



About the Cisco 4x InfiniBand Host Channel Adapter Expansion Card for IBM BladeCenter


This chapter describes the hardware and InfiniBand host drivers associated with the Cisco 4x InfiniBand Host Channel Adapter Expansion Card for IBM BladeCenter. This chapter contains the following sections:

- [Introduction, page 1-1](#)
- [HCA Expansion Cards and Server Switch Modules, page 1-2](#)
- [Supported Protocols, page 1-3](#)

Introduction

The Cisco 4x InfiniBand Host Channel Adapter Expansion Card for IBM BladeCenter (referred to as the HCA Expansion Card) provides InfiniBand I/O capability to processor blades in the IBM BladeCenter unit. The host channel adapter adds two InfiniBand ports to the CPU blade cards to create an IB-capable high density cluster. [Figure 1-1](#) shows the HCA Expansion Card.

Figure 1-1 Cisco 4x InfiniBand Host Channel Adapter Expansion Card for IBM BladeCenter



HCA Expansion Cards communicate with one another through the Cisco 4x InfiniBand Switch Module for IBM BladeCenter. For information on the server switch module, see the *Cisco 4x InfiniBand Switch Module for IBM BladeCenter User Guide*.

Features

The HCA Expansion Card includes the following features:

- PCI-Express Interface to dual 4x InfiniBand bridge
- Two 4x InfiniBand interfaces, line rate of the interfaces are 10 Gbps per link, theoretical maximum
- 128 MB table memory (133 MHz DDR SDRAM)
- I2C serial EEPROM holding system Vital Product Data (VPD)
- IBM proprietary blade daughter card form factor
- Identical operation and interfaces for existing Cisco host channel adapter drivers
- Forced air cooling compatible for highly reliable operation
- Port-to-port failover for select protocols (see the [“Supported Protocols”](#) section on page 1-3 for details)

Supported Blade Servers

[Table 1-1](#) lists the IBM blade servers that support the HCA Expansion Card, and includes the maximum number of HCA Expansion Cards that the blade supports.

Table 1-1 Supported Blade Servers

Blade Server	Maximum Number of Supported HCA Expansion Cards
HS21 8853	1
LS21 7971	1
LS41 7972	1

HCA Expansion Cards and Server Switch Modules

Within the BladeCenter chassis, Server Switch Modules manage traffic to and from HCA Expansion Cards on the BladeCenter hosts. Each host channel adapter expansion card adds two IB ports to a BladeCenter host. Each host channel adapter port connects through the BladeCenter unit backplane to a particular Server Switch Module bay. With the Server Switch Module and HCA Expansion Cards, you can create a nonredundant, single-switch topology or a redundant, dual-switch topology.

Single-Switch Topology

When you populate just one BladeCenter module bay with a Server Switch Module, you create a bisectonal bandwidth topology. However, this topology does not provide redundant links from the HCA Expansion Cards to the Server Switch Module. We strongly recommend that you implement a dual-switch topology (see the [“Dual-Switch Topology”](#) section on page 1-3) to avoid single points of failure.

Dual-Switch Topology

To enable IB redundancy on the BladeCenter unit, you must install one Server Switch Module in each available bay. HCA Expansion Cards do not support redundant links to a single switch bay. When you add a second Server Switch Module to the BladeCenter unit, each port of each HCA Expansion Card connects to one Server Switch Module.

Supported Protocols

The Cisco Host Channel Adapter Expansion Card supports the following protocols:

- Internet Protocol over InfiniBand (IPoIB)
- Socket Direct Protocol (SDP)
- User Direct Access Programming Library (uDAPL)
- SCSI RDMA Protocol (SRP)
- Message Passing Interface (MPI)

**Note**

Cisco supports all protocols for Linux. As of this release, Cisco only supports IPoIB and SRP for Windows.

IPoIB

IPoIB allows IP networks to utilize the InfiniBand fabric. SDP and uDAPL use IPoIB to resolve IP addresses. You configure IPoIB as you would a standard Ethernet interface. IPoIB automatically adds IB interface names to the network configuration. The interface names correspond to the ports on the host channel adapter.

**Note**

IPoIB supports port-to-port failover on the host channel adapter daughter card.

SDP

SDP provides a high-performance, zero-copy data transfer protocol used for stream-socket networking over an InfiniBand fabric. You can configure the Cisco driver to automatically translate TCP to SDP based on source IP, destination, or program name.

uDAPL

The User Direct Access Programming Library (uDAPL) defines a set of APIs that leverages InfiniBand's remote direct memory access (RDMA) capabilities. uDAPL is installed transparently with the Cisco driver library. Your application must explicitly support uDAPL. This library requires no manual configuration. However, if your application supports uDAPL, it may require additional configuration changes. For more information see your application documentation.

SRP

SRP runs SCSI commands across RDMA-capable networks so that InfiniBand hosts can communicate with Fibre Channel storage devices. This information is used to assign devices and mount file systems so that the data on those file systems is accessible to the host. The SRP driver is installed as part of the Cisco driver package and is loaded automatically upon host reboot. This protocol requires that a Fibre Channel gateway be present in the IB fabric.

**Note**

SRP supports port-to-port failover on the host channel adapter daughter card.

MPI

MPI is available from the Cisco support web site at <http://www.cisco.com/public/sw-center/>. The following apply to MPI:

- There is no restriction on which host channel adapter port is used.
- Support for Opteron 64-bit operation is provided.