
Before Reading This Manual

Thank you for purchasing the PRIMERGY 1Gbit/s Ethernet I/O Module (hereinafter referred to as this product or the card).

The card can be installed in the expansion card slot of the Fujitsu server blade to configure the Local Area Network (LAN) system. This manual explains the 1Gbit/s Ethernet I/O Modules and the LAN driver (for Windows). Read this manual carefully to handle the product correctly.

For details about the LAN driver, refer to manuals supplied with the server blade or the Fujitsu PRIMERGY website:

(<http://primergy.fujitsu.com>)

May, 2007

For Your Safety

This manual contains important information, required to operate this product safely.

Thoroughly review the information in this manual before using this product. Especially note the points under "Safety", and only operate this product with a complete understanding of the material provided.

This manual should be kept in an easy-to-access location for quick reference when using this product.



High Safety

The Products are designed, developed and manufactured as contemplated or general use, including without limitation, general office use, personal use, household use, and ordinary industrial use, but are not designed, developed and manufactured as contemplated for use accompanying fatal risks or dangers that, unless extremely high safety is secured, could lead directly to death, personal injury, severe physical damage, or other loss (hereinafter "High Safety Required Use"), including without limitation, nuclear reaction control in nuclear facility, aircraft flight control, air traffic control, mass transport control, medical life support system, missile launch control in weapon system. You shall not use this Product without securing the sufficient safety required for the High Safety Required Use. If you wish to use this Product for High Safety Required Use, please consult with our sales representatives in charge before such use.




Remarks

Warning Descriptions

Various symbols are used throughout this manual. These are provided to emphasize important points for your safety and that of others. The symbols and their meanings are as follows. Make sure to fully understand these before reading this manual.



 WARNING	Ignoring this symbol could be potentially lethal.
 CAUTION	Ignoring this symbol may lead to injury and/or damage this product.

The following symbols are used to indicate the type of warning or cautions being described.

	The triangle mark emphasizes the urgency of the WARNING and CAUTION. Details are described next to the triangle.
	A barred circle (⊘) warns against certain actions (Do Not). Details are described next to the circle.
	A black circle indicates actions that must be taken. Details are described next to the black circle.

Symbols

The following are symbols used throughout this manual.

Symbols	Definition
 IMPORTANT	These sections explain prohibited actions and points to note when using this product. Make sure to read these sections.
 POINT	These sections explain information needed to operate the hardware and software properly. Make sure to read these sections.
→	This mark indicates reference pages or manuals.

Entering commands (Keys)

CD-ROM drive names are shown as [CD-ROM drive]. Enter your drive name according to your environment.

[CD-ROM drive]:\Setup.exe

Abbreviations

The following expressions and abbreviations are used to describe the product names used in this manual.

Product names	Expressions and abbreviations		
1Gbit/s Ethernet I/O Module (PG-LND201)	this product		
PG-LND201 LAN Driver	LAN driver or the driver		
Microsoft® Windows Server® 2003, Standard Edition	Windows Server 2003	Windows	
Microsoft® Windows Server® 2003, Enterprise Edition			
Microsoft® Windows Server® 2003 R2, Standard Edition			
Microsoft® Windows Server® 2003 R2, Enterprise Edition			
Microsoft® Windows Server® 2003, Standard x64 Edition	Windows Server 2003 x64		
Microsoft® Windows Server® 2003, Enterprise x64 Edition			
Microsoft® Windows Server® 2003 R2, Standard x64 Edition			
Microsoft® Windows Server® 2003 R2, Enterprise x64 Edition	Windows 2000 Server		
Microsoft® Windows® 2000 Server			
Microsoft® Windows® 2000 Advanced Server			
Red Hat® Enterprise Linux® ES (v.3 for x86)	Linux		
Red Hat® Enterprise Linux® AS (v.3 for x86)			
Red Hat® Enterprise Linux® AS (v.4 for x86)			
Red Hat® Enterprise Linux® ES (v.4 for x86)			
Red Hat® Enterprise Linux® AS (v.4 for EM64T)			
Red Hat® Enterprise Linux® ES (v.4 for EM64T)			
Red Hat® Enterprise Linux® 5 (for x86)			
Red Hat® Enterprise Linux® 5 (for Intel64)			
PRIMERGY BX620 S3 Server Blade	BX620 S3 Server Blade		Server Blade
PRIMERGY BX620 S4 Server Blade	BX620 S4 Server Blade		
PRIMERGY BX600 S3 Blade Server System Unit	Chassis		
PRIMERGY BX600 S2 Blade Server System Unit			
Switch Blade	Switch Blade		
Catalyst Blade Switch 3040			
FC Pass-Thru Blade	FC Pass-Thru Blade		
FC Switch Blade (PG-FCS103/PG-FCS102)	FC Switch Blade		

Safety

For safe use of this product, it is vital that the following warnings are heeded.

Handling this product

WARNING

Electric Shock



- Do not tinker with the product. Doing so may cause fire or electric shock.
- Keep this product away from water. Failure to do so may cause fire or electric shock.
- When there is lightning nearby, unplug all power cords and external connecting cords from this product. Failure to do so may cause destruction of the devices and fire.

CAUTION



- Since this product is delicate, avoid using or storing it under extreme conditions, such as excessively high or low temperature, high humidity, or in direct sunlight. Do not bend or damage the card or subject it to extreme shock. Doing so may cause failure or fire.

Recycle

When scrapping this product, contact an office listed in "Appendix A Contact Information" (→pg.195). This product must be disposed of as industrial waste.

Checking the Items Supplied

Before using the product, check that no supplied or attached items are missing.

If any items are missing, contact an office listed in "Appendix A Contact Information" (→pg.195).

- **1Gbit/s Ethernet I/O Module**
- **User's Guide (this manual)**
- **Screw (3 screws)**

Intel is a registered trademark of Intel Corporation in the USA.

Microsoft, Windows, and Windows Server are trademarks or registered trademarks of Microsoft Corporation in the United States and other countries.

Linux is a trademark or registered trademark of Linus Torvalds in the United States and other countries.

Red Hat and all Red Hat-based trademarks and logos are trademarks or registered trademarks of Red Hat, Inc. in the United States and other countries.

Other product names used are trademarks or registered trademarks of their respective manufacturers.

Other products are copyrights of their respective manufacturers.

All Rights Reserved, Copyright© FUJITSU LIMITED 2007

Contents

1	1 Gbit/s Ethernet I/O Module	106
1.1	Overview	106
1.2	Specifications	106
1.3	Functionality and Features	107
2	Installing a 1Gbit/s Ethernet I/O Module	115
2.1	Installing in a Server Blade	116
3	Installing the LAN Driver	118
3.1	Installing the Driver Software	118
4	Broadcom Gigabit Ethernet Teaming Services	119
4.1	Broadcom Gigabit Ethernet Overview	119
4.2	Glossary	119
4.3	Teaming Concepts	121
4.4	Teaming Mechanisms	130
4.5	Types of Teams	133
4.6	Attributes of the Features Associated with Each Type of Team	136
4.7	Teaming and Other Advanced Networking Properties	138
4.8	General Network Considerations	140
4.9	Event Log Messages	151
5	Broadcom Advanced Control Suite 2 (BACS2)	156
5.1	BACS2 Overview	156
5.2	Installing the BACS2	158
5.3	Starting BACS2	162
5.4	Setting of BACS2	162
5.5	Configuring Teaming	176
Appendix A	Contact Information	195

1 1 Gbit/s Ethernet I/O Module

This chapter explains the features and specifications of this product.

1.1 Overview

This product is an LAN expansion board exclusive to BX620 S3 Server Blade and BX620 S4 Server Blade.

This product has 2 LAN controllers, and provides LAN connection completely separate/ independent of the onboard LAN. The external access is performed via the Switch Blade installed to the network blade slot 3 or 4 (NET3 or NET4) on the chassis.



- ▶ A server blade installed with this products and a server blade installed with a fibre channel expansion board cannot be installed to the same chassis.
- ▶ When a Switch Blade is installed to network blade slot 3 or 4 (NET3 or NET4), a FC Pass-Thru Blade or FC Switch Blade cannot be installed at the same time.

1.2 Specifications

Item		Specifications
Product ID		PG-LND201
Host bus specifications	Interface	PCI-Express (x4)
	Data transfer rate	Max. 1 Gbps
	Data transfer system	Bus master, DMA
External interface		1000BASE-T Ethernet (SerDes) Interface
Controller LSI		BCM 5708S
Applicable model		BX620 S3 Server Blade BX620 S4 Server Blade

1.3 Functionality and Features

1.3.1 Functional Description

This product is a new class of Gigabit Ethernet (GbE) converged network interface controller (C-NIC) that can simultaneously perform accelerated data networking, storage networking, and high-performance clustering on a standard Ethernet network. The C-NIC offers acceleration for all popular protocols used in the data center, such as:

- TCP Offload Engine (TOE) for accelerating TCP

POINT

- ▶ Offloading technologies are supported when this product is installed in a system with Windows Server 2003 with Scalable Networking Pack (SNP).

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. For details about teaming, refer to "Broadcom Advanced Server Program (BSAP) Overview" (→pg.109) in "1.3.3 Teaming Function" (→pg.109). For a description of VLANs, refer to "1.3.4 Virtual LAN Function" (→pg.112). For instructions on configuring teaming and creating VLANs on Windows operating systems, refer to "1.3.3 Teaming Function" (→pg.109).

1.3.2 Features

The following is a list of this product features:

- TCP Offload Engine (TOE)
- Single-chip solution
 - Standard Ethernet frame size (1518 bytes)
 - PCI Express v1.0A, x4
 - Full fast-path TCP offload
 - Zero copy capable hardware
- Other performance features
 - TCP, IP, UDP checksum
 - TCP segmentation
- Manageability
 - Broadcom Advanced Control Suite2 (BACS2) diagnostic and configuration software suite
 - Supports PXE 2.0 specification
 - Statistics for SNMP MIB II, Ethernet-like MIB, and Ethernet MIB (IEEE Std 802.3z, Clause 30)
 - IPMI support
- Advanced network features
 - Jumbo frames (up to 9 KB)
 - Virtual LANs
 - Flow Control (IEEE Std 802.3x)
 - LiveLink™ (supported in both the 32-bit and 64-bit Windows operating systems)
 - Logical Link Control (IEEE Std 802.2)
- Layer-2 Priority Encoding (IEEE 802.1p)

- High-speed on-chip RISC processor
- Up to 4 classes of service (CoS)
- Integrated 96 KB frame buffer memory
- Support for multicast addresses via 128 bits hashing hardware function
- EM64T processor support

TCP Offload Engine (TOE)

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

This product's network adapter's TOE function allows simultaneous operations of up to 1024 fully offloaded TCP connections. The TOE support on the adapter significantly reduces the host CPU utilization while preserving the implementation of the operating system stack.

Broadcom Advanced Control Suite 2 (BACS2)

Broadcom Advanced Control Suite 2 (BACS2), a component of the Broadcom teaming software, is an integrated utility that provides information about each network adapter that is installed in your system.

The BACS2 also enables you to perform detailed tests, diagnostics, and analysis on each adapter, as well as to modify property values and view traffic statistics for each adapter. BACS2 is used on a Windows operating systems to configure teaming and to add VLANs.

1.3.3 Teaming Function

Broadcom Advanced Server Program (BSAP) Overview

Broadcom Advanced Server Program (BSAP) is the Broadcom teaming software for Windows Server 2003 and Windows 2000 Server operating systems. BSAP runs within the Broadcom Advanced Control Suite 2 (BACS2) utility.

BSAP supports four types of teams for Layer 2 teaming:

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

POINT

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

For more information on network adapter teaming concepts, refer to "4.3 Teaming Concepts" (→pg.121).

Load Balancing and Fault Tolerance

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

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

Types of Teams

The available types of teams for the Windows Server 2003/Windows 2000 Server operating systems are:

- Smart Load Balancing and Failover
→ "Smart Load Balancing™ and Failover"(pg.110)
- Link Aggregation (802.3ad)
→ "Link Aggregation (802.3ad)"(pg.110)
- Generic Trunking (FEC/GEC)/802.3ad-Draft Static
→ "Generic Trunking (FEC/GEC)/802.3ad-Draft Static"(pg.110)
- SLB (Auto-Fallback Disable)
→ "SLB (Auto-Fallback Disable)"(pg.111)

IMPORTANT

- ▶ Link aggregation is not supported in the Blade Server.

E

Smart Load Balancing™ and Failover

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

POINT

- ▶ 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.
- ▶ IPX balances only on the transmitting side of the team; other protocols are limited to the primary adapter.
- ▶ 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.

IMPORTANT

- ▶ Link aggregation (802.3ad) is not supported in the Blade Server.

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. However, the Generic Trunking (FEC/GEC)/802.3ad-Draft Static type of team, 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 exception of when the standby member is active, if a primary member comes back on line, the team continues using the standby member, rather than switching back to the primary member.

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

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

Limitations of Smart Load Balancing™ and Failover/SLB (Auto-Fallback Disable) Types of Teams

Smart Load Balancing™ (SLB) is a protocol-specific scheme. The state of support for IP, IPX, and NetBEUI protocols is listed below.

