p>
English	<p>This product has been determined to be in compliance with 2006/95/EC (Low Voltage Directive), 2004/108/EC (EMC Directive), and amendments of the European Union.</p> <p>A "Declaration of Conformity" in accordance with the preceding directives and standards has been made and is on file at Broadcom Corporation, 190 Mathilda Place, Sunnyvale, California 94086, USA.</p> <p>European Union, Class B This Broadcom device is classified for use in a typical Class B domestic environment.</p> <p>European Union, Class A WARNING: This is a Class A product. In a domestic environment this product may cause radio frequency interference in which case the user may be required to take adequate measures.</p>

EESTLANE Estonian	<p>Antud toode vastab direktiividele 2006/95/EÜ (Madalpinge direktiiv), 2004/108/EÜ (EMC direktiiv) ja ELi parandustele.</p> <p>Vastavalt ülalloodud direktiividele ja standarditele on koostatud „Vastavusdeklaratsioon”, mis on arvel ettevõttes Broadcom Corporation, 190 MathildaPlace, Sunnyvale, California 94086, USA.</p> <p>Euroopa Liit, Klass B</p> <p>Antud Broadcom toode on klassifitseeritud kasutamiseks tüüpilises B-klassi koduses keskkonnas.</p> <p>Union européenne, classe A</p> <p>AVERTISSEMENT : Ce produit est un produit de classe A. Dans un environnement résidentiel, ce produit peut provoquer des perturbations radioélectriques, auquel cas l'utilisateur peut se voir obligé de prendre les mesures appropriées.</p>
Finnish	<p>Tämä tuote täyttää Euroopan unionin direktiivin 2006/95/EY (pienjännitedirektiivi) ja direktiivin 2004/108/EY (sähkömagneettisesta yhteensopivuudesta annettu direktiivi), sellaisina kuin ne ovat muutettuina, vaatimukset.</p> <p>Yllä mainittujen direktiivien ja standardien mukainen vaatimustenmukaisuusvakuutus on tehty, ja sitä säilyttää Broadcom Corporation, 190 Mathilda Place, Sunnyvale, California 94086, USA.</p> <p>Euroopan unioni, luokka B</p> <p>Tämä Broadcom-laite on luokiteltu käytettäväksi tyypillisessä luokan B kotiympäristössä.</p> <p>Euroopan unioni, Luokka A</p> <p>VAROITUS: Tämä on Luokan A tuote. Asuinympäristössä tämä laite saattaa aiheuttaa radiotaajuushäiriöitä, mikä saattaa edellyttää toimia laitteen käyttäjältä.</p>
FRANÇAIS French	<p>Ce produit a été déclaré conforme aux directives 2006/95/EC (Directive sur la faible tension), 2004/108/EC (Directive EMC) et aux amendements de l'Union européenne.</p> <p>Une « Déclaration de Conformité » relative aux normes et directives précédentes a été rédigée et est enregistrée auprès de Broadcom Corporation, 190 Mathilda Place, Sunnyvale, California 94086, USA.</p> <p>Union européenne, classe B</p> <p>Cet appareil Broadcom est classé pour une utilisation dans un environnement résidentiel classique (classe B).</p> <p>Union européenne, classe A</p> <p>AVERTISSEMENT : Ce produit est un produit de classe A. Dans un environnement résidentiel, ce produit peut provoquer des perturbations radioélectriques, auquel cas l'utilisateur peut se voir obligé de prendre les mesures appropriées.</p>
DEUTSCH German	<p>Es ist befunden worden, dass dieses Produkt in Übereinstimmung mit 2006/95/EC (Niederspannungs-Richtlinie), 2004/108/EC (EMV-Richtlinie) und Ergänzungen der Europäischen Union steht.</p> <p>Eine Konformitätserklärung in Übereinstimmung mit den oben angeführten Normen ist abgegeben worden und kann bei Broadcom Corporation, 190 Mathilda Place, Sunnyvale, California 94086, USA.</p> <p>Europäische Union, Klasse B</p> <p>Dieses Gerät von Broadcom ist für die Verwendung in einer typisch häuslichen Umgebung der Klasse B vorgesehen.</p> <p>Europäische Union, Klasse A</p> <p>WARNUNG: Dies ist ein Produkt der Klasse A. In einer häuslichen Umgebung kann dieses Produkt Hochfrequenzstörungen verursachen. In diesem Fall muss der Benutzer die entsprechenden Maßnahmen treffen.</p>
ΕΛΛΗΝΙΚΟΣ Greek	<p>Το προϊόν αυτό συμμορφώνεται με τις οδηγίες 2006/95/ΕΕ (Οδηγία περί χαμηλής τάσης), 2004/108/ΕΕ (Οδηγία περί ηλεκτρομαγνητικής συμβατότητας), και τροποποιήσεις τους από την Ευρωπαϊκή Ένωση.</p> <p>Μία «Δήλωση Συμμόρφωσης» σύμφωνα με τις προηγούμενες οδηγίες και πρότυπα υπάρχει και είναι αρχαιοθετημένη στο Broadcom Corporation, 190 Mathilda Place, Sunnyvale, California 94086, USA.</p> <p>Ευρωπαϊκή Ένωση, Κατηγορία Β</p> <p>Αυτή η συσκευή Broadcom είναι κατάλληλη για χρήση σε ένα σύνθετες οικιακό περιβάλλον κατηγορίας Β.</p> <p>Ευρωπαϊκή Ένωση, Κατηγορία Α</p> <p>ΠΡΟΕΙΔΟΠΟΙΗΣΗ: Αυτό είναι ένα προϊόν κατηγορίας Α. Σε οικιακό περιβάλλον, αυτό το προϊόν μπορεί να προκαλέσει παρεμβολές ραδιοσυχνότητας (RF), στην οποία περίπτωση μπορεί να απαιτηθεί η λήψη κατάλληλων μέτρων από τον χρήστη.</p>

