

Broadcom[®] NetXtreme[®] BCM57XX User Guide

Last revised: April 2017

2CS57XX-CDUM514-R

Broadcom, the pulse logo, Connecting everything, Avago, Avago Technologies, and the A logo are among the trademarks of Broadcom and/or its affiliates in the United States, certain other countries and/or the EU.

Copyright © 2017 by Broadcom. All Rights Reserved.

The term “Broadcom” refers to Broadcom Limited and/or its subsidiaries. For more information, please visit www.broadcom.com.

Broadcom 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 is believed to be accurate and reliable. However, Broadcom 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 nor the rights of others.

Last revised: April 2017

2CS57XX-CDUM514-R

Table of Contents

- Section 1: Installing the Hardware** 10
 - Safety Precautions** 10
 - Preinstallation Checklist** 11
 - Installing the Adapter** 11
 - Connecting the Network Cables** 12
 - Copper 12

- Section 2: Functionality and Features** 13
 - Functional Description** 13
 - Features** 13
 - Power Management 14
 - Adaptive Interrupt Frequency 14
 - Dual DMA Channels 14
 - ASIC with Embedded RISC Processor 15
 - Broadcom Advanced Control Suite 15
 - Supported Operating Environments** 15
 - Network Link and Activity Indication** 15

- Section 3: Windows Driver and Management Application Installation** 17
 - Installing the Driver Software** 18
 - Using Silent Installation 19
 - Modifying the Driver Software** 19
 - Repairing or Reinstalling the Driver Software** 20
 - Removing the Device Drivers** 20
 - Viewing or Changing the Properties of the Adapter** 21
 - Setting Power Management Options** 21
 - Configuring the Communication Protocol To Use With BACS4** 22
 - Using WS-MAN 22
 - WS-MAN Windows Server Configuration 22
 - WS-MAN Windows Client Installation 29
 - Using WMI 31
 - Step 1: Set up Namespace Security Using WMI Control 31

Step 2: Grant DCOM Remote Launch and Activate Permission.....	31
Special Configuration for WMI on Different Systems.....	33
Section 4: Linux Driver and Management Application Installation	34
Packaging	34
Installing TG3 Driver Software.....	35
Installing the Source RPM Package	35
Building the Driver from the Source TAR File	36
Network Installations	36
Unloading/Removing the TG3 Driver.....	36
Unloading/Removing the Driver from an RPM Installation.....	36
Removing the Driver from a TAR Installation.....	36
Driver Messages.....	37
Teaming with Channel Bonding.....	37
Linux Management Application Installation	38
Overview	38
Communication Protocols.....	38
Installing WS-MAN or CIM-XML on Linux Server	39
Step 1: Install OpenPegasus	39
Step 2: Start CIM Server on the Server	41
Step 3: Configure OpenPegasus on the Server	41
Step 4: Install Broadcom CMPI Provider	42
Step 5: Perform Linux Firewall Configuration, If Required.....	43
Step 6: Install BACS and Related Management Applications.....	44
Installing WS-MAN or CIM-XML on Linux Client.....	45
Configure HTTPS on Linux Client.....	45
Installing the Broadcom Advanced Control Suite.....	47
Section 5: VMware Driver Software	48
Packaging	48
Drivers	48
Download, Install, and Update Drivers.....	48
Driver Parameters.....	48
Driver Parameters.....	49
Driver Defaults	49
Driver Messages	49

Section 6: Using Broadcom Advanced Control Suite 4	51
Broadcom Advanced Control Suite Overview	51
Starting Broadcom Advanced Control Suite	52
BACS Interface	52
Explorer View Pane	53
Context View Selector	54
Filter View.....	54
Context View Pane	54
Menu Bar	54
Description Pane	55
Configuring Preferences in Windows	55
Connecting to a Host	56
Managing the Host	57
Information Tab: Host Information	57
Managing the Network Adapter	59
Viewing Adapter Information	59
Viewing Driver Information	60
Viewing Resource Information.....	61
Viewing Hardware Information	62
Testing the Network.....	63
Running Diagnostic Tests	64
Analyzing Cables	65
Setting Adapter Properties	66
Viewing Statistics	68
General Statistics	69
Configuring Teaming	69
Team Types	70
Using the Broadcom Teaming Wizard	70
Using Expert Mode	83
Creating a Team	83
Modifying a Team	86
Adding a VLAN	87
Viewing VLAN Properties and Statistics and Running VLAN Tests	88
Deleting a VLAN	89

Configuring LiveLink for a Smart Load Balancing and Failover and SLB (Auto-Fallback Disable) Team.....	90
Saving and Restoring a Configuration	91
Viewing BASP Statistics	92
Configuring With the Command Line Interface Utility.....	93
Managing VLANs.....	94
Overview	94
Adding VLANs to Teams.....	95
Troubleshooting BACS.....	96
Section 7: Teaming	97
Overview	98
Load Balancing and Fault Tolerance.....	98
Types of Teams	98
Smart Load Balancing™ and Failover	99
Link Aggregation (802.3ad).....	99
Generic Trunking (FEC/GEC)/802.3ad-Draft Static	99
SLB (Auto-Fallback Disable).....	100
Limitations of Smart Load Balancing and Failover/SLB (Auto-Fallback Disable) Types of Teams.....	100
LiveLink™ Functionality	101
Teaming and Large Send Offload/Checksum Offload Support.....	101
Section 8: Broadcom Gigabit Ethernet Teaming Services	102
Introduction	103
Teaming Glossary	103
Teaming Concepts	104
Network Addressing.....	104
Teaming and Network Addresses.....	105
Description of Teaming Types	105
Software Components	109
Hardware Requirements	110
Ethernet Switch.....	110
Router	110
Supported Features by Team Type	110
Selecting a Team Type	112
Teaming Mechanisms	113

Architecture	113
Outbound Traffic Flow	114
Inbound Traffic Flow (SLB Only)	114
Protocol Support.....	115
Performance.....	115
Driver Support by Operating System.....	116
Supported Teaming Speeds.....	117
Teaming and Other Advanced Networking Properties	118
Checksum Offload	119
IEEE 802.1p QoS Tagging	119
Large Send Offload	119
Jumbo Frames.....	119
IEEE 802.1Q VLANs	119
Wake on LAN	120
Preboot Execution Environment (PXE)	120
General Network Considerations	121
Teaming Across Switches	121
Switch-Link Fault Tolerance	121
Spanning Tree Algorithm.....	123
Topology Change Notice (TCN).....	124
Port Fast/Edge Port.....	124
Teaming with Microsoft NLB/WLBS	125
Application Considerations	126
Teaming and Clustering—Microsoft Cluster Software.....	126
Teaming and Network Backup	126
Load Balancing and Failover.....	127
Fault Tolerance	128
Troubleshooting Teaming Problems	130
Teaming Configuration Tips	130
Troubleshooting Guidelines.....	131
Frequently Asked Questions	132
Event Log Messages	135
Windows System Event Log Messages	135
Base Driver (Physical Adapter/Miniport).....	135
Intermediate Driver (Virtual Adapter/Team).....	137

Section 9: iSCSI Protocol and Broadcom Boot Agent Software	139
iSCSI Boot	139
Supported Operating Systems for iSCSI Boot	139
iSCSI Boot Setup	139
Configuring the iSCSI Target.....	139
Configuring iSCSI Boot Parameters	140
MBA Boot Protocol Configuration	141
iSCSI Boot Configuration	141
Enabling CHAP Authentication	145
Configuring the DHCP Server to Support iSCSI Boot.....	145
DHCP iSCSI Boot Configurations for IPv4.....	145
DHCP iSCSI Boot Configuration for IPv6	148
Configuring the DHCP Server	148
Preparing the iSCSI Boot Image.....	149
Booting.....	152
Other iSCSI Boot Considerations	152
Changing the Speed & Duplex Settings in Windows Environments	152
Locally Administered Address	152
Virtual LANs.....	153
Troubleshooting iSCSI Boot.....	153
Broadcom Boot Agent Driver Software	154
Overview	154
Setting Up MBA in a Client Environment	154
Configuring the MBA Driver	154
Setting Up the BIOS	156
Section 10: Manageability	157
CIM	157
SNMP	158
BASP Subagent.....	158
BASP Extensible-Agent	158
Section 11: Specifications	160
10/100/1000BASE-T Cable Specifications	160
Performance Specifications	160

Section 12: Troubleshooting	161
Hardware Diagnostics	161
BACS Diagnostic Tests Failures	161
BACS Network Test Failures	162
Troubleshooting Checklist	163
Checking for Network Link and Activity	163
Checking if Current Drivers Are Loaded	164
Windows	164
Linux	164
Running a Cable Length Test	164
Testing Network Connectivity	165
Windows	165
Linux	165
Broadcom Boot Agent	165
Broadcom Advanced Server Program (BASP)	166
Kernel Debugging over Ethernet	166
Miscellaneous	166
Section 13: Regulatory Information	167
FCC Class B Notice	167
VCCI Class B Notice	168
VCCI Class B Statement (Japan)	168
CE Notice	168
Canadian Regulatory Information (Canada Only)	172
Industry Canada, Class B	172
Industry Canada, classe B	172
MIC Notice (Republic of Korea Only)	173
B CLASS Device	173
BSMI	174

Section 1: Installing the Hardware

- [Safety Precautions](#)
- [Preinstallation Checklist](#)
- [Installing the Adapter](#)
- [Connecting the Network Cables](#)



Note: This section applies only to add-in NIC models of Broadcom NetXtreme Gigabit Ethernet adapters.

Safety Precautions



Caution! The adapter is being installed in a system that operates with voltages that can be lethal. Before you remove the cover of your system, you must 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 non-conducting tools.
- Verify that the system is powered OFF and 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 server is using the latest BIOS.
2. If your system is booted to an operating system, gracefully power down the OS.
3. When system shutdown is complete, turn off the power and unplug the power cord.
4. Holding the adapter card by the edges, remove it from its shipping package and place it on an antistatic surface.
5. Check the adapter for visible signs of damage, particularly on the card edge connector. Never attempt to install any damaged adapter.

Installing the Adapter

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

1. Review the [Safety Precautions](#) and [Preinstallation Checklist](#). Before installing the adapter, ensure the system power is OFF and unplugged from the power outlet, and that proper electrical grounding procedures have been followed.
2. Open the system case, and select any empty PCI Express slot.
3. Remove the blank cover-plate from the slot that you selected.
4. Align the adapter connector edge with the 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

Copper

The Broadcom NetXtreme Gigabit Ethernet adapter has one or more RJ-45 connectors used for attaching the system to an Ethernet copper-wire segment.



Note: The Broadcom NetXtreme Gigabit Ethernet 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: “10/100/1000BASE-T Cable Specifications”](#) lists the cable requirements for connecting to 10/100/1000BASE-T ports:

Table 1: 10/100/1000BASE-T Cable Specifications

Port Type	Connector	Media	Maximum Distance
10BASE-T	RJ-45	Category 3, 4, or 5 UTP	100 meters (328 feet)
100/1000BASE-T ¹	RJ-45	Category 5 ² UTP	100 meters (328 feet)

¹1000BASE-T signaling requires four twisted pairs of Category 5 balanced cabling, as specified in ISO/IEC 11801:1995 and EIA/TIA-568-A (1995) and tested using procedures defined in TIA/EIA TSB95.

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

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



Note: After the cable is properly connected at both ends, the port LEDs on the adapter should be functional. See [Table 1: “10/100/1000BASE-T Cable Specifications,” on page 12](#) for a description of network link and activity indications

Section 2: Functionality and Features

- [Functional Description](#)
- [Features](#)
- [Supported Operating Environments](#)
- [Network Link and Activity Indication](#)

Functional Description

Broadcom NetXtreme Gigabit Ethernet adapters connect a PCI Express™ compliant system to a Gigabit Ethernet network. Broadcom NetXtreme Gigabit Ethernet adapters incorporate a technology that transfers data at a maximum rate of 1 gigabit per second—10 times the rate of Fast Ethernet adapters.

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 [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 Gigabit Ethernet adapter features for all supported operating systems:

- Integrated quad 10/100/1000BASE-T and quad 1000BASE-X/SGMII 1.25 Gbaud SerDes transceivers
- Energy Efficient Ethernet™ compliant with IEEE Std 802.3az-2010
- IEEE 802.3ap Clause 73 auto-negotiation
- Quad 10/100/1000BASE-T full-duplex/half-duplex MACs
- Quad 1000BASE-X/SGMII full-duplex/half-duplex MACs
- Automatic MDI crossover
- x2 PCI Express v2.0 at 5 GT/s or 2.5 GT/s
- MSI and MSI-X capabilities—up to 17 MSIX vectors
- I/O Virtualization support for VMware NetQueue and Microsoft VMQ
 - 17 receive queues and 16 transmit queues
 - 17 MSI-X vectors supporting per queue interrupt to host
- Flexible MSI-X vector to transmit/receive queue association
- Function Level Reset
- Receive Side Scaling (RSS) with per queue MSI-X vector support and support for UDP RSS hash type
- Transmit Side Scaling (TSS) and multi-Tx queue with per queue MSI-X vector support
- Jumbo frame support for up to 9000-byte payload

- Virtual LAN (VLAN) support— IEEE 802.1q VLAN tagging
- TCP, IP, UDP checksum offload
- Large Send Offload (LSO), TCP Segmentation Offload (TSO)
- Hardware assist for IEEE 1588 and IEEE 802.1AS time synchronization implementations
- IEEE 802.3x flow control
- SMBus 2.0 Interface
- Statistics for SNMP MIB II, Ethernet-like MIB and Ethernet MIB (IEEE 802.3z, Clause 30)
- ACPI power management compliance
- Advanced power management by a Central Power Management Unit (CPMU)
- Efficient integrated switching regulator controller
- On-chip temperature monitor
- PCI Express CLKREQ support
- Power Management Offload (PM Offload)
- Serial flash and EEPROM NVRAM support; flash auto-configure
- ECC error detection and correction on internal SRAM
- JTAG boundary scan support

Power Management

Wake on LAN (Magic Packet, Wake Up Frame, specific pattern) is supported.



Note: Adapter speed connection when the system is down waiting for a wake-up signal is either 10 Mbps or 100 Mbps, but can return to 1000 Mbps 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.

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.

Dual DMA Channels

The PCIe interface on Broadcom NetXtreme Gigabit Ethernet adapters contains two independent DMA channels for simultaneous read and write operations.

ASIC with Embedded RISC Processor

The core control for Broadcom NetXtreme Gigabit Ethernet 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.

Broadcom NetXtreme manageability operations such as DMTF, SMASH, DASH, and NC-SI pass-through run on a high-performance application processor engine (APE), which is a separate from the traditional network processing engine.

Broadcom Advanced Control Suite

Broadcom Advanced Control Suite (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 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. BACS is used on Windows operating systems to configure teaming and to add VLANs. See [Using Broadcom Advanced Control Suite](#) for detailed information and instructions.

Supported Operating Environments

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

- Microsoft® Windows® (32-bit and 64-bit extended) Server 2008 or later
- Linux® (64-bit extended) RHEL 6.9, 7.1 or later; SLES 11 SP4, SLES 12 SP4 or later
- VMware ESXi 6.0 or later
- Oracle Solaris

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 2: “Network Link and Activity Indicated by RJ-45 Port LEDs,” on page 15](#). Broadcom Advanced Control Suite also provides information about the status of the network link and activity (see [Viewing Adapter Information](#)).

Table 2: Network Link and Activity Indicated by RJ-45 Port LEDs

Port LED	LED Appearance	Network State
Link LED	Off	No link (cable disconnected)
	Continuously illuminated	Link (Green—1 Gb, Amber —10 or 100 Mb)
Activity LED	Off	No network activity
	Blinking	Network activity

Section 3: Windows Driver and Management Application Installation

- [Installing the Driver Software](#)
- [Modifying the Driver Software](#)
- [Repairing or Reinstalling the Driver Software](#)
- [Removing the Device Drivers](#)
- [Viewing or Changing the Properties of the Adapter](#)
- [Setting Power Management Options](#)
- [Configuring the Communication Protocol To Use With BACS4](#)

Installing the Driver Software



Note: These instructions are based on the assumption that your Broadcom NetXtreme 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 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.



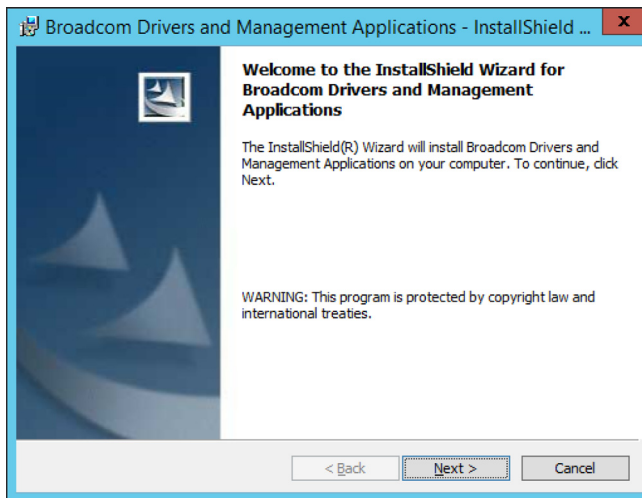
Notes:

- 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 installed before the Broadcom NetXtreme Gigabit Ethernet adapter can be used with your Windows operating system. Drivers are located on the installation CD.
- BACS is not supported on the Server Core installation option for Microsoft Windows Server 2008 R2.

Download the driver installer from the following link:

<http://www.dell.com/support>

Use the self-extracting zip file.



Using Silent Installation

Refer to the readme.txt file in the installation folder for command line instructions.

**Notes:**

- All commands are case sensitive.
- For detailed instructions and information about unattended installs, refer to the Silent.txt file in the Driver_Management_Apps_Installer folder.



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.

Modifying the Driver Software

To modify the driver software

1. In Control Panel, double-click **Add or Remove Programs**.
2. Click **Broadcom Drivers and Management Applications**, and then click **Change**.
3. Click **Next** to continue.
4. Click **Modify, Add, or Remove** to change program features. This option does not install drivers for new adapters. For information on installing drivers for new adapters, see [Repairing or Reinstalling the Driver Software](#).
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. The installer will determine if a system restart is necessary. Follow the on-screen instructions.

Repairing or Reinstalling the Driver Software

To repair or reinstall the driver software

1. In Control Panel, double-click **Add or Remove Programs**.
2. Click **Broadcom Drivers and Management Applications**, and then click **Change**.
3. Click **Next** to continue.
4. Click **Repair or Reinstall** to repair errors or install drivers for new adapters.
5. Click **Next** to continue.
6. Click **Install**.
7. Click **Finish** to close the wizard.
8. The installer will determine if a system restart is necessary. Follow the on-screen instructions.

Removing the Device Drivers

When removing the device drivers, any management application that is installed is also removed.



Note: Windows Server 2008 and later 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 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 Drivers and Management Applications**, and then click **Remove**. Follow the on-screen prompts.
3. Reboot your system to completely remove the drivers. If you fail to reboot your system, you will not be able to successfully install the drivers.

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 4**.
2. Click the Advanced section of the **Configurations** tab.

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.

Configuring the Communication Protocol To Use With BACS4

There are two main components of the BACS4 management application: the provider component and the client software. A provider is installed on a server, or “managed host”, that contains one or more NICs. The provider collects information on the NICs and makes it available for retrieval from a management PC on which the client software is installed. The client software enables viewing information from the providers and configuring the NICs. The BACS client software includes a graphical user interface (GUI) and a command line interface (CLI).

A communication protocol enables communication between the provider and the client software. Depending on the mix of operating systems (Linux, Windows, or both) on the clients and managed hosts in your network, you can choose an appropriate communication protocol to use. See [Linux Management Application Installation](#) for a description of the available communication protocols for each network configuration.

The instructions in this chapter address only the scenario where Windows managed hosts are communicating with Windows clients. In these scenarios, you can use either the WMI or the WS-MAN (WinRM) communication protocols. When you use the driver installer described in this chapter to install both the driver and the management applications, the provider for both WMI and WS-MAN is installed on the managed host. Additionally, the BACS4 utility is installed on the client. The following sections provide additional configuration steps for the communication protocol you select.

For Linux installations, the driver is installed separately from the management applications. See [Linux Driver Software and Management Application Installation](#) for related instructions.

Using WS-MAN

To use the WS-MAN communication protocol, follow the instructions in the following sections:

- [WS-MAN Windows Server Configuration](#)
- [WS-MAN Windows Client Installation](#)

WS-MAN Windows Server Configuration

Step 1: Install the WinRM Software Component on Server

On the following operating systems, WinRM 2.0 is preinstalled:

- Windows 7
- Windows 8
- Windows 8.1
- Windows Server 2008 R2
- Windows Server 2012
- Windows 2012 R2
- Windows Server 2016

For Windows Server 2008, install Windows Management Framework Core, which includes WinRM 2.0 and Windows Powershell 2.0, from the following link:

<http://www.microsoft.com/download/en/details.aspx?displaylang=en&id=11829>

Step 2: Perform Basic Configuration on the Server

The Windows firewall must be enabled for WinRM to work properly. For detailed information about firewall configuration, see [Step 7: Additional Server Configuration](#). After the firewall is configured, open a command prompt and run the following command to enable the remote management on the Windows server:

```
winrm quickconfig
```

You can use the following command to view the configuration information for the service:

```
winrm get winrm/config
```

Step 3: Perform User Configuration on the Server

To connect to WinRM, the account must be a member of the local administrators group on the local or remote computer. The output of the `get winrm/config` command will be as follows:

```
RootSDDL = O:NSG:BAD:P(A;;GA;;;BA)S:P(AU;FA;GA;;;WD)(AU;SA;GWGX;;;WD)
```

BA stands for BUILTIN\Administrators.

To add another user group to the WinRM allowed connect list, you can modify the RootSDDL to include the new user group. You will need the SDDL ID for the new group. For example, the following command adds the new user group with SDDL ID S-1-5-21-1866529496-2433358402-1775838904-1021.

```
winrm set winrm/config/Service @{RootSDDL="O:NSG:BAD:P(A;;GA;;;BA)(A;;GA;;;S-1-5-21-1866529496-2433358402-1775838904-1021)S:P(AU;FA;GA;;;WD)(AU;SA;GWGX;;;WD)"} 
```

Step 4: Perform HTTP Configuration on the Server

To use the BACS GUI, you must configure the HTTP protocol, as follows:



Note: The default HTTP port is 5985 for WinRM 2.0.

1. Click **Start** (or press the Windows logo key) and select **Run**.
2. Enter **gpedit.msc** to open the local Group Policy editor.
3. Under **Computer Configuration**, open the **Administrative Templates** folder and then open the **Windows Components** folder.
4. Select **Windows Remote Management (WinRM)**.
5. Under **Windows Remote Management (WinRM)**, select **WinRm Client**.
6. Under **WinRM Client**, double-click **Trusted Hosts**.
7. In the **TrustedHostsList**, enter the host names of the clients. If all clients are trusted then enter an asterisk (*) only.
8. Select **WinRM Service**.
9. Enable **Allow Basic Authentication**.