Operating System	Failover/ Fallback - All Broadcom			Failover/ Fallback - Multivendor			Load Balance - All Broadcom			Load Balance - Multivendor		
	IP	IPX	Net BE UI	IP	IPX	Net BE UI	IP	IPX	Net BE UI	IP	IPX	Net BE UI
Windows Server 2003	Y	Y	N	Y	N	N	Y	Y	N	Y	N	N
Windows Server 2003 with SNP	Y	Y	N	Y	N	N	Y	Y	N	Y	N	N
Windows 2000 Server	Y	Y	Y	Y	N	N	Y	Y	N	Y	N	N

Y = supported, N = not supported

The Smart Load Balancing™ type of team works with all Ethernet switches without having to configure the switch ports to any special trunking mode. Only IP traffic is load-balanced in both inbound and outbound directions. IPX traffic is load-balanced in the outbound direction only. Other protocol packets are sent and received through one primary interface only. Failover for non-IP traffic is supported only for 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.

POINT

- ▶ 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.

E

LiveLink™

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

LiveLink™ is supported in both 32-bit and 64-bit Windows operating systems. Refer to the Channel Bonding documentation for similar functionality in Linux Channel Bonding (refer to <http://www.redhat.com/docs/manuals/enterprise/RHEL-3-Manual/ref-guide/s1-modules-ethernet.html>).

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.

1.3.4 Virtual LAN Function


VLAN Overview

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

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

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

- Example of Servers Supporting Multiple VLANs with Tagging



"• Example of Servers Supporting Multiple VLANs with Tagging" (→pg.113)" 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 "• Example VLAN Network Topology" (→pg.113).

- Example VLAN Network Topology

Component	Description
VLAN #1	An IP subnet consisting of the Main Server, PC #3, and PC #5. This subnet represents an engineering group.
VLAN #2	Includes the Main Server, PCs #1 and #2 via shared media segment, and PC #5. This VLAN is a software development group.
VLAN #3	Includes the Main Server, the Accounting Server and PC #4. This VLAN is an accounting group.
Main Server	A high-use server that needs to be accessed from all VLANs and IP subnets. The Main Server has a Broadcom adapter installed. All three IP subnets are accessed via the single physical adapter interface. The server is attached to one of the switch ports, which is configured for VLANs #1, #2, and #3. Both the adapter and the connected switch port have tagging turned on. Because of the tagging VLAN capabilities of both devices, the server is able to communicate on all three IP subnets in this network, but continues to maintain broadcast separation between all of them.
Accounting Server	Available to VLAN #3 only. The Accounting Server is isolated from all traffic on VLANs #1 and #2. The switch port connected to the server has tagging turned off.
PCs #1 and #2	Attached to a shared media hub that is then connected to the switch. PCs #1 and #2 belong to VLAN #2 only, and are logically in the same IP subnet as the Main Server and PC #5. The switch port connected to this segment has tagging turned off.

Component	Description
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.

POINT

- ▶ 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, refer to "5.5 Configuring Teaming" (→pg.176).

POINT

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

2 Installing a 1Gbit/s Ethernet I/O Module

This chapter explains the installation procedure in the server blade.

WARNING

Electric Shock



- When installing or removing this product, make sure to remove the server blade from the chassis. Failure to do so may cause electric shock. For details on how to remove the server blade from the chassis, refer to "Blade Server System Unit Hardware Guide".

CAUTION



- The circuit boards and soldered parts of internal options are exposed. They can be damaged by static electricity. Before handling them, first touch a metal part of the server blade to discharge static electricity from your body.
- Do not touch the circuitry on boards or soldered parts. Hold the metallic areas or the edges of the circuit boards.
- These products are susceptible to static electricity. Place them on conductive pads or keep them in their packaging as long as they are not necessary.


2.1 Installing in a Server Blade

POINT

- ▶ To connect a 1Gbit/s Ethernet I/O Module to the external LAN (device), it is necessary to install a Switch Blade or GbE Pass-Thru Blade to network blade slot 3 or network blade slot 4 (NET3 or NET4) of the chassis.

2.1.1 Installation Position of the 1Gbit/s Ethernet I/O Module

Install the 1Gbit/s Ethernet I/O Module in the expansion board slot in the BX620 S3/BX620 S4 Server Blade.



2.1.2 Installation Procedure for the 1Gbit/s Ethernet I/O Module


- 1 Turn off the server blade where the 1Gbit/s Ethernet I/O Module will be installed.**
→ "3.3 Turning Off the Server " in "Blade Server System Unit Hardware Guide"
- 2 Touch a metal part of the chassis to discharge static electricity from your body.**
- 3 Remove the server blade from the chassis.**
→ "4.2 Installing Server Blades" in "Blade Server System Unit Hardware Guide"

4 Remove the top cover.

→"7.2 Removing and Attaching the Top Cover" in "Server Blade User's Guide"


5 Install the 1Gbit/s Ethernet I/O Module.

Make sure the 1Gbit/s Ethernet I/O Module is securely on the slot.



6 Secure the 1Gbit/s Ethernet I/O Module with the screws.

Secure the 1Gbit/s Ethernet I/O Module with the three screws included with this product.



7 Attach the top cover.

→"7.2 Removing and Attaching the Top Cover" in "Server Blade User's Guide"

8 Install the server blade to the chassis.

→"4.2 Installing Server Blades" in "Blade Server System Unit Hardware Guide"

3 Installing the LAN Driver

This chapter explains how to install the LAN driver.

3.1 Installing the Driver Software

POINT

- ▶ Get the LAN driver from the ServerStart CD-ROM included with the BX620 S4 server blade and install.
If using a LAN driver, refer to the operation manual or help file included with the driver.
- ▶ Do not use LAN drivers downloaded from Broadcom's online service in the PRIMERGY server.
- ▶ Before installing the driver software, verify that the Windows operating system has been upgraded to the latest version with the latest service pack applied.
- ▶ To use TCP/IP Offload Engine (TOE), you must have Windows Server 2003 with Scalable Networking Pack (SNP)

3.1.1 Installing the LAN Drivers

The installer will detect if SNP for Windows Server 2003 is installed on your machine. If it is, the installer will install the NDIS 5.2 driver, which is necessary in order to use TOE. If SNP for Windows Server 2003 is not installed on your machine, the installer will install the NDIS 5.1 driver and you will not be able to use TOE.

1 Insert the ServerStart Disc1 CD-ROM.

2 Execute the following EXE file.

- For Windows Server 2003 x64
 \DRIVERS\LAN\Broadcom\Ext\W2K3x64\DrvInst\setup.exe
- For Windows Server 2003
 \DRIVERS\LAN\Broadcom\Ext\W2K3\DrvInst\setup.exe
- For Windows Server 2000
 \DRIVERS\LAN\Broadcom\Ext\W2K\DrvInst\setup.exe

3 Click [Next].

The installation of the LAN driver is started.

4 Click [Finish].

5 Eject the CD-ROM, and restart the server blade.

4 Broadcom Gigabit Ethernet Teaming Services

This chapter explains the technology and considerations when working with the network teaming services.

4.1 Broadcom Gigabit Ethernet Overview

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 manual is provided to assist IT professionals during the deployment and troubleshooting of system applications that require network fault tolerance and load balancing.

4.2 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)

Item	Definition
PXE	Preboot Execution Environment
RAID	Redundant Array of Inexpensive Disks
Smart Load Balancing™ and Failover	Switch-independent failover type of team in which the primary team member handles all incoming and outgoing traffic while the standby team member is idle until a failover event (for example, loss of link occurs). The intermediate driver (BASP) manages incoming/outgoing traffic.
Smart Load Balancing (SLB)	Switch-independent load balancing and failover type of team, in which the intermediate driver manages outgoing/incoming traffic.
TCP	Transmission Control Protocol
UDP	User Datagram Protocol
WINS	Windows name service
WLBS	Windows Load Balancing Service

4.3 Teaming Concepts

Storage devices use RAID technology to group individual hard drives. Switch ports can be grouped together using technologies such as Cisco Gigabit EtherChannel, IEEE 802.3ad Link Aggregation, Bay Network Multilink Trunking, and Extreme Network Load Sharing. Network interfaces on servers can be grouped together into a team of physical ports called a virtual adapter.

4.3.1 Teaming and Network Addresses

Network Addressing

To understand how teaming works, it is important to understand how node communications work in an Ethernet network. This Section is based on the assumption that the reader is familiar with the basics of IP and Ethernet network communications. The following information provides the concepts of network addressing used in an Ethernet network.

Every Ethernet network interface in a host platform, such as a computer system, requires a globally unique Layer 2 address and at least one globally unique Layer 3 address. Layer 2 is the Data Link Layer, and Layer 3 is the Network layer as defined in the OSI model. The Layer 2 address is assigned to the hardware and is often referred to as the MAC address or physical address. This address is pre-programmed at the factory and stored in NVRAM on a network interface card or on the system motherboard for an embedded LAN interface. The Layer 3 addresses are referred to as the protocol or logical address assigned to the software stack. IP and IPX are examples of Layer 3 protocols. In addition, Layer 4 (Transport Layer) uses port numbers for each network upper level protocol such as Telnet or FTP. These port numbers are used to differentiate traffic flows across applications. The combination of the IP address and the TCP port number is called a socket.

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

Teaming and Network Addresses

A team of adapters will function as a single virtual network interface, and do not appear to be any different than a non-teamed adapter to other network devices. 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.

E

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

4.3.2 Types of Teams

"Available Teaming Types" shows a summary of the teaming types and their classification.

Available Teaming Types

Teaming Type	Switch-Dependent* ¹	Link Aggregation Control Protocol Support Required on the Switch	Load Balancing	Failover
Smart Load Balancing and Failover	–	–	○	○
SLB (Auto-Fallback Disable)	–	–	○	○
Link Aggregation (802.3ad) * ²	○	○	○	○
Generic Trunking (FEC/GEC)/802.3ad-Draft Static	○	–	○	○

*1: Switch must support specific type of team.

*2: Link aggregation is not supported in the Blade Server.

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 2-8 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. The 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 send 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.

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.

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.



- ▶ Link aggregation is not supported in the Blade Server.

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.

4.3.3 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 (refer to "• Process for Selecting a Team Type" (→pg.129)). The miniport driver controls the host LAN controller directly to enable functions such as send, receive, and interrupt processing. The intermediate driver fits between the miniport driver and the protocol layer multiplexing several miniport driver instances, and creating a virtual adapter that looks like a single adapter to the NDIS layer. NDIS provides a set of library functions to enable the communications between either miniport drivers or intermediate drivers and the protocol stack. The protocol stack implements IP, IPX and ARP. A protocol address such as an IP address is assigned to each miniport device instance, but when an Intermediate driver is installed, the protocol address is assigned to the virtual team adapter and not to the individual miniport devices that make up the team.

The Broadcom supplied teaming support is provided by three individual software components that work together and are supported as a package. When one component is upgraded, all the other components must be upgraded to supported versions. "• Broadcom Teaming Software Component" (→pg.126) describes the three software components and their associated files for supported operating systems.

- Broadcom Teaming Software Component

Software Component	Broadcom Name	Windows File Name
–	Virtual Bus Driver (VBD)	bxvbdx.sys bxvbda.sys*
Miniport Driver	Broadcom Base Driver	bxnd50x.sys bxnd51x.sys bxnd51a.sys* bxnd52x.sys bxnd52a.sys*
Intermediate Driver	Broadcom Advanced Server Program (BASP)	Baspxp32.sys Baspw2k.sys
Configuration User Interface	Broadcom Advanced Control Suite 2 (BACS2)	BACS2

*: For 64-bit systems

4.3.4 Hardware Requirements

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

- **Repeater Hub**

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

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

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

- **Switching Hub**

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

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

- **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.

4.3.5 Configuring Teaming

The Broadcom Advanced Control Suite 2 (BACS2) utility is used to configure teaming in the supported operating system environments.

The BACS2 utility is designed to run in one of the following 32-bit and 64-bit Windows operating systems: Microsoft Windows 2000 Server and Windows Server 2003. BACS2 is used to configure load balancing and fault tolerance teaming, and VLANs. In addition, it displays the MAC address, driver version, and status information. BACS2 also includes a number of diagnostics tools such as hardware diagnostics, cable testing, and a network topology test.

4.3.6 Supported Features by Team Type

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

- 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 ^(*)	SLB	Generic Trunking	Link Aggregation
Number of ports per team (same broadcast domain)	2-8	2-8	2-8	2-8
Number of teams	4	4	4	4
Adapter fault tolerance	Yes	Yes	Yes	Yes
Switch link fault tolerance (same broadcast domain)	Yes	Yes	Switch-dependent	Switch-dependent
TX load balancing	No	Yes	Yes	Yes
RX load balancing	No	Yes	Yes (performed by the switch)	Yes (performed by the switch)
Requires compatible switch	No	No	Yes	Yes
Heartbeats to check connectivity	No	No	No	No

Type of Team	Fault Tolerance	Load Balancing	Switch-Dependent Static Trunking	Switch-Independent Dynamic Link Aggregation (IEEE 802.3ad)
Function	SLB with Standby ^(*1)	SLB	Generic Trunking	Link Aggregation
Mixed media (adapters with different media)	Yes	Yes	Yes (switch-dependent)	–
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 ^(*2)	Yes ^(*2)	Yes ^(*2)	Yes ^(*2)
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

*1: SLB with one primary and one standby member.