MAGYAR Hungarian	<p>A termék megfelel a 2006/95/EGK (alacsony feszültségű eszközökre vonatkozó irányelv), a 2004/108/EGK (EMC irányelv) és az Európai Unió ajánlásainak.</p> <p>Az előbbieken ismertetett irányelvek és szabványok szellemében „Megfelelési nyilatkozat” készült, amely az írországi Broadcom Corporation, 190 Mathilda Place, Sunnyvale, California 94086, USA.</p> <p>Európai Unió, „B” osztály Ez a Broadcom eszköz „B” osztályú besorolást kapott, tipikus lakossági környezetben való használatra alkalmas.</p> <p>Európai Unió, „A” osztály FIGYELEM! „A” osztályba sorolt termék. Lakóhelyi környezetben ez a termék rádiófrekvenciás (RF) interferenciát okozhat, ebben az esetben a felhasználónak gondoskodnia kell a szükséges ellenintézkedésekről.</p>
PORTUGUES Iberian Portuguese	<p>Este produto está em conformidade com 2006/95/EC (Directiva de baixa tensão), com 2004/108/EC (Directiva de compatibilidade electromagnética) e com as alterações da União Europeia.</p> <p>Foi elaborada uma “declaração de conformidade” de acordo com as normas e directivas anteriores, encontrando-se arquivada na Broadcom Corporation, 190 Mathilda Place, Sunnyvale, California 94086, USA.</p> <p>União Europeia, Classe B Este dispositivo Broadcom está classificado para utilização num ambiente doméstico típico Classe B.</p> <p>União Europeia, Classe A ADVERTÊNCIA: Este é um produto Classe A. Num ambiente doméstico, este produto pode provocar interferências de frequência de rádio, podendo ser necessário que o utilizador adopte as medidas adequadas.</p>
ITALIANO Italian	<p>Il presente prodotto è stato determinato essere conforme alla 2006/95/CE (Direttiva Bassa Tensione), alla 2004/108/CE (Direttiva CEM) e a rettifiche da parte dell'Unione Europea.</p> <p>Una “Dichiarazione di conformità” secondo gli standard e le direttive precedenti è stata emessa e registrata presso Broadcom Corporation, 190 Mathilda Place, Sunnyvale, California 94086, USA.</p> <p>Unione Europea, Classe B Il presente dispositivo Broadcom è classificato per l'uso nel tipico ambiente domestico di Classe B.</p> <p>Unione Europea, Classe A AVVERTENZA: Questo prodotto è classificato come Classe A. In un ambiente domestico il presente prodotto potrebbe provocare interferenze di radiofrequenza, nel qual caso potrebbe essere richiesto all'utente di adottare misure adeguate.</p>
LATVISK Latvian	<p>Šis izstrādājums atbilst direktīvām 2006/95/EK (Direktīva par zemsprieguma iekārtām), 2004/108/EK (Direktīva par elektromagnētisko saderību) un to labojumiem Eiropas Savienības ietvaros.</p> <p>“Atbilstības deklarācija”, kas ir saskaņā ar iepriekšminētajām direktīvām un standartiem, ir sastādīta un tiek glabāta firmā Broadcom Corporation, 190 Mathilda Place, Sunnyvale, California 94086, USA.</p> <p>Eiropas Savienība, klase B Šī firmas Broadcom ražotā ierīce ir atzīta par derīgu darbam B klasei atbilstošos mājas apstākļos.</p> <p>Eiropas Savienība, A klase BRĪDINĀJUMS. Šis A klases izstrādājums. Izmantojot šo izstrādājumu mājas apstākļos, tas var radīt radiotraucējumus; šajā gadījumā lietotājam var būt nepieciešams veikt atbilstošus pasākumus.</p>
Lithuanian	<p>Buvo nustatyta, kad šis produktas atitinka direktyvą 73/23/EEB (žemos įtampos direktyva), 89/336/EEB (elektromagnetinio suderinamumo direktyva) ir Europos Sąjungos pataisas.</p> <p>Atitikties deklaracija pagal visas galiojančias direktyvas ir standartus yra sudaryta ir saugoma įrašyta faile Broadcom Corporation, 190 Mathilda Place, Sunnyvale, California 94086, USA.</p> <p>Europos Sąjunga, B klasė Šis „Broadcom“ prietaisas yra klasifikuotas naudoti įprastose B klasės gyvenamosiose aplinkose.</p> <p>Europos Sąjunga, A klasė ĮSPĖJIMAS. Tai yra A klasės produktas. Gyvenamosiose aplinkose šis produktas gali kelti radijo dažnių trikdžius. Tokiu atveju naudotojui gali reikėti imtis atitinkamų priemonių.</p>