10. Enable **Allow unencrypted traffic**.
11. Close the **Group Policy** window.
12. From the command prompt, run the following command to configure WinRM with default settings:
`winrm qc` or `winrm quickconfig`
13. When the tool displays “**Make these changes[y/n]?**”, enter “**y**”.
14. Enter one of the following commands to check whether an HTTP listener is created:
`winrm enumerate winrm/config/listener`
or
`winrm e winrm/config/Listener`
15. Enter the following command from the command prompt to test locally.
`winrm id`

Step 5: Perform HTTPS Configuration on the Server (to use HTTPS rather than HTTP)

This step consists of two distinct processes: generating a self-signed certificate, if certificate does not exist, and importing it to a Windows server. If one does not already exist, you must configure a self-signed certificate on the Windows server to enable HTTPS/SSL communication with the BACS GUI on the Windows client. The Windows client also must be configured with the self-signed certificate. See [Perform HTTPS Configuration \(if you plan to use HTTPS\)](#).



Note: The self-signed certificate can be created on any Windows server. The server does not require BACS to be installed. The self-signed certificate generated on any Windows server should be copied on the local drive of client.

1. Click **Start** (or press the Windows logo key) and select **Run**.
2. Enter `gpedit.msc` to open the local Group Policy editor.
3. Under **Computer Configuration**, open the **Administrative Templates** folder and then open the **Windows Components** folder.
4. Select **Windows Remote Management (WinRM)**.
5. Under **Windows Remote Management (WinRM)**, select **WinRm Client**.
6. Under **WinRM Client**, double-click **Trusted Hosts**.
7. In the **TrustedHostsList**, enter the host names of the clients. If all clients are trusted then enter an asterisk (*) only.
8. Select **WinRM Service**.
9. Enable **Allow Basic Authentication**.

To generate a self-signed certificate for the Windows Server:

Openssl on Windows can be used to generate the self-signed certificate, as follows:

1. Enter the following command to generate a private key:
`openssl genrsa -des3 -out server.key 1024`
2. You are prompted to enter a passphrase. Be sure to remember the passphrase.

3. Use the following steps to generate a Certificate Signing Request (CSR).

During the generation of the CSR, you are prompted for several pieces of information. When prompted for the “Common Name”, enter the Windows Server host name or IP address.

Enter the following command (sample responses are shown):

```
openssl req -new -key server.key -out server.csr
```

If this command does not work, try the following:

```
openssl req -new -key server.key -out server.csr -config openssl.cnf
```

The openssl.cnf file should be placed in the same directory where openssl is placed. Openssl.cnf is located in the folder C:\Program Files (x86)\GnuWin32\share.

The following information is requested:

- Country Name (2 letter code) []: **US**
- State or Province Name (full name) []: **California**
- Locality Name (e.g., city) []: **Irvine**
- Organization Name (e.g., company) []: **Broadcom Corporation**
- Organizational Unit Name (e.g., section) []: **Engineering**
- Common Name (e.g., YOUR name) []: Enter the host name or IP address of the Windows server. For IPv6, enter the Common Name in the format [xyxy:xxx:.....:xxx], **including the brackets []**.
- (Optional) Email Address []:

Enter the following additional attributes to be sent with your certificate request:

- A challenge password []: **password1**
- An optional company name []:

4. Remove the passphrase from the key.

Enter the following commands:

```
cp server.key server.key.org  
openssl rsa -in server.key.org -out server.key
```

5. Generate a self-signed certificate:

To generate a self-signed certificate which is active for 365 days, enter the following command:

```
openssl x509 -req -days 365 -in server.csr -signkey server.key -out server.crt
```

The following output displays:

```
Signature ok  
subject=/C=US/ST=California/L=Irvine/O=Broadcom Corporation/OU=Engineering/CN=MGMTAPP-  
LAB3/emailAddress=  
Getting Private key
```

6. Enter the following command to verify the generated self-signed certificate.

```
openssl verify server.crt
```

The following output displays:

```
server.crt:/C=US/ST=California/L=Irvine/O=Broadcom Corporation/OU=Engineering/  
CN=MGMTAPP- LAB3/emailAddress=  
error 18 at 0 depth lookup:self signed certificate  
OK
```

Ignore the error message “error 18 at 0 depth lookup:self signed certificate”. This error indicates that this is a self-signed certificate.

7. Convert the certificate from “crt” to “pkcs12” format, as follows:

For a Windows server, the certificate should be in pkcs12 format. Enter the following command:

```
openssl pkcs12 -export -in server.crt -inkey server.key -out hostname.pfx
```

You will be prompted for the following:

Enter Export Password:

Verifying - Enter Export Password:

Enter the password and be sure to remember it. The password is required when importing the certificate on the Windows server and client.

8. Make a copy of the certificate file `server.crt` and place it on the server where BACS will be installed, so that it can be imported. If you plan to use a Windows client to connect to the server running BACS, then the certificate also needs to be transferred (copied and pasted) to the client system.

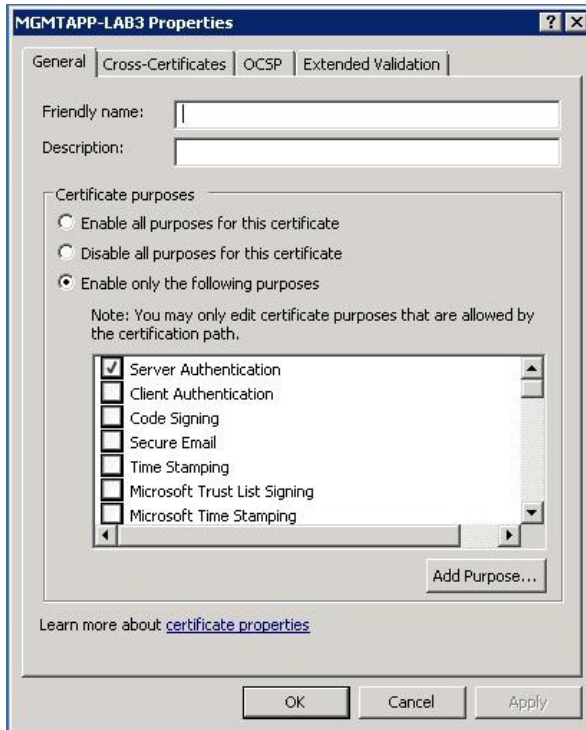


Note: A separate certificate must be generated for an IPv4 address, IPv6 address, and Hostname.

To install the self-signed certificate on Windows server:

Transfer the file `hostname.pfx` you generated on the Windows server before you install the certificate:

1. Click **Start** (or press the Windows logo key) and select **Run**.
2. Enter **MMC** and click **OK**.
3. Click **File > Add/Remove Snap-in**.
4. Click **Add**.
5. Select **Certificates** and click **Add**.
6. Select **Computer account**.
7. Click **Next** and then click **Finish**.
8. Click **Close**, then click **OK**.
9. Open the **Certificates (Local Computer)** folder and then open the **Personal** folder.
10. Right-click **Certificates**, select **All Tasks** and then click **Import**.
11. Click **Next** to begin the Certificate Import Wizard.
12. Browse to select **hostname.pfx**.
13. When you are prompted for the password for the private key, enter the same password you created in [To generate a self-signed certificate for the Windows Server](#).
14. Follow the instructions, select the defaults, and continue.
The certificate is shown as installed on the right side of the window. The name will be the name you specified while creating a self-signed certificate.
15. Right-click on the certificate and select **Properties**.
A dialog box displays, as follows:



16. Ensure that only **Server Authentication** is enabled, as shown in the figure.

17. Open **Trusted Root Certification Authorities** and then open **Certificates**.

18. Follow the instructions from [Step 11 on page 26](#) to [Step 17 on page 27](#)



Note: See [Perform HTTPS Configuration \(if you plan to use HTTPS\)](#) for instructions on importing the self-signed certificate on a client.

Step 6: Configure WinRM HTTPS/SSL on the Server

1. Create WinRM Listener, as follows:

- a. Click **Start** (or press the Windows logo key) and select **Run**.
- b. Enter **MMC** and click **OK**.
- c. Select the self-signed certificate from the Personal store.
For example, if the certificate is created with a host name, the host name will appear.
- d. Double-click the certificate to open it.
- e. Click the **Details** tab.
- f. Scroll down and select the **Thumbprint** field.
- g. Select and copy the thumbprint in the **Details** window so you can insert it in the next step.
- h. Return to the command prompt.
- i. Enter the following command:

```
winrm create winrm/config/Listener?Address=*&Transport=
HTTPS @{Hostname="<HostName or IPAddress>";
CertificateThumbprint="<paste from the previous step and remove the spaces>"}
```

**Notes:**

- If the certificate was generated using the host name, enter the host name. If it was generated using the IP address, enter the IP address. For an IPv6 address, use brackets [] around the address.
- If HTTPS is configured in your system, the listener must be deleted before creating a new HTTPS listener. Use the following command:

```
winrm delete winrm/config/Listener?Address=*&Transport=HTTPS
```

- a. The above command creates a listener on the HTTPS port (5986) using any/all network address of the server, and my SelfSSL generated certificate.
- b. You can use the `winrm` command to modify or set the HTTPS listener, as WinRM listeners can be configured on any user defined port.
- c. From command prompt, run the following command to verify that the listener(s) that have been configured:

```
winrm e winrm/config/listener
```

1. Test HTTPS/SSL connection on the server.

- a. At the command prompt on the server, enter the following command:

```
winrs -r:https://yourserver:5986 -u:username -p:password hostname
```

- b. If set up correctly, the output of the command shows the server host name.
- c. To check WinRM Service Configuration, run the following command:

```
winrm get winrm/config/service
```

Step 7: Additional Server Configuration

If necessary, modify the firewall rules as follows:

Windows Server 2008 R2

1. From the **Administrative Tools** menu, open **Windows Firewall with Advanced Security**.
2. Right-click **Inbound Rules** and select **New Rule**.
The new rule wizard opens.
3. Select **Port** and click **Next**.
4. On the **Protocol and Ports** screen, select **TCP** and enter the specific port, for example, 5985 for HTTP or 5986 for HTTPS.
5. Click **Next**.
6. On the **Action** screen, select **Allow the connection** and click **Next**.
7. For **Profile**, you can select all three profiles if your server is in a workgroup.
8. Specify a name for the rule and click **Finish**.
9. Ensure that the new rule and is enabled (the green check box is selected).

Windows XP

1. Click **Start > Control Panel**, and then double-click **Windows Firewall**.
2. Click the **Exceptions** tab

3. Click **Add Port**.
4. Enter a meaningful **Name**, for example “WinRM rule” and port number, for example, 5985 for HTTP or 5986 for HTTPS.
5. Click **OK**.

Useful WinRM Commands

Command	Description
<code>winrm quickconfig</code> or <code>winrm qc</code>	Configures WinRM with default settings
<code>winrm enumerate winrm/config/Listener</code> or <code>winrm e winrm/config/Listener</code>	Helps to check which service listener are enabled and listening on which port and IP Address.
<code>winrm get winrm/config/Service</code>	Checks WinRM Service Configuration.
<code>winrm delete winrm/config/Listener?Address=*&Transport=HTTPS</code>	Deletes a Listener (in this case deleting a HTTPS listener).

Useful WinRM Websites

- <http://msdn.microsoft.com/en-us/library/aa384372%28v=vs.85%29.aspx>
- <http://technet.microsoft.com/en-us/library/cc782312%28WS.10%29.aspx>
- <http://msdn.microsoft.com/en-us/library/aa384295%28v=VS.85%29.aspx>
- The following articles on “<http://support.microsoft.com>”:
 - “Configuring WINRM for HTTPS”
 - “Windows Management Framework (Windows PowerShell 2.0, WinRM 2.0, and BITS 4.0)”

WS-MAN Windows Client Installation

On the Windows client, perform following configuration steps.

1. Perform HTTP Configuration (if you plan to use HTTP)
 - a. Click **Start** (or press the Windows logo key) and select **Run**.
 - b. Enter `gpedit.msc` to open the local Group Policy editor.
 - c. Under **Computer Configuration**, open the **Administrative Templates** folder and then open the **Windows Components** folder.
 - d. Select **Windows Remote Management (WinRM)**.
 - e. Under **Windows Remote Management (WinRM)**, select **WinRm Client**.
 - f. Under **WinRM Client**, double-click **Trusted Hosts**.
 - g. In the **TrustedHostsList**, enter the host names of the clients and click **OK**. If all clients are trusted then enter “*” only.
 - h. Select **WinRM Service**.
 - i. Enable **Allow Basic Authentication** and click **OK**.
 - j. Run the following command from the command prompt to test the connection:


```
winrm id -remote:<remote machine Hostname or IP Address>
```
2. Perform HTTPS Configuration (if you plan to use HTTPS)

After you generate a self-signed certificate, as described in [To generate a self-signed certificate for the Windows Server](#), you can import the certificate on the client to facilitate a connection between server and client. Ensure that all steps mentioned in section [To generate a self-signed certificate for the Windows Server](#) are completed, including copying *hostname.pfx* at the location from where client can access it, before you proceed with the following steps.

- a. Click **Start** (or press the Windows logo key) and select **Run**.
- b. Enter **MMC** and click **OK**.
- c. Click **File** and select **Add/Remove Snap-in**.
- d. Click **Add**.
- e. Select **Certificates** and click **Add**.
- f. Select **Computer account** and click **Next**.
- g. Click **Finish**.
- h. Click **Close** and then click **OK**.
- i. Under **Certificates (Local Computer)**, right-click on **Trusted Root Certification Authorities**, select **All Tasks**, and select **Import**.
- j. Click **Next** to begin the Certificate Import Wizard.
- k. Browse to select the .pfx file you generated in [To generate a self-signed certificate for the Windows Server](#). Change the selection in the **Files of type** list to **Personal Information Exchange (*.pfxas, *.p12)**, select the *hostname.pfx* file and click **Open**.
- l. Enter the password you assigned to the private key and click **Next**.

3. Configure WinRM HTTPS/SSL

You can run `winrm` from a client to retrieve information from the server using WinRM HTTPS connection. Use the following steps to test the WinRM HTTPS/SSL connection from client:

- a. To retrieve the server operating system information, enter the following command.

```
winrm e wmi/root/cimv2/Win32_OperatingSystem -r:https://yourservername -u:username -p:password -skipCAcheck
```
- b. To retrieve the server WinRM identity information, enter the following command.

```
winrm id -r:https://yourservername -u:username -p:password -skipCAcheck
```
- c. To enumerate Windows services on the server, enter the following command.

```
winrm e wmicimv2/Win32_service -r:https://yourservername -u:username -p:password -skipCAcheck
```



Note: It is important to use `-skipCAcheck` switch in the `winrm` command line testing, as the certificate is self-generated and not imported on the client. Otherwise, the following error message displays:
WSManFault.

Using WMI

No special configuration is required to use WMI on the Windows client. Perform the steps in the following sections to configure WMI on the Windows server.

Step 1: Set up Namespace Security Using WMI Control

The WMI Control provides one way to manage namespace security. You can start the WMI Control from the command prompt using this command:

```
wiimgmt
```

On Windows 9x or Windows NT4 computers that have WMI installed, use this command instead:

```
wbemctl.exe
```

Alternatively, you can access the WMI Control and the Security tab as follows:

1. Right-click on **My Computer** and click **Manage**.
2. Double-click **Services and Applications** and then double-click **WMI Control**.
3. Right-click **WMI Control** and then click **Properties**.
4. In WMI Control Properties, click the **Security** tab.
5. A folder named Root with a plus sign (+) next to it should now be visible. Expand this tree as necessary to locate the namespace for which you want to set permissions.
6. Click **Security**.

A list of users and their permissions appears. If the user is on the list, modify the permissions as appropriate. If the user is not on the list, click **Add** and add the user from the location (local machine, domain, etc.) where the account resides.



Note: You can add these exports at the end of the `.bash_profile`. This file is located in the `/root` directory.

- In order to view and set namespace security, the user must have Read Security and Edit Security permissions. Administrators have these permissions by default, and can assign the permissions to other user accounts as required.
- If this user needs to access the namespace remotely, you must select the Remote Enable permission.
- By default, user permissions set on a namespace apply only to that namespace. If you want the user to have access to a namespace and all subnamespaces in the tree below it, or in subnamespaces only, click **Advanced**. Click **Edit** and specify the scope of access in the dialog box that displays.

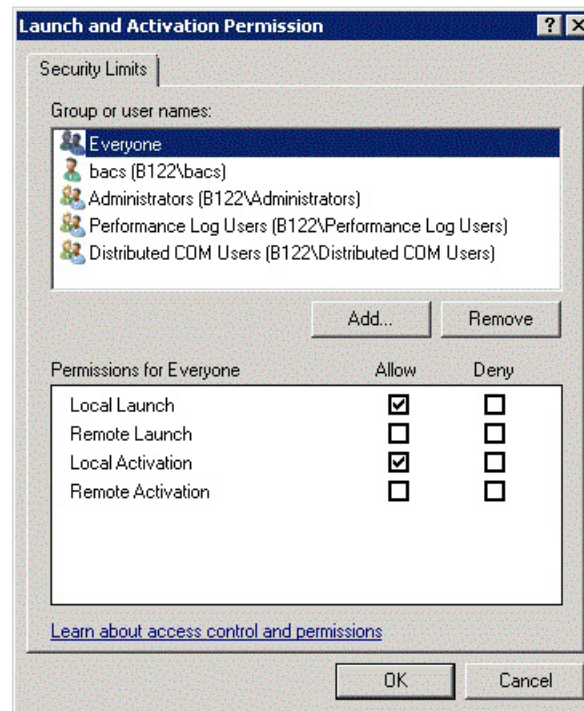
Step 2: Grant DCOM Remote Launch and Activate Permission

In the Windows domain environment, the Domain Administrator account has the necessary privilege level to access the WMI component for BACS management and, therefore, no special configuration is needed. In a large enterprise, however, a user who is accessing the local or remote host using the BACS4 client GUI may not always have the domain administrator account privilege. It is necessary to configure WMI security access on the remote host to allow the user to connect to it using the BACS4 client GUI.

This configuration can be easily done using the following procedure. If you do not have sufficient privileges to configure security for WMI access, contact your Network Administrator.

1. Click **Start**, click **Run**, type **DCOMCNFG**, and then click **OK**.
2. The Component Services dialogue box displays.
3. Open **Component Services** and then open **Computers**.
4. Right-click **My Computer** and click **Properties**.
5. In **My Computer Properties**, click the **COM Security** tab.
6. Under **Launch and Activation Permissions**, click **Edit Limits**.
7. Follow these steps if your name or your group does not appear in the **Groups or user names** list.
 - a. In the Launch Permission dialog box, click **Add**.
 - b. In the Select Users, Computers, or Groups dialog box, add your name and the group in the **Enter the object names to select** box, and then click **OK**.
 - c. In the Launch Permission dialog box, select your user and group in the **Group or user names** list.
 - d. In the **Permissions for User** area, select **Allow** for **Remote Launch** and **Remote Activation**, and then click **OK**.

Figure 1: Launch and Activation Permission



For more information, see [Securing a Remote WMI Connection](#) on the Microsoft Developer Network site.

Special Configuration for WMI on Different Systems

In Windows Vista and Windows 7, in order to let all users in the administrator group connect using the WMI namespace, the user might need to change the LocalAccountTokenFilterPolicy as needed.

Section 4: Linux Driver and Management Application Installation

- [Packaging](#)
- [Installing TG3 Driver Software](#)
- [Network Installations](#)
- [Unloading/Removing the TG3 Driver](#)
- [Driver Messages](#)
- [Teaming with Channel Bonding](#)
- [Linux Management Application Installation](#)

Packaging

The Linux TG3 driver is released in the following packaging formats (file names):

- Source RPM (*tg3-version.3dkms.src.rpm*)
- Supplemental (*tg3_sup-version.tar.gz*)
- Compressed tar (*tg3-version.tar.gz*)

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

Installing TG3 Driver Software

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

Installing the Source RPM Package

Prerequisites:

- Linux kernel source
- C compiler

Procedure:

1. Install the source RPM package.

```
rpm -ivh tg3-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/tg3.spec or rpmbuild -bb SPECS/tg3.spec  
rpmbuild -bb SPECS/tg3.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/SOURCE
```

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.i386.rpm
```

Complete the installation of the source RPM.

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

```
rpm -ivh RPMS/x86_64/tg3-version.i386.rpm
```

Depending on the kernel, the driver is installed to the following path:

2.6.x kernels:

```
/lib/modules/kernel_version/kernel/drivers/net/tg3.ko
```

4. Load the driver.

```
modprobe tg3
```

To configure the network protocol and address, refer to the Linux version-specific documentation.

Building the Driver from the Source TAR File

1. Create a directory (*tg3-version*) and extract the TAR files to the directory.

```
tar xvzf tg3-version.tgz
```

2. Build the driver *tg3.o* as a loadable module for the running kernel.

```
CD tg3-version
make clean
make; make install
```

3. Test the driver by loading it.