*2: Requires at least one Broadcom adapter in the team.




▶ Link aggregation is not supported in the Blade Server.

4.3.7 Selecting a Team Type

The following flow chart provides the decision flow when planning for Layer 2 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: IEEE 802.3ad 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, however, then SLB is the only choice (refer to "• Process for Selecting a Team Type" (→pg.129)).

- Process for Selecting a Team Type

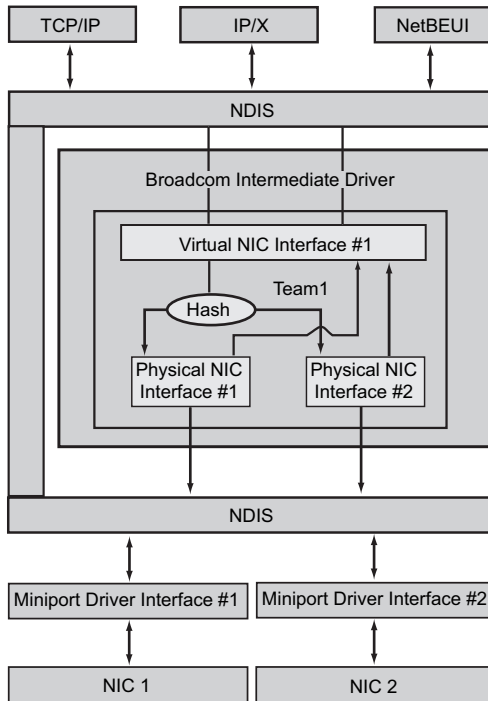


4.4 Teaming Mechanisms

4.4.1 Architecture

The Broadcom Advanced Server Program (BASP) is implemented as an NDIS intermediate driver (refer to "• Teaming Across Switches Without an Interswitch Link" (→pg.141)). It operates below protocol stacks such as TCP/IP and IPX and appears as a virtual adapter. This virtual adapter inherits the MAC Address of the first port initialized in the team. A Layer 3 address must also be configured for the virtual adapter. The primary function of BASP is to balance inbound (for SLB) and outbound traffic (for all teaming modes) among the physical adapters installed on the system selected for teaming. The inbound and outbound algorithms are independent and orthogonal to each other. The outbound traffic for a particular session can be assigned to a given port while its corresponding inbound traffic can be assigned to a different port.

- Intermediate Driver



4.4.2 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.

4.4.3 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. However, the processing of ARP packets, both inbound and outbound, 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.

4.4.4 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 a 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-----> PhysAdapter1 -----> 10.0.0.1
```

Or the streams may look like this:

```
IGMP -----> PhysAdapter1 -----> 10.0.0.1
ICMP -----> PhysAdapter1 -----> 10.0.0.1
TCP-----> PhysAdapter1 -----> 10.0.0.1
UDP-----> PhysAdapter1 -----> 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.

4.4.5 Performance

Modern network interface cards provide many hardware features that reduce CPU utilization by offloading certain CPU intensive operations (→"4.7 Teaming and Other Advanced Networking Properties"(pg.138)). 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. In such cases, the application may be better suited to taking 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.

4.5 Types of Teams

4.5.1 Switch-Independent

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

Smart Load Balancing and Failover

SLB provides for switch-independent, bidirectional, fault-tolerant teaming and load balancing. Switch independence implies that there is no specific support for this function required in the switch, allowing SLB to be compatible with all switches. Under SLB, all adapters in the team have separate MAC addresses. The load-balancing algorithm operates on Layer 3 addresses of the source and destination nodes, which enables SLB to load balance both incoming and outgoing traffic. The BASP intermediate driver continually monitors the physical ports in a team for link loss. In the event of link loss on any port, traffic is automatically diverted to other ports in the team. The SLB teaming mode supports switch fault tolerance by allowing teaming across different switches-provided the switches are on the same physical network or broadcast domain.

• Network Communications

The following are the key attributes of SLB:

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

- **Applications**

The SLB algorithm is most appropriate in home and small business environments where cost is a concern, or with commodity switching equipment. SLB teaming works with unmanaged Layer 2 switches and is a cost-effective way of getting redundancy and link aggregation at the server.

Smart Load Balancing also supports teaming physical adapters with differing link capabilities. In addition, SLB is recommended when switch fault tolerance with teaming is required.

- **Configuration Recommendations**

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

4.5.2 Switch-Dependent

Generic Static Trunking

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

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

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

- **Network Communications**

The following are the key attributes of Generic Static Trunking:

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

- **Applications**

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

- **Configuration Recommendations**

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

Dynamic Trunking (IEEE 802.3ad Link Aggregation)

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

- **Network Communications**

The following are the key attributes of Dynamic Trunking:

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

- **Applications**

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

- **Configuration Recommendations**

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

4.6 Attributes of the Features Associated with Each Type of Team

The attributes of the features associated with each type of team are summarized below.

- Smart Load Balancing™

Feature	Attribute
User interface	Broadcom Advanced Control Suite 2 (BACS2)
Number of teams	Maximum 4
Number of adapters per team	Maximum 8
Hot replace	Yes
Hot add	Yes
Hot remove	Yes
Link speed support	Different speeds
Frame protocol	IP
Incoming packet management	BASP
Outgoing packet management	BASP
LiveLink support	Yes
Failover event	Loss of link
Failover time	<500 ms
Fallback time	1.5 s ^(*1) (approximate)
MAC address	Different
Multivendor teaming	Yes

*1: Make sure that Port Fast or Edge Port is enabled.

- Generic Trunking

Feature	Attribute
User interface	Broadcom Advanced Control Suite 2 (BACS2)
Number of teams	Maximum 4
Number of adapters per team	Maximum 8
Hot replace	Yes
Hot add	Yes
Hot remove	Yes
Link speed support	Different speeds ^(*1)
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 ^(*2) (approximate)

Feature	Attribute
MAC address	Same for all adapters
Multivendor teaming	Yes

*1: Some switches require matching link speeds to correctly negotiate between trunk connections.

*2: Make sure that Port Fast or Edge Port is enabled.

- Dynamic Trunking

Feature	Attribute
User interface	Broadcom Advanced Control Suite 2 (BACS2)
Number of teams	Maximum 4
Number of adapters per team	Maximum 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 ^(*) (approximate)
MAC address	Same for all adapters
Multivendor teaming	Yes

*1: Make sure that Port Fast or Edge Port is enabled.

Speeds Supported for Each Type of Team

The various link speeds that are supported for each type of team are listed in "• Link Speeds in Teaming" (→pg.137). Mixed speed refers to the capability of teaming adapters that are running at different link speeds.

- 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

4.7 Teaming and Other Advanced Networking Properties

Before creating a team, adding or removing team members, or changing advanced settings of a team member, make sure each team member has been configured in a similar way. Settings to check include VLANs and QoS Packet Tagging, Jumbo Frames, and the various offloads. Advanced adapter properties and teaming support are listed below.

- Advanced Adapter Properties and Teaming Support

Adapter Properties	Supported by Teaming Virtual Adapter
Checksum Offload	Yes
IEEE 802.1p QoS Tagging	No
Large Send Offload	Yes ^(*1)
Jumbo Frames	Yes ^(*2)
IEEE 802.1Q VLANs	Yes ^(*3)
Wake on LAN	No
Preboot Execution environment (PXE)	Yes ^(*4)

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

*2: Must be supported by all adapters in the team.

*3: Only for Broadcom adapters.

*4: As a PXE server only, not as a client.

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

POINT

- ▶ All adapters on the team must support the property listed in "• Advanced Adapter Properties and Teaming Support" (→pg.138) in order for the team to support the property.

Checksum Offload

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

IEEE 802.1p QoS Tagging

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

Large Send Offload

Large Send Offload (LSO) is a feature provided by Broadcom network adapters that prevents an upper level protocol such as TCP from breaking a large data packet into a series of smaller packets with headers appended to them. The protocol stack need only generate a single header for a data packet as large as 64 KB, and the adapter hardware breaks the data buffer into appropriately-sized Ethernet frames with the correctly sequenced header (based on the single header originally provided).

Jumbo Frames

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

IEEE 802.1Q VLANs

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

Preboot Execution Environment

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

4.8 General Network Considerations

4.8.1 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.

4.8.2 Switch-Link Fault Tolerance

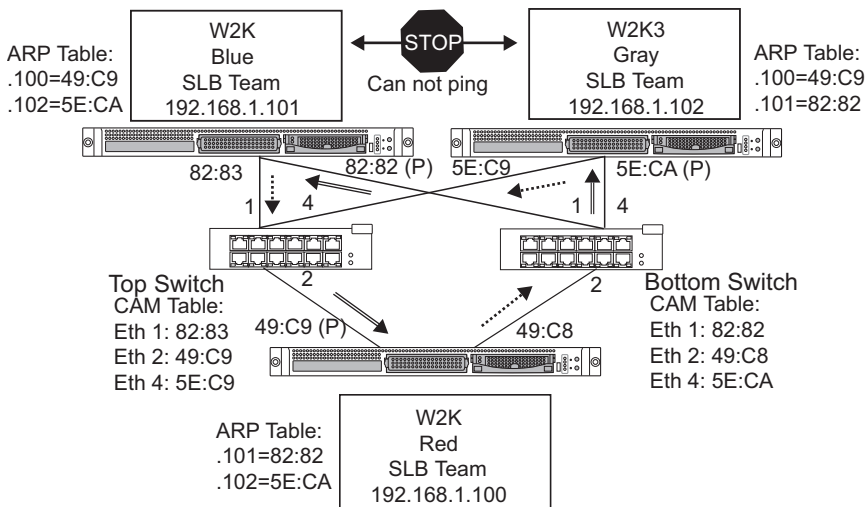
The diagrams below describe the operation of an SLB team in a switch fault tolerant configuration. We show the mapping of the ping request and ping replies in an SLB team with two active members. All servers (Blue, Gray and Red) have a continuous ping to each other. "• Teaming Across Switches Without an Interswitch Link" (→pg.141) is a setup without the interconnect cable in place between the two switches. "• Teaming Across Switches With Interconnect" (→pg.142) has the interconnect cable in place, and "• Failover Event" (→pg.142) 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 (arrow of a dotted line), while the primary team member receives the respective ICMP echo replies (arrow of a double line). 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.

- 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.

• Teaming Across Switches With Interconnect



"• Teaming Across Switches With Interconnect" (→pg.142) represents a failover event in which the cable is unplugged on the Top Switch port 4. This is a successful failover with all stations pinging each other without loss of connectivity.

• Failover Event



4.8.3 Spanning Tree Algorithm

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.

4.8.4 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. For example, 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.

4.8.5 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.

4.8.6 Layer 3 Routing/Switching

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

4.8.7 Teaming with Hubs (for troubleshooting purposes only)

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

4.8.8 Hub Usage in Teaming Network Configurations

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

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


4.8.9 SLB Teams

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

SLB Team Connected to a Single Hub

SLB teams configured as shown in "• Team Connected to a Single Hub" (→pg.146) maintain their fault tolerance properties. Either server connection could potentially fail, and network functionality is maintained. Clients could be connected directly to the hub, and fault tolerance would still be maintained; server performance, however, would be degraded.

- Team Connected to a Single Hub



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

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

Teaming with Microsoft NLB/WLBS

It is known that the SLB mode of teaming does not work in an NLB unicast environment. It is not known, however, why the SLB mode should not work in a NLB multicast environment. The SLB teaming algorithm is mutually exclusive with the NLB unicast mechanism.

4.8.10 Troubleshooting Teaming Problems

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

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

Teaming Configuration Tips

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

- Although mixed-speed SLB teaming is supported, it is recommended that all adapters in a team be the same speed (either all Gigabit Ethernet or all Fast Ethernet).
- 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.
- All switches that the team is directly connected to must have the same hardware revision, firmware revision, and software revision to be supported.
- 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.
- Do not use the Locally Administered Address on any physical adapter that is a member of a team.
- Remove any static IP address from the individual physical team members before the team is built.
- 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.
- 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.
- 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.
- Verify the base (Miniport) and team (intermediate) drivers are from the same release package. The mixing of base and teaming drivers from different CD releases is not supported.
- Test the connectivity to each physical adapter prior to teaming.
- Test the failover and fallback behavior of the team before placing into a production environment.
- When moving from a nonproduction network to a production network, it is strongly recommended to test again for failover and fallback.
- Test the performance behavior of the team before placing 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.