Maltese	<p>Ġie stabbilit li dan il-prodott hu konformi ma' 2006/95/KE (Direttiva dwar il-Vultagġ Baxx), 2004/108/KE (Direttiva EMC), u emendi ta' l-Unjoni Ewropea.</p> <p>Saret "Dikjarazzjoni ta' Konformità" b'konformità mad-direttivi u ma' l-istandards imsemmijin qabel, u din tinsab iffajljata għand Broadcom Corporation, 190 Mathilda Place, Sunnyvale, California 94086, USA.</p> <p>Unjoni Ewropea, Klassi B</p> <p>Dan it-tagħmir Broadcom hu kklassifikat għall-użu f' ambjent residenzjali tipiku ta' Klassi B.</p> <p>Unjoni Ewropea, Klassi A</p> <p>TWISSIJA: Dan huwa prodott ta' Klassi A. F'ambjent domestiku dan il-prodott jista' jikkawża interferenza tal-frekwenza tar-radju (RF), f'liema każ l-utent jista' jkun mefġieg li jieffu miżuri adegwati.</p>
POLSKI Polish	<p>Niniejszy produkt został określony jako zgodny z dyrektywą niskonapięciową 2006/95/WE i dyrektywą zgodności elektromagnetycznej 2004/108/WE oraz poprawkami do nich.</p> <p>Zgodnie ze stosownymi dyrektywami i normami została sporządzona „Deklaracja zgodności”, która jest dostępna w aktach firmy Broadcom Corporation, 190 Mathilda Place, Sunnyvale, California 94086, USA.</p> <p>Unia Europejska, klasa B</p> <p>Niniejsze urządzenie firmy Broadcom zostało zakwalifikowane do klasy B, do użytku w typowych środowiskach domowych.</p> <p>Unia Europejska, klasa A</p> <p>OSTRZEŻENIE: Urządzenie to jest urządzeniem klasy A. W warunkach domowych urządzenie to może wywoływać zakłócenia o częstotliwości radiowej, wymagające od użytkownika podjęcia odpowiednich działań zaradczych.</p>
ROMÂN Romanian	<p>S-a stabilit că acest produs respectă cerințele Directivei 2006/95/CE privind echipamentele de joasă tensiune, ale Directivei 2004/108/CE (Directiva EMC) privind compatibilitatea electromagnetică și ale amendamentelor Uniunii Europene.</p> <p>Conform directivei și standardelor de mai sus, a fost emisă o „Declarație de Conformitate”, arhivată la sediul Broadcom Corporation, 190 Mathilda Place, Sunnyvale, California 94086, USA.</p> <p>Uniunea Europeană, Clasa B</p> <p>Acest echipament Broadcom este clasificat pentru utilizare într-un mediu casnic tipic de Clasă B.</p> <p>Uniunea Europeană, Clasa A</p> <p>AVERTISMENT: Acesta este un produs din Clasa A. În mediul casnic, acest produs poate cauza interferențe radio, caz în care utilizatorul trebuie să ia măsurile necesare.</p>
SLOVENSKÝ Slovakian	<p>Tento výrobok vyhovuje požiadavkám smernice 2006/95/EC (smernica o nízkom napätí), 2004/108/EC (smernica o elektromagnetickej kompatibilite) a neskorším zmenám a doplnkom Európskej.</p> <p>„Vyhlásenie o zhode“ vydané v súlade s predchádzajúcimi smernicami a štandardmi sa nachádza v spoločnosti Broadcom Corporation, 190 Mathilda Place, Sunnyvale, California 94086, USA.</p> <p>Európska únia, Trieda B</p> <p>Toto zariadenie Broadcom triedy B je určené pre domáce prostredie.</p> <p>Európska únia, Trieda A</p> <p>VAROVANIE: Toto je zariadenie triedy A. V domácom prostredí môže tento produkt spôsobovať rušenie rádiových frekvencií. V takom prípade musí používateľ prijať príslušné opatrenia.</p>
Slovenian	<p>Ta izdelek je v skladu z 2006/95/ES (Direktiva o nizki napetosti), 2004/108/ES (Direktiva o elektromagnetni združljivosti) in dopolnili Evropske unije.</p> <p>«Izjava o skladnosti» je bila sprejeta v skladu s predhodnimi direktivami in standardi in je shranjena na naslovu Broadcom Corporation, 190 Mathilda Place, Sunnyvale, California 94086, USA.</p> <p>Evropska unija, razred B</p> <p>Ta Broadcomova naprava je razvrščena za uporabo v značilnem bivalnem okolju razreda B.</p> <p>Evropska unija, razred A</p> <p>OPOZORILO: To je izdelek razreda A. V domačem okolju lahko ta izdelek povzroča motnje radijskih frekvenc, v tem primeru mora uporabnik ustrezno ukrepati.</p>