```
rmmod tg3
modprobe tg3
```

No message should be returned if this command runs properly.



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

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

Network Installations

For network installations through NFS, FTP, or HTTP (using a network boot disk or PXE), use the *tg3* driver that is part of the Linux operating system distribution.

Unloading/Removing the TG3 Driver

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

Unloading/Removing the Driver from an RPM Installation

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

```
rmmod tg3
```

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

```
rpm -e tg3-<version>
```

Removing the Driver from a TAR Installation

If the driver was installed using **make install** from the tar file, the *tg3.o* driver file has to be manually deleted from the operating system. See [Installing the Source RPM Package](#) for the location of the installed driver.

If there is an interface configuration that is related to the *tg3* driver, then bring the interface down first by using **ifconfig ethx down** and then **rmmod tg3**.

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.

Driver Sign on

```
tg3.c:version (date)
```

NIC Detected

```
eth#: Tigon3 [partno (BCM95xxx) rev 4202 PHY (57xx) (PCI Express) 10/100/1000BaseT
Ethernet :00:xx:xx:xx:xx:xx
eth#: RXcsuns [1] LinkChg REG [0] MIirq [0] ASF [0] Split [0] Wirespeed [1]TSOcap [1]
eth#: dma_rwctrl [76180000]
ACPI : PCI interrupt 0000:02:02.0 [A] -> GSI 26 (level,low) -> IRQ 233
```

Flow Control

```
tg3: eth#: Flow control is configured for TX and for RX.
```

Link Up and Speed Indication

```
tg3: eth#: Link is up at 1000 Mbps, full duplex.
```

Link Down Indication

```
tg3: eth#: Link is down.
```

Teaming with Channel Bonding

With the TG3 driver, you can team adapters together using the bonding kernel module and a channel bonding interface. Refer to your Linux documentation for more information on Linux Channel Bonding.

Linux Management Application Installation

- [Overview](#)
- [Installing WS-MAN or CIM-XML on Linux Server](#)
- [Installing WS-MAN or CIM-XML on Linux Client](#)
- [Installing the Broadcom Advanced Control Suite](#)

Overview

The Broadcom Advanced Control Suite version 4 (BACS4) is a management application for configuring the NetXtreme I families of adapters. BACS4 GUI software operates on Windows and BACS CLI is used on Linux server operating systems.

This chapter describes how to install the BACS4 management application on Linux systems. For Windows systems, an installation program is provided which installs both the Windows drivers and the management applications, including BACS4 (see [Windows Driver and Management Application Installation](#) for instructions).

There are two main components of the BACS4 utility: the provider component and the client software. A provider is installed on a server, or “managed host”, that contains one or more NICs. The provider collects information on the NICs and makes it available for retrieval from a management PC on which the client software is installed. The client software enables viewing information from the providers and configuring the NICs. The BACS client software includes a command line interface (CLI).

Communication Protocols

A communication protocol enables exchanging information between provider and the client software. These are proprietary or open-source implementations of the Web-Based Enterprise Management (WBEM) and Common Information Model (CIM) standards from the Distributed Management Task Force (DMTF). Network administrators can choose the best option based on the prevailing standard on their network.

The following table shows the available options based on the operating systems installed on the managed host and the client.

<i>If the client uses:</i>	<i>And the managed host uses:</i>	<i>BACS can use these communication protocols:</i>
Windows	Windows	WMI WS-MAN (WinRM)
Windows	Linux	CIM-XML (OpenPegasus) WS-MAN (OpenPegasus)
Linux	Windows	WS-MAN (WinRM)
Linux	Linux	CIM-XML (OpenPegasus) WS-MAN (OpenPegasus)

<i>If the client uses:</i>	<i>And the managed host uses:</i>	<i>BACS can use these communication protocols:</i>
<ul style="list-style-type: none"> • WMI = Windows Management Instrumentation. • WS-MAN = Web Service-Management. WinRM is a Windows-based implementation and OpenPegasus is an open-source implementation of the that operates on Linux. • CIM-XML = An XML-based version of OpenPegasus. 		

If your network includes a mix of Windows and Linux clients accessing Windows and Linux servers, then WS-MAN is a suitable choice. If Linux is the only OS installed on the servers, then CIM-XML is an option. If the network includes only Windows servers and clients, WMI is an option. WMI is very simple to configure but is supported only on the Windows OS. (See [Windows Driver and Management Application Installation](#) for instructions on installing and configuring the Windows protocols.)

BACS installation includes installing the provider component on the managed host and the client software on the management station. The installation process differs based on the combination of operating systems installed on the client and managed host and on the selected communication protocol.

Installing WS-MAN or CIM-XML on Linux Server

Step 1: Install OpenPegasus

On the Red Hat Linux OS, two installation options are available:

- [From the Inbox RPM \(Red Hat Only\)](#)
- [From Source \(Red Hat and SuSE\)](#)

On the SUSE Linux Enterprise Server 11 (SLES11) OS, you must use the source RPM.



Note: The Inbox RPM does not support the WS-MAN communication protocol. To use WS-MAN, you must install OpenPegasus from source.

From the Inbox RPM (Red Hat Only)

In Red Hat Linux, an Inbox OpenPegasus RPM is available as `tog-pegasus-<version>.<arch>.rpm`.

1. Use the following command to install `tog-pegasus`:

```
rpm -ivh tog-openpegasus-<version>.<arch>.rpm
```
2. Use the following command to start Pegasus:

```
/etc/init.d/tog-pegasus start
```



Note: On SuSE Linux, the Inbox OpenPegasus RPM is not available. OpenPegasus must be installed from source, as described in the following procedure.

Note that in inbox Pegasus, HTTP is not enabled by default. After Inbox OpenPegasus is installed successfully, if no further configuration is required, then follow the instructions in [Step 4: Install Broadcom CMPI Provider](#). To enable HTTP, see [Enable HTTP](#).

From Source (Red Hat and SuSE)

The OpenPegasus source can be downloaded from www.openpegasus.org.



Note: If not already installed, download and install the openssl and libopenssl-devel rpm. This step is optional and required only if you are planning to use HTTPS to connect the client to the managed host.

Set the Environment Variables

Set the environment variables for building OpenPegasus as follows.

<i>Environment Variable</i>	<i>Description</i>
PEGASUS_ROOT	The location of the Pegasus source tree
PEGASUS_HOME	The location for the built executable, repository; e.g., \$PEGASUS_HOME/bin, PEGASUS_HOME/lib, \$PEGAUS_HOME/repository, and \$PEGASUS_HOME/mof subdirectories.
PATH	\$PATH:\$PEGASUS_HOME/bin
PEGASUS_ENABLE_CMPI_PROVIDER_MANAGER	True
PEGASUS_CIM_SCHEMA	"CIM222"
PEGASUS_PLATFORM	For Linux 32 bit systems: "LINUX_IX86_GNU" For Linux 64 bit systems: "LINUX_X86_64_GNU"
PEGASUS_HAS_SSL	Optional. Set to "true" for HTTPS support.
PEGASUS_ENABLE_PROTOCOL_WSMAN	Optional. Set to "true" for WSMAN protocol support.

Additional Settings

The \$PEGASUS_HOME variable must be set up in the shell environment, and \$PEGASUS_HOME/bin needs to be appended to the \$PATH environment.

Examples

- export PEGASUS_PLATFORM="LINUX_X86_64_GNU"
- export PEGASUS_CIM_SCHEMA="CIM222"
- export PEGASUS_ENABLE_CMPI_PROVIDER_MANAGER=true
- export PEGASUS_ROOT="/share/pegasus-2.10-src"
- export PEGASUS_HOME="/pegasus"
- export PATH=\$PATH:\$PEGASUS_HOME/bin

For SSL Support, add the following environment variable:

- export PEGASUS_HAS_SSL=true

For WS-MAN Support, add the following environment variable:

- export PEGASUS_ENABLE_PROTOCOL_WSMAN=true

CIM-XML and WSMAN in OpenPegasus use the same ports for HTTP or HTTPS. The default port numbers for HTTP and HTTPS are 5989 and 5989, respectively.



Note: You can add these exports at the end of the `.bash_profile`. This file is located in the `/root` directory.

- The environment variables will be set when a user logs in using PuTTY.
- On the Linux system itself, for each terminal where the environment variables are not set, run the following command:

```
source /root/.bash_profile
```
- When you logout and login, the environment variables will be set.

Build and install OpenPegasus

From `$PEGASUS_ROOT` (the location of the Pegasus source root directory), run the following:

```
make clean
make
make repository
```



Note: Whenever OpenPegasus is built from source, all configurations are reset to the default values. If you are rebuilding OpenPegasus, you must redo the configuration as mentioned in [Step 3: Configure OpenPegasus on the Server](#).

Step 2: Start CIM Server on the Server

Use the `cimserver` command to start CIM server. To stop CIM server, use the command `cimserver -s`.

To check whether OpenPegasus has been installed properly, enter the following command:

```
cimcli ei -n root/PG_Interop PG_ProviderModule
```



Note: For OpenPegasus compiled from source, `PEGASUS_HOME` must be defined when you start CIM server. Otherwise, CIM server will not load the repository properly. Consider setting `PEGASUS_HOME` in the `“.bash_profile”` file.

Step 3: Configure OpenPegasus on the Server

Use the `cimconfig` command to configure OpenPegasus, as shown in the following table:

Command	Description
<code>cimconfig -l</code>	List all valid property names.
<code>cimconfig -l -c</code>	List all valid property names and its value
<code>cimconfig -g <property name></code>	Query a particular property.
<code>cimconfig -s <property name>=<value> -p</code>	Set a particular property.
<code>cimconfig --help</code>	Find out more about the command.

CIM server must be started before running `cimconfig`, and must be restarted for configuration changes to take effect.

Enable Authentication

The following OpenPegasus properties have to be set as described in this section. Otherwise, the Broadcom CIM Provider will not work properly. Ensure the following are set before launching BACS and connecting to the provider.

Start CIM server if it is not already started. Then, set the following:

- `cimconfig -s enableAuthentication=true -p`
- `cimconfig -s enableNamespaceAuthorization=false -p`
- `cimconfig -s httpAuthType=Basic -p`
- `cimconfig -s passwordFilePath=cimserver.passwd -p`
- `cimconfig -s forceProviderProcesses=false -p`

If you want root user to connect remotely:

- `cimconfig -s enableRemotePrivilegedUserAccess=true -p`

User configuration with privilege: The Linux system users are used for OpenPegasus authentication. The systems users have to be added to OpenPegasus using `cimuser` to connect via BACS:

- `cimuser -a -u <username> -w <password>`

Example: `cimuser -a -u root -w linux1`

Enable HTTP

1. If CIM server is not started, start it.
2. Use the following command to set up an HTTP port (optional):
`cimconfig -s httpPort=5988 -p`
This property is not available for Inbox OpenPegasus.
3. Use the following command to enable HTTP connection:
`cimconfig -s enableHttpConnection=true -p`
4. Use the `cimserver -s` and `cimserver` commands, respectively, to stop and restart CIM server for the new configuration to take effect.

Enable HTTPS

1. If CIM server is not started, start it.
2. Set up HTTPS port with the following command (optional):
`cimconfig -s httpsPort=5989 -p`
This property is not available for inbox OpenPegasus.
3. Enable HTTPS connection with 'the following command:
`cimconfig -s enableHttpsConnection=true -p`
4. Use the `cimserver -s` and `cimserver` commands, respectively, to stop and restart CIM server for the new configuration to take effect.

Step 4: Install Broadcom CMPI Provider

Ensure that OpenPegasus is installed properly before installing CMPI Provider.

Install

Enter following command to install Broadcom CMPI Provider.

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

Uninstall

Enter following command to uninstall Broadcom CMPI Provider:

```
% rpm -e BRCM_CMPIProvider
```

Step 5: Perform Linux Firewall Configuration, If Required

Follow these procedures to open the appropriate ports in the firewall:

Red Hat

1. Click **System**, select **Administration**, and then select **Firewall**.
2. Select **Other Ports**.
3. In the Port and Protocol Dialog box, select **User Defined**.
4. In the **Port/Port Range** field, add the port number.
5. In the **Protocol** field, add the protocol as TCP or UDP, etc.
6. Click **Apply** for the firewall rules to take effect.

Example:

- For CIM-XML over HTTP, the port number is 5988 and protocol is TCP.
- For CIM-XML over HTTPS, the port number is 5989 and protocol is TCP.

SuSE

1. Click **Compute** and then click **YaST**.
2. Select **Security & Users** on the left pane.
3. On the right pane, double-click **Firewall**.
4. Select **Custom Rules** on the left pane.
5. On the right pane click **Add**.
6. Enter the following values:
 - **Source Network:** 0/0 (means all)
 - **Protocol:** TCP (or the appropriate protocol)
 - **Destination Port:** <Port Number> or <Range of Port Numbers>
 - **Source Port:** Leave blank.
7. Click **Next** and then click **Finish** for the firewall rules to take effect.

Example:

For CIM-XML, use the following values:

- **Source Network:** 0/0 (means all)
- **Protocol:** TCP
- **Destination Port:** 5988:5989
- **Source Port:** Leave blank.

Step 6: Install BACS and Related Management Applications

See [Installing the Broadcom Advanced Control Suite](#).

Installing WS-MAN or CIM-XML on Linux Client

No special software components are required on the Linux client system to use the HTTP except installing the BACS management application. However, for WS-MAN installations, you can optionally configure the HTTPS protocol for use with BACS.

Configure HTTPS on Linux Client

Follow these steps if you want to use HTTPS rather than HTTP (WS-MAN only):

Generate a Self-Signed Certificate for Windows/Linux Server

Openssl on Linux or Windows can be used to generate the self-signed certificate, as follows:



Note: You can download and install openssl from <http://gnuwin32.sourceforge.net/packages/openssl.htm>.

1. Enter the following command to generate a private key:

```
openssl genrsa -des3 -out server.key 1024
```

2. You are prompted to enter a passphrase. Be sure to remember the passphrase.
3. Use the following steps to generate a Certificate Signing Request (CSR).

During the generation of the CSR, you are prompted for several pieces of information. When prompted for the “Common Name”, enter the Windows Server host name or IP address.

Enter the following command (sample responses are show):

```
openssl req -new -key server.key -out server.csr
```

If this command does not work, try the following:

```
openssl req -new -key server.key -out server.csr -config openssl.cnf
```

The openssl.cnf file should be placed in the same directory where openssl is placed. Openssl.cnf is located in the folder C:\Program Files (x86)\GnuWin32\share.

The following information is requested:

- Country Name (2 letter code) []: **US**
- State or Province Name (full name) []: **California**
- Locality Name (e.g., city) []: **Irvine**
- Organization Name (e.g., company) []: **Broadcom Corporation**
- Organizational Unit Name (e.g., section) []: **Engineering**
- Common Name (e.g., YOUR name) []: Enter the host name or IP address of the Windows server. For IPv6, enter the Common Name in the format [xyxy:xxx:.....:xxx], **including the brackets []**.
- (Optional) Email Address []:

Enter the following additional attributes to be sent with your certificate request:

- A challenge password []: **linux1**
- An optional company name []:

4. Remove the passphrase from the key.

Enter the following commands:

```
cp server.key server.key.org
```

```
openssl rsa -in server.key.org -out server.key
```

5. Generate a self-signed certificate:

To generate a self-signed certificate which is active for 365 days, enter the following command:

```
openssl x509 -req -days 365 -in server.csr -signkey server.key -out server.crt
```

The following output displays:

```
Signature ok
subject=/C=US/ST=California/L=Irvine/O=Broadcom Corporation/OU=Engineering/CN=MGMTAPP-
LAB3/emailAddress=
Getting Private key
```

6. Enter the following command to verify the generated self-signed certificate.

```
openssl verify server.crt
```

The following output displays:

```
server.crt:/C=US/ST=California/L=Irvine/O=Broadcom Corporation/OU=Engineering/
CN=MGMTAPP- LAB3/emailAddress=
error 18 at 0 depth lookup:self signed certificate
OK
```

Ignore the error message “error 18 at 0 depth lookup:self signed certificate”. This error indicates that this is a self-signed certificate.

7. Convert the certificate from “crt” to “pkcs12” format, as follows:

For a Windows server, the certificate should be in pkcs12 format. Enter the following command:

```
openssl pkcs12 -export -in server.crt -inkey server.key -out hostname.pfx
```

You will be prompted for the following:

```
Enter Export Password:
Verifying - Enter Export Password:
```

Enter the password and be sure to remember it. The password is required when importing the certificate on the Windows server and client.

8. Make a copy of the certificate file server.crt and place it on the server where BACS will be installed, so that it can be imported. If you plan to use a Windows or Linux client to connect to the server running BACS, then the certificate also needs to be transferred (copied and pasted) to the client system.

In Linux, the certificate should have the extension “.pem”. The extension “.crt” and “.pem” are the same, so there is no need to use the `openssl` command to convert from .crt to .pem. You can simply copy the file as-is.



Note: A separate certificate must be generated for an IPv4 address, IPv6 address, and Hostname.

Import Self-Signed Certificate on Linux Client

On Linux distributions, note the following certificate directory:

- For all SuSE versions, the certificate directory is `/etc/ssl/certs`.
- For Red Hat, the certificate directory can be different for each version. For some versions, it is `/etc/ssl/certs` or `/etc/pki/tls/certs`. For other versions, find out the certificate directory.

Copy *hostname.pem*, which you created in [Generate a Self-Signed Certificate for Windows/Linux Server](#), into the certificate directory of the Linux client. For example, if the certificate directory is `/etc/ssl/certs`, copy *hostname.pem* to `/etc/ssl/certs`.

1. Change directory to `/etc/ssl/certs`.
2. Create a hash value by running the following command.

```
openssl x509 -noout -hash -in hostname.pem
```

A value such as the following will be returned.
100940db
3. Create a symbolic link to the hash value by running the following command:

```
ln -s hostname.pem 100940db.0
```

Test HTTPS/SSL Connection from Linux Client

Use the following command to test whether the certificate is installed correctly on Linux:

```
# curl -v --capath /etc/ssl/certs https://Hostname or IPAddress:5986/wsman
```

If this fails, then the certificate is not installed correctly and an error message displays, indicating to take corrective action.

Installing the Broadcom Advanced Control Suite

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

Before you begin:

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

To install BACS CLI

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

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

To use BACS CLI, refer to the file `BACSCLI_Readme.txt` provided with the release files.

To remove BACS CLI

To uninstall the RPM package, use the following command:

```
% rpm -e BACScLi
```

Section 5: VMware Driver Software

- [Packaging](#)
- [Drivers](#)

Packaging

The VMware driver is released in the following packaging format.

Table 3: VMware Driver Packaging

<i>Format</i>	<i>Drivers</i>
VMware VIB	vmware-esx-drivers-net-tg3-version.x86_64.vib

Drivers

Download, Install, and Update Drivers

To download, install, or update the VMware ESX/ESXi driver for NetXtreme I GbE network adapters, see <http://www.vmware.com/support>.

Driver Parameters

NetQueue

The optional parameter **force_netq** can be used to set the number of Rx and Tx net queues. BCM57XX devices that support NetQueue are the BCM5718, BCM5719, and BCM5720.

Allowed values for x are -1 to 15:

- 1–15 will force the number of NetQueues for the given NIC.
- 0 disables NetQueue.
- -1 specifies to use the default driver NetQueue value.

The number of “x” entries can go up to 32, which means the maximum supported NICs = 32.

Example usage:

```
esxcfg-module -s force_netq=-1,0,1,2 tg3]
```

- tg3 NIC 0: Use the default number of NetQueues.
- tg3 NIC 1: Disable the NetQueue feature.
- tg3 NIC 2: Use 1 NetQueue.
- tg3 NIC 3: Use 2 NetQueues.

Note that the NIC # above does not correspond to the vmnic<#>. The NIC number is the system vmnic probe order number. Optimally, the number of NetQueues matches the number of CPUs on the machine.

Driver Parameters

Several optional parameters can be supplied as a command line argument to the vmkload_mod command. These parameters can also be set via the esxcfg-module command. See the man page for more information.

Driver Defaults

Table 4: VMware Driver Defaults

Parameter	Default Value
Speed	Autonegotiation with all speeds advertised
Flow Control	Autonegotiation with rx and tx advertised
MTU	1500 (range 46–9000)
Rx Ring Size	200 (range 0–511). Some chips are fixed at 64.
Rx Jumbo Ring Size	100 (range 0–255). Not all chips support the jumbo ring and some chips that support jumbo frames do not use the jumbo ring.
Tx Ring Size	511 (range (MAX_SKB_FRAGS+1) – 511). 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	20 (range 0–1023)
Coalesce RX Microseconds irq	20 (range 0–255)
Coalesce rx frames	5 (range 0–1023)
Coalesce rx frames irq	5 (range 0–255)
Coalesce TX Microseconds	72 (range 0–1023)
Coalesce tx usecs irq	20 (range 0–255)
Coalesce tx frames	53 (range 0–1023)
Coalesce tx frames irq	5 (range 0–255)
Coalesce stats usecs	1000000 (approx. 1 sec.). Some coalescing parameters are not used or have different defaults on some chips.
MSI	Enabled (if supported by the chip and passed the interrupt test).
WoL	Disabled

Driver Messages

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

Driver Sign On

```
tg3.c:v3.118g (Jan 4, 2012)
```

NIC Detected

```
vmnic0#: Tigon3 [partno (BCM95xxx) rev 4202 PHY (57xx) (PCI Express) 10/100/1000BaseT  
Ethernet :00:xx:xx:xx:xx:xx  
vmnic0#: RXchecksums [1] LinkChg REG [0] MIirq [0] ASF [0] Split [0] Wirespeed [1]TSOcap [1]  
vmnic0#: dma_rwctrl [76180000]  
ACPI : PCI interrupt 0000:02:02.0 [A] -> GSI 26 (level,low) -> IRQ 233
```

Link Up and Speed Indication

```
tg3: vmnic0: Link is up at 1000 Mbps, full duplex.  
tg3: vmnic0: Flow control is on for TX and on for RX.
```

Link Down Indication