- Make sure the link light is ON for every adapter and all the cables are attached.
- Check that the matching base and intermediate drivers belong to the same release and are loaded correctly.
- Check for a valid IP Address using the Windows ipconfig command.
- 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.
- Check that the adapters and the switch are configured identically for link speed and duplex.
- If possible, break the team and check for connectivity to each adapter independently to confirm that the problem is directly associated with teaming.
- Check that all switch ports connected to the team are on the same VLAN.
- 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	Answer
Under what circumstances is traffic not load balanced? Why is all traffic not load balanced evenly across the team members?	The bulk of traffic does not use IP/TCP/UDP or the bulk of the clients are in a different network. The receive load balancing is not a function of traffic load, but a function of the number of clients that are connected to the server.
What network protocols are load balanced when in a team?	Broadcom's teaming software only supports IP/TCP/UDP traffic. All other traffic is forwarded to the primary adapter.
Which protocols are load balanced with SLB and which ones are not?	Only IP/TCP/UDP protocols are load balanced in both directions: send and receive. IPX is load balanced on the transmit traffic only.
Can I team a port running at 100 Mbps with a port running at 1000 Mbps?	Mixing link speeds within a team is only supported for Smart Load Balancing™ teams and 802.3ad teams.
Can I team a fiber adapter with a copper Gigabit Ethernet adapter?	Yes with SLB, and yes if the switch allows for it in FEC/GEC and 802.3ad.
What is the difference between adapter load balancing and Microsoft's Network Load Balancing (NLB)?	Adapter load balancing is done at a network session level, whereas NLB is done at the server application level.

Question	Answer
Can I connect the teamed adapters to a hub?	Teamed ports can be connected to a hub for troubleshooting purposes only. However, this practice is not recommended for normal operation because the performance would be degraded due to hub limitations. Connect the teamed ports to a switch instead.
Can I connect the teamed adapters to ports in a router?	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.
Can I use teaming with Microsoft Cluster Services?	Yes. Teaming is supported on the public network only, not on the private network used for the heartbeat link.
Can PXE work over a virtual adapter (team)?	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.
What is the maximum number of ports that can be teamed together?	Up to eight ports can be assigned to a team.
What is the maximum number of teams that can be configured on the same server?	Up to four teams can be configured on the same server.
Why does my team loose 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.
Can I connect a team across multiple switches?	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.
How do I upgrade the intermediate driver (BASP)?	The intermediate driver cannot be upgraded through the Local Area Connection Properties. It must be upgraded using the Setup installer.

Question	Answer
Can I configure NLB and teaming concurrently?	Yes, but only when running NLB in a multicast mode (NLB is not supported with MS Cluster Services).
Should both the backup server and client servers that are backed up be teamed?	Because the backup server is under the most data load, it should always be teamed for link aggregation and failover. A fully redundant network, however, requires that both the switches and the backup clients be teamed for fault tolerance and link aggregation.
Is there any special configuration required in the tape backup software or hardware to work with adapter teaming?	No special configuration is required in the tape software to work with teaming. Teaming is transparent to tape backup applications.
How do I know what driver I am currently using?	In all operating systems, the most accurate method for checking the driver revision is to physically locate the driver file and check the properties.
Can SLB detect a switch failure in a Switch Fault Tolerance configuration?	No. SLB can only detect the loss of link between the teamed port and its immediate link partner. SLB cannot detect link failures on other ports.
Why does my team lose connectivity for the first 30 to 50 seconds after the primary adapter is restored (fall-back after a failover)?	During a fall-back event, link is restored causing Spanning Tree Protocol to configure the port for blocking until it determines that it can move to the forwarding state. You must enable Port Fast or Edge Port on the switch ports connected to the team to prevent the loss of communications caused by STP.
Where do I monitor real time statistics for an adapter team in a Windows server?	Use the BACS2 to monitor general, IEEE 802.3 and custom counters.

4.9 Event Log Messages

4.9.1 Windows System Event Log messages

The known base and intermediate Windows System Event Log status messages for this product are listed in "4.9.2 Base Driver (Physical Adapter/Miniport)" (→pg.151) and "4.9.3 Intermediate Driver (Virtual Adapter/Team)" (→pg.153). 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).

4.9.2 Base Driver (Physical Adapter/Miniport)

"• Base Driver Event Log Messages" (→pg.151) lists the event log messages supported by the base driver, explains the cause for the message, and provides the recommended action.

- 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: reseal 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.

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 BACS2 diagnostics; check that the network cable is good.

Message Number	Message	Cause	Corrective Action
17	Unknown PHY detected. Using a default PHY initialization routine.	The driver could not read the PHY ID.	Replace the adapter.
18	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.
19	Driver initialization failed.	Unspecified failure during driver initialization.	Reinstall the driver, update to a newer driver, run BACS2 diagnostics, or replace the adapter.

4.9.3 Intermediate Driver (Virtual Adapter/Team)

The intermediate driver is identified by BLFM regardless of the base driver revision. "• Intermediate Driver Event Log Messages" (→pg.153) lists the event log messages supported by the intermediate driver, explains the cause for the message, and provides the recommended action.

- Intermediate Driver Event Log Messages

Message Number	Message	Cause	Corrective Action
1	Unable to register with NDIS.	The driver cannot register with the NDIS interface.	Unload any 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 "Blf".
4	Broadcom Advanced Server Program Driver has started.	The driver has started.	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.

Message Number	Message	Cause	Corrective Action
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.
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.

Message Number	Message	Cause	Corrective Action
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 deactivated and is no longer participating in network traffic.	The driver does not recognize the installed adapter.	Informational message only. No action is required.

5 Broadcom Advanced Control Suite 2 (BACS2)

This chapter explains the Broadcom Advanced Control Suite 2 (BACS2).

5.1 BACS2 Overview

BACS2 is an integrated utility that provides useful information about each network adapter that is installed in your system. BACS2 also enables you to perform detailed tests, diagnostics, and analyses on each adapter, as well as to view and modify property values and view traffic statistics for each adapter.

BACS2 contains three panes:

- In the Information/Task pane, users can view available information and perform certain tests, diagnostics, and analysis on a selected device by clicking a specific tab.
- To the left of the Information/Task pane is the Device Name pane, which lists the names of the individual network adapters and the individual members of teams that have been created.
- A third pane contains the Menu bar.

Broadcom Advanced Server Program (BASP), which runs within Broadcom Advanced Control Suite 2 (BACS2), is used in Expert Mode to configure teams for load balancing, fault tolerance, and virtual local area networks (VLANs). The Teaming Wizard can also be used to configure teams and VLANs. BASP functionality is available only on systems that use at least one network adapter.




- ▶ Link aggregation is not supported in the Blade Server.

5.1.1 Types of Information Provided by BACS2

BACS2 lists all of the network adapters in your system and provides the following information (if available) about each adapter:

- Driver Status
- IP Address
- Speed
- Driver Name
- Firmware Version
- Bus No.
- Interrupt Request
- VLAN Name
- MAC Address
- Memory Address
- Duplex
- Driver Version
- ASIC Version
- Device No.
- Team Name
- VLAN ID
- Link Status
- Physical Address
- Slot No.
- Driver Date
- Bus Type
- Function No.
- Team Type
- Properties

The details of a function provided by BACS2 are shown below.

Function	Details
Vital Sign	At-a-glance information on all of the LAN adapters in your system. →"5.4.1 Vital Sign"(pg.162)
Resources	Shows the resource settings for the selected adapter. →"5.4.2 Resources"(pg.165)
Hardware	Shows the hardware information for the selected adapter. →"5.4.3 Hardware"(pg.166)
Advanced	Shows the available properties and their values for the selected adapter. →"5.4.4 Advanced"(pg.167)
Network Test	Confirms network connectivity to a remote station. →"5.4.5 Network Test"(pg.170)
Diagnostics	Performs comprehensive diagnostics. →"5.4.6 Diagnostics"(pg.170)
Statistics	Provides detailed performance statistics for the selected adapter. →"5.4.7 Statistics"(pg.172)
Resource Allocations	Displays a pie chart of the allocated TOE, iSCSI, and RDMA connections, as well as the unallocated and unlicensed resources. Only available with this product.  ► This setting is not supported.
Licenses	Displays licensing information for the TOE, iSCSI, and RDMA technologies. Only available with this product. →"5.4.8 Licenses"(pg.175)

5.2 Installing the BACS2

If [Broadcom Control Suite 2] is not displayed in the "Control Panel", install BACS2 according to the following installation procedures:

POINT

- ▶ Ensure that this product is installed in the Server Blade before installing BACS2.
- ▶ Before you begin the installation, close all applications, windows, or dialog boxes.
- ▶ To use TCP/IP Offload Engine (TOE), you must have Windows Server 2003 with Scalable Networking Pack (SNP)

IMPORTANT


- ▶ The installer checks if SNP for Windows Server 2003 is installed on the machine. If it is installed, the installer installs the NDIS 5.2 driver, which is necessary in order to use TOE. If SNP for Windows Server 2003 is not installed on the machine, the installer will install the NDIS 5.1 driver, a user cannot use TOE.
- ▶ Get the Broadcom Control Suite from the ServerStart CD-ROM included with the BX620 S4 server blade and install.
If using Broadcom, refer to the operation manual or help file included with the driver.

For Windows Server 2003 x64

1 Start the following EXE file from the ServerStart Disc1 CD-ROM attached to the BX620 S4.

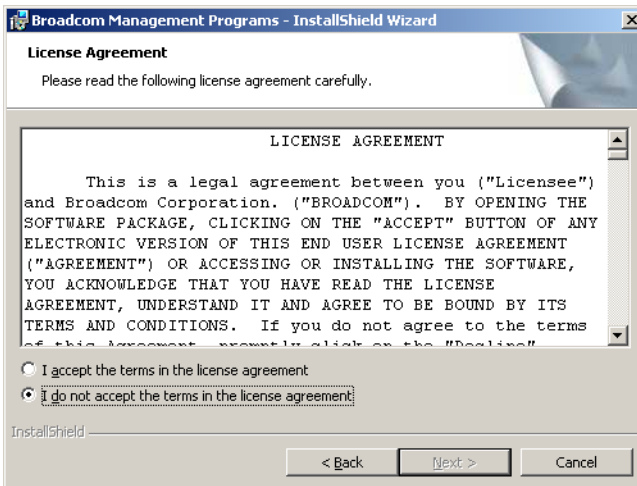
[CD-ROM drive]: \PROGRAMS\GENERAL\Broadcom\MgmtApps_x64\setup.exe

The installer starts up.



2 Click [Next].

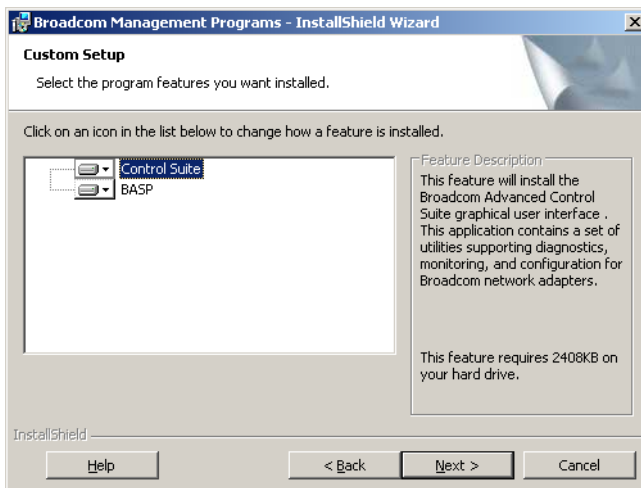
License agreement window appears



E

3 Click "I accept the terms in the license agreement" and click [Next].

Custom Set up window appears.



4 Click [Next].

Proceed the installation by following the window instructions.

For Windows Server 2003, Windows 2000 Server

POINT

- ▶ If the OS is installed using ServerStart, "BACS2" is already installed with the driver. If the OS is installed manually, BACS2 will not be installed.

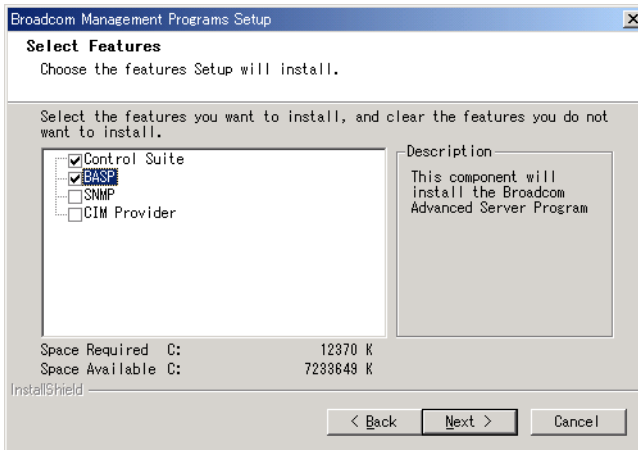
1 Start the following EXE file from the ServerStart Disc1 CD-ROM attached to the BX620 S4.

[CD-ROM drive]: \PROGRAMS\GENERAL\Broadcom\MgmtApps\setup.exe

The installer starts up.

2 Proceed the installation by following the window instructions.

When the window below appears during the installation procedure, check [BASP] and click [Next].



5.3 Starting BACS2

To start BACS2, click "Control Panel" → "Broadcom Control Suite 2".