<p>ESPAÑOL Spanish</p>	<p>Este producto se ha fabricado de conformidad con la Directiva para bajo voltaje 2006/95/EC (Low Voltage Directive), la Directiva para compatibilidad electromagnética 2004/108/EC (EMC Directive) y las enmiendas de la Unión Europea. Se ha realizado una "Declaración de conformidad" de acuerdo con las directivas y estándares anteriores y está archivada en Broadcom Corporation, 190 Mathilda Place, Sunnyvale, California 94086, USA. Unión Europea, Clase B Este dispositivo Broadcom está clasificado para ser utilizado en un entorno doméstico convencional de Clase B. Unión Europea, Clase A ADVERTENCIA: éste es un producto de Clase A. En un entorno doméstico, este producto puede causar interferencia de radio frecuencia, en cuyo caso el usuario debe tomar las medidas oportunas.</p>
<p>SVENSK Swedish</p>	<p>Denna produkt överensstämmer med EU-direktivet 2006/95/EC (lågspänningsdirektivet), 2004/108/EC (EMC direktivet), och andra ändringar enligt den Europeiska unionen. En "Försäkran om överensstämmelse" i enlighet med de föregående direktiven och standarderna har framställts och finns registrerad hos Broadcom Corporation, 190 Mathilda Place, Sunnyvale, California 94086, USA. Europeiska unionen, klass B Den här Broadcom-enheten är klassificerad för användning i vanlig klass B-bostadsmiljö. Europeiska unionen, klass A VARNING: Detta är en klass A-produkt. I en bostadsmiljö kan denna produkt orsaka störningar i radiofrekvenser, så att användaren får vidtaga lämpliga åtgärder.</p>
<p>TÜRK Turkish</p>	<p>Bu ürünün 2006/95/EC (Düşük Voltaj Direktifi), 2004/108/EC (EMC Direktifi), ve Avrupa Birliği'nin ilavelerine uygun olduğu belirlenmiştir. Yukarıda belirtilen direktifler ve standartlara uygun olarak, bir "Uygunluk Beyanı" hazırlanmıştır, ve Broadcom Corporation, 190 Mathilda Place, Sunnyvale, California 94086, USA. Avrupa Birliği, B Sınıfı Bu Broadcom cihazı, tipik bir B sınıfı, ev içi ortamda kullanılmak üzere sınıflandırılmıştır. Avrupa Birliği, A Sınıfı UYARI: Bu bir A sınıfı üründür. Bu ürün, mahalli bir çevrede (yada ev içinde) radyo frekans kargaşasına sebep olabilir, bu durumda ise kullanıcının gerekli önlemleri alması zorunlu olabilir.</p>

CANADIAN REGULATORY INFORMATION (CANADA ONLY)

INDUSTRY CANADA, CLASS B

Broadcom NetXtreme II Gigabit Ethernet Controller

Broadcom Corporation
190 Mathilda Place
Sunnyvale, California 94086 USA

This Class B digital apparatus complies with Canadian ICES-003.

Notice: The Industry Canada regulations provide that changes or modifications not expressly approved by Broadcom could void your authority to operate this equipment.