```
tg3: vmnic0: Link is down.
```

Section 6: Using Broadcom Advanced Control Suite 4

- [Broadcom Advanced Control Suite Overview](#)
- [Starting Broadcom Advanced Control Suite](#)
- [BACS Interface](#)
- [Configuring Preferences in Windows](#)
- [Connecting to a Host](#)
- [Managing the Host](#)
- [Managing the Network Adapter](#)
- [Viewing Statistics](#)
- [Configuring Teaming](#)
- [Configuring With the Command Line Interface Utility](#)
- [Managing VLANs](#)
- [Troubleshooting BACS](#)

Broadcom Advanced Control Suite Overview

Broadcom Advanced Control Suite (BACS) is an integrated utility that provides useful information about each network adapter that is installed in your system. BACS also enables you to perform detailed tests, diagnostics, and analyses on each adapter, as well as to view and modify property values and view traffic statistics for network objects. BACS operates on Windows and Linux operating systems. BACS supports Broadcom NetXtreme 1 and Netlink Gigabit controllers. It does not support Broadcom NetXtreme-E 10G/25G controllers.

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



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

The BACS application includes a graphical user interface and a command line interface (BACSCLI). BACS GUI and BACS CLI can operate on the following operating system families:

- Windows
- Windows Server

For information on the latest supported OS versions, refer to the release documentation provided with your software distribution.

Starting Broadcom Advanced Control Suite

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

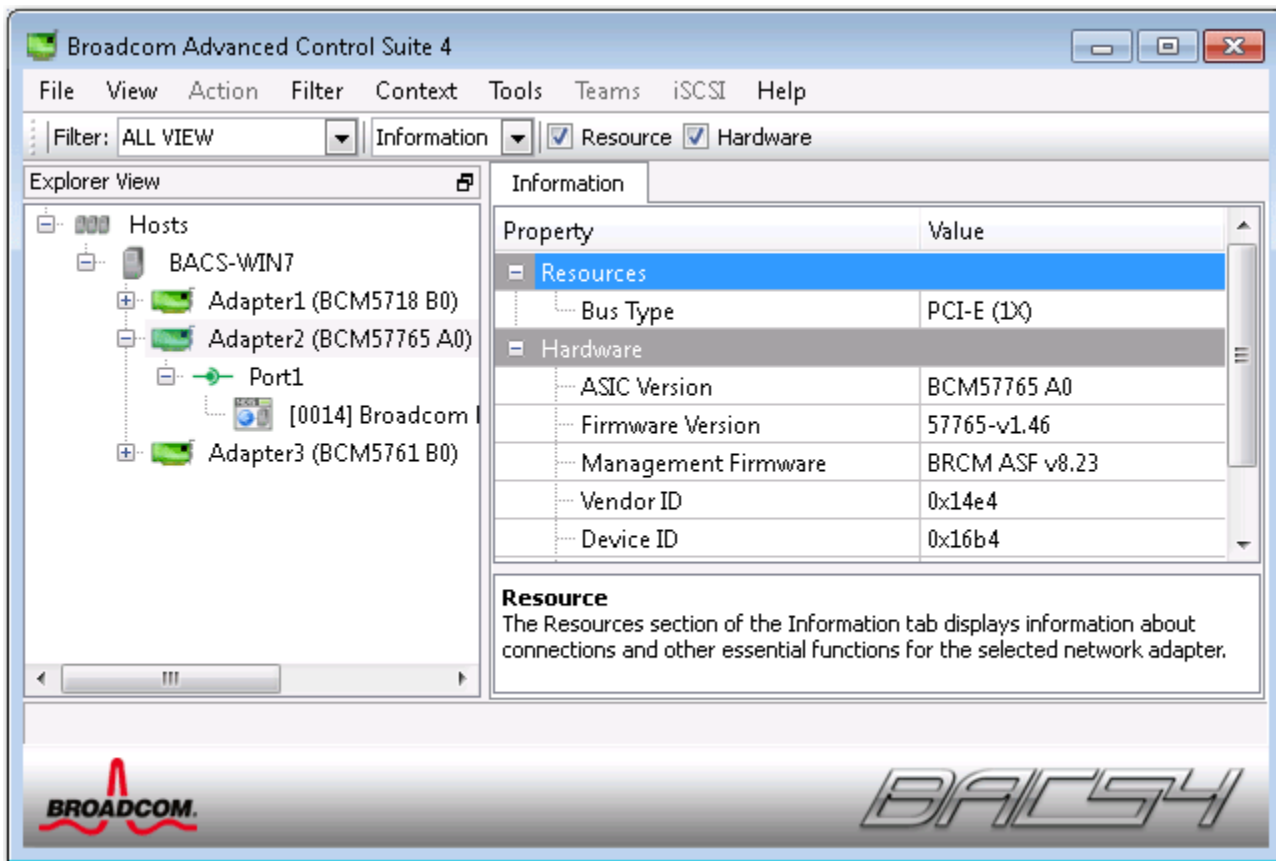
On Linux systems, you can double-click the BACS4 desktop icon, or access the BACS program from the task bar under **System Tools**. (If you are having difficulty launching BACS on a Linux system, see the related topic in [Troubleshooting BACS](#).)

BACS Interface

The BACS interface is comprised of the following regions:

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

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



Explorer View Pane

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

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

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

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

- **X.** A red “X” that appears on the device’s icon indicates the device is currently not connected to the network.
- **Greyed out.** A device icon that appears greyed out indicates the device is currently disabled.

Context View Selector

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

Filter View

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

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

Context View Pane

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

Menu Bar

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

File menu

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

Action menu

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

View menu

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

Tools menu

- Options: Used for configuring BACS preferences.

Teams (Windows only)

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

Description Pane

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

Configuring Preferences in Windows

To enable or disable the BACS tray icon in Windows

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

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

Setting the teaming mode in Windows

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

Setting the Explorer View refresh time in Windows

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

Connecting to a Host

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

To add a local host

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

To add a remote host

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

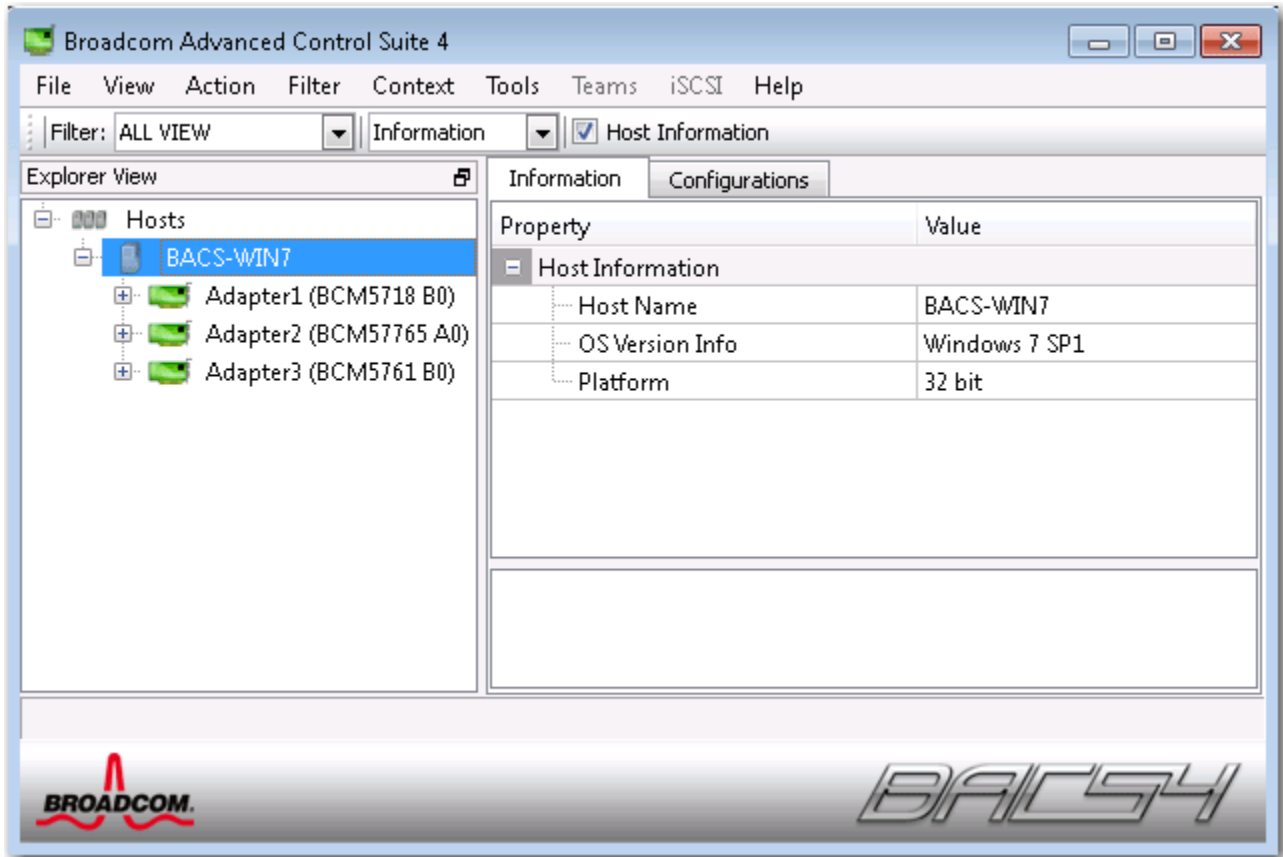
Managing the Host

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

- Information
- Configuration

To view host information

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



Information Tab: Host Information

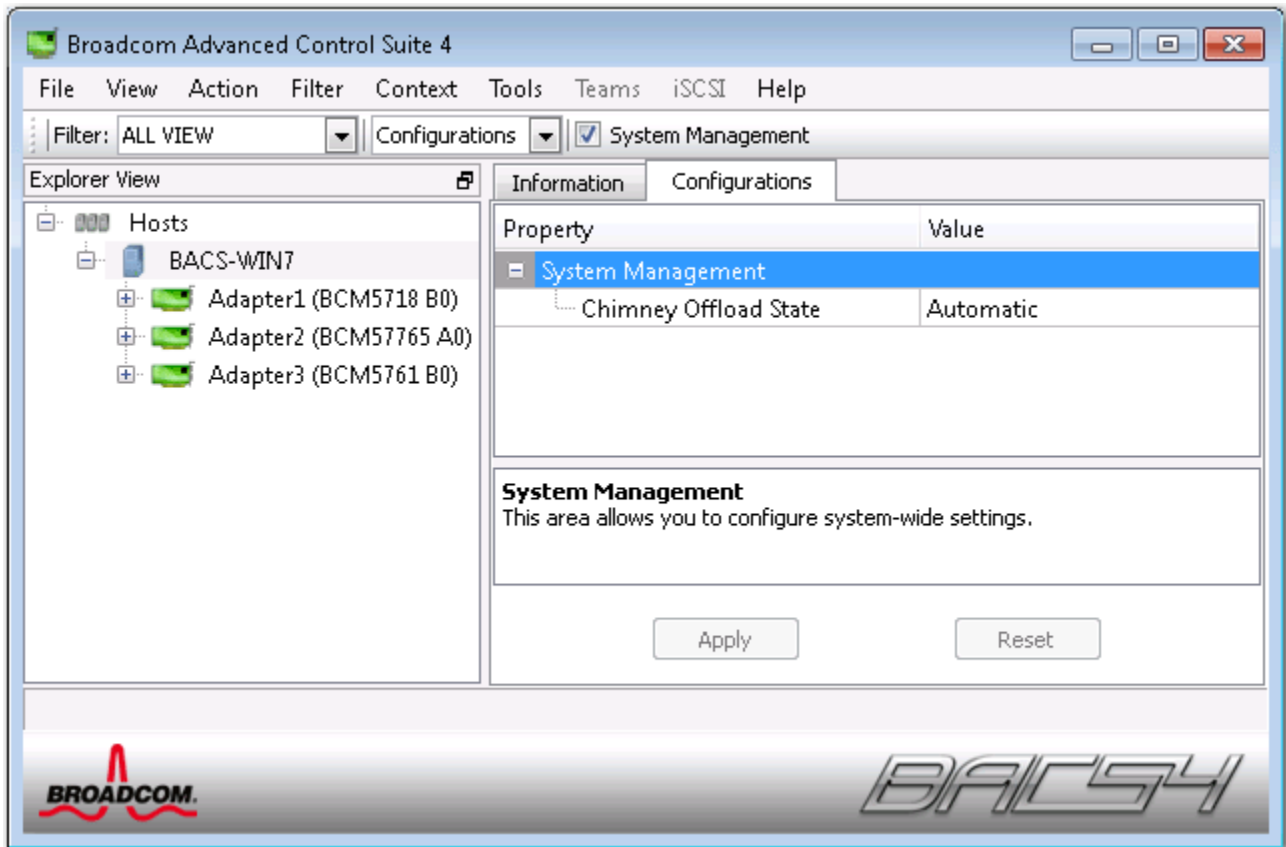
Host Name. Displays the name of the host.

OS Version Info. Displays the operating system, including the version.

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

To configure the host

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



Managing the Network Adapter

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

- Information
- Configuration

Viewing Adapter Information

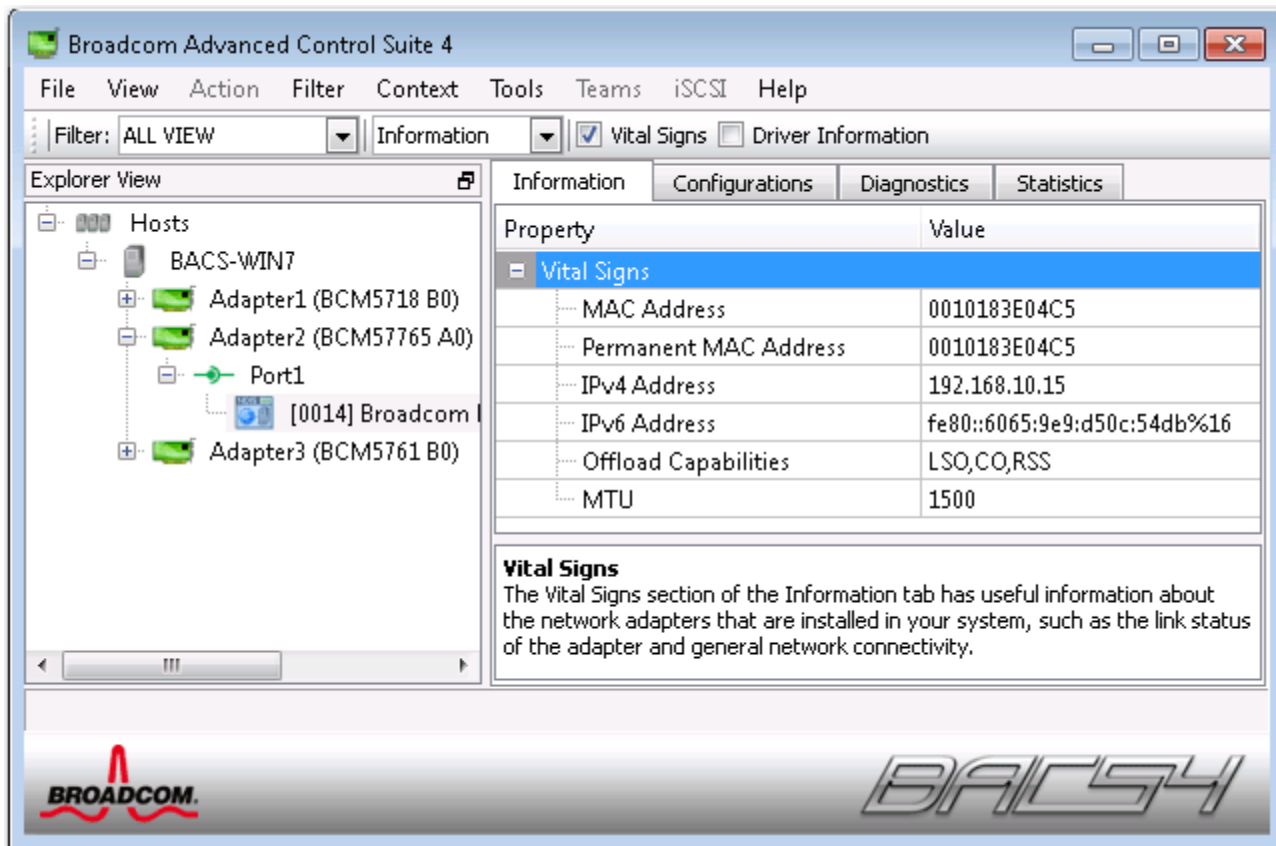
The **Vital Signs** section of the **Information** tab has useful information about the network adapters that are installed in your system, such as the link status of the adapter and general network connectivity.

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



Notes:


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



Viewing Driver Information

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 Explorer View pane, then click the **Information** tab.



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.

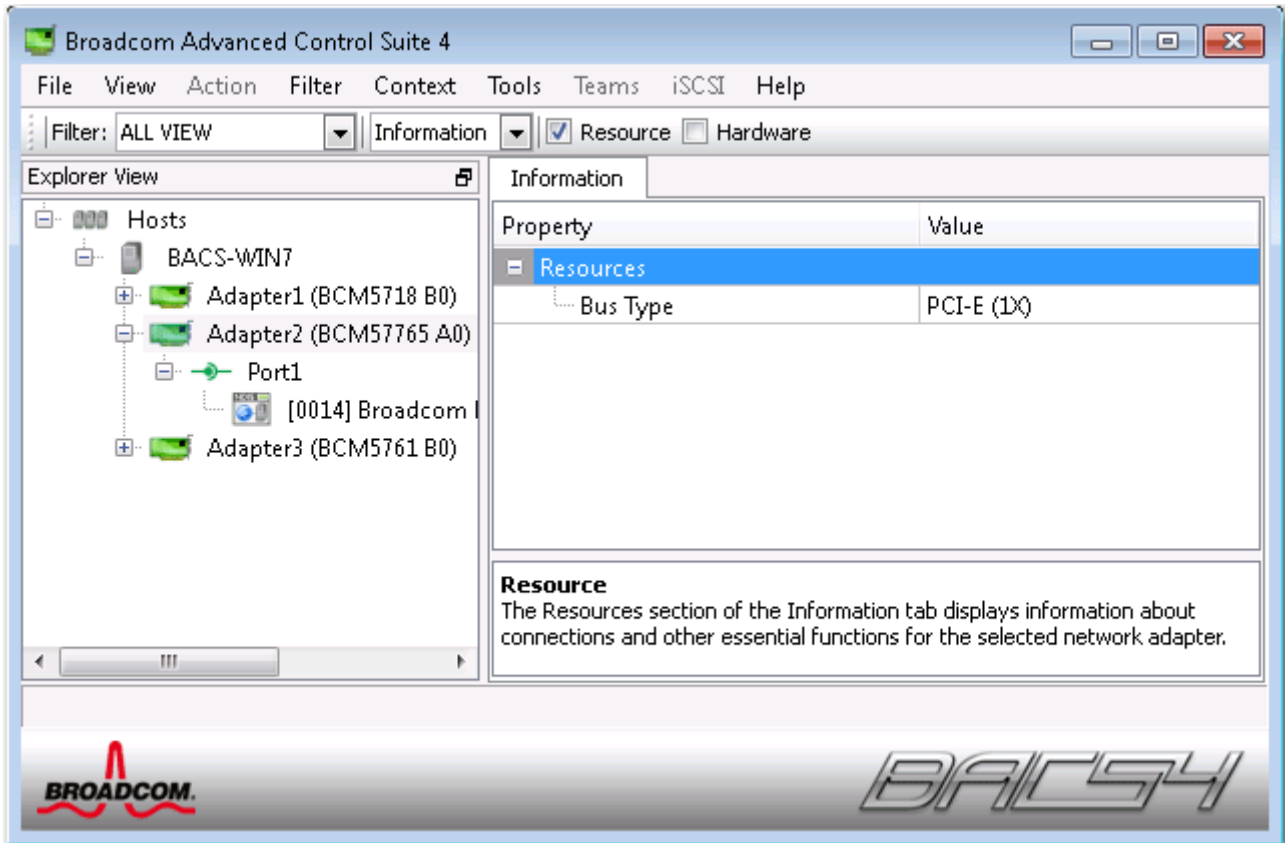
Viewing Resource Information

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 Explorer View pane, then click the **Information** tab.



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.

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.

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 Explorer View pane, then click the Information tab.



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

The screenshot shows the Broadcom Advanced Control Suite 4 application window. The Explorer View pane on the left displays a tree structure under 'Hosts' for 'BACS-WIN7', listing three network adapters: Adapter1 (BCM5718 B0), Adapter2 (BCM57765 A0), and Adapter3 (BCM5761 B0). The 'Information' tab is active, and the 'Hardware' section is expanded, showing a table of properties and values for the selected adapter.

Property	Value
Hardware	
ASIC Version	BCM57765 A0
Firmware Version	57765-v1.46
Management Firmware	BRCM ASF v8.23
Vendor ID	0x14e4
Device ID	0x16b4
Sub-System Vendor ID	0x14e4
Sub System ID	0x16b4

Below the table, there is a section titled **Hardware** with the following text: "The Hardware section of the Information tab displays information about the hardware settings for the selected network adapter."

ASIC Version. The chip version of the Broadcom adapter (this information is not available for adapters made by others).

Firmware Version. The firmware version of the Broadcom adapter (this information is not available for adapters made by others). This information is only available for Broadcom NetXtreme adapters.

Vendor ID. The vendor ID.

Device ID. The adapter ID.

Subsystem Vendor ID. The subsystem vendor ID.

Subsystem ID. The subsystem ID.

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.



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 Explorer View pane.
2. 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**.
3. To change the destination IP address, select **IP address to ping**. In the Network Test window, enter a Destination IP address, then click **OK**.
4. Click **Test**.

The results of the network test are displayed in the **Status** 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 tests are performed continuously, then the number of passes and fails in the **Result** field for each test increments every time the tests are performed. For example, if a test is performed four times and there are no fails, the value in the **Result** field for that test is 4/0. However, if there were 3 passes and 1 fail, the value in the **Result** field is 3/1.



Notes:

- You must have administrator privileges to run diagnostic tests.
- The network connection is temporarily lost while these tests are running.
- Not all Broadcom adapters support each test.

To run the diagnostic tests once

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

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

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

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

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

On-Chip CPU. This test verifies the operation of the internal CPUs in the adapter.

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

LoopBack MAC. This test verifies that the NDIS driver is able to send packets to and receive packets from the adapter.

LoopBack PHY. This test verifies 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.
- For Broadcom NetXtreme adapters, the cable analysis test can only run for gigabit link speed connections and when there is no connection.
- This option is not available for all Broadcom network adapters.

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 Explorer View pane.
3. 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**.
4. Click **Run**.
5. In the error message window that warns of the network connection being temporarily interrupted, click **Yes**.

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

Status. This shows the type of link on this cable 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 duplex 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 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.

To set adapter properties

1. Click the name of the adapter in the Explorer View pane, and click the **Configurations** tab.
2. From the **Advanced** section, select the property you want to set.
3. 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).
4. Click **Apply** to confirm the changes to all properties. Click **Reset** to return the properties to their original values.

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 of 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.

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.

- **Auto** (default). PAUSE frame receipt and transmission are optimized.
- **Disable.** PAUSE frame receipt and transmission are disabled.
- **Rx PAUSE.** PAUSE frame receipt is enabled.
- **Rx/Tx PAUSE.** PAUSE frame receipt and transmission are enabled.
- **Tx PAUSE.** PAUSE frame transmission is enabled.

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.
- **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 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.

Wake Up Capabilities. Enables the network adapter to wake up from a low-power mode when it receives a network wake-up frame. Two types of wake-up frames are possible: Magic Packet and Wake Up Frame.

This property is only available for Broadcom NetXtreme adapters.

- **Both** (default). Selects both Magic Packet and Wake Up Frame as wake-up frames.
- **Magic Packet.** Selects Magic Packet as the wake-up frame.
- **None.** Selects no wake-up frame.
- **Wake Up Frame.** Selects Wake Up Frame as the wake-up frame and allows the network adapter to wake the system when an event such as a ping or an Address Resolution Protocol (ARP) request is received. This option works in conjunction with the operating system power mode saving and does not work if the Power Save setting does not enable WOL.

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

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



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

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

Risk Assessment of VLAN Tagging through the NDIS Miniport Driver

Broadcom's NDIS 6.0 miniport driver provides the means to allow a system containing a Broadcom adapter to connect to a tagged VLAN. Unlike BASP, however, the NDIS 6 driver's support for VLAN participation is only for a single VLAN ID.

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

A properly configured network that has multiple VLANs should maintain separate IP segments for each VLAN. This is necessary since outbound traffic relies on the routing table to identify which adapter (virtual or physical) to pass traffic through and does not determine which adapter based on VLAN membership. Since support for VLAN tagging on Broadcom's NDIS 6.0 driver is limited to transmit (Tx) traffic only, there is a risk of inbound traffic (Rx) from a different VLAN being passed up to the operating system. However, based on the premise of a properly configured network above, the IP segmentation and/or the switch VLAN configuration may provide additional filtration to limit the risk.

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

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

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 Explorer View pane, then click the Statistics tab.

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



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.

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 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.
- 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.
- 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 WOL 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.

Team Types

You can create four types of load balance teams:

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

For a description of these types, see “Load Balancing and Fault Tolerance” in the *Broadcom® NetXtreme® BCM57XX User Guide*.

Using the Broadcom Teaming Wizard

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

1. Create or edit a team:


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

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



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

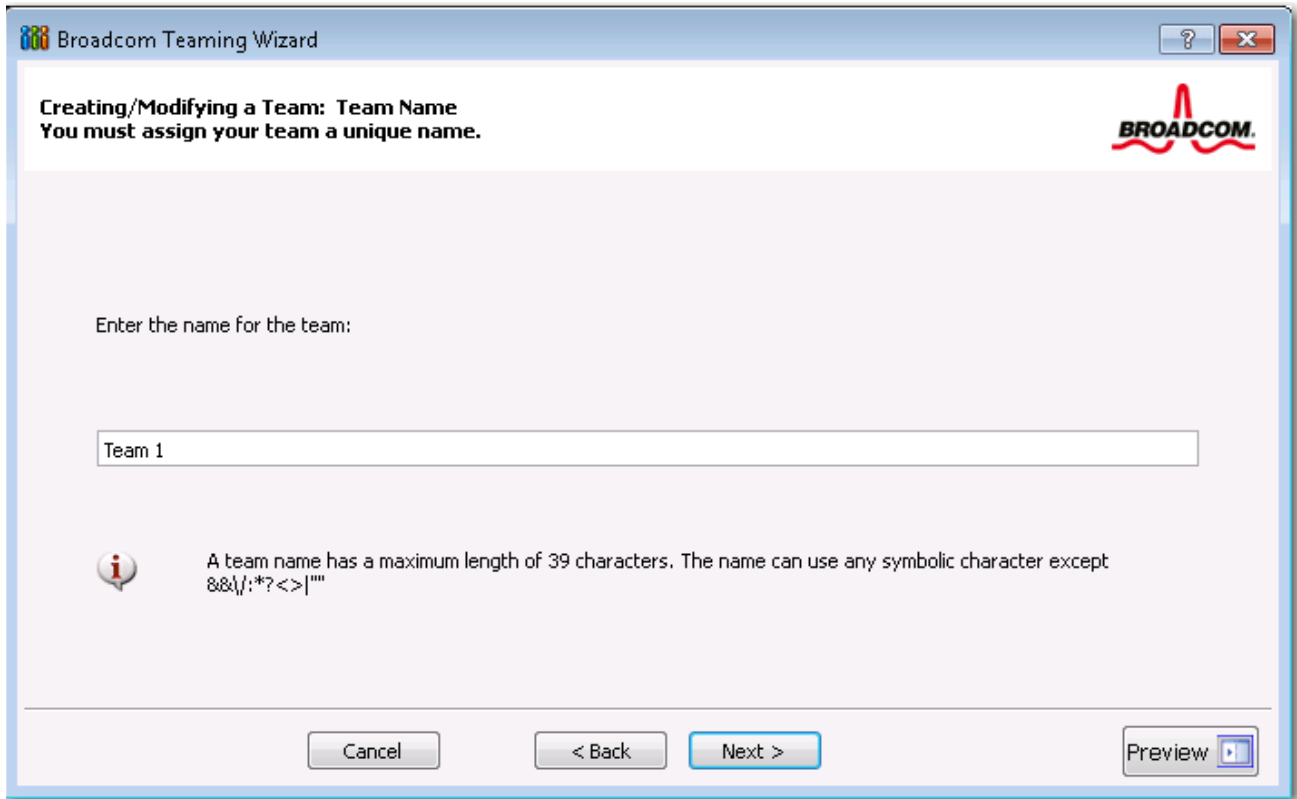
2. To continue using the wizard, click **Next**.



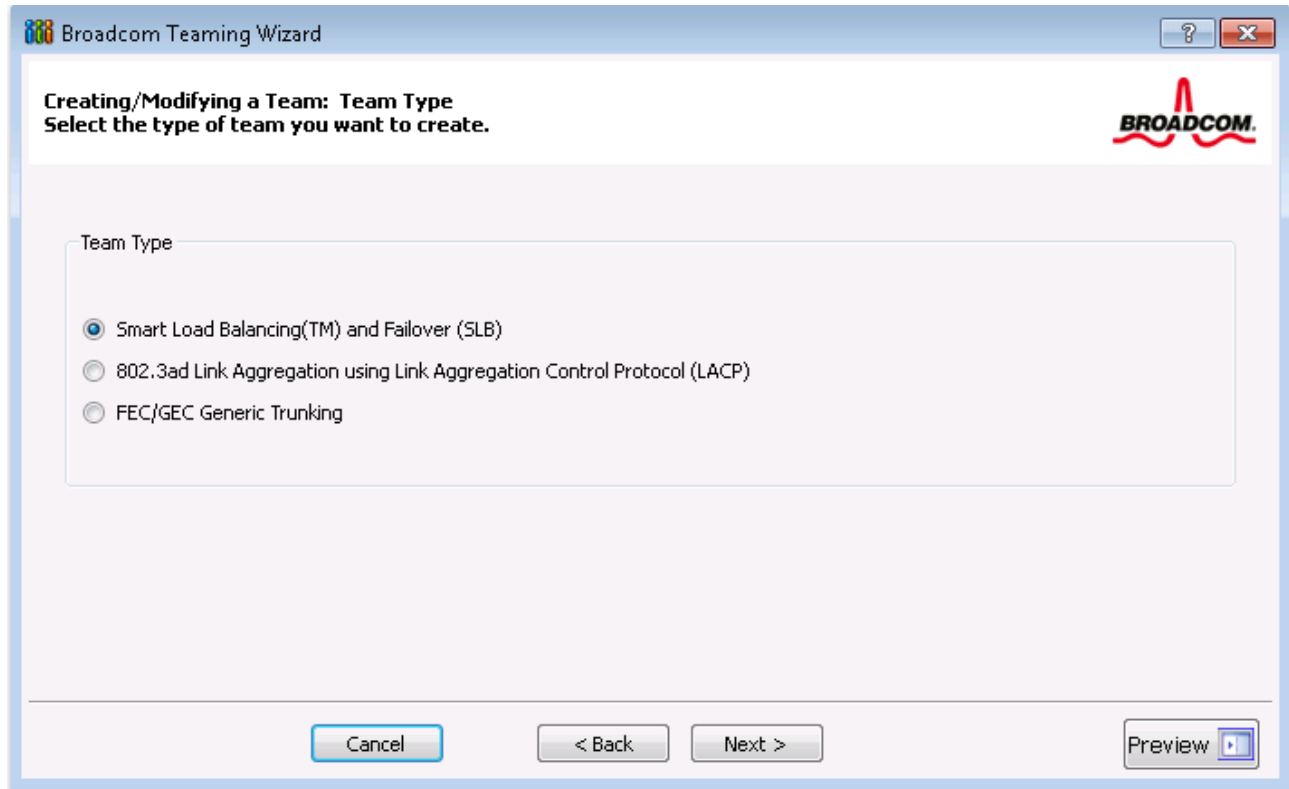
3. Type the team name and then click **Next**. If you want to review or change any of your settings, click **Back**. Click **Cancel** to discard your settings and exit the wizard.



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



4. 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.



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




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


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



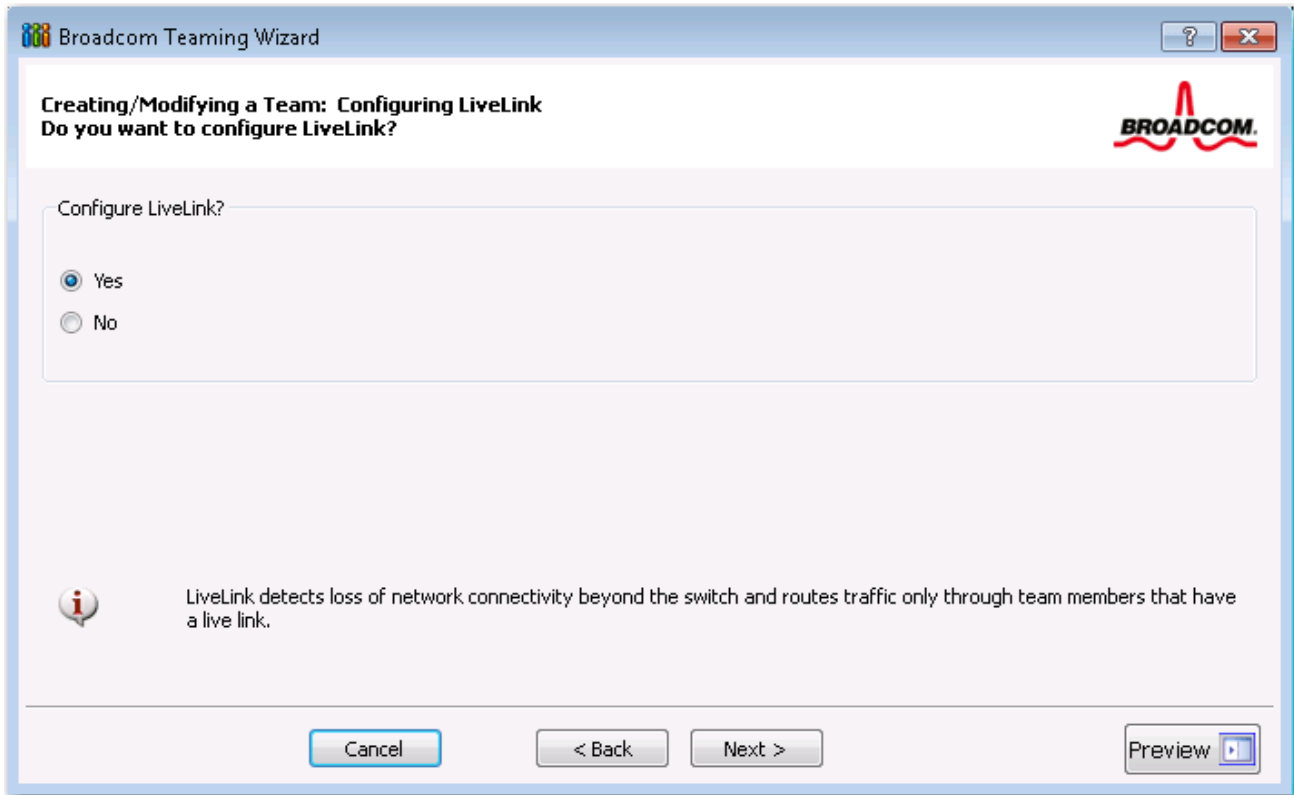
Note: Adding a network adapter to a team where its driver is disabled may negatively affect the offloading capabilities of the team. This may have an impact on the team's performance. Therefore, it is recommended that only driver-enabled network adapters be added as members to a team.



- 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**.



8. If you want to configure LiveLink, select **Yes**, otherwise select **No**, then click **Next**.



9. 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).

10. 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 16 on page 79](#)).

11. 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.

12. 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**.



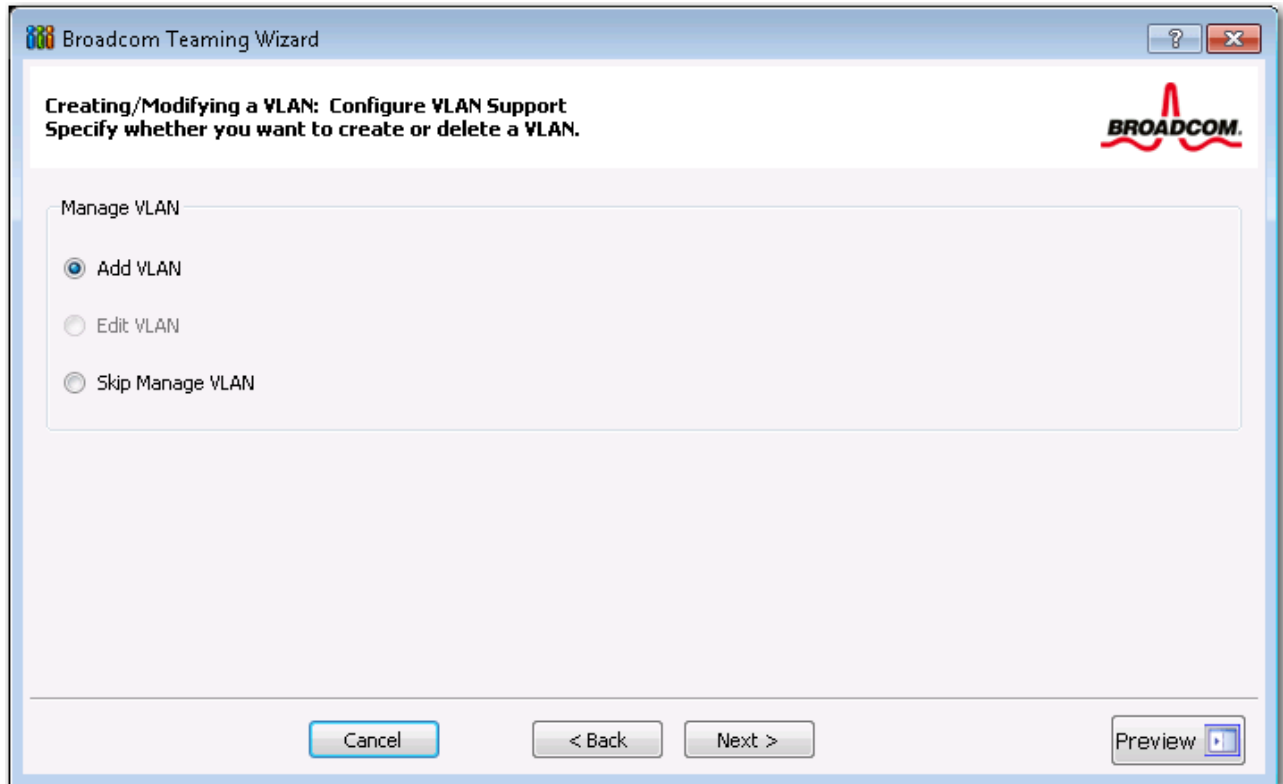
Note: All of the member IP addresses must be in the same subnet as the subnet of the probe targets.

13. 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 18 on page 80](#) 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.



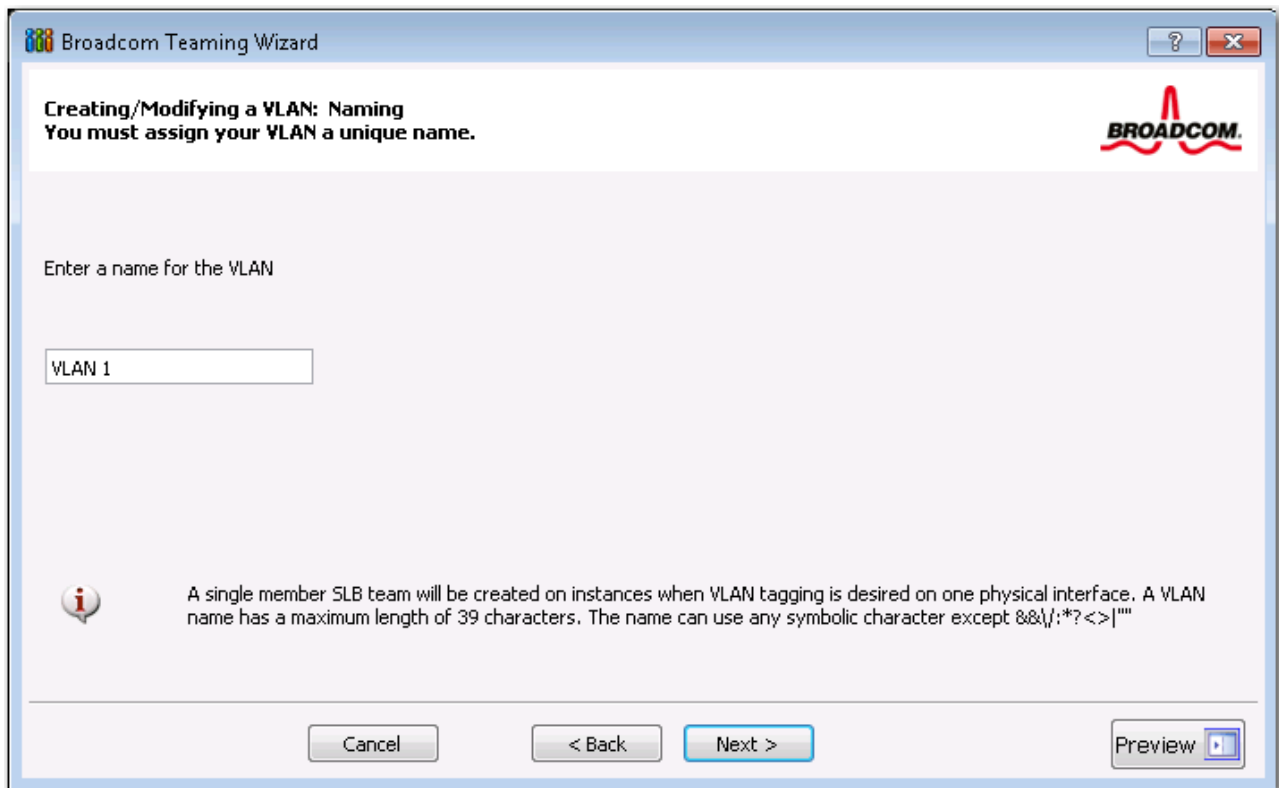
Note: VLANs can only be created when all team members are Broadcom adapters.




14. Type the VLAN name and then click **Next**.



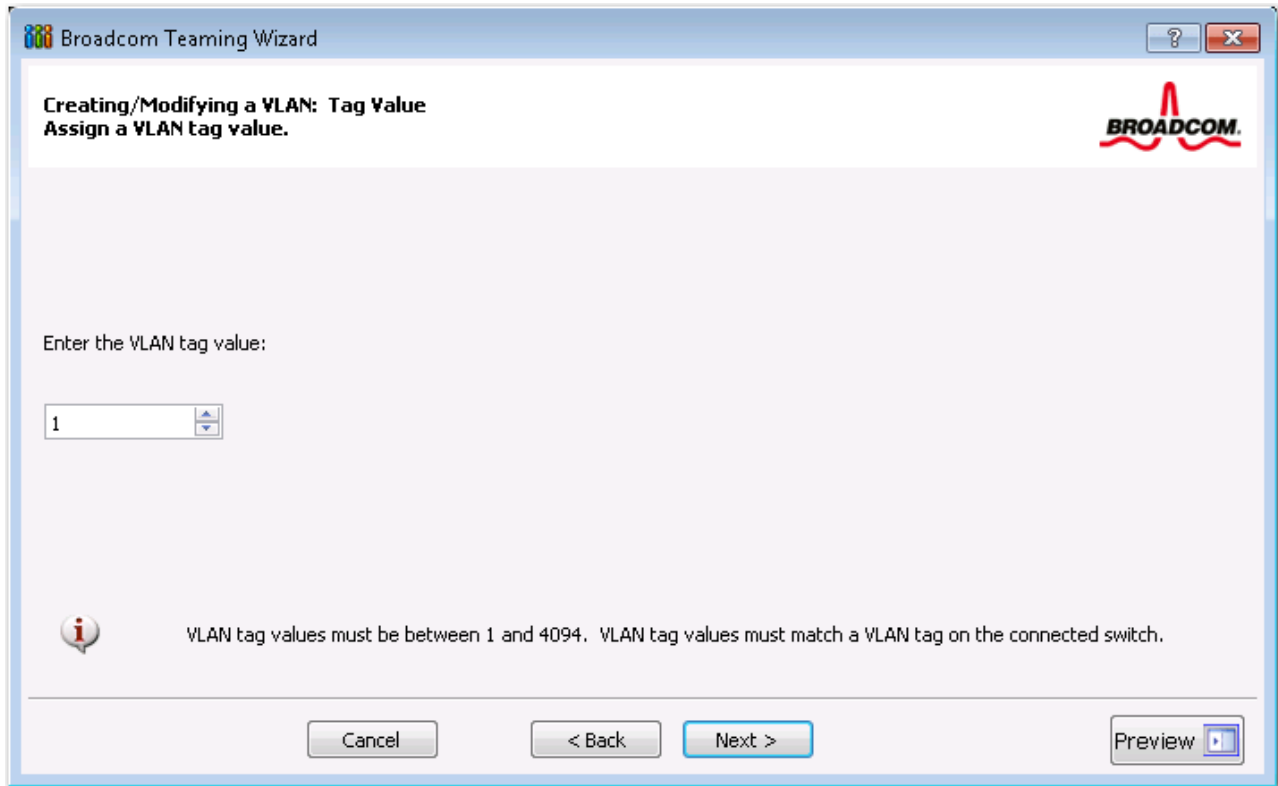
Note: The team name cannot exceed 39 characters, cannot begin with spaces, and cannot contain any of the following characters: & \ / : * ? < > |



15. 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 17 on page 80](#) of this procedure).



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




17. 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.

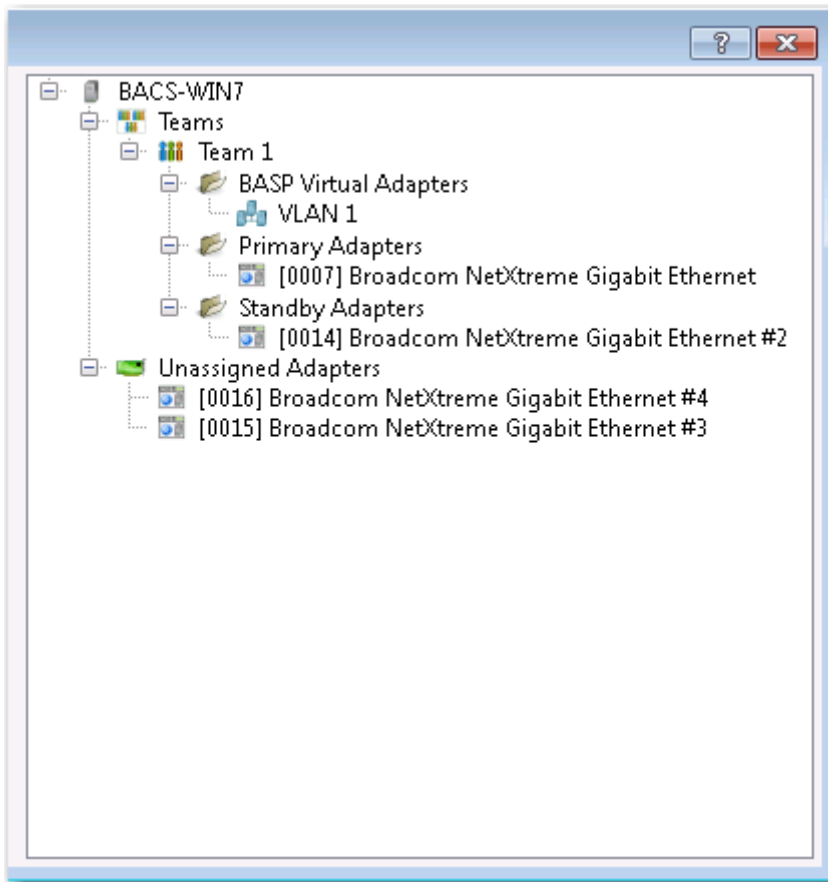


Note: You can define up to 64 VLANs per team (63 VLANs that are tagged and 1 VLAN that is not tagged). Adding several VLANS may slow down the reaction time of the Windows interface due to memory and processor time usage for each VLAN. The degree to which Windows performance may suffer depends on system configuration.

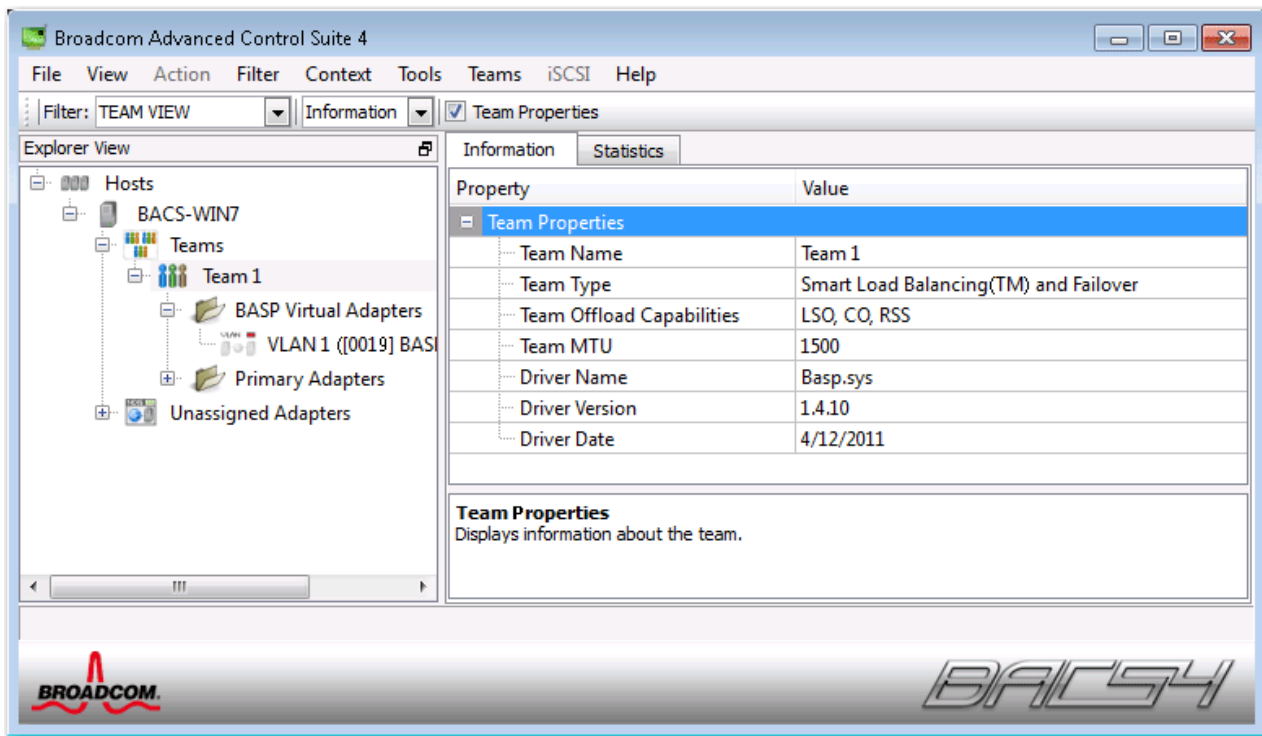
18. To apply and commit the changes to the team, select **Commit changes to system and Exit the wizard**. To apply your changes but continue using the wizard, select **Save changes and continue to manage more teams**. Click Finish.



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.



19. 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.




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, then select **Expert Mode** or **Wizard Mode** (the default is Wizard Mode).

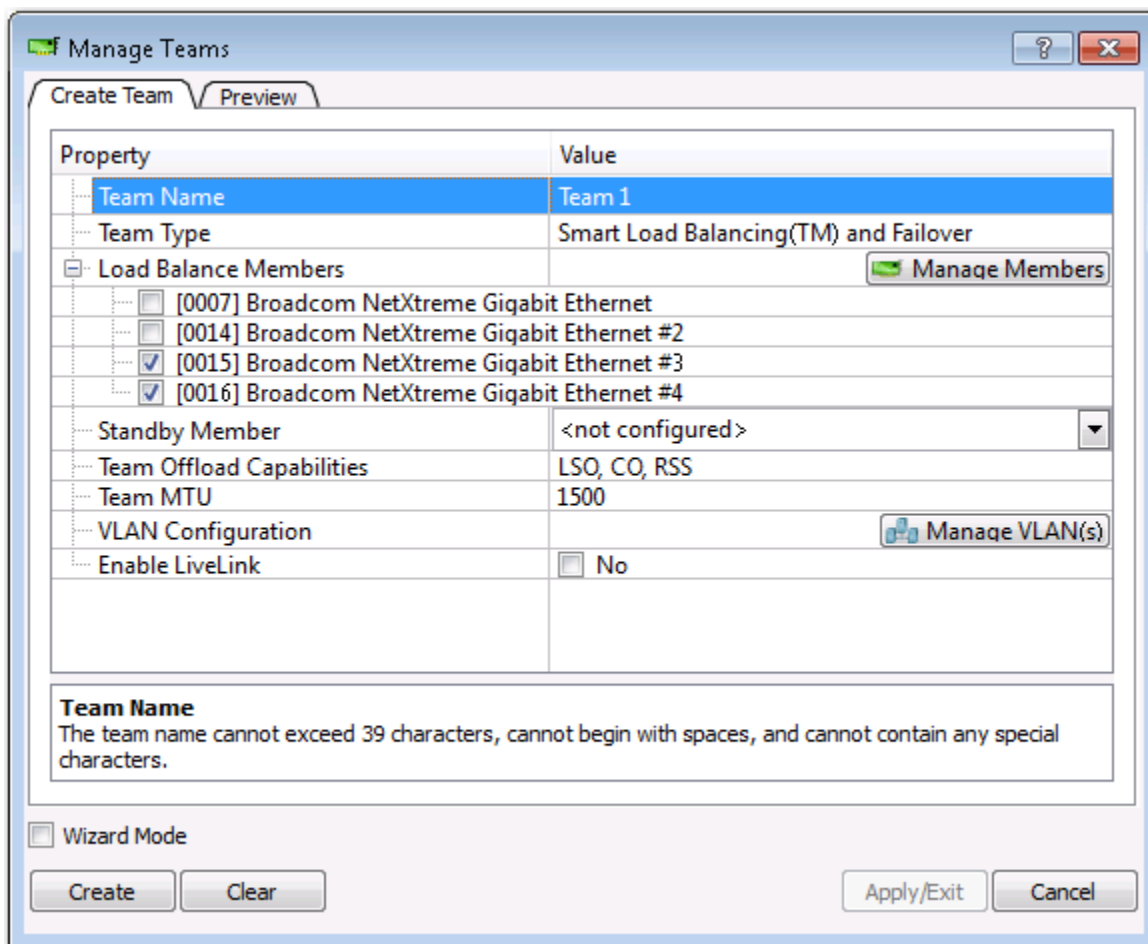
Creating a Team


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

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


 **Note:** If you want to always use Expert Mode to create a team, click **Default to Expert Mode on next start**.

3. Click the **Create Team** tab.



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

4. Click the **Team Name** field to enter a team name.
5. Click the **Team Type** field to select a team type.
6. Assign any available adapter or adapters to the team by selecting the adapter from the **Load Balance Members** list. There must be at least one adapter selected in the **Load Balance Members** list.
7. You can assign any other available adapter to be a standby member by selecting it from the **Standby Member** list.

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

The Large Send Offload (LSO), Checksum Offload (CO), and RSS indicate if the LSO, CO, and/or RSS properties are supported for the team. The LSO, CO, and RSS properties are enabled for a team only when all of the members support and are configured for the feature.



Note: 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.

8. Type the value for **Team MTU**.
9. Click **Create** to save the team information.
10. Repeat steps [Step 4](#) through [Step 9](#) 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.
11. Click **Apply/Exit** to create all the teams you have defined and exit the Manage Teams window.
12. 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 reenabling it, you must first disable and reenabling all team members before you reenabling 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.
1. Configure the team IP address.
 - a. From **Control Panel**, double-click **Network Connections**.
 - b. Right-click the name of the team to be configured, and then click **Properties**.
 - c. On the **General** tab, click **Internet Protocol (TCP/IP)**, and then click **Properties**.
 - d. Configure the IP address and any other necessary TCP/IP configuration for the team, and then click **OK** when finished.

Modifying a Team

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

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


To modify a team

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



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

3. Click the **Edit Team** tab.



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


Adding a VLAN

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

You can define up to 64 VLANs per team (63 VLANs that are tagged and 1 VLAN that is not tagged). VLANs can only be created when all teams members are 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. From the **Teams** menu, select **Add VLAN**.
2. The Welcome screen appears.
3. Click **Expert Mode**.
4. On the **Create Team** tab of the **Manage Teams** window, click **Manage VLAN(s)**.
5. Type the VLAN name, then select the type and ID.
6. Click **Create** to save the VLAN information. As VLANs are defined, they can be selected from the Team Name list, but they have not yet been created.
7. Continue this process until all VLANs are defined, then click **OK** to create them.



8. Click **Yes** when the message is displayed indicating that the network connection will be temporarily interrupted.



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

Viewing VLAN Properties and Statistics and Running VLAN Tests

To view VLAN properties and statistics and to run VLAN tests

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

Deleting a VLAN

The procedure below applies when you are in Expert Mode.