Click the tab that provides the information of interest or from which to perform a desired test, diagnostic, analysis, or set adapter properties.

To create a team, from the "Tools" menu, click [Create a Team], which starts the Teaming Wizard.


5.4 Setting of BACS2

5.4.1 Vital Sign

The "Vital Sign" tab shows useful information about 1Gbit/s Ethernet I/O Module, Onboard LAN, and other network adapters that are installed in your system. Such information includes the link status of the adapter and network connectivity. To view this information for each network adapter, click the name of the adapter listed in the "Name" pane.

POINT

- ▶ Information about Intel LAN card is less comprehensive than the information provided for this product.
- ▶ Some information may not be available for all network adapters.

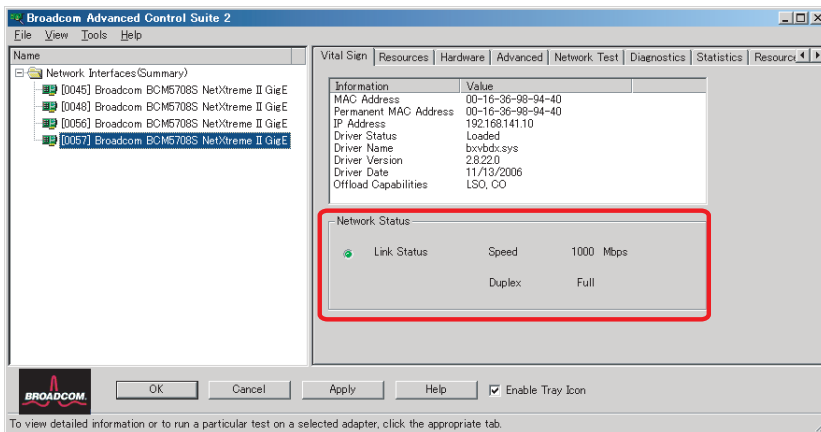


Item	Description
MAC Address	This is the physical MAC address that is assigned to the adapter by the manufacturer. The physical address is never all 0s.
Permanent MAC Address	The unique hardware address assigned to the network adapter.
IP Address	The network address that is associated with the adapter. If the IP address is all 0s, this means that the associated driver has not been bound with Internet Protocol (IP).

Item	Description
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/Version/Date	The file name of the adapter driver.
LiveLink IP Address	The network address of the LiveLink enabled adapter.
Offload Capabilities	The offload capabilities supported by the adapter.
LSO	Large Send Offload prevents an upper level protocol such as TCP from breaking a large data packet into a series of smaller packets with headers appended to them.
CO	Checksum Offload allows the TCP/IP/UDP checksums for send and receive traffic to be calculated by the adapter hardware rather than by the host CPU.
BASP State	Information about the status of the BASP application. This information is displayed only when there is a team (→"1.3.3 Teaming Function"(pg.109)).

Network Status

The following network status information is provided.

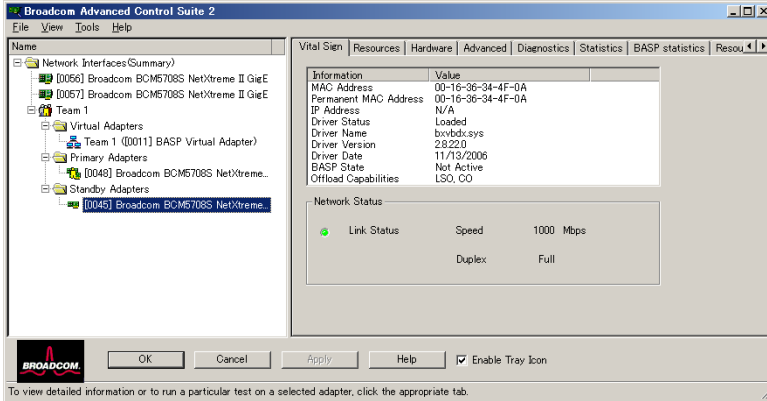


Item	Description
Link Status	The indicator is green if a link is established. A red indicator means that a link is not established.
Speed	The link speed of the adapter
Duplex	The duplex mode in which the adapter is operating

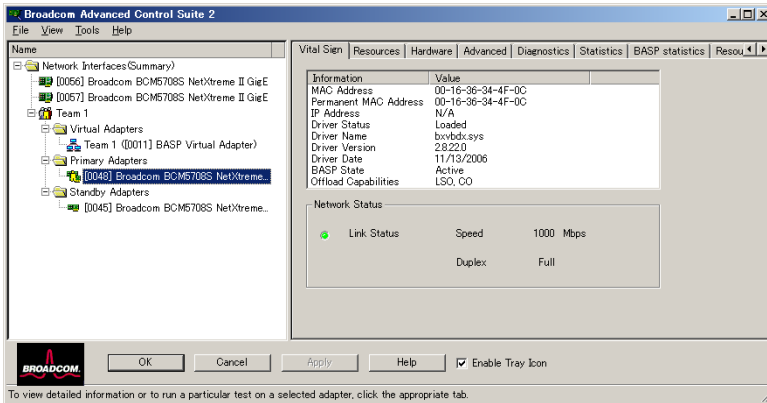
E

Team Status

The team status is indicated by the appearance of the icons representing the team and the team members. If the adapter icon shows only the adapter, the adapter is connected to a network, but is not correctly participating in the team, which corresponds to a Not Active BASP state. This condition does not apply to an adapter that is a standby member of a team. The standby adapter may be working properly even though it is Not Active.



If the adapter icon shows a superimposed running yellow figure, the adapter is connected and participating in the team correctly, which corresponds to an Active BASP state.




If the adapter icon shows a superimposed red letter X, the adapter is not connected to the network.

5.4.2 Resources

The following information can be checked on the "Resources" tab.

POINT

- ▶ Some information may not be available for all network adapters.



Item	Description
Bus Type	The type of input/output (I/O) interconnect used by the adapter
Bridge	The bridge type, which is the PCI-E to PCI-X bridge.
Bridge Lanes	The number of PCI-E lanes connected to the bridge.
Bridge Speed	The clock speed on PCI-E bus.
Slot No	The slot number on the system board occupied by the adapter.
Bus Speed (MHz)	The bus clock signal frequency used by the adapter.
Bus Width (bit)	The number of bits that the bus can transfer at a single time to and from the adapter.
Bus No	Indicates the number of the bus in which 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.

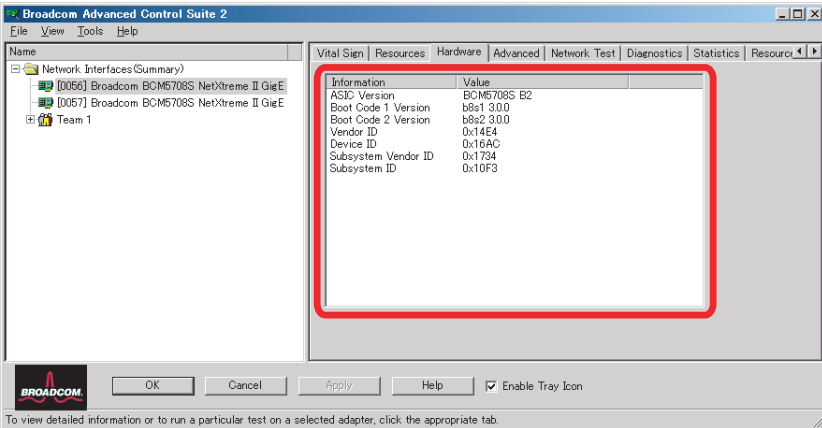
E

5.4.3 Hardware

The following information can be checked on the "Hardware" tab.




- ▶ Some information may not be available for all network adapters.



Item	Description
ASIC Version	The chip version of the Broadcom adapter (this information is not available for adapters made by others).
Boot Code 1 Version	The version of boot code 1.
Boot Code 2 Version	The version of boot code 2.
Vendor ID	The vendor ID.
Device ID	The adapter ID.
Subsystem Vendor ID	The subsystem vendor ID.
Subsystem ID	The subsystem ID.

5.4.4 Advanced

The following information can be checked on the "Advanced" tab.



The "Advanced" tab allows 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. To view the value of a property, click the name of the property in the "Property" list. The property value is displayed in the "Value" box. To change the value, click an item in the "Value" list or type a new value, as appropriate (selection options are different for different properties).

POINT

- ▶ 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 network adapters.

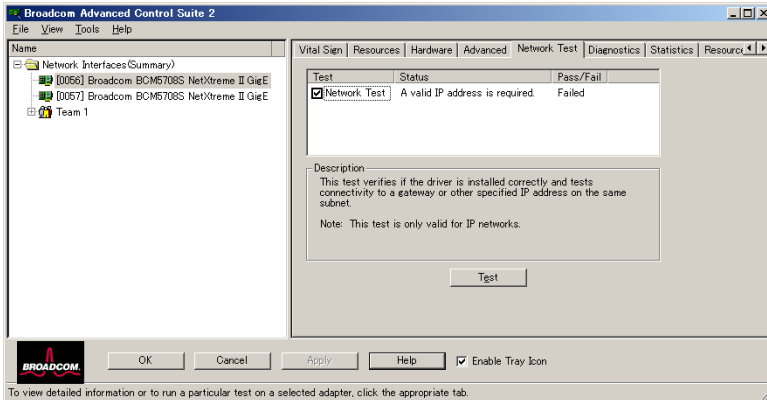
Item	Description
802.1p QoS	The 802.1p QoS property enables quality of service, which is an Institute of Electrical and Electronics Engineering (IEEE) specification that treats different types of network traffic differently to ensure required levels or reliability and latency according to the type of traffic. This property is disabled by default. Unless the network infrastructure supports QoS, do not enable QoS. Otherwise, problems may occur.

Item	Description
Checksum Offload	<p>Normally, the checksum function is computed by the protocol stack. When you select one of the Checksum Offload property values (other than None), the checksum can be computed by the network adapter.</p> <ul style="list-style-type: none"> • Rx TCP/IP Checksum Enables receive TCP/IP/UDP checksum offload. • Tx TCP/IP Checksum Enables transmit TCP/IP/UDP checksum offload. • Tx/Rx TCP/IP Checksum (default) Enables transmit and receive TCP/IP/UDP checksum offload. • None Disables checksum offload.
Flow Control	<p>The Flow Control property enables or disables the receipt or transmission of PAUSE frames. PAUSE frames enable the network adapter and a switch to control the transmit rate. The side that is receiving the PAUSE frame momentarily stops transmitting.</p> <ul style="list-style-type: none"> • 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.
Large Send Offload	<p>Normally the TCP segmentation is done by the protocol stack. When you enable the Large Send Offload property, the TCP segmentation can be done by the network adapter.</p> <ul style="list-style-type: none"> • Disable Disables Large Send Offload • Enable. (default) Enables Large Send Offload

Item	Description
Locally Administered Address	<p>The Locally Administered Address is a user-defined MAC address that is used in place of the MAC address originally assigned to the network adapter. Every adapter in the network must have its own unique MAC address. This locally administered address consists of a 12-digit hexadecimal number.</p> <ul style="list-style-type: none"> • Value Assigns a unique node address for the adapter • Not Present (default) Uses the factory-assigned node address on the adapter <p>The appropriate assigned ranges and exceptions for the locally administered address include the following:</p> <ul style="list-style-type: none"> • The range is 00:00:00:00:00:01 to FF:FF:FF:FF:FF:FD. • Do not use a multicast address (least significant bit of the high byte = 1). • Do not use all 0s or all F's.
Maximum Transfer Unit	<p>The Maximum Transfer Unit property allows the network adapter to transmit and receive oversized Ethernet frames that are greater than 1514 bytes, but less than or equal to 9000 bytes in length. This property requires the presence of a switch that is able to process jumbo frames. Frame size is set at 1500 bytes by default. To increase the size of the received frames, increment the byte quantity in 500-byte increments.</p>
Speed & Duplex	<p>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.</p> <ul style="list-style-type: none"> • 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 Sets the speed and mode for optimum network connection

5.4.5 Network Test

On the "Network Test" tab, you can 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. Network Test uses TCP/IP. The network test sends ICMP packets to remote systems and waits for a response. If a gateway is configured, the test automatically sends packets to that system. If a gateway is not configured or if the gateway is unreachable, the test prompts you for a destination IP address.

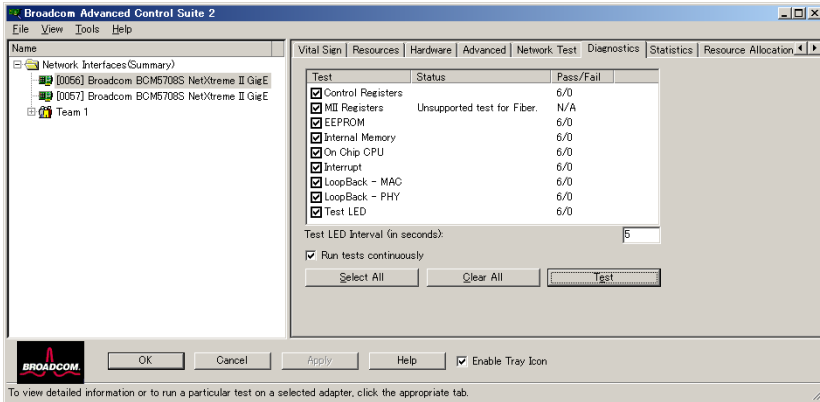


5.4.6 Diagnostics

On the "Diagnostics" tab, you can perform diagnostic tests on the physical components of a network adapter. The tests are continuously performed. The number of passes and fails in the "Pass/Fail" column increments each time the tests are performed. For example, if a test is performed four times and there are no fails, the value in the "Pass/Fail" column is 4/0. If there were 3 passes and 1 fail, however, the value in the "Pass/Fail" column is 3/1.