INDUSTRY CANADA, CLASS A

Broadcom NetXtreme II Gigabit Ethernet Controller

Broadcom NetXtreme II 10 Gigabit Ethernet Controller

Broadcom Corporation
190 Mathilda Place
Sunnyvale, California 94086 USA

This Class A digital apparatus complies with Canadian ICES-003.

Notice: The Industry Canada regulations provide that changes or modifications not expressly approved by Broadcom could void your authority to operate this equipment.

INDUSTRY CANADA, CLASSE B

Broadcom NetXtreme II Gigabit Ethernet Controller

Broadcom Corporation
190 Mathilda Place
Sunnyvale, California 94086 USA

Cet appareil numérique de la classe B est conforme à la norme canadienne ICES-003.

Avis : Dans le cadre des réglementations d'Industry Canada, vos droits d'utilisation de cet équipement peuvent être annulés si des changements ou modifications non expressément approuvés par Broadcom y sont apportés.

INDUSTRY CANADA, CLASSE A

Broadcom NetXtreme II Gigabit Ethernet Controller

Broadcom NetXtreme II 10 Gigabit Ethernet Controller

Broadcom Corporation
190 Mathilda Place
Sunnyvale, California 94086 USA

Cet appareil numérique de classe A est conforme à la norme canadienne ICES-003.

Avis : Dans le cadre des réglementations d'Industry Canada, vos droits d'utilisation de cet équipement peuvent être annulés si des changements ou modifications non expressément approuvés par Broadcom y sont apportés.

KOREA COMMUNICATIONS COMMISSION (KCC) NOTICE (REPUBLIC OF KOREA ONLY)

B CLASS DEVICE

B급 기기 (가정용 방송통신기기)	이 기기는 가정용(B급)으로 전자파적합등록을 한 기기로서 주로 가정에서 사용하는 것을 목적으로 하며, 모든 지역에서 사용할 수 있습니다.
-----------------------	--

Broadcom NetXtreme II Gigabit Ethernet Controller

Broadcom Corporation
190 Mathilda Place
Sunnyvale, California 94086 USA

Note that this device has been approved for non-business purposes and may be used in any environment, including residential areas.

A CLASS DEVICE

A급 기기 (업무용 방송통신기기)	이 기기는 업무용(A급)으로 전자파적합등록을 한 기기이오니 판매자 또는 사용자는 이 점을 주의하시기 바라며, 가정 외의 지역에서 사용하는 것을 목적으로 합니다.
-----------------------	---

Broadcom NetXtreme II Gigabit Ethernet Controller

Broadcom NetXtreme II 10 Gigabit Ethernet Controller

Broadcom Corporation
190 Mathilda Place
Sunnyvale, California 94086 USA

User Diagnostics: Broadcom NetXtreme II™ Network Adapter User Guide

- [Introduction](#)
- [System Requirements](#)
- [Performing Diagnostics](#)
- [Diagnostic Test Descriptions](#)

INTRODUCTION

Broadcom NetXtreme II User Diagnostics is an MS-DOS based application that runs a series of diagnostic tests (see [Table 3](#)) on the Broadcom NetXtreme II 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.

SYSTEM REQUIREMENTS

Operating System: MS-DOS 6.22

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

PERFORMING DIAGNOSTICS

At the MS-DOS prompt, type `uxdiag` (for BCM5706/BCM5708/BCM5709 network adapters) or `uediag` (for BCM57710 network adapters) followed by the command options. The `uxdiag` command options are shown in [Table 1](#) and the `uediag` command options are shown in [Table 2](#). 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 1: uxdiag Command Options

Command Options	Description
<code>uxdiag</code>	Performs all tests on all Broadcom NetXtreme II adapters in your system.
<code>uxdiag -c <devnum></code>	Specifies the adapter (devnum) to test. Use all in place of a specific device number to test all adapters.
<code>uxdiag -cof</code>	Allows tests to continue after detecting a failure.
<code>uxdiag -F</code>	Forces an upgrade of the image without checking the version.
<code>uxdiag -fbc <bc_image></code>	Specifies the bin file to update the bootcode.
<code>uxdiag -fib <ib_image></code>	Specifies the bin file for iSCSI boot.
<code>uxdiag -fbc</code>	Programs the iSCSI configuration block. Used only with <code>-fib <ib_image></code> .
<code>uxdiag -fibp</code>	Programs the iSCSI configuration software. Used only with <code>-fib <ib_image></code> .
<code>uxdiag -fipmi <ipmi_image></code>	Specifies the bin file to update IPMI firmware.
<code>uxdiag -fmba <mba_image></code>	Specifies the bin file to update the MBA.
<code>uxdiag -fncsi <ncsi_image></code>	Specifies the bin file to update the NCSI firmware.
<code>uxdiag -fnvm <raw_image></code>	Programs the raw image into NVM.
<code>uxdiag -fump <ump_image></code>	Specifies the bin file to update UMP firmware.
<code>uxdiag -help</code>	Displays the Broadcom NetXtreme II User Diagnostics (uxdiag) command options.
<code>uxdiag -l <iteration num></code>	Specifies the number of iterations to run on the selected tests.
<code>uxdiag -idmatch</code>	Enables matching of VID, DID, SVID, and SSID from the image file with device IDs. Used only with <code>-fnvm <raw_image></code> .
<code>uxdiag -log <file></code>	Logs the test results to a specified log file.
<code>uxdiag -mba <1/0></code>	Enables/disables Multiple Boot Agent (MBA) protocol. 1 = Enable 0 = Disable
<code>uxdiag -mbap <n></code>	Sets the MBA boot protocol. 0 = PXE 1 = RPL 2 = BOOTP 3 = iSCSI_Boot