To delete a VLAN

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



Note: If you delete a team, any VLANs configured for that team 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. From the **Teams** menu, select **Edit Team**.
2. Click Expert Mode (to configure LiveLink using the Teaming Wizard, see [Using the Broadcom Teaming Wizard](#)).
3. In the Manage Members window, click the **Edit Team** tab.
4. Select **Enable LiveLink**. The LiveLink Configuration options appear below.
5. It is recommended to accept the default values for **Probe interval** (the number of seconds between each retransmission of a link packet to the probe target) and **Probe maximum retries** (the number of consecutively missed responses from a probe target before a failover is triggered). To specify different values, click the desired probe interval in the **Probe interval (seconds)** list and click the desired maximum number of probe retries in the **Probe maximum retries** list.
6. Set the **Probe VLAN ID** to correspond with the VLAN where the probe target(s) resides. This will apply the appropriate VLAN tag to the link packet based on the shared configuration of the attached switch port(s).



Note: Each LiveLink enabled team can only communicate with Probe Targets on a single VLAN. Also, VLAN ID 0 is equivalent to an untagged network.

7. Select **Probe Target 1** and type the target IP address for one or all probe targets.



Note: Only the first probe target is required. You can specify up to 3 additional probe targets to serve as backups by assigning IP addresses to the other probe targets.

8. Select one of the listed team members and type the member IP address.



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

9. Click **Update**. Repeat these steps for each of the other listed team members.
10. Click **Apply/Exit**.

Saving and Restoring a 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.
5. If a configuration is already loaded, a message is displayed that asks if you want to save your current configuration. Click **Yes** to save the current configuration. Otherwise, the configuration data that is currently loaded is lost.



Note: The team may take a very long time to restore if the team is configured with many VLANs and a static IP address.

Viewing BASP Statistics

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

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

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

Tx. Packet. This is the number of packets transmitted.

Tx. Packet Discarded. This is the number of packets discarded.

Tx. Packet Queued. This is the number of packets queued.

Rx. Packet. This is the number of packets received.

Rx. Packet Discarded. This is the number of packets discarded.

Probes Retried. This is the number of consecutively missed responses from a probe target before a failover is triggered.

Configuring With the Command Line Interface Utility

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

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

BACSCLI is installed when BACS is installed with the installer.

Managing VLANs

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 2: Example of Servers Supporting Multiple VLANs with Tagging




Figure 2 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 5:

Table 5: 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.

Table 5: Example VLAN Network Topology (Cont.)

Component	Description
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). 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 a third-party NIC, if a third-party NIC is a member of a failover team, VLANs cannot be configured for that team.

Troubleshooting BACS

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

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

Solution: This message displays if you try to run a second instance of BACS. If you receive this message but are certain that no instance of BACS is currently running, a previous instance of BACS may have quit unexpectedly. To clear that instance, remove the file “/dev/shm/sem.Global-BACS-{C50398EE-84A7-4bc3-9F6E-25A69603B9C0}.”

Section 7: Teaming

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



Note: See [Broadcom Gigabit Ethernet Teaming Services](#) for detailed information on the following topics:

- Glossary of Terms and Acronyms
- Teaming Concepts
- Software Components
- Hardware Requirements
- Supported Features by Team Type
- Selecting a Team Type
- Teaming Mechanisms
- Architecture
- Types of Teams
- Driver Support by Operating System
- Supported Teaming Speeds
- Teaming and Other Advanced Networking Features
- General Network Considerations
- Application Considerations
- Troubleshooting Teaming Problems
- Frequently-Asked Questions
- Event Log Messages

Overview

Adapter teaming allows you to group network adapters together to function as a team. The benefits of teaming include allowing membership to VLANs, providing load balancing between adapters, and offering fault tolerance. These benefits can be combined such that you can couple the functionality of load balancing for the load balance members and the capability of employing a failover with having the team participate on different VLANs.

Broadcom Advanced Server Program (BASP) is the Broadcom teaming software. For Windows operating systems, BASP is configured through the [Broadcom Advanced Control Suite \(BACS\)](#) utility. For Linux operating systems, teaming is done with channel bonding (see [Teaming with Channel Bonding](#)).

BASP supports four types of load balancing teams:

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

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 adapters are installed in the same system, they can be grouped with up to 16 teams.

Each team can consist of up to eight adapters, with one adapter used as a standby for Smart Load Balancing and Failover (SLB) or SLB (Auto-Fallback Disabled) team types. If traffic is not identified on any of the adapter team member connections due to failure of the adapter, cable, or switch, the load will be distributed to the remaining team members with an active connection. In the event that all primary adapters fail, traffic will be distributed to the standby adapter. Existing sessions are maintained with no impact on the user.

Types of Teams

The available types of teams for the supported operating systems are shown in the following table:

Table 6: Types of Teams

Operating System	Available Types of Teams
Windows Server 2008 and Windows Server 2012	Smart Load Balancing and Failover Link Aggregation (802.3ad) Generic Trunking (FEC/GEC)/802.3ad-Draft Static SLB (Auto-Fallback Disable)
	NOTE: Windows Server 2012 provides built-in teaming support, called NIC Teaming. It is not recommended that users enable teams through NIC Teaming and BASP at the same time on the same adapters.

Table 6: Types of Teams

Operating System	Available Types of Teams
Linux	Team adapters using the bonding kernel module and a channel bonding interface. See your Linux documentation for more information.

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 the Layer 3 protocol (IP); rather, it works with existing Layer 2 and Layer 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) at the switch or port is recommended. This minimizes the downtime due to spanning tree loop determination when failing over. LiveLink mitigates such issues.
- If a team member is linked at 1000 Mbit/s and another team member is linked at 100 Mbit/s, most of the traffic is handled by the 1000 Mbit/s team member.

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.

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.

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.

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.

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.

Table 7: Smart Load Balancing

Operating System	Failover/Fallback — All Broadcom	Failover/Fallback — Multivendor
Protocol	IP	IP
Windows Server 2008	Y	Y
Windows Server 2008 R2	Y	Y
Windows Server 2012	Y	Y
Operating System	Load Balance — All Broadcom	Load Balance — Multivendor
Protocol	IP	IP
Windows Server 2008	Y	Y
Windows Server 2008 R2	Y	Y
Windows Server 2012	Y	Y
Windows Server 2012 R2	Y	Y

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. 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) at the switch is recommended. This minimizes the downtime due to the spanning tree loop determination when failing over. LiveLink mitigates such issues.

LiveLink™ Functionality

LiveLink™ functionality is a feature of BASP that is available only for the Smart Load Balancing™ and Failover type of teaming. The purpose of LiveLink is to detect network connectivity beyond the switch and to route traffic only through team members that have a live link. This function is accomplished through the teaming software (see [Configuring LiveLink for a Smart Load Balancing and Failover and SLB \(Auto-Fallback Disable\) Team](#)). The teaming software periodically probes (issues a link packet from each team member) one or more specified target network adapter(s). The probe target(s) responds when it receives the link packet. If a team member does not detect a response within a specified amount of time after a specified number of retries, 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, refer to Channel Bonding in your Linux documentation.

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.

Section 8: Broadcom Gigabit Ethernet Teaming Services

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

Introduction

- [Teaming Glossary](#)
- [Teaming Concepts](#)
- [Software Components](#)
- [Hardware Requirements](#)
- [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 systems. 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.

Teaming Glossary

Table 8: Teaming Glossary

Item	Definition
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
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 Balance 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.

Table 8: Teaming Glossary

Item	Definition
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
UDP	User Datagram Protocol
WINS	Windows name service
WLBS	Windows Load Balancing Service

Teaming Concepts

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

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 is an example of a Layer 3 protocol. 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. 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 9](#) shows a summary of the teaming types and their classification.

Table 9: Available Teaming Types

Teaming Type	Switch-Dependent (Switch must support specific type of team)	Link Aggregation Control Protocol support is 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)	•	•	•	•

Table 9: Available Teaming Types

Teaming Type	Switch-Dependent (Switch must support specific type of team)	Link Aggregation Control Protocol support is required on the switch	Load Balancing	Failover
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.

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.

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 3](#)). 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. 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 10](#) describes the three software components and their associated files for supported operating systems.

Table 10: Broadcom Teaming Software Component

Software Component	Broadcom Name	Windows	Linux
Miniport Driver	Broadcom Base Driver	b57nd60X.sys	tg3
Intermediate Driver	Broadcom Advanced Server Program (BASP)	Basp.sys	bonding
Configuration User Interface	Broadcom Advanced Control Suite (BACS)	BACS	BACS CLI
NDIS 6 Driver	Windows Vista and later x86 driver	b57nd60x.sys	N/A
	Windows Vista and later x64 driver	b57nd60a.sys	

The Broadcom Advanced Control Suite (BACS) utility is designed to run in 32-bit and 64-bit Windows Server operating system. BACS 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 also includes a number of diagnostics tools such as hardware diagnostics, cable testing, and a network topology test.

Hardware Requirements

- [Ethernet Switch](#)
- [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.

Ethernet Switch

Ethernet switches allow an Ethernet network to be broken into multiple broadcast 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.

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.

Supported Features by Team Type

[Table 11](#) provides a feature comparison across the team types supported by Broadcom NICs. 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 16 teams in a single system. These teams can be any combination of the supported teaming types, but each team must be on a separate network or subnet.

Table 11: Comparison of Team Types

Type of Team	Fault Tolerance	Load Balancing	Switch-Dependent Static Trunking	Switch-Independent Dynamic Link Aggregation (IEEE 802.3ad)
Function	SLB with Standby ^a	SLB	Generic Trunking	Link Aggregation
Number of ports per team (same broadcast domain)	2–8	2–8	2–8	2–8
Number of teams	16	16	16	16
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

Table 11: Comparison of Team Types (Cont.)

Type of Team	Fault Tolerance	Load Balancing	Switch-Dependent Static Trunking	Switch-Independent Dynamic Link Aggregation (IEEE 802.3ad)
Heartbeats to check connectivity	No	No	No	No
Mixed media (adapters with different media)	Yes	Yes	Yes (switch-dependent)	Yes
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

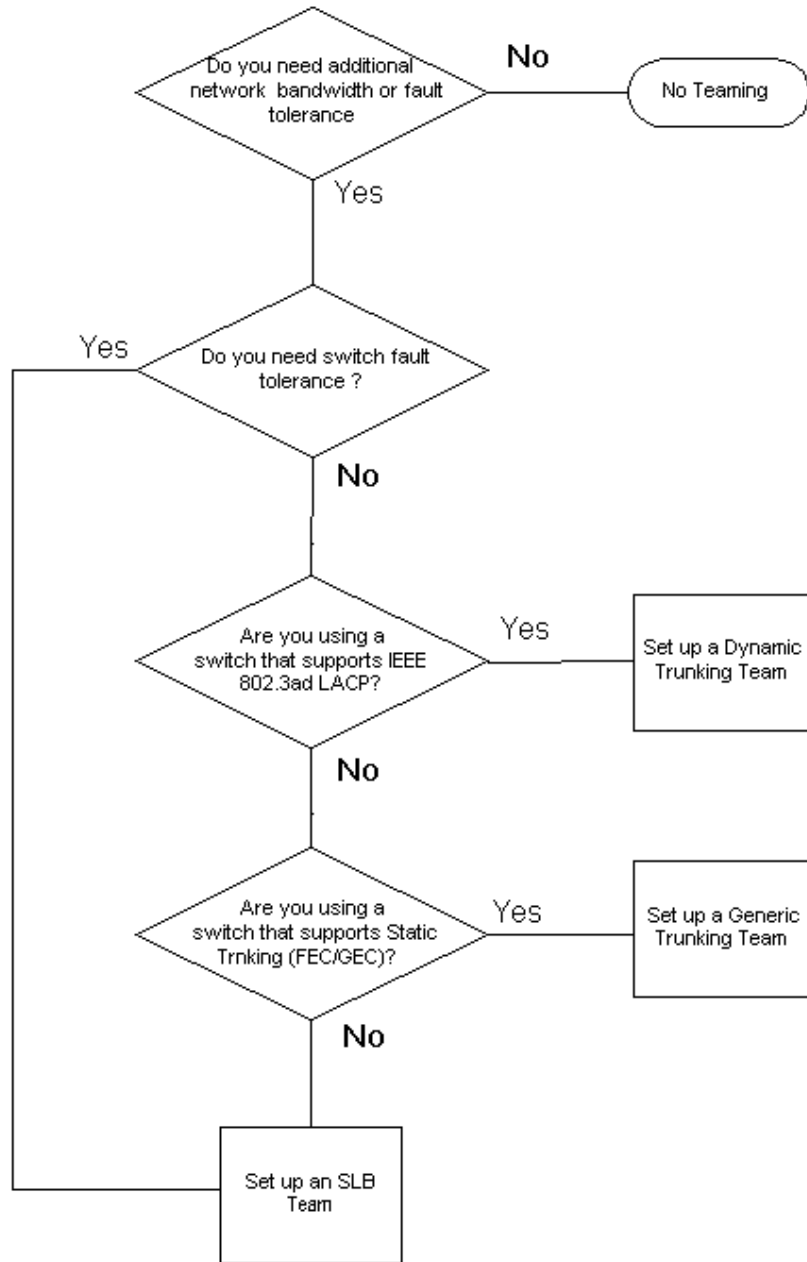
^a SLB with one primary and one standby member.

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

Selecting a Team Type

The following flow chart provides the decision flow when planning for teaming. 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 3](#)).

Figure 3: Process for Selecting a Team Type




Teaming Mechanisms

- [Architecture](#)
- [Driver Support by Operating System](#)
- [Supported Teaming Speeds](#)

Architecture

The Broadcom Advanced Server Program is implemented as an NDIS intermediate driver (see [Figure 4](#)). It operates below protocol stacks such as TCP/IP 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 4: 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.

Driver Support by Operating System

As previously noted, BASP is supported in Windows Server 2008 and 2012 operating system environments.

The various teaming mode features are summarized in the table below.

Table 12: Teaming Mode Features

Features	Windows Support
Smart Load Balancing™	
User interface	BACS ^a
Number of teams	16
Number of adapters per team	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
Failover event	Loss of link or LiveLink event
Failover time	<500 ms
Fallback time	1.5 s ^b (approximate)
LiveLink support	Yes
MAC address	Different
Multivendor teaming	Yes
Generic Trunking	
User interface	BACS
Number of teams	16
Number of adapters per team	8
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
Dynamic Trunking	

Table 12: Teaming Mode Features (Cont.)

Features	Windows Support
User interface	BACS
Number of teams	16
Number of adapters per team	8
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 Broadcom Advanced Control Suite

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

Supported Teaming Speeds

The various link speeds that are supported for each type of team are listed in [Table 13](#). Mixed speed refers to the capability of teaming adapters that are running at different link speeds.

Table 13: Link Speeds in Teaming

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

Teaming and Other Advanced Networking Properties

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

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. The advanced adapter properties and teaming support are listed in [Table 14](#).

Table 14: Advanced Adapter Properties and Teaming Support

Adapter Property	Supported by Teamed Virtual Adapter
Checksum Offload	Yes
IEEE 802.1p QoS Tagging	No
Large Send Offload	Yes ^a
Jumbo Frames	Yes ^b
IEEE 802.1Q VLANs	Yes
Wake on LAN	No
Preboot Execution environment (PXE)	Yes ^c

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

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

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

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; because the intermediate driver is located directly between the protocol layer and the miniport driver, the protocol layer is not able to offload any checksums.

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).

Jumbo Frames

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

The IEEE 802.3ac standard 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 (PXE)

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 Across Switches](#)
- [Spanning Tree Algorithm](#)
- [Teaming with Microsoft NLB/WLBS](#)

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 5](#) is a setup without the interconnect cable in place between the two switches. [Figure 6](#) has the interconnect cable in place, and [Figure 7](#) 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 5: 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 6: Teaming Across Switches With Interconnect





Figure 7 represents a failover event in which the cable is unplugged on the Top Switch port 4. This is a successful failover with all stations pinged each other without loss of connectivity.

Figure 7: 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.

Teaming with Microsoft NLB/WLBS

The SLB mode of 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.

Application Considerations

- [Teaming and Clustering—Microsoft Cluster Software](#)
- [Teaming and Network Backup](#)

Teaming and Clustering—Microsoft Cluster Software

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 subnet. 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.



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

Teaming and Network Backup

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

When you perform network backups in a nonteam 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 fibre 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.

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 8](#) 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:

1. Client-Server ARP cache; which points to the backup server MAC address. This is determined by the Broadcom intermediate driver inbound load balancing algorithm.
2. 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 \(SLB Only\)](#)).

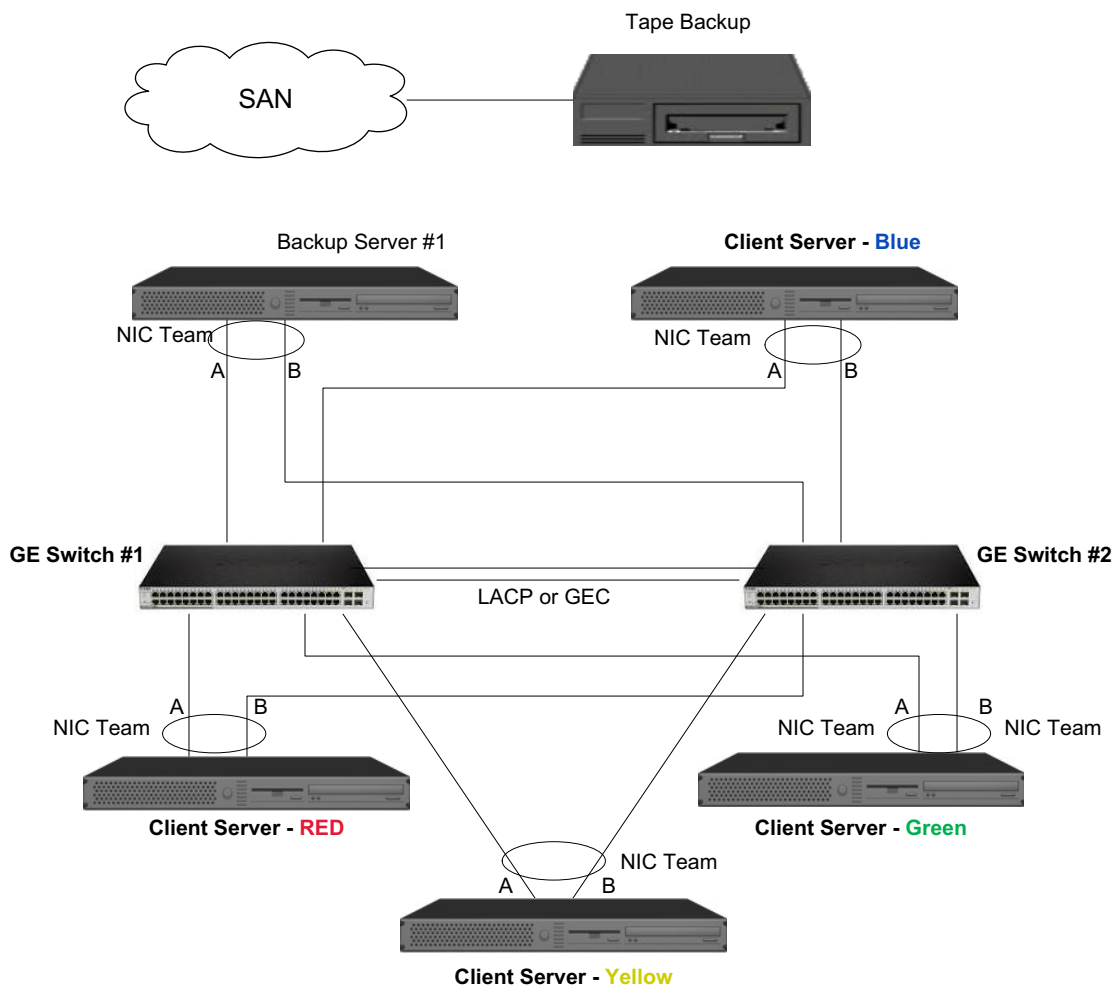
The teamed interface on the backup server transmits a gratuitous address resolution protocol (G-ARP) to Client-Server Red, which in turn, causes the client server ARP cache to get updated with the Backup Server MAC address. The load balancing mechanism within the teamed interface determines the MAC address 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-Server Red, the backup server invokes its SLB algorithm to *inform* Client-Server Blue (thru 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 8](#). 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 8: 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:

1. Mirror all uplink ports from the team at the switch.
2. If the team spans two switches, mirror the interlink trunk as well.
3. Sample all mirror ports independently.
4. 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. It is recommended that, for an SLB team, all adapters be the same link 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 enter a multicast or broadcast address in the Locally Administered Address field.
6. Do not assign a Locally Administered Address on any physical adapter that is a member of a team.
7. Verify that power management is disabled on all physical members of any team (the **Allow the computer to turn off this device to save power** box on the **Power Management** tab in adapter **Properties** should be cleared—see [Setting Power Management Options](#) in “Windows Driver and Application Installation”).
8. Remove any static IP address from the individual physical team members before the team is built.
9. 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.
10. 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.

11. 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.