POINT

- ▶ You must have Windows administrator privileges to perform diagnostics.
- ▶ The network connection is temporarily lost when these tests are running.




Item	Description
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 for configuring network speeds such as 1000 Mbit/s.
EEPROM	This test verifies the content of the electrically erasable programmable read-only memory (EEPROM) by reading a portion of the EEPROM and computing the checksum. The test fails if the computed checksum is different from the checksum stored in the EEPROM. An EEPROM image upgrade does not require a code change for this test.
Internal Memory	This test verifies that the internal memory of the adapter is functioning properly. The test writes patterned values to the memory and reads back the results. The test fails if an erroneous value is read back. The adapter cannot function if its internal memory is not functioning properly.
On-Chip CPU	This test verifies the operation of the internal CPUs in the adapter.
Interrupt	This test verifies that the Network Device Driver Interface Specification (NDIS) driver is able to receive interrupts from the adapter.
Loopback MAC and Loopback PHY	These tests verify that the NDIS driver is able to send packets to and receive packets from the adapter.
Test LED	This test causes all of the port LEDs to blink 5 times for the purpose of identifying the adapter.

5.4.7 Statistics

On the "Statistics" tab, you can view traffic statistics for both Broadcom network adapters and network adapters made by others. Statistical information and coverage are more comprehensive for Broadcom adapters.

POINT

- ▶ Some statistics may not be available for all network adapters.



General Statistics

Item	Description
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, or 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 therefore, are excluded.
Directed Frames Rx.	A count of directed data frames that were successfully received.

Item	Description
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, or 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, or 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.

IEEE 802.3 Statistics

Item	Description
Frames Rx. with Alignment Error	A count of the frames that were not an integral number of octets in length and do not pass the FCS check. This counter is incremented when the receive status is reported as Alignment Error.
Frames Tx. with One Collision	A count of the frames that were involved in a single collision and were subsequently transmitted successfully. This counter is incremented when the result of a transmission is reported as Transmit OK, and the attempt value is 2.
Frames Tx. with more than One Collision	A count of the frames that were involved in more than one collision and were subsequently transmitted successfully. This counter is incremented when the transmit status is reported as Transmit OK, and the value of the attempts variable is greater than 2 and less than or equal to the attempt limit.
Frames Tx. after Deferral	A count of the frames that were delayed being transmitted on the first attempt because the medium was busy. The frames involved in any collision are not counted.

Custom Statistics

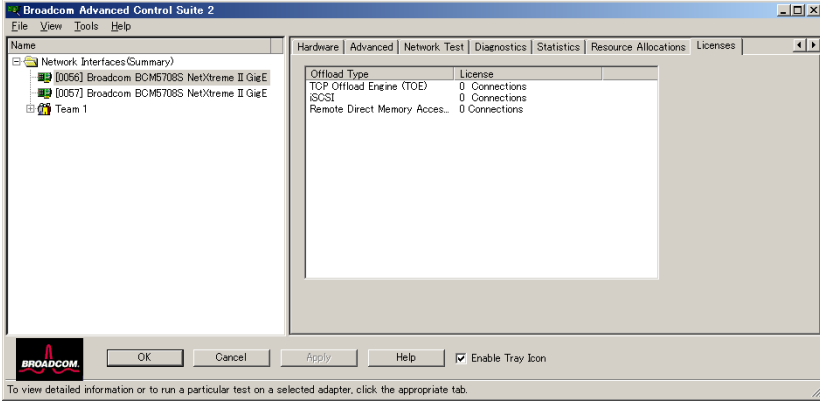
Item	Description
Out of Recv. Buffer	The number of times the adapter ran out of Receive Buffer Descriptors.
Frames size less than 64-byte with bad FCS	The number of frames with a size less than 64 bytes with bad FCS.
MAC Rx w/ Pause Command and Length = 0	MAC control frames with the pause command and a length equal to 0.
MAC Rx w/ Pause Command and Length greater than 0	MAC control frames with the pause command and a length greater than 0.
MAC Rx w/ no Pause Command	MAC control frames with no pause command.
MAC Sent X-on	MAC Transmit with X-on was on.
MAC Sent X-off	MAC Transmit with X-on was off.
Large Send Offload Transmit Requests	The number of times the adapter was requested to transmit a packet performing TCP segmentation.

5.4.8 Licenses

On the "Licenses" tab, you can view the number of connections available for TOE, iSCSI, and RDMA offload technologies. You can also upgrade your license for additional connections.

POINT

- ▶ The "Licenses" tab is only available for this product.
- ▶ Not all offload technologies are available with all adapters.



5.5 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 LAN (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 →"4 Broadcom Gigabit Ethernet Teaming Services"(pg.119)

Teaming can be accomplished by either of the following methods:

- Using the Broadcom Teaming Wizard
→"5.5.2 Creating and Modifying a Team Using the Teaming Wizard"(pg.177)
- Using Expert Mode
→"5.5.3 Using Expert Mode"(pg.186)

POINT

- ▶ 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 does not support Microsoft Network Load Balancing (NLB).
- ▶ 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.

For more details about Teaming, refer to "1.3.3 Teaming Function" (→pg.109).

IMPORTANT

- ▶ Link aggregation is not supported in the Blade Server.

5.5.1 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.

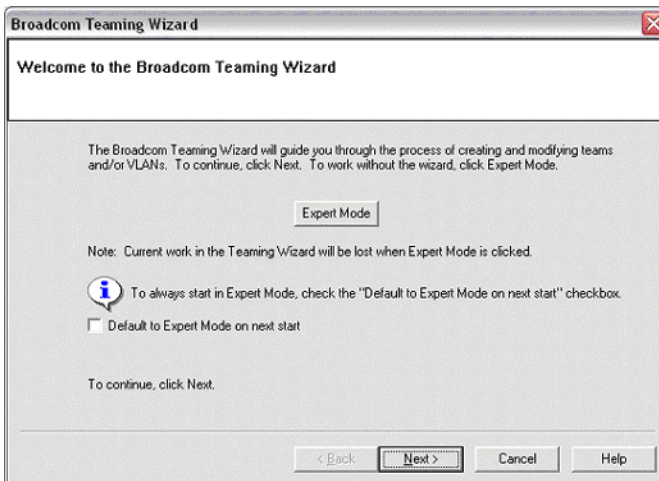
5.5.2 Creating and Modifying a Team Using the Teaming Wizard

1 On the BACS2 "Tools" menu, click "Create a Team".

POINT

- ▶ If you prefer to work without the wizard for now, click [Expert Mode] and then click [Next]. If you want to always use Expert Mode to create a team, select "Default to Expert Mode on next start" and then click [Next]. Refer to "5.5.3 Using Expert Mode" (→pg.186).

2 To continue using the wizard, click [Next].

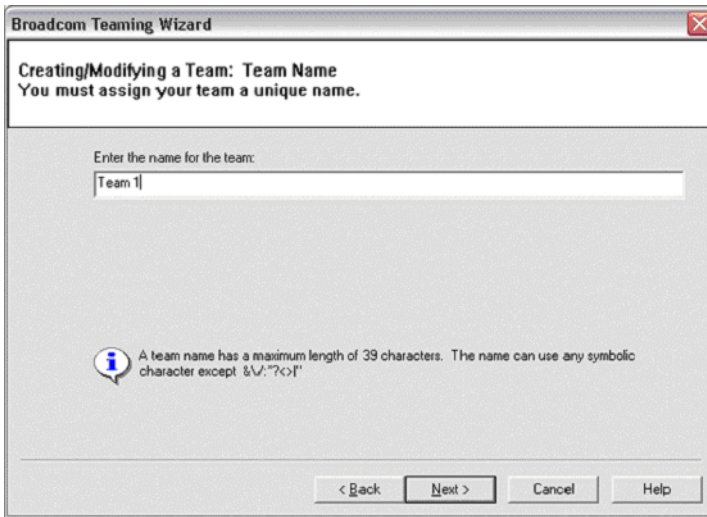


3 Type the team name and then click [Next].

- If you want to review or change any of your settings, click [Back].
- Click [Cancel] to discard your settings and exit the wizard.

POINT

- ▶ 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 and then click [Next].

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].

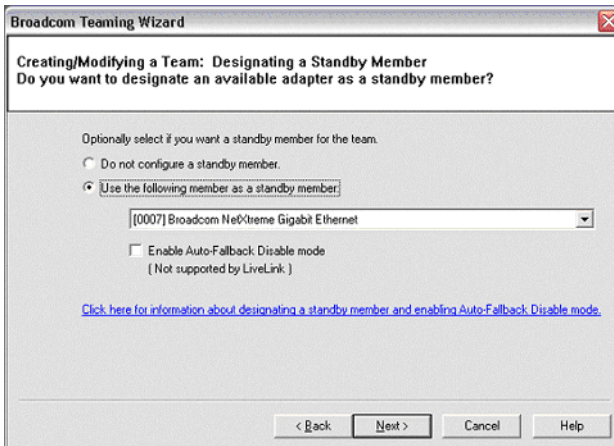
POINT

- ▶ There must be at least one 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 (Y) or not supported (N) 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.



- 6** If you want to designate one of the adapters as a standby member (optional), then click "Use the following member as a standby member".

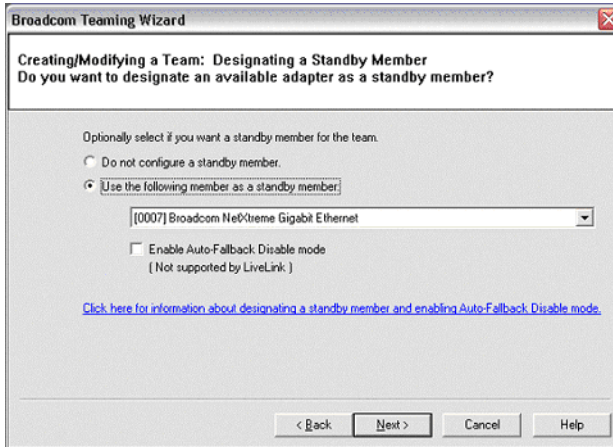


7 Select the standby member from the list of adapters.

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.

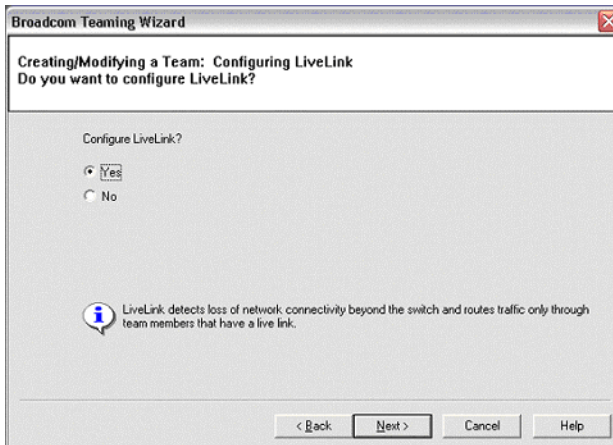
To enable this feature, click "Enable Auto-Fallback Disable mode" and then click [Next].

The Auto-Fallback Disable mode feature is enabled, LiveLink cannot be used.



8 If you want to configure LiveLink, click [Yes] and then click [Next].

If you do not want to configure LiveLink, click [No], and click [Next].



9 Set 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 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].**

POINT

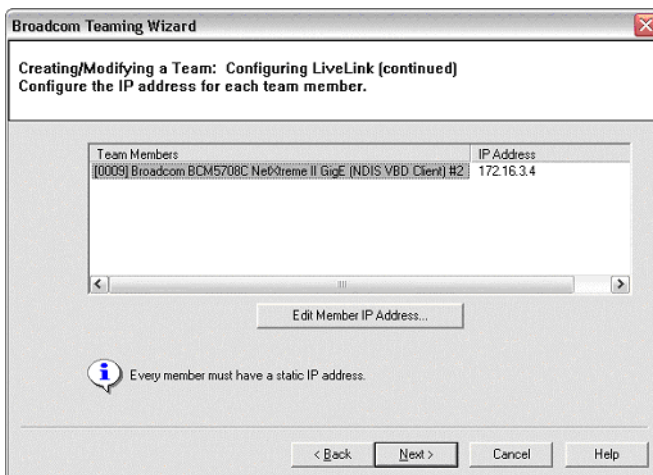
- ▶ Only the first probe target is required. You can specify up to three additional probe targets to serve as backups by assigning IP addresses to the other probe targets.