Table 1: uxdiag Command Options (Cont.)

Command Options	Description
uxdiag -mbas <n>	Sets the MBA/PXE speed. 0 = Auto 1 = 10H 2 = 10F 3 = 100H 4 = 100F 6 = 1000F
uxdiag -mbav <1 0>	Enables/disables MBA VLAN. 1 = Enable 0 = Disable
uxdiag -mbavval <n>	Sets MBA VLAN (<65536).
uxdiag -mfw <1/0>	Enables/disables management firmware. 1 = Enable 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.
uxdiag -wol <1/0>	Enables/disables Magic Packet WOL. 1 = Enable 0 = Disable

Table 2: uediag Command Options

Command Options	Description
uediag	Performs all tests on all Broadcom NetXtreme II 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 -l <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. 1 = Enable 0 = Disable
uediag -mbap <n>	Sets the MBA boot protocol. 0 = PXE 1 = RPL 2 = BOOTP 3 = iSCSI_Boot
uediag -mbav <1/0>	Enables/disables MBA VLAN. 1 = Enable 0 = Disable
uediag -mbavval <n>	Sets MBA VLAN (<65536).
uediag -mfw <1/0>	Enables/disables management firmware. 1 = Enable 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.

Table 2: uediag Command Options (Cont.)

Command Options	Description
uediag -wol <1/0>	Enables/disable Magic Packet WOL. 1 = Enable 0 = Disable

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 3](#).

Table 3: Diagnostic Tests

Test		Description
Number	Name	
Group A: Basic Functional Tests		
A1	Register	Verifies that registers accessible through the PCI/PCI-E 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 Broadcom NetXtreme II 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		



Table 3: Diagnostic Tests (Cont.)

Test		Description
Number	Name	
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 events occur, 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.



Table 3: Diagnostic Tests (Cont.)

Test		Description
Number	Name	
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.

Troubleshooting: Broadcom NetXtreme II™ Network Adapter User Guide

- [Hardware Diagnostics](#)
- [Checking Port LEDs](#)
- [Troubleshooting Checklist](#)
- [Checking if Current Drivers are Loaded](#)
- [Running a Cable Length Test](#)
- [Testing Network Connectivity](#)
- [Microsoft Virtualization with Hyper-V](#)
- [Removing the Broadcom NetXtreme II Device Drivers](#)
- [Upgrading Windows 2000 Server to Windows Server 2003](#)
- [Upgrading Windows Operating Systems](#)
- [Preparing an Answer File for Unattended Installation](#)
- [Broadcom Boot Agent](#)
- [Broadcom Advanced Server Program \(BASP\)](#)
- [Linux](#)
- [Miscellaneous](#)

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](#)).

CHECKING PORT LEDs

See [Network Link and Activity Indication](#) to check the state of the network link and activity.

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 Broadcom NetXtreme II 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.

CHECKING IF CURRENT DRIVERS ARE LOADED

WINDOWS

See [Viewing Vital Signs](#) to view vital information about the adapter, link status, and network connectivity.

NETWARE

To verify that the driver is loaded properly, type

```
LOAD BX2.LAN FRAME_ETHERNET_II NAME=BX2_1_EII
```

This command automatically verifies if the link is active. If the link is active, the command returns Link is up.

From the command line, type `config` then press **ENTER**. The following status information is displayed:

```
Broadcom NetXtreme II Gigabit Ethernet Adapter Version:  
Hardware Setting:  
Node Address:  
Frame Type:  
Board Name:
```

LAN Protocol: ARP (see note)

LAN Protocol: IP Addr: (see note)



NOTE: The LAN protocol status is displayed after an IP address is assigned to the adapter.