12. Verify the base (Miniport) and team (intermediate) drivers are from the same release package.
13. Test the connectivity to each physical adapter prior to teaming.
14. Test the failover and fallback behavior of the team before placing into a production environment.
15. When moving from a nonproduction network to a production network, it is strongly recommended to test again for failover and fallback.
16. Test the performance behavior of the team before placing it into a production environment.

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 Ethernet 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 package and are loaded correctly.
3. Check for a valid IP address using the **ipconfig** command for Windows.
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 system.
<hr/>	
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.
<hr/>	
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.
<hr/>	
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, as stated earlier.
<hr/>	
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.
<hr/>	
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 system application level.
<hr/>	
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.
<hr/>	
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.
<hr/>	
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 8 ports can be assigned to a team.

Question:	What is the maximum number of teams that can be configured on the same system?
Answer:	Up to 16 teams can be configured on the same system.

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, click the BASP 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 system and client systems that are backed up be teamed?
Answer:	Because the backup system 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 detect only the loss of link between the teamed port and its immediate link partner. SLB cannot detect link failures on other ports. For more information, see [LiveLink™ Functionality](#).

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

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

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 driver and intermediate driver Windows System Event Log status messages for the Broadcom NetXtreme Gigabit Ethernet adapters are listed in the following section. 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/Miniport)

[Table 15](#) lists the event log messages supported by the base driver, explains the cause for the message, and provides the recommended action.

Table 15: Base Driver Event Log Messages

Message Number	Message	Cause	Corrective Action
1	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	Failed to allocate map registers	The driver cannot allocate map registers from the operating system.	Unload other drivers that may allocate map registers.
3	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	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	The network link is up.	The adapter has established a link.	Informational message only. No action is required.
6	Network controller configured for 10Mb half-duplex link.	The adapter has been manually configured for the selected line speed and duplex settings.	Informational message only. No action is required.

Table 15: Base Driver Event Log Messages (Cont.)

Message Number	Message	Cause	Corrective Action
7	Network controller configured for 10Mb full-duplex link.	The adapter has been manually configured for the selected line speed and duplex settings.	Informational message only. No action is required.
8	Network controller configured for 100Mb half-duplex link.	The adapter has been manually configured for the selected line speed and duplex settings.	Informational message only. No action is required.
9	Network controller configured for 100Mb full-duplex link.	The adapter has been manually configured for the selected line speed and duplex settings.	Informational message only. No action is required.
10	Network controller configured for 1Gb half-duplex link.	The adapter has been manually configured for the selected line speed and duplex settings.	Informational message only. No action is required.
11	Network controller configured for 1Gb full-duplex link.	The adapter has been manually configured for the selected line speed and duplex settings.	Informational message only. No action is required.
12	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.
13	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.
14	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.
15	Driver initialized successfully.	The driver has successfully loaded.	Informational message only. No action is required.
16	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 diagnostics; check that the network cable is good.
18	Unknown PHY detected. Using a default PHY initialization routine.	The driver could not read the PHY ID.	Replace the adapter.
19	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 15: Base Driver Event Log Messages (Cont.)

Message Number	Message	Cause	Corrective Action
20	Driver initialization failed.	Unspecified failure during driver initialization.	Reinstall the driver, update to a newer driver, run Broadcom Advanced Control Suite diagnostics, or replace the adapter.
21	Ethernet@WireSpeed is enabled and could not negotiate maximum link speed.	Potential faulty cable or connection.	Reconnect the cable or change the cable.
22	Unable to install device driver for obsolete network controller for this Operating System.	The latest outbox driver doesn't support obsolete device any more.	Use OS inbox driver, or replace the device with latest one.
256	Not enough contiguous physical memory for coalescing pool.	Driver can't allocate enough shared memory for coalescing packet buffers.	Remove/Disable other adapter from the system or increase system memory.

Intermediate Driver (Virtual Adapter/Team)

Table 16 lists the event log messages supported by the intermediate driver, explains the cause for the message, and provides the recommended action.

Table 16: Intermediate Driver Event Log Messages

System Event Message Number	Message	Cause	Corrective Action
1	Unable to register with NDIS.	The driver cannot register with the NDIS interface.	Unload other NDIS drivers.
2	Unable to instantiate the management interface.	The driver cannot create a device instance.	Reboot the operating system.
3	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>Blf</i> .
4	Broadcom Advanced Server Program Driver has started.	Another driver has created a conflicting device name.	Informational message only. No action is required.
5	Broadcom Advanced Server Program Driver has stopped.	The driver has stopped.	Informational message only. No action is required.
6	Could not allocate memory for internal data structures.	The driver cannot allocate memory from the operating system.	Close running applications to free memory
7	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.
8	Successfully bind to adapter.	The driver successfully opened the physical adapter.	Informational message only. No action is required.

Table 16: Intermediate Driver Event Log Messages (Cont.)

System Event Message Number	Message	Cause	Corrective Action
9	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.
10	Network adapter is connected.	The physical adapter is connected to the network (it has established link).	Informational message only. No action is required.
11	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.
12	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.
13	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.
14	Network adapter is enabled via management interface.	The driver has successfully enabled a physical adapter through the management interface.	Informational message only. No action is required.
15	Network adapter is disabled via management interface.	The driver has successfully disabled a physical adapter through the management interface.	Informational message only. No action is required.
16	Network adapter is activated and is participating in network traffic.	A physical adapter has been added to or activated in a team.	Informational message only. No action is required.
17	Network adapter is de-activated and is no longer participating in network traffic.	The driver does not recognize the installed adapter.	Informational message only. No action is required.

Section 9: iSCSI Protocol and Broadcom Boot Agent Software

- [iSCSI Boot](#)
- [Broadcom Boot Agent Driver Software](#)

iSCSI Boot

Broadcom NetXtreme 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 both Windows and Linux operating systems, iSCSI boot can be configured to boot with the general parameters shown in [Table 17](#).

Supported Operating Systems for iSCSI Boot

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

- Windows Server operating system
- Enterprise Linux distribution

iSCSI Boot Setup

This section applies to BIOS-mode iSCSI boot. Refer to the system documentation for UEFI iSCSI boot setup.

iSCSI boot is not supported in BIOS mode when local storage when local storage (especially RAID) is present due to EBDA memory constraints.

The iSCSI boot setup consists of:

- [Configuring the iSCSI Target](#)
- [Configuring iSCSI Boot Parameters](#)
- [Preparing the iSCSI Boot Image](#)
- [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.
6. 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 iSCSI Boot Parameters

Configure the Broadcom iSCSI boot software for either static or dynamic configuration. Refer to [Table 17](#) for configuration options available from the General Parameters screen.

[Table 17](#) 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 17: 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.
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.

Table 17: Configuration Options (Cont.)

Option	Description
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.

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**.



Note: For platforms on which the boot protocol is set via the BIOS, see your system documentation for more information.

4. Select **iSCSI Boot Configuration** from **Main Menu**.



Note: If iSCSI boot firmware is not programmed in the NetXtreme network adapter, selecting **iSCSI Boot Configuration** will not have any effect.

iSCSI Boot Configuration

- [Static iSCSI Boot Configuration](#)
- [Dynamic iSCSI Boot Configuration](#)

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 17](#).

To configure the iSCSI boot parameters using static configuration

1. From the **General Parameters Menu** screen, set the following:
 - **TCP/IP parameters via DHCP:** Disabled. (For IPv4.)
 - **IP Autoconfiguration:** Disabled. (For IPv6)

- **iSCSI parameters via DHCP:** Disabled
 - **CHAP Authentication:** Disabled
 - **Boot to iSCSI target:** Disabled
 - **DHCP Vendor ID:** BRCM ISAN
 - **Link Up Delay Time:** 0
 - **Use TCP Timestamp:** Enabled (for some targets such as the Dell/EMC AX100i, it is necessary to enable **Use TCP Timestamp**)
 - **Target as First HDD:** Disabled
 - **LUN Busy Retry Count:** 0
 - **IP Version:** IPv6. (For IPv6)
2. Select **ESC** to return to the **Main** menu.
 3. From the **Main** menu, select **Initiator Parameters**.
 4. From the **Initiator Parameters** screen, type values for the following:
 - IP Address (unspecified IPv4 and IPv6 addresses should be “0.0.0.0” and “::”, respectively)
 - 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.

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, 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
8. Select **ESC** to return to the **Main** menu.
9. Select **ESC** and select **Exit and Save Configuration**.
10. Select **F4** to save your MBA configuration.

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 17](#).

**Notes:**

- 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 **General Parameters Menu** screen, set the following:
 - **TCP/IP parameters via DHCP:** Enabled. (For IPv4.)
 - **IP Autoconfiguration:** Enabled. (For IPv6)
 - **iSCSI parameters via DHCP:** Enabled
 - **CHAP Authentication:** Disabled
 - **Boot to iSCSI target:** Disabled
 - **DHCP Vendor ID:** BRCM ISAN
 - **Link Up Delay Time:** 0
 - **Use TCP Timestamp:** Enabled (for some targets such as the Dell/EMC AX100i, it is necessary to enable **Use TCP Timestamp**)
 - **Target as First HDD:** Disabled
 - **LUN Busy Retry Count:** 0
 - **IP Version:** IPv6. (For IPv6)
2. Select **ESC** to return to the **Main** menu.



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

3. Select **Exit and Save Configurations**.

Enabling CHAP Authentication

Ensure that CHAP authentication is enabled on the target.

To enable CHAP authentication

1. From the **General Parameters** screen, set **CHAP Authentication** to Enabled.
2. 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)
3. Select **ESC** to return to the **Main** menu.
4. From the **Main** menu, select **1st Target Parameters**.
5. 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)
6. Select **ESC** to return to the **Main** menu.
7. 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 18: 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 19: 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 as 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 20: DHCP Option 17 Suboption Definition

<i>Suboption</i>	<i>Definition</i>
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 20](#), 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.

Preparing the iSCSI Boot Image

- [Windows Server 2008 R2 and SP2 iSCSI Boot Setup](#)
- [Windows Server 2012 iSCSI Boot Setup](#)
- [Linux iSCSI Boot Setup](#)

Windows Server 2008 R2 and SP2 iSCSI Boot Setup

Windows Server 2008 R2 and Windows Server 2008 SP2 support iSCSI booting. The following procedure references Windows Server 2008 R2 but is common to both the Windows Server 2008 R2 and SP2.

Required CD/ISO image:

- Windows Server 2008 R2 x64 with the Broadcom drivers injected. Refer to the Microsoft knowledge base topic KB974072 at support.microsoft.com.



Note: Refer to the *silent.txt* file for the specific driver installer application for instructions on how to extract the individual Windows NetXtreme drivers.

Other software required:

- Bindview.exe (Windows Server 2008 R2 only; see KB976042)

Procedure:

1. Remove any local hard drives on the system to be booted (the “remote system”).
2. Load the latest Broadcom MBA and iSCSI boot images onto NVRAM of the adapter.
3. Configure the BIOS on the remote system to have the Broadcom MBA as the first bootable device, and the DVDROM as the second device.
4. Configure the iSCSI target to allow a connection from the remote device. Ensure that the target has sufficient disk space to hold the new O/S installation.
5. Boot up the remote system. When the Preboot Execution Environment (PXE) banner displays, press **Ctrl+S** to enter the PXE menu.
6. At the PXE menu, set **Boot Protocol** to **iSCSI**.
7. Enter the iSCSI target parameters.
8. In General Parameters, set the **Boot to Target** parameter to **One-Time Disabled**.
9. Save the settings and reboot the system.
The remote system should connect to the iSCSI target and then boot from the DVDROM device.
10. Boot to DVD and begin installation.
11. Answer all the installation questions appropriately (specify the Operating System you want to install, accept the license terms, etc.).
When the **Where do you want to install Windows?** window appears, the target drive should be visible. This is a drive connected via the iSCSI boot protocol, located in the remote iSCSI target.
12. Select **Next** to proceed with Windows Server 2008 R2 installation.

A few minutes after the Windows Server 2008 R2 DVD installation process starts, a system reboot will follow. After the reboot, the Windows Server 2008 R2 installation routine should resume and complete the installation.

13. Following another system restart, check and verify that the remote system is able to boot to the desktop.
14. After Windows Server 2008 R2 is booted up, load the driver and run Bindview.exe.
 - a. Select **All Services**.
 - b. Under **WFP Lightweight Filter** you should see **Binding paths** for the AUT. Right-click and disable them. When done, close out of the application.
15. Verify that the OS and system are functional and can pass traffic by pinging a remote system's IP, etc.

Windows Server 2012 iSCSI Boot Setup

Windows Server 2012 supports iSCSI booting and installation. Broadcom requires the use of a “slipstream” DVD with the latest Broadcom drivers injected. See the Microsoft knowledge base topic KB974072 at support.microsoft.com.

The following procedure prepares the image for installation and booting:

1. Remove any local hard drives on the system to be booted (the “remote system”).
2. Load the latest Broadcom MBA and iSCSI boot images into the NVRAM of the adapter.
3. Configure the BIOS on the remote system to have the Broadcom MBA as the first bootable device and the DVDROM as the second device.
4. Configure the iSCSI target to allow a connection from the remote device. Ensure that the target has sufficient disk space to hold the new O/S installation.
5. Boot up the remote system. When the Preboot Execution Environment (PXE) banner displays, press **Ctrl+S** to enter the PXE menu.
6. At the PXE menu, set **Boot Protocol** to **iSCSI**.
7. Enter the iSCSI target parameters.
8. In General Parameters, set the **Boot to Target** parameter to **One-Time Disabled**.
9. Save the settings and reboot the system.

The remote system should connect to the iSCSI target and then boot from the DVDROM device.
10. Boot from DVD and begin installation.
11. Answer all the installation questions appropriately (specify the Operating System you want to install, accept the license terms, etc.).

When the **Where do you want to install Windows?** window appears, the target drive should be visible. This is a drive connected via the iSCSI boot protocol, located in the remote iSCSI target.
12. Select **Next** to proceed with Windows 2012 installation.

A few minutes after the Windows 2012 DVD installation process starts, a system reboot will occur. After the reboot, the Windows 2012 installation routine should resume and complete the installation.
13. Following another system restart, check and verify that the remote system is able to boot to the desktop.

14. After Windows 2012 boots to the OS, Broadcom recommends running the driver installer to complete the Broadcom driver and application installation.

Linux iSCSI Boot Setup

Linux iSCSI boot is supported on Red Hat Enterprise Linux 6.9, 7.1 and later and SUSE Linux Enterprise Server (SLES) 11 SP4, or SLES 12 SP2 and later.

1. For driver update, obtain the latest Broadcom Linux driver CD.
2. Configure the iSCSI Boot Parameters for DVD direct install to target by disabling the Boot from target option on the network adapter.
3. Change the boot order as follows:
 - a. Boot from the network adapter.
 - b. Boot from the CD/DVD driver.
4. Reboot the system.
5. System will connect to iSCSI target, then will boot from CD/DVD drive.
6. Follow the corresponding OS instructions.
 - a. RHEL 6.9 — Type “linux dd” at “boot:” prompt and press enter
 - b. SuSE 11.X — Choose **installation** and type **withiscsi=1 netsetup=1** at the boot option. If driver update is desired, choose **YES** for the F6 driver option.
7. If driver update is desired, follow the instructions to load the driver CD; otherwise skip this step.
8. At the “networking device” prompt, choose the desired network adapter port and press **OK**.
9. At “configure TCP/IP prompt”, configure the way the system acquire IP address and press **OK**.
10. If static IP was chosen, you need to enter IP information for iscsi initiator.
11. (RHEL) Choose to “skip” media testing.
12. Continue installation as desired. A drive will be available at this point. After file copying is done, remove CD/DVD and reboot the system.
13. When the system reboots, enable “boot from target” in iSCSI Boot Parameters and continue with installation until it is done.

At this stage, the initial installation phase is complete. The rest of the procedure pertains to creating a new customized initrd for any new components update:

14. Update iscsi initiator if desired. You will first need to remove the existing initiator using **rpm -e**.
15. Make sure all runlevels of network service are on:

```
chkconfig network on
```
16. Make sure 2,3 and 5 runlevels of iscsi service are on.

```
chkconfig -level 235 iscsi on
```
17. For Red Hat 6.0, make sure Network Manager service is stopped and disabled.
18. Install iscsiui0 if desired (not required for SuSE 10).
19. Install linux-nx2 package if desired.

20. Install bibt package.
21. Remove ifcfg-eth*.
22. Reboot.
23. For SUSE 11.1, follow the remote DVD installation workaround shown below.
24. After the system reboots, log in, change to the /opt/bcm/bibt folder, and run iscsi_setup.sh script to create the initrd image.
25. Copy the initrd image to the /boot folder.
26. Change the grub menu to point to the new initrd image.
27. To enable CHAP, you need to modify iscsid.conf (Red Hat only).
28. Reboot and change CHAP parameters if desired.
29. Continue booting into the iSCSI Boot image and select the image you created.
30. For IPv6, you can now change the IP address for both the initiator and the target to the desired IPv6 address in the NVRAM configuration.

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 or Linux over the network and operate as if it were a local disk drive.

1. Reboot the server.
2. Select **CTRL+S**.
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](#)).

Other iSCSI Boot Considerations

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

Changing the Speed & Duplex Settings in Windows Environments

Booting via the NDIS path is supported. The Speed & Duplex settings can be changed using the BACS management utility for iSCSI boot via the NDIS path.

Locally Administered Address

A user-defined MAC address assigned through the Locally Administered Address property of the Advanced section of the BACS Configurations tab is not supported on iSCSI boot-enabled devices.

Virtual LANs

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

Troubleshooting iSCSI Boot

The following troubleshooting tips are useful for iSCSI boot.

Problem: The Broadcom iSCSI Crash Dump utility will not work properly to capture a memory dump when the link speed for iSCSI boot is configured for 10 Mbps or 100 Mbps.

Solution: The iSCSI Crash Dump utility is supported when the link speed for iSCSI boot is configured for 1 Gbps. 10 Mbps or 100 Mbps is not supported.

Problem: An iSCSI target is not recognized as an installation target when you try to install Windows Server 2008 by using an IPv6 connection.

Solution: This is a known third-party issue. See Microsoft Knowledge Base KB 971443, <http://support.microsoft.com/kb/971443>.

Problem: The iSCSI configuration utility will not run.

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

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.

Problem: Unable to update inbox driver if a non-inbox hardware ID present.

Solution: Create a custom slipstream DVD image with supported drivers present on the install media.

Broadcom Boot Agent Driver Software

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

Overview

Broadcom NetXtreme Gigabit Ethernet adapters support Preboot Execution Environment (PXE), Remote Program Load (RPL), iSCSI boot, and Bootstrap Protocol (BootP). Multi-Boot Agent (MBA) is a software module that allows your networked system to boot with the images provided by remote systems 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/system environment. A network consists of one or more boot systems that provide boot images to multiple systems 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.
- **Windows Deployment Service (WDS).** Broadcom PXE client can be used to install Windows operating systems using WDS.

Setting Up MBA in a Client Environment

Use the following procedure for add-in NICs. For LOMs, refer to your computer's system guide.

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

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

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.

Using CCM

1. Restart your system.
2. Press **CTRL+s** within 4 seconds after you are prompted to do so. A list of adapters displays.
 - a. Select the adapter to configure and press **Enter**. The Main Menu displays.
 - b. Select **MBA Configuration** to display the MBA Configuration menu.