Target	IP Address
Probe Target 1	172.16.1.2
Probe Target 2	
Probe Target 3	
Probe Target 4	

- 11 Click 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].**

POINT

- ▶ All of the member IP addresses must be in the same subnet as the subnet of the probe targets.



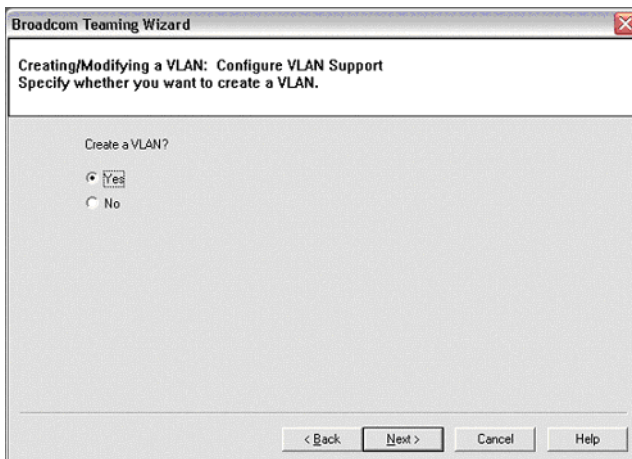
12 If you want to create a VLAN on the team, click [Yes] and then click [Next].

If you do not want to create a VLAN to the team, click [No], click [Next], and continue with the wizard from the "Finish" screen (refer to Step 17 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.

POINT

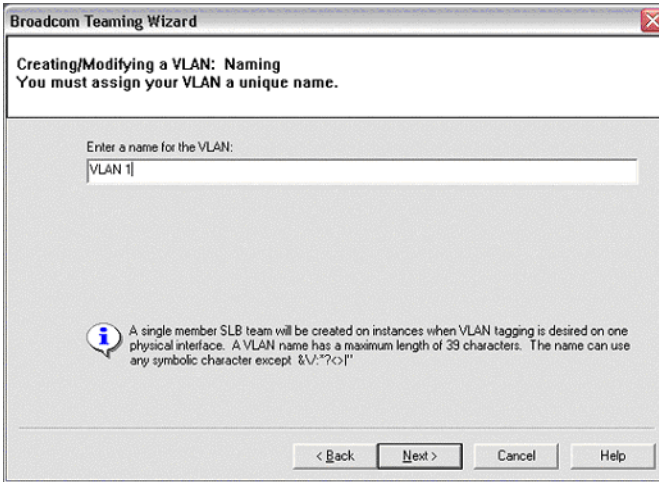
- ▶ VLANs can only be created when all team members are Broadcom adapters.



13 Type the VLAN name and then click [Next].

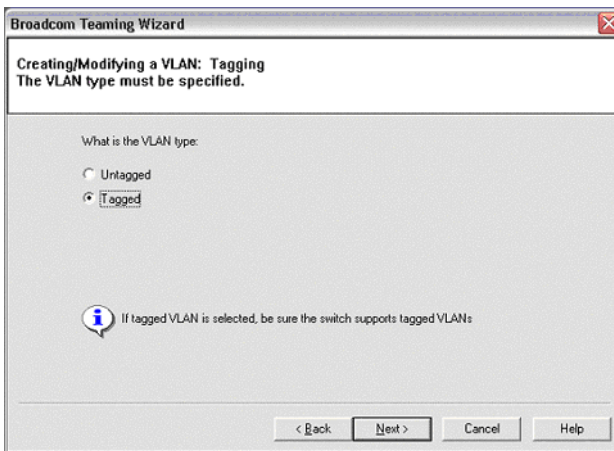
POINT

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




14 To tag the VLAN, click "Tagged" and then click [Next].

Otherwise, click "Untagged", click [Next], and continue with the wizard to add additional VLANs (refer to Step 16 of this procedure).



15 Type the VLAN tag value and then click [Next].

The value must be between 1 and 4094.

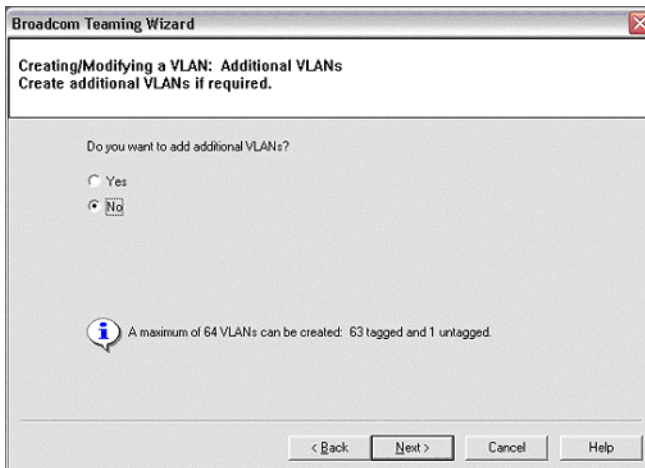


16 Click [Yes] to add another VLAN and then click [Next].

Repeat until you do not want to add any additional VLANs.

POINT

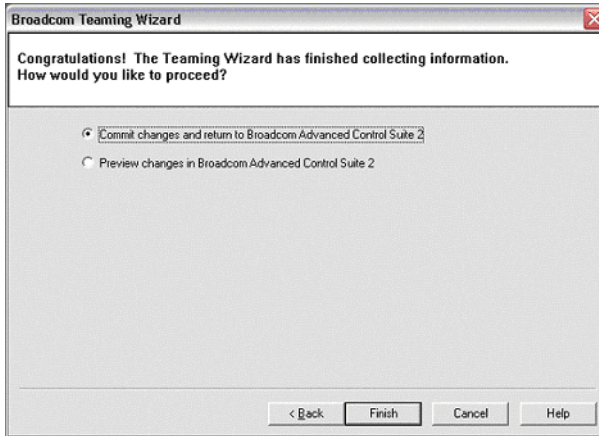
- ▶ You can define up to 64 VLANs per team (63 VLANs that are tagged and 1 VLAN that is not tagged). Adding several VLANs may slow down the reaction time of the Windows interface due to memory and processor time usage for each VLAN. The degree to which Windows performance may suffer depends on system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



17 To apply and commit the changes to the team, click "Commit changes and return to Broadcom Advanced Control Suite 2" and then click [Finish].

To preview the changes to the team in BACS2, click "Preview changes in Broadcom Advanced Control Suite 2" and then click [Finish].

The wizard exits and BACS2 opens.

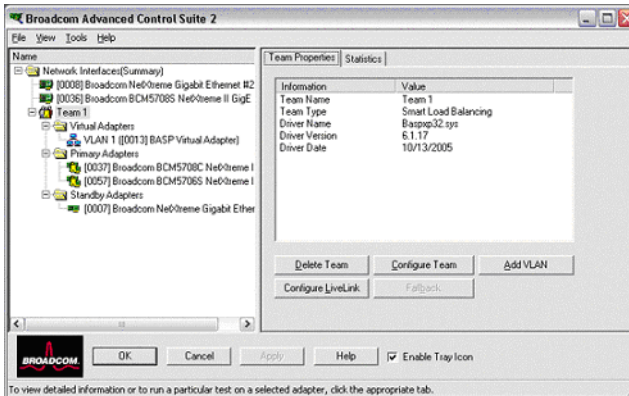


18 Click [Finish] to commit the changes.

Click [Cancel] to discard the changes.

Click the team to view the team's statistics from the "Statistics" tab and the team properties from the "Team Properties" tab.

Click the VLAN to view the properties of the VLAN from the "Vital Signs" tab.



5.5.3 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 team. To create a team using the wizard, refer to "5.5.2 Creating and Modifying a Team Using the Teaming Wizard" (→pg.177). To switch to the Teaming Wizard from the BACS2 "Tools" menu, click "Customize". Select the "Default Teaming Mode" tab and select "Wizard Mode".

Creating a Team

POINT

- ▶ Enabling Dynamic Host Configuration Protocol (DHCP) is not recommended for members of an SLB type of team.

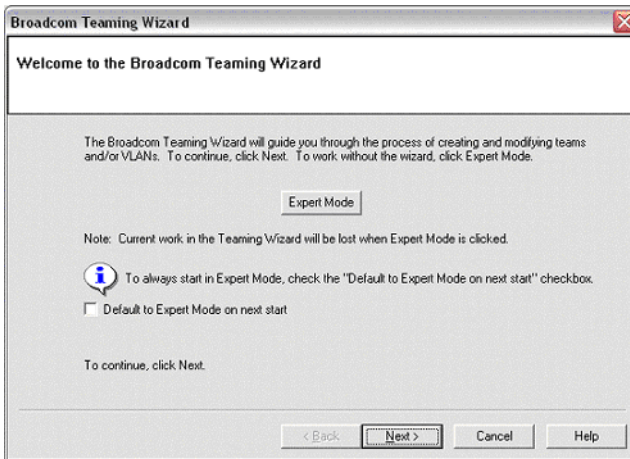
1 Click the BACS2 "Tools" menu → "Create a Team".

The wizard's Welcome screen appears.

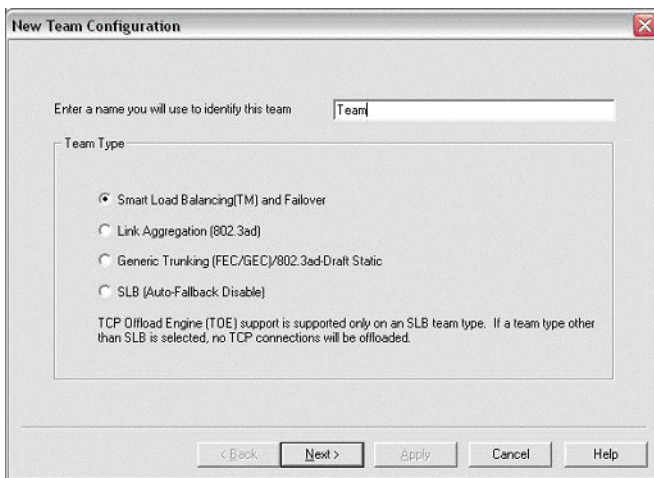
2 To work without the wizard, click "Expert Mode".

POINT

- ▶ If you want to always use Expert Mode to create a team, check "Default to Expert Mode on next start". Click [Next].




- 3 Type a team name in the "Enter a name you will use to identify this team" box, and click the type of team, and then click [Next].



- 4 Assign any available adapter or adapters to the team by moving the adapter from the "Available Adapters" list to the "Load Balance Members" list.
- 5 Assign any other available adapter or adapters to the team by moving the adapter from the "Available Adapters" list to the "Standby Member" list.

POINT

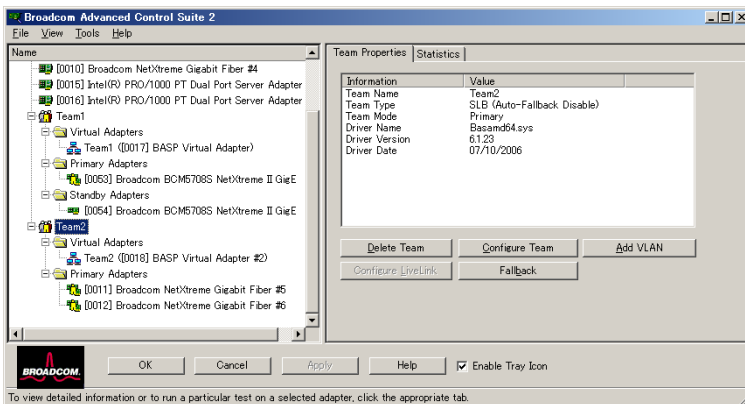
- ▶ There must be at least one network adapter assigned to the team.



6 Click [Yes] when the message is displayed indicating that the network connection will be temporarily interrupted.

POINT

- ▶ 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.
- 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.



7 Configure the team IP address.

1. Click "Control Panel" → "Network Connections".
2. Right-click the name of the team to be configured, and then click "Properties".
3. On the "General" tab, click "Internet Protocol (TCP/IP)", and then click "Properties".
4. 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)

The following shows how to modify a team.

- 1 Click the BACS2 "Tools" menu → "Create a Team".**
The wizard Welcome screen appears.
- 2 Click [Next] to continue modifying a team using the wizard or click [Expert Mode] to work in Expert Mode.**
- 3 Make the desired changes, and then click [OK].**
- 4 Click [Apply].**
- 5 Click [Yes] when the message is displayed indicating that the network connection will be temporarily interrupted.**

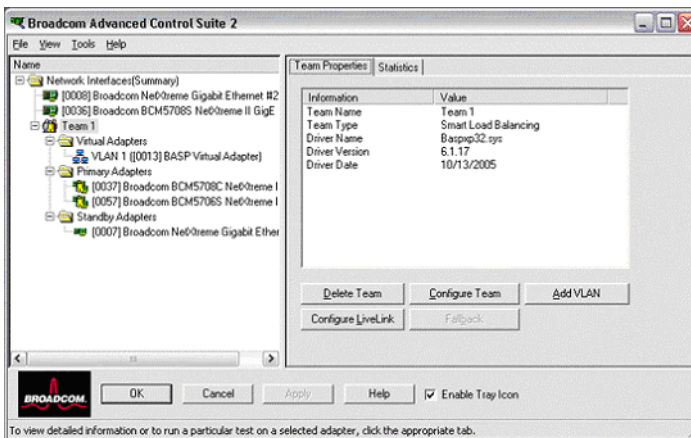
Adding a VLAN

You also 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.

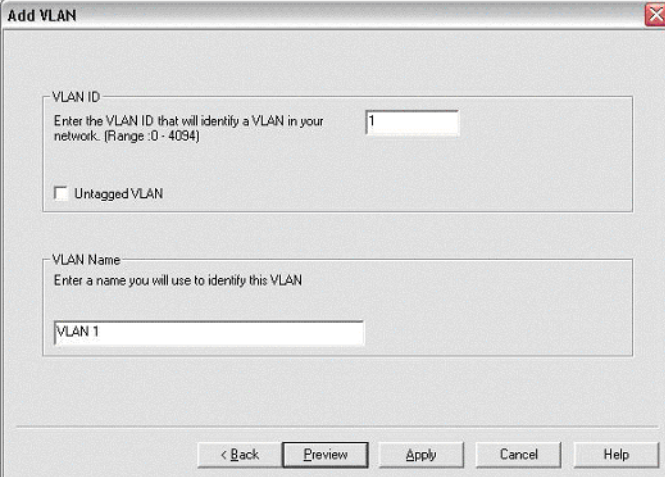
You can define up to 64 VLANs per team (63 VLANs that are tagged and 1 VLAN that is not tagged). VLANs can only be created when all teams members are Broadcom adapters. If you try to create a VLAN with a non-Broadcom adapter, an error message is displayed.

The following shows how to configure a team with a VLAN.

- 1 Click the name of the team you want to configure.**
- 2 From the "Team Properties" tab, click [Add VLAN].**



3 Type the VLAN ID and VLAN name, and click [Apply].



The screenshot shows a dialog box titled "Add VLAN". It has two main sections. The first section is labeled "VLAN ID" and contains the text "Enter the VLAN ID that will identify a VLAN in your network. (Range 0 - 4094)". Below this text is a text input field containing the number "1". There is also a checkbox labeled "Untagged VLAN" which is currently unchecked. The second section is labeled "VLAN Name" and contains the text "Enter a name you will use to identify this VLAN". Below this text is a text input field containing "VLAN 1". At the bottom of the dialog box, there are five buttons: "< Back", "Preview", "Apply", "Cancel", and "Help".

POINT

- ▶ If you type a VLAN name or ID and the name already exists, an error message is displayed.

4 Click [Yes] when the message is displayed indicating that the network connection will be temporarily interrupted.

POINT

- ▶ 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 Running VLAN Tests

The following shows how to view VLAN properties and to run VLAN tests.

- 1 Click the name of the VLAN adapter of interest.
- 2 Click the "Vital Sign" tab to view the properties of the VLAN adapter.
- 3 Click the "Network Test" tab to run a network test on the VLAN adapter.

Deleting a VLAN

To delete a VLAN, right-click the VLAN and select "Remove VLAN".

1 Click the "Tools" menu → "Configure a Team".

If there is more than one team, click the name of the team that has the VLAN you want to delete, and then click [OK].

2 Click [Remove VLAN].

3 Click [Apply].

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

POINT

- ▶ If you delete a team, any VLANs configured for that team are also deleted.

Configuring LiveLink for a Smart Load Balancing and Failover Team

Read the following notes before you attempt to configure LiveLink.

POINT

- ▶ 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. (→"4.8.10 Troubleshooting Teaming Problems"(pg.147).
- ▶ 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.

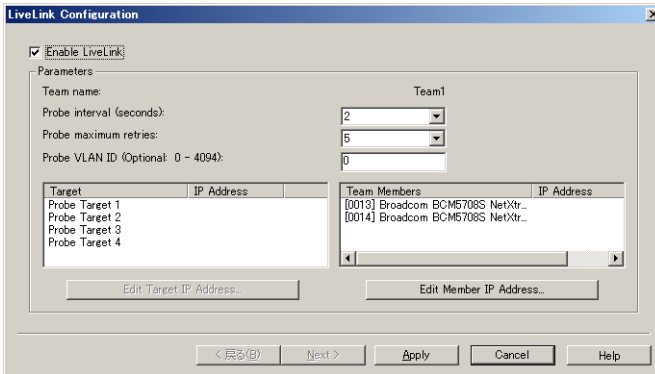
The following shows how to configure a LiveLink.

1 Right-click the name of the Smart Load Balance and Failover (SLB) team, and then click "Configure LiveLink".

2 Check the "Enable LiveLink" box.

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).

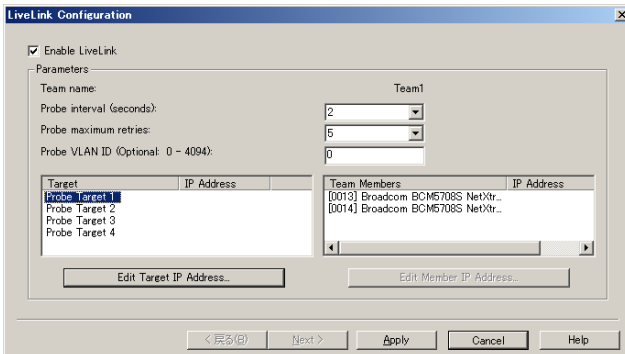
To specify different values, click the desired probe interval in the "Probe interval (seconds)" list and click the desired maximum number of probe retries in the "Probe maximum retries" list.



3 Click the probe target at the top of the list, click "Edit Target IP Address", type the target IP address for one or all probe targets in the "IP Address" box, and then click [OK].

POINT

- ▶ 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.



- 4 Click one of the listed team members, click "Edit Member IP Address", type the member IP address in the "IP Address" box, and then click [OK].**

 **POINT**

- ▶ All of the member IP addresses must be in the same subnet as the subnet for the probe targets.

- 5 Repeat step 4 for each of the other listed team members.**
- 6 Click [Apply].**

Configuring LiveLink in VLAN-tagged Environments

 **CAUTION**



- For the teams with VLANs (on which LiveLink is enabled): to be able to communicate with the probe target, both the probe target and the team must be on an untagged VLAN (VLAN ID 0). Otherwise, the team loses connectivity.

- 1 Ensure that the team has an untagged VLAN (VLAN ID 0).**
- 2 Ensure there is network connectivity between the team and the probe target on the untagged VLAN.**
- 3 Right-click the name of the SLB team with VLAN(s), and then click "Configure LiveLink".**
- 4 Click the "Enable LiveLink" box.**
- 5 Click the "desired probe interval" (the number of seconds between each retransmission of the link packet to the probe target) in the "Probe interval (seconds)" list.**
- 6 Click the desired maximum number of probe retries in the "Probe maximum retries" list.**
- 7 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, and then click [OK].**

 **POINT**

- ▶ It is not necessary to specify more than one probe target. If you do want to specify more than one, for each additional probe target (up to a total of 4), click the next probe target in the list, type the target IP address in the "IP Address" box, and then click [OK].

- 8 Click one of the listed team members, click "Edit Member IP Address", type the member IP address in the "IP Address" box, and then click [OK].**

 **POINT**

- ▶ All of the member IP addresses must be in the same subnet as the subnet for the probe targets.

- 9 Repeat step 8 for each of the other listed team members.**
- 10 Click [Apply].**

Viewing the Team Properties and Statistics

The following shows how to view the team properties and statistics.

- 1 Click the name of the newly created team.**
- 2 Click the "Statistics" tab to view the team statistics.**

Saving and Restoring a Configuration

The following shows how to save a configuration.

- 1 Click the "File" menu → "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.

The following shows how to restore a configuration.

- 1 Click the "File" menu → "Team Restore".**
- 2 Click the name of the file to be restored, and then click [Open].**

 **POINT**

- ▶ If necessary, go to the folder where the file is located.

- 3 Click [Apply].**
- 4 Click [Yes] when the message is displayed indicating that the network connection will be temporarily interrupted.**

If a configuration is already loaded, a message is displayed that asks if you want to save your current configuration.

Click [Yes] to save the current configuration. Otherwise, the configuration data that is currently loaded is lost.

Appendix A Contact Information

- **Australia:**
Fujitsu Australia Limited
Tel: +61-2-9776-4555
Fax: +61-2-9776-4556
Address: 2 Julius Avenue (Cnr Delhi Road) North Ryde, Australia N.S.W. 2113
- **China:**
Fujitsu (China) Holdings Co., Ltd.
Tel: +86-21-5292-9889
Fax: +86-21-5292-9566
Address: 18F, Citic Square, 1168 West Nanjing Road Shanghai, China 200041
- **Hong Kong:**
Fujitsu Hong Kong Limited
Tel: +852-2827-5780
Fax: +852-2827-4724
Address: 10/F., Lincoln House, 979 King's Road Taikoo Place, Island East, Hong Kong
- **Indonesia:**
PT. Fujitsu Systems Indonesia Offices Headquarters
Tel: +62-21-570-9330 (Hunting)
Fax: +62-21-573-5150
Address: Wisma Kyoei Prince 10th Floor Jl. Jend. Sudirman Kav 3-4 Jakarta, Indonesia 10220
- **Korea:**
Fujitsu Korea Ltd.
Tel: +82-2-3787-6000
Fax: +82-2-3787-6066
Address: Susong Tower Building, 83-1 Susong-Dong Jongno-Gu, Seoul, Republic of Korea 110-140
- **Malaysia:**
Fujitsu (Malaysia) Sdn. Bhd.
Tel: +60-3-8318-3700
Fax: +60-3-8318-8700
Address: 1st Floor, No.3505 Jalan Technokrat 5 63000 Cyberjaya, Selangor Darul Ehsan Malaysia
- **Philippines:**
Fujitsu Philippines, Inc.
Tel: +63-2-812-4001
Fax: +63-2-817-7576
Address: 2nd Floor, United Life Building, A. Arnaiz Legaspi Village, Makati, Metro Manila Philippines

- Singapore:
 Fujitsu Asia Pte. Ltd.
 Tel: +65-6777-6577
 Fax: +65-6771-5502
 Address: 20, Science Park Road, #03-01 TeleTech Park, Singapore Science Park II,
 Singapore 117674
- Taiwan:
 Fujitsu Taiwan Limited
 Tel: +886-2-2311-2255
 Fax: +886-2-2311-2277
 Address: 19F, No.39, Section 1, Chung hwa Road Taipei, Taiwan
- Thailand:
 Fujitsu Systems Business (Thailand) Ltd.
 Tel: +66-2-500-1500
 Fax: +66-2-500-1555
 Address: 12th Floor, Olympia Thai Tower, 444 Rachadapisek Road Samsennok, Huaykwang,
 Bangkok, Thailand 10310
- Vietnam:
 Fujitsu Vietnam Limited
 Tel: +84-4-831-3895
 Fax: +84-4-831-3898
 Address: Unit 802-8th floor, Fortuna Tower Hanoi 6B Lang ha Street, Ba dinh District, Hanoi
 Socialist Republic of Vietnam
- United States:
 Fujitsu Computer Systems Corporation
 Tel: +1-800-831-3183
 Fax: +1-408-496-0575
 Address: 1250 East Arques Avenue, Sunnyvale, CA USA 94088-3470

For the latest information, refer to the Fujitsu PRIMERGY website (<http://primergy.fujitsu.com>).

PRIMERGY

LAN 拡張ボード (1Gbps)
(PG-LND201)
取扱説明書
1Gbit/s Ethernet I/O Module
(PG-LND201)
User's Guide

B7FY-2091-01-00

発行日 2007年5月
発行責任 富士通株式会社

Issued on May, 2007
Issued by FUJITSU LIMITED

Printed in Japan

- 本書の内容は、改善のため事前連絡なしに変更することがあります。
- 本書に記載されたデータの使用に起因する、第三者の特許権およびその他の権利の侵害については、当社はその責を負いません。
- 無断転載を禁じます。
- 落丁、乱丁本は、お取り替えいたします。
- The contents of this manual may be revised without prior notice.
- Fujitsu assumes no liability for damages to third party copyrights or other rights arising from the use of any information in this manual.
- No part of this manual may be reproduced in any form without the prior written permission of Fujitsu.
- Any manual which has missing pages or which is incorrectly collated will be replaced.

FUJITSU



古紙パルプ配合率100%再生紙を使用

このマニュアルはリサイクルに配慮して製本されています。
不要になった際は、回収・リサイクルに出してください。



Free Manuals Download Website

<http://myh66.com>

<http://usermanuals.us>

<http://www.somanuals.com>

<http://www.4manuals.cc>

<http://www.manual-lib.com>

<http://www.404manual.com>

<http://www.luxmanual.com>

<http://aubethermostatmanual.com>

Golf course search by state

<http://golfingnear.com>

Email search by domain

<http://emailbydomain.com>

Auto manuals search

<http://auto.somanuals.com>

TV manuals search

<http://tv.somanuals.com>