LINUX

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

```
lsmod
```

If the driver is loaded, a line similar to the one below is displayed, where *size* is the size of the driver in bytes, and *n* is the number of adapters configured.

<i>Module</i>	<i>Size</i>	<i>Used by</i>
BCM5706	<i>size</i>	<i>n</i>
BCM5708	<i>size</i>	<i>n</i>

RUNNING A CABLE LENGTH TEST

For Windows operating systems, see [Analyzing Cables](#) for information on running a cable length test. Cable analysis is not available for NetXtreme II 10 GbE network adapters.

TESTING NETWORK CONNECTIVITY



NOTE: When using forced link speeds, verify that both the adapter and the switch are forced to the same speed, or that at least one link partner is configured for auto-negotiation.

WINDOWS

Network connectivity can be tested using the [Testing the Network](#) feature in Broadcom Advanced Control Suite 3.

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.

NETWARE

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.

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.

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 NetXtreme II 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 1 identifies Hyper-V supported features that are configurable for NetXtreme II network adapters. This table is not an all-inclusive list of Hyper-V features.

Table 1: Configurable Network Adapter Hyper-V Features

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

SINGLE NETWORK ADAPTER

Windows Server 2008

When configuring a NetXtreme II 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.
- A TOE-enabled network adapter that is bound to a Hyper-V virtual network will report TOE as an offload capability in BACS; however, TOE will not work. This is a limitation of Hyper-V. Hyper-V does not support TOE.
- 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

When configuring a NetXtreme II 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.

TEAMED NETWORK ADAPTERS

[Table 2](#) identifies Hyper-V supported features that are configurable for NetXtreme II teamed network adapters. This table is not an all-inclusive list of Hyper-V features.

Table 2: Configurable Teamed Network Adapter Hyper-V Features

Feature	Supported in Windows Server 2008	Supported in Windows Server 2008 R2	Comments/Limitation
Smart Load Balancing and Failover (SLB) team type	Yes	Yes	Multi-member SLB team allowed with latest BASP6 version.
Link Aggregation (IEEE 802.3ad LACP) team type	Yes	Yes	–
Generic Trunking (FEC/GEC) 802.3ad Draft Static team type	Yes	Yes	–
Failover	Yes	Yes	–
LiveLink	Yes	Yes	–
Large Send Offload (LSO)	Limited*	Yes	*Conforms to miniport limitations outlines in Table 1 .
Checksum Offload (CO)	Limited*	Yes	*Conforms to miniport limitations outlines in Table 1 .
TCP Offload Engine (TOE)	Limited*	Yes	*Conforms to miniport limitations outlines in Table 1 .
Hyper-V VLAN over an adapter	Yes	Yes	–
Hyper-V VLAN over a teamed adapter	Yes	Yes	–
Hyper-V VLAN over a VLAN	Limited*	Limited*	Only an untagged VLAN.
Hyper-V virtual switch over an adapter	Yes	Yes	–
Hyper-V virtual switch over a teamed adapter	Yes	Yes	–
Hyper-V virtual switch over a VLAN	Yes	Yes	–
iSCSI boot	No	No*	*Remote boot to SAN is supported.

Windows Server 2008

When configuring a team of NetXtreme II 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 TOE-enabled team that is bound to a Hyper-V virtual network will report TOE as an offload capability in BACS; however, TOE will not work. This is a limitation of Hyper-V. Hyper-V does not support TOE.
- 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 NetXtreme II 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. However, the VLAN tagging capability through Hyper-V cannot be enabled with this configuration. Hyper-V VLAN tagging will only work if bound to an untagged BASP virtual adapter.
- 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.

REMOVING THE BROADCOM NETXTREME II DEVICE DRIVERS

Uninstall the Broadcom NetXtreme II 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 Broadcom NetXtreme II device drivers, see [Removing the Device Drivers](#).

UPGRADING WINDOWS 2000 SERVER TO WINDOWS SERVER 2003

Prior to performing an OS upgrade from Windows 2000 Server to Windows Server 2003 when a Broadcom NetXtreme II adapter is installed in your system, Broadcom recommends that you uninstall the adapter's driver. This is to ensure that no errors are encountered in the installation of the Windows Server 2003 drivers for the NetXtreme II adapters. The device drivers for Windows 2000 and Windows Server 2003 are not the same and require the installation of the Windows 2003 drivers to be performed after the OS upgrade is completed.

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

Prior to performing an OS upgrade when a Broadcom NetXtreme II 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 Broadcom adapter drivers and the BACS application.