```

Comprehensive Configuration Management v7.8.10
Copyright (C) 2000-2013 Broadcom Corporation
All rights reserved.

MBA Configuration Menu

Option ROM           : Enabled
Boot Protocol        : iSCSI
Boot Strap Type      : Auto
Hide Setup Prompt    : Disabled
Setup Key Stroke     : Ctrl-S
Banner Message Timeout : 5 Seconds
Link Speed           : 1Gbps
Pre-boot Wake On LAN : Disabled
VLAN Mode            : Disabled
VLAN ID              : 1
Boot Retry Count     : 0

Enable/Disable Option ROM
[←|→|Enter|Space]:Toggle Value; [↑|↓]:Next Entry; [ESC]:Quit
Current Adapter:Primary, Bus=03 Device=00 Func=00, MAC=00:10:18:A7:19:10

```

- 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: For some, but not all, iSCSI boot-capable LOMs, the boot protocol is set via the BIOS. See your system documentation for more information.



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.

- 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.
- Press **F4** to save your settings.
- Press **ESC** when you are finished.

Using UEFI

- Restart your system.
- Enter the System Setup or Device Setting configuration menu.
- Select the device on which you want to change MBA settings.
- Select **MBA Configuration Menu**.

5. Use the drop-down menu to select the boot protocol of choice, if boot protocols other than Preboot Execution Environment (PXE) are available. If available, other boot protocols include iSCSI and Bootstrap Protocol (BOOTP).



Note: For iSCSI boot-capable LOMs, the boot protocol is set via the BIOS. See your system documentation for more information.

6. 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.
7. Select **Back** to go to Main menu
8. Select **Finish** to save and exit.

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.

Section 10: Manageability

- [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 platforms such as Windows Server 2008. Broadcom will support CIM on Windows Server 2008, 2012, and 2012 R2 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 both Broadcom and other vendors' controllers. `BRCM_ExtraCapacityGroup` class provides team configuration for the Broadcom Advanced Server Program (BASP) 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 2008, 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 Windows Server 2008 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
- baspmat.mib
- basptrap.mib

BASP Extensible-Agent

The Broadcom NetXtreme Gigabit Ethernet Controller Extended Information SNMP extensible-agent, bcmif.dll, is designed for Windows Server SNMP service.

The extensible-agent allows the SNMP manager software to actively monitor the configurations of the Broadcom NetXtreme 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 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 Broadcom NetXtreme adapter installation CD.

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

Section 11: Specifications

- [10/100/1000BASE-T Cable Specifications](#)
- [Performance Specifications](#)

10/100/1000BASE-T Cable Specifications

Table 21: 10/100/1000BASE-T Cable Specifications

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

¹1000BASE-T signaling requires 4 twisted pairs of Category 5 balanced cabling, as specified in ISO/IEC 11801:1995 and ANSI/EIA/TIA-568-A (1995) and tested for additional performance using testing procedures defined in TIA/EIA TSB95.

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

Performance Specifications

Table 22: Performance Specifications

Feature	Specification
PCI Express™ Type Controllers (BCM57XX Controllers)	
PCI Express Interface	x1, x2, x4 link width
PCI Express Aggregated Bandwidth (transmit and receive)	2.5 Gbps or 5.0 Gbps
10/100/1000BASE-T	10/100/1000 Mbps (full-duplex)

Section 12: Troubleshooting

- [Hardware Diagnostics](#)
- [Troubleshooting Checklist](#)
- [Checking for Network Link and Activity](#)
- [Checking if Current Drivers Are Loaded](#)
- [Running a Cable Length Test](#)
- [Testing Network Connectivity](#)
- [Broadcom Boot Agent](#)
- [Broadcom Advanced Server Program \(BASP\)](#)
- [Kernel Debugging over Ethernet](#)
- [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 Windows environments, see [Running Diagnostic Tests](#)).

BACS Diagnostic Tests Failures

If any of the following tests fail while running the diagnostic tests from the [Running Diagnostic Tests](#) tab in BACS, this may indicate a hardware issue with the NIC or LOM that is installed in the system.

- Control Registers
- MII Registers
- EEPROM
- Internal Memory
- On-Chip CPU
- Interrupt
- Loopback - MAC
- Loopback - PHY
- Test LED

Below are troubleshooting steps that may help correct the failure.

1. Remove the failing device and reseal it in the slot, ensuring the card is firmly seated in the slot from front to back.
2. Rerun the test.
3. If the card still fails, replace it with a different card of the same model and run the test. If the test passes on the known good card, contact your hardware vendor for assistance on the failing device.

4. Power down the machine, remove AC power from the machine, and then reboot the system.
5. Remove and re-install the diagnostic software.
6. Contact your hardware vendor.

BACS Network Test Failures

Typically, the BACS [Testing the Network](#) failures are the result of a configuration problem on the network or with the IP addresses. Below are common steps when troubleshooting the network.

1. Verify that the cable is attached and you have proper link.
2. Verify that the drivers are loaded and enabled.
3. Replace the cable that is attached to the NIC/LOM.
4. Verify that the IP address is assigned correctly using the command "ipconfig" or by checking the OS IP assigning tool.
5. Verify that the IP address is correct for the network to which the adapter(s) is connected.

Troubleshooting Checklist



Caution! Before you open the case of your system, review [Safety Precautions](#).

The following checklist provides recommended actions to take to resolve problems installing the Broadcom NetXtreme Gigabit Ethernet 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 [Installing the Hardware](#). 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 system 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 for Network Link and Activity

See [Testing Network Connectivity](#) or [Viewing Adapter Information](#) to check the state of the network link and activity as indicated by the port LEDs.

Checking if Current Drivers Are Loaded

Windows

See [Viewing Adapter Information](#) to view useful information about the adapter, its link status, and network connectivity.

Linux

To verify that the TG3 Linux driver is loaded properly, run:

```
lsmod | grep tg3
```

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.

Table 23: Linux Driver

<i>Module</i>	<i>Size</i>	<i>Used by</i>
TG3	<i>size</i>	<i>n</i>

Running a Cable Length Test

In Windows environments, a cable length test can be run. See [Analyzing Cables](#) for information about running a cable length test.

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 both sides are configured for auto-negotiation.

Windows

Use the ping command to determine if the network connection is working.



Note: Network connectivity can also be tested using the [Testing the Network](#) feature in Broadcom Advanced Control Suite 2.

1. Verify that the drivers are loaded and enabled.
2. Verify that the cable is attached and you have a proper link.
3. Click **Start**, and then click **Run**.
4. Type **cmd** in the **Open** box, and then click **OK**.
5. Type **ipconfig /all** to view the network connection to be tested.
6. Verify that the IP address is correct for the network to which the adapter(s) is connected.
7. Type **ping IP address**, and then press ENTER.

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

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. Go to [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 the network connection is working or not.

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

Solution: To physically remove a teamed NIC from a system, you must first delete the NIC from the team. Not doing this before shutting down could result in breaking the team on a subsequent reboot, which may result in unexpected team behavior.

Problem: The teaming changes I made when I modified my team using INETCFG did not take effect.

Solution: When you modify a team using INETCFG, you may need to reboot after reinitialization for the changes to the team to take effect.

Kernel Debugging over Ethernet

Problem: When attempting to perform kernel debugging over an Ethernet network on a Windows 8.0 or Windows Server 2012 system, the system will not boot. This problem may occur with some adapters on systems where the Windows 8.0 or Windows Server 2012 OS is configured for UEFI mode. You may see a firmware error on screen, indicating that a Non Maskable Interrupt exception was encountered during the UEFI pre-boot environment.

Solution: Refer to the Microsoft knowledge base topic number 2920163, "[Non Maskable Interrupt error during boot on a system which has been configured for kernel debugging over Ethernet](#)".

Miscellaneous

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.

Section 13: Regulatory Information

- [FCC Class B Notice](#)
- [VCCI Class B Notice](#)
- [CE Notice](#)
- [Canadian Regulatory Information \(Canada Only\)](#)
- [MIC Notice \(Republic of Korea Only\)](#)
- [BSMI](#)

FCC Class B Notice

Broadcom NetXtreme Gigabit Ethernet Controller
BCM95721A211
BCM95722A2202

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 you change or modify the adapter without permission of Broadcom, you may void your authority to operate the equipment.

Broadcom Corporation
190 Mathilda Place
Sunnyvale, California 94086 USA

VCCI Class B Notice

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、受信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。

CE Notice

BЪЛГАРСКИ Bulgarian	<p>Този продукт отговаря на 2006/95/EC (Нисковолтова директива), 2004/108/EC (Директива за електромагнитна съвместимост) и измененията на Европейския съюз.</p> <p>Европейски съюз, Клас B</p> <p>Това устройство на Broadcom е класифицирано за използване в типичната за Клас B жилищна среда.</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>
ČESKY Czech	<p>Bylo ustanoveno, že tento produkt splňuje směrnici 2006/95/EC (nizkonapěťová směrnice), směrnici 2004/108/EC (směrnice EMC) a dodatky Evropské unie.</p> <p>Evropská unie, třída B</p> <p>Toto zařízení společnosti Broadcom je klasifikováno pro použití v obvyklém prostředí domácnosti (třída B).</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>
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>Den Europæiske Union, Klasse B</p> <p>Denne Broadcom-enhed er klassificeret til anvendelse i et typisk Klasse B-hjemligt miljø.</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>
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>Europese Unie/Klasse B</p> <p>Dit Broadcom-apparaat is geclassificeerd voor gebruik in een typische klasse B woonomgeving.</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>
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>European Union, Class B</p> <p>This Broadcom device is classified for use in a typical Class B domestic environment.</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>
EESTLANE Estonian	<p>Antud toode vastab direktiividele 2006/95/EU (Madalpinge direktiiv), 2004/108/EU (EMC direktiiv) ja ELi parandustele.</p> <p>Euroopa Liit, Klass B</p> <p>Antud Broadcom toode on klassifitseeritud kasutamiseks tüüpilises B-klassi kodus keskkonnas.</p> <p>Vastavalt ülaltoodud direktiividele ja standarditele on koostatud „Vastavusdeklaratsioon“, mis on arvel ettevõttes Broadcom Corporation, 190 Mathilda Place, Sunnyvale, California 94086, USA.</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, vaatimuksia.</p> <p>Euroopan unioni, luokka B</p> <p>Tämä Broadcom-laite on luokiteltu käytettäväksi tyypillisessä luokan B kotiympäristössä.</p> <p>Yllä mainittujen direktiivien ja standardien mukainen vaatimustenmukaisuusvakuutus on tehty, ja sitä säilyttää Broadcom Corporation, 190 Mathilda Place, Sunnyvale, California 94086, USA.</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>Union européenne, classe B</p> <p>Cet appareil Broadcom est classé pour une utilisation dans un environnement résidentiel classique (classe B).</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>

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>Europäische Union, Klasse B Dieses Gerät von Broadcom ist für die Verwendung in einer typisch häuslichen Umgebung der Klasse B vorgesehen. 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>
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>Ευρωπαϊκή Ένωση, Κατηγορία Β Αυτή η συσκευή Broadcom είναι κατάλληλη για χρήση σε ένα σύνθετες οικιακό περιβάλλον κατηγορίας Β. Μία «Δήλωση Συμμόρφωσης» σύμφωνα με τις προηγούμενες οδηγίες και πρότυπα υπάρχει και είναι αρχειοθετημένη στο Broadcom Corporation, 190 Mathilda Place, Sunnyvale, California 94086, USA.</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>Európai Unió, „B” osztály Ez a Broadcom eszköz „B” osztályú besorolást kapott, tipikus lakossági környezetben való használatra alkalmas. Az előbbiakban ismertetett irányelvek és szabványok szellemében „Megfelelőségi nyilatkozat” készült, amely az irországi Broadcom Corporation, 190 Mathilda Place, Sunnyvale, California 94086, USA.</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>União Europeia, Classe B Este dispositivo Broadcom está classificado para utilização num ambiente doméstico típico Classe B. 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>
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>Unione Europea, Classe B Il presente dispositivo Broadcom è classificato per l'uso nel tipico ambiente domestico di Classe B. 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>
LATVISKS Latvian	<p>Sis 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>Eiropas Savienība, klase B Šī firmas Broadcom ražotā ierīce ir atzīta par derīgu darbam B klasei atbilstošos mājas apstākļos. “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>
Lithuanian	<p>Buvo nustatyta, kad šis produktas atitinka direktyvą 73/23/EEB (žemos įtampos direktyvą), 89/336/EEB (elektromagnetinio suderinamumo direktyvą) ir Europos Sąjungos pataisas.</p> <p>Europos Sąjunga, B klasė Šis „Broadcom“ prietaisas yra klasifikuotas naudoti įprastose B klasės gyvenamosiose aplinkose. Atitikties deklaracija pagal visas galiojančias direktyvas ir standartus yra sudaryta ir saugoma įrašyta faile Broadcom Corporation, 190 Mathilda Place, Sunnyvale, California 94086, USA.</p>

Maltese	<p>Gie stabbilit li dan il-prodott hu konformi ma' 2006/95/KE (Direttiva dwar il-Vultaġġ Baxx), 2004/108/KE (Direttiva EMC), u emendi ta' l-Unjoni Ewropea.</p> <p>Unjoni Ewropea, Klassi B</p> <p>Dan it-tagħmir Broadcom hu kklassifikat għall-użu f'ambjent residenzjali tipiku ta' Klassi B. 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>
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>Unia Europejska, klasa B</p> <p>Niniejsze urządzenie firmy Broadcom zostało zakwalifikowane do klasy B, do użytku w typowych środowiskach domowych.</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>
ROMAN 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>Uniunea Europeană, Clasa B</p> <p>Acest echipament Broadcom este clasificat pentru utilizare într-un mediu casnic tipic de Clasă B. Conform directivelor și standardelor de mai sus, a fost emisă o „Declarația de Conformitate”, arhivată la sediul Broadcom Corporation, 190 Mathilda Place, Sunnyvale, California 94086, USA.</p>
SLOVENSKY Slovakian	<p>Tento výrobok vyhovuje požiadavkám smernice 2006/95/EC (smernica o nízkom napätí), 2004/108/EC (smernica o elektromagnetickej kompatibilitate) a neskorším zmenám a doplnkom Európskej.</p> <p>Európska únia, Trieda B</p> <p>Toto zariadenie Broadcom triedy B je určené pre domáce prostredie.</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>
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>Evropska unija, razred B</p> <p>Ta Broadcomova naprava je razvrščena za uporabo v značilnem bivalnem okolju razreda B.</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>
ESPAÑOL Spanish	<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 emiendas de la Unión Europea.</p> <p>Unión Europea, Clase B</p> <p>Este dispositivo Broadcom está clasificado para ser utilizado en un entorno doméstico convencional de Clase B.</p> <p>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.</p>
SVENSK Swedish	<p>Denna produkt överensstämmer med EU-direktivet 2006/95/EC (lågspänningsdirektivet), 2004/108/EC (EMC direktivet), och andra ändringar enligt den Europeiska unionen.</p> <p>Europeiska unionen, klass B</p> <p>Den här Broadcom-enheten är klassificerad för användning i vanlig klass B-bostadsmiljö.</p> <p>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.</p>
TURK Turkish	<p>Bu ürünün 2006/95/EC (Düşük Voltaj Direktifi), 2004/108/EC (EMC Direktifi), ve Avrupa Birliği'nin ilavelerine uygun olduğu belirlenmiştir.</p> <p>Avrupa Birliği, B Sınıfı</p> <p>Bu Broadcom cihazı, tipik bir B sınıfı, ev içi ortamda kullanılmak üzere sınıflandırılmıştır. Yukarıda belirtilen direktifler ve standartlara uygun olarak, bir "Uygunluk Beyanı" hazırlanmıştır, ve Broadcom Corporation, 190 Mathilda Place, Sunnyvale, California 94086, USA.</p>

Canadian Regulatory Information (Canada Only)

Industry Canada, Class B

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, classe B

Cet appareil numérique de la classe B est conforme à la norme NMB-003 du Canada.

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.

MIC Notice (Republic of Korea Only)

B CLASS Device

Broadcom NetXtreme Gigabit Ethernet Controller
 BCM95721A211
 BCM95722A2202

기종별	사용자안내문
B급 기기 (가정용)	이 기기는 가정용으로 전자파적합등록을 한 기기로서 주거지역에서는 물론 모든 지역에서 사용할 수 있습니다.



- 1. 기기의 명칭(모델명) : BCM95721A211
- 2. 인증번호 : E-G021-04-2613(B)
- 3. 인증받은 자의 상호 : Broadcom
- 4. 제조년월일 : 5/12/2004
- 5. 제조자/제조국가 : Foxconn/China



- 1. 기기의 명칭(모델명) : BCM95722A2202G
- 2. 인증번호 : BCM-BCM95722A2202G (B)
- 3. 인증받은 자의 상호 : BROADCOM
- 4. 제조년월일 : 04/30/2007
- 5. 제조자/제조국가 : Foxconn/China

Note that this device has been approved for non-business purposes and may be used in any environment, including residential areas.

BSMI

BSMI 通告 (僅限於台灣)

大多數的 Dell 電腦系統被 BSMI (經濟部標準檢驗局) 劃分為乙類數位裝置。但是, 使用某些選件會使有些組態的等級變成甲類。若要確定您的電腦系統適用等級, 請檢查所有位於電腦底部或背面板、擴充卡安裝托架, 以及擴充卡上的 BSMI 註冊標籤。如果其中有一甲類標籤, 即表示您的系統為甲類數位裝置。如果只有 BSMI 的檢磁號碼標籤, 則表示您的系統為乙類數位裝置。

一旦確定了系統的 BSMI 等級, 請閱讀相關的 BSMI 通告。請注意, BSMI 通告規定凡是未經 Dell Inc. 明確批准的擅自變更或修改, 將導致您失去此設備的使用權。

此裝置符合 BSMI (經濟部標準檢驗局) 的規定, 使用時須符合以下兩項條件:

- 此裝置不會產生有害干擾。
- 此裝置必須能接受所接收到的干擾, 包括可能導致無法正常作業的干擾。

乙類

此設備經測試證明符合 BSMI (經濟部標準檢驗局) 之乙類數位裝置的限制規定。這些限制的目的是為了在住宅區安裝時, 能防止有害的干擾, 提供合理的保護。此設備會產生、使用並散發射頻能量; 如果未遵照製造廠商的指導手冊來安裝和使用, 可能會干擾無線電通訊。但是, 這並不保證在個別的安裝中不會產生干擾。您可以透過關閉和開啓此設備來判斷它是否會對廣播和電視收訊造成干擾; 如果確實如此, 我們建議您嘗試以下列一種或多種方法來排除干擾:

- 重新調整天線的接收方向或重新放置接收天線。
- 增加設備與接收器的距離。
- 將設備連接至不同的插座, 使設備與接收器連接在不同的電路上。
- 請向經銷商或有經驗的無線電 / 電視技術人員查詢, 以獲得幫助。