

# Cisco Server Fabric Switch InfiniBand Fabric

## Introduction

Cisco® Server Fabric Switches (SFSs) facilitate utility computing by dramatically simplifying the data center architecture, creating a unified, “wire-once” fabric that aggregates I/O and server resources.

This document provides an overview of InfiniBand essentials and then focuses on the use and functions of the Ethernet gateway.

## InfiniBand Overview

InfiniBand is a standards-based network technology that Cisco uses to build a server area network. InfiniBand is connection oriented as opposed to broadcast Carrier Sense Multiple Access/Collision Detect (CSMA/CD) Ethernet.

The InfiniBand protocol specifies a scalable interconnect:

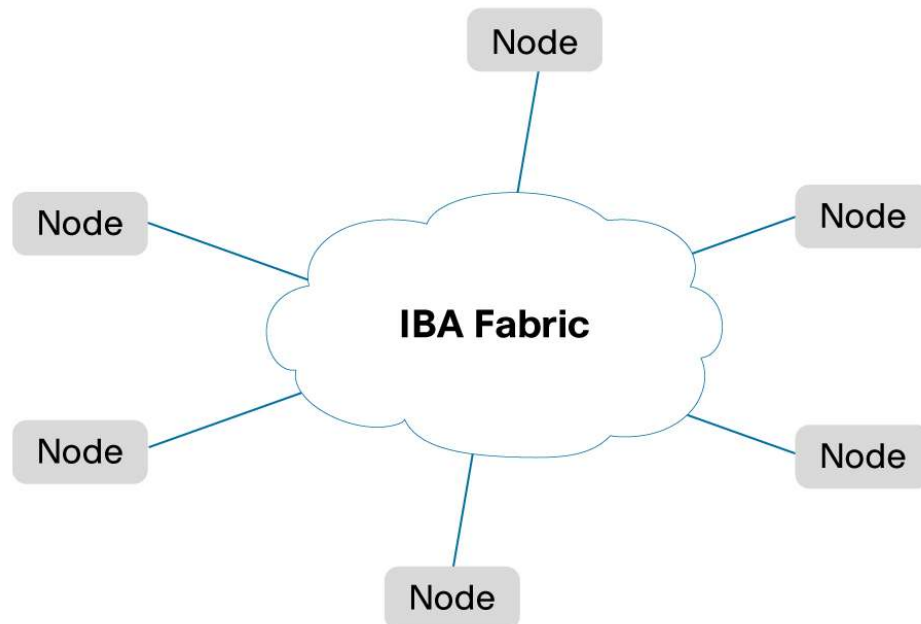
1X = 2.5 Gbps

4X = 10 Gbps

12X = 30 Gbps

At a high level, InfiniBand is the interconnect for end nodes, as illustrated in Figure 1.

**Figure 1.** InfiniBand Architecture Fabric Interconnect



## Terminology and Concepts

### Channel Adapter

A channel adapter is a device that terminates a link and performs transport-level functions. A channel adapter can be either a host channel adapter (HCA) or a target channel adapter (TCA).

### **Subnet Manager**

The subnet manager is the governing entity for the InfiniBand fabric. Each InfiniBand subnet has at least one subnet manager. Each subnet manager resides on a port of a channel adapter, router, or switch and can be implemented either in hardware or software. When there are multiple subnet managers on a subnet, one subnet manager will be the master subnet manager. The remaining subnet managers must be standby subnet managers.

The master subnet manager is a crucial element in initializing and configuring an InfiniBand subnet. The master subnet manager is selected as part of the initialization process for the subnet and is responsible for these tasks:

- Discovering the physical topology of the subnet
- Assigning local identifiers (LIDs) to the end nodes, switches, and routers
- Establishing possible paths among the end nodes
- Sweeping the subnet, discovering topology changes, and managing changes as nodes are added and deleted.

### **Local Identifier (LID)**

A local identifier (LID) is a nonpersistent address that is assigned to a port by the subnet manager. The LID is unique within the subnet and is used for directing packets within the subnet. The source and destination LIDs are present in the local route header.

### **Subnet Prefix**

The subnet prefix is a 0 to 64-bit (a function of scope) identifier that is used to uniquely identify a set of end ports, which are managed by a common subnet manager.

### **Unicast Identifier**

A unicast identifier is the value that specifies a single end port. A packet sent to a unicast identifier is delivered to the end port, which is located using that identifier.

InfiniBand Architecture (IBA) standards define the following:

- Two unicast identifiers (a global identifier [GID]), which may be unique across subnets
- A LID, which is unique only within a subnet

### **Extended Universal Identifier 64 (EUI-64)**

The extended universal identifier (EUI) is a defined 64-bit identifier that is assigned to a device. This identifier is created by concatenating a 24-bit company\_id value with a 40-bit extension identifier. The company\_id value is assigned by the IEEE Registration Authority; the extension identifier is assigned by the organization with the assigned company\_id.

Each HCA, TCA, switch, and router is assigned a EUI-64 global unique identifier (GUID) by the manufacturer.

### **Global Unique Identifier (GUID)**

A global unique identifier (GUID) is a globally unique identifier that is compliant with EUI-64.

### **Global Identifier (GID)**

A global identifier (GID) is a 128-bit unicast or multicast identifier that is used to identify an end port or a Multicast group. It's a catenation of the subnet prefix and the GUID in an InfiniBand fabric.

### **Partition**

A partition (P\_KEY) is a collection of channel adapter ports that are allowed to communicate with one another. Ports may be members of multiple partitions simultaneously. Ports in different partitions are unaware of each others' presence. A partition key value is carried in packets and stored in channel adapters.

### **Simplifying the Fabric with the Server Switch**

The server switch simplifies the fabric by connecting every server with a single high-bandwidth, low-latency network cable. Rather than requiring multiple adapters and cables for each server, the unified fabric aggregates Ethernet, Fibre Channel, and clustering interconnects into a 10-Gbps InfiniBand cable.

The server switch connects servers to a pool of shared Fibre Channel and Ethernet ports over line-rate gateways and creates virtual I/O subsystems on each host, including virtual host bus adapters (HBAs) and virtual IP interfaces.

Servers can then share a centralized pool of Ethernet and Fibre Channel ports in the same way that a storage area network (SAN) creates a pool of shared storage that can be managed independently of the servers themselves.

### **Maximizing Data Center Investments**


Consolidating resources over the unified fabric eliminates the costs of underused Fibre Channel HBAs and network interface devices (NICs) as well as the associated cabling complexity. Instead of designing a data center to accommodate bandwidth peaks using a dedicated switch port for each host, a data center can now share remote Fibre Channel and Gigabit Ethernet ports—which allows networks to be designed based on the average load across multiple servers.

Designing for actual use can save up to 50 percent of the cost of the I/O associated with a server. In addition, eliminating multiple adapters and local storage and introducing a single high-bandwidth, low-latency connection means that the size of the server is restricted only by CPU and memory requirements.

This maximization of resources often results in a reduction in the size and cost of the server as well as in its space, power, and cooling needs; the result is an immediate return-on-investment savings of up to 50 percent.

Figure 2 shows a unified server fabric.

**Figure 2.** Unified Server Fabric



## Virtual I/O

Virtual I/O has two major components:

- Virtual IP interfaces and an InfiniBand-to-Ethernet gateway
- Virtual HBA and an InfiniBand-to-Fibre Channel gateway.

An InfiniBand driver package is installed on the host, which includes an IP-over-InfiniBand (IPoIB) driver, Small Computer System Interface (SCSI) driver (called the SCSI Remote Direct Memory Access (RDMA) Protocol [SRP]), and other RDMA protocols. The server uses the IP and SCSI drivers to communicate through the gateways, bridging IP subnets and allowing hosts to access Fibre Channel attached storage.

### Virtual IP Interfaces

Loading IPoIB drivers and creating virtual IP interfaces on the InfiniBand-capable servers allows them to communicate directly with existing IP servers.

To configure the IPoIB driver on the host, the administrator configures an interface on the ib0 and ib1 ports on the HCA, similar to the way eth0 and eth1 ports are configured on an Ethernet NIC. In Linux, the administrator configures the interface with ifconfig, and in Windows, the administrator can configure the interfaces in the device driver dialog box in the Control Panel. The InfiniBand partition is also configured at this location.

Although the server may not be physically connected directly to the LAN or WAN, this interface can transparently communicate to other Ethernet-attached hosts using the InfiniBand-to-Ethernet gateway.

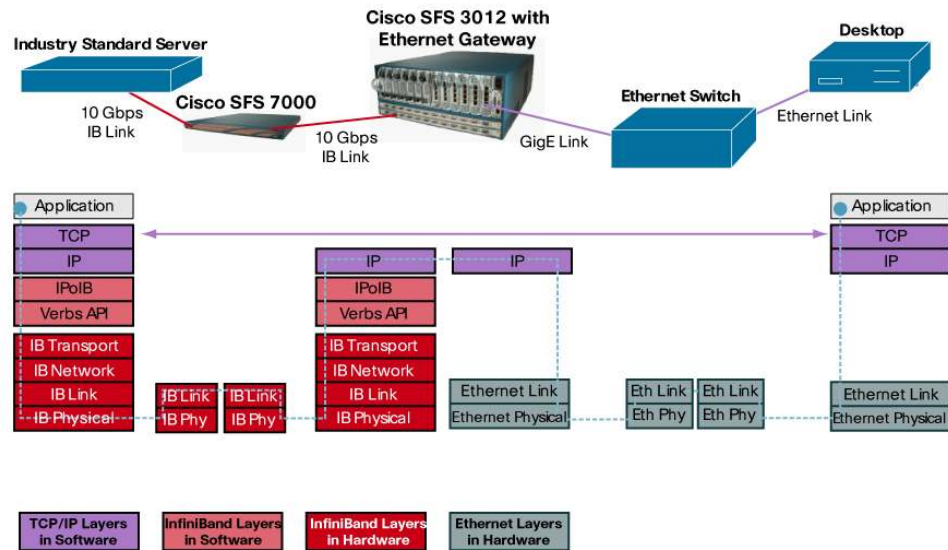
### IPoIB

Using IPoIB, existing IP-based applications and tools work using existing socket libraries. Networking diagnostic tools such as ping and traceroute work, and server switches can integrate with existing network management tools such as CiscoWorks, Tivoli, Unicenter, and OpenView using Simple Network Management Protocol (SNMP).

IPoIB is used within the InfiniBand fabric for standard IP-based communications, as well as address lookups. IPoIB is also translated across the Ethernet gateway to IP over Ethernet (IPoE)

to provide Layer 2 bridging based on the IP address (not the MAC address). IP addresses are lookup keys in the forwarding tables, and IP addresses are used to make forwarding decisions. Administrators create bridge groups for bridging between InfiniBand and Ethernet subnets, which translates the IPoIB frames to IPoE frames (Figure 3).

**Figure 3.** IPoIB Translation



### Ethernet Gateway Architecture

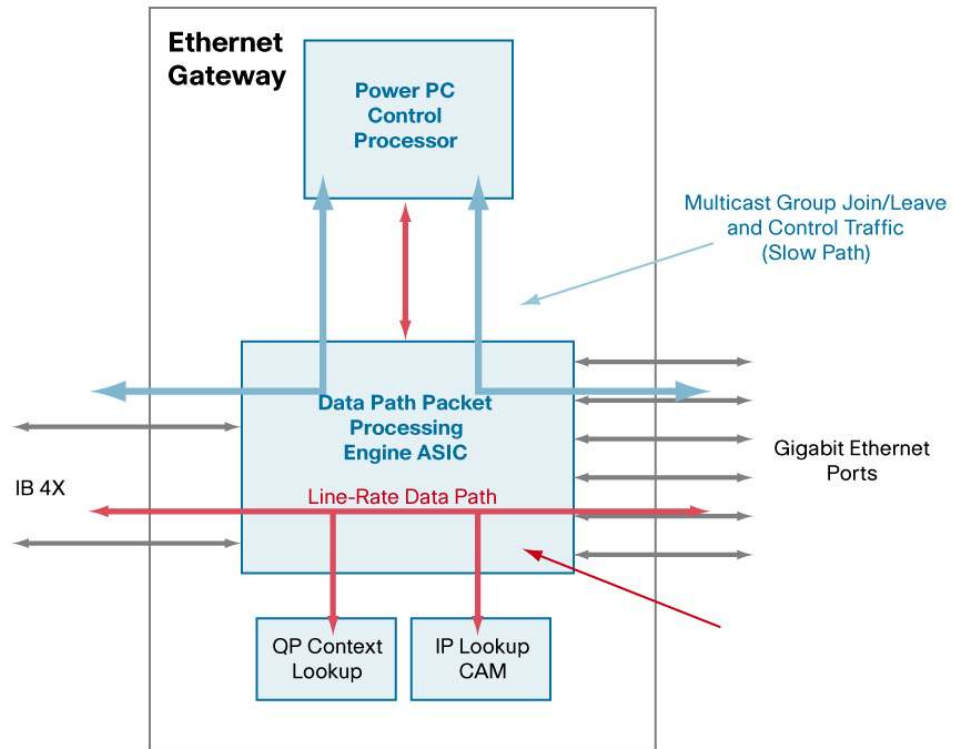
The InfiniBand-to-Ethernet gateway is based on specialized chipsets that perform bridging at a cut-through line rate on all six Gigabit Ethernet ports.

### Data Paths

The gateway has two distinct paths for data and control (Figure 4):

- *Slow path* PowerPC processor—Multicast joins and leaves are handled in the slow path.
- *Fast path* hardware packet processing engine—The fast path handles multicast packet processing by looking up IP-to-InfiniBand multicast group mappings and changing packet headers to InfiniBand multicast addresses, which are also handled in hardware. Multicast joins and leaves are handled in the slow path.

**Figure 4.** Data Paths within the Ethernet Gateway




### Internal Gateway Ports and External Ethernet Ports

Internally, the Ethernet gateway connects to the InfiniBand network using two 10-Gbps gateway ports (Figure 5). When configuring the gateway (or internal) ports, a user can specify the slot number alone or the slot number and a specific internal port number.

Externally, there are six 1-Gbps ports that can be connected to the access or the distribution layer. All six external ports can be bonded together to form an aggregate, providing higher throughput.