PREPARING AN ANSWER FILE FOR UNATTENDED INSTALLATION

When creating an answer file for an unattended installation or for the System Preparation Tool (Sysprep) utility, the following sections must be included under the [Unattend] section:

```
OemPreinstall=Yes  
OemPnpDriversPath=Drivers\NIC
```



NOTES:

- This does not include an unattended installation when performed from a CD-ROM.
- The path shown in OemPnpDriversPath can be appended with the path to other applicable drivers.

For an unattended installation, place the Broadcom NetXtreme II driver files in the \$OEM\$\\$1\Drivers\NIC directory. For Sysprep, the drivers are located in Drivers\NIC at the root of the system drive. The Broadcom NetXtreme II driver files are listed below according to the operating system to be installed:

Windows 2000: bxvbd.inf, bxvbdx.sys, bxvbd.cat, bxnd.inf, bxnd.cat, bxnd50x.sys, bxndcox.dll, bxdiag.cat, bxdiag.inf, bxdiagx.sys, wdfcoinstaller01005.dll, and wUDFUpdate_01005.dll

Windows Server 2003 (ia32): bxvbd.inf, bxvbdx.sys, bxvbd.cat, bxnd.inf, bxnd.cat, bxnd52x.sys, bxndcox.dll, bxdiag.cat, bxdiag.inf, bxdiagx.sys, wdfcoinstaller01005.dll, and wUDFUpdate_01005.dll

Windows Server 2003 (x64): bxvbd.inf, bxvbdx.sys, bxvbd.cat, bxnd.inf, bxnd.cat, bxnd52a.sys, bxndcoa.dll, bxdiag.cat, bxdiag.inf, bxdiagx.sys, wdfcoinstaller01005.dll, and wUDFUpdate_01005.dll

When applying network properties through an answer file for a NetXtreme II device where the PnP ID is being used as the identifier, show as follows:

For a BCM5708C NetXtreme II GigE - InfId = "b06bdrv\2nd&pci_164c14e4"

For a BCM5708S Netxtreme II GigE - InfId = "b06bdrv\2nd&pci_16ac14e4"

For a BCM5706C Netxtreme II GigE - InfId = "b06bdrv\2nd&pci_164a14e4"

For a BCM5706S Netxtreme II GigE - InfId = "b06bdrv\2nd&pci_16aa14e4"

Currently, the PCI location (PCI bus, device, function numbers) method is not supported for the NetXtreme II device as an identifier or the adapter in the answer file due to a limitation with the Windows operating system.

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.

BROADCOM ADVANCED SERVER PROGRAM (BASP)

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 NetXtreme II adapter to a team may cause unpredictable results.

Solution: Prior to creating the team, unbind NLB from the NetXtreme II adapter, create the team, and then bind NLB to the team.



Problem: A BASP team in Windows 2003 may not function properly if a team member driver property is modified.

Solution: Due to a limitation in Windows 2003, the features of team members should remain static during the entire life of the team. To change the characteristics of a team member, remove the team member from the team, modify the team member, and then re-add the adapter to the team. The limitation does not exist in Windows 2008 and later.

Problem: A system containing an 802.3ad team causes a Netlogon 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 239924 (<http://support.microsoft.com/kb/239924/>) 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.

Problem: Errors occur when an RSS-enabled adapter is included as a member of a team on Windows Server 2003 and later.

Solution: Disable RSS from all members of the team and rebuild the team.

LINUX

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 NetXtreme II 10 GbE network adapters installed in Linux systems.

Solution: For NetXtreme II 10 GbE 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 NetXtreme II 10 GbE network adapter.

Problem: On a NetXtreme II 1 GbE 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.

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.

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 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: Although installed, the Broadcom Advanced Control Suite 3 (BACS) application does not start and an error message appears.

Solution: .NET Framework 2.0 is required for BACS 3 to operate. Install .NET Framework 2.0.

Problem: Broadcom Advanced Control Suite 3 takes a long time to start up.

Solution: This is a limitation of .NET Framework 2.0. Apply .NET Framework 2.0 SP1 and reboot your system.

Problem: The Broadcom NetXtreme II adapter may not perform at optimal level 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: Although the Broadcom 5708S SerDes adapter is capable of connecting at speeds up to 2.5 Gbps when licensed and configured, Windows Task Manager incorrectly reports the speed at 2 Gbps.

Solution: This reporting error is a known Microsoft issue. Locate the actual link speed from [Viewing Vital Signs](#) in BACS.

Problem: Cannot configure Resource Reservations in BACS 3 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: TOE performance is more susceptible to packet loss when flow control is disabled.

Solution: Enable flow control to reduce the number of packets lost.

Problem: A DCOM error message (event ID 10016) appears in the System Event Log during the installation of the Broadcom adapter 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>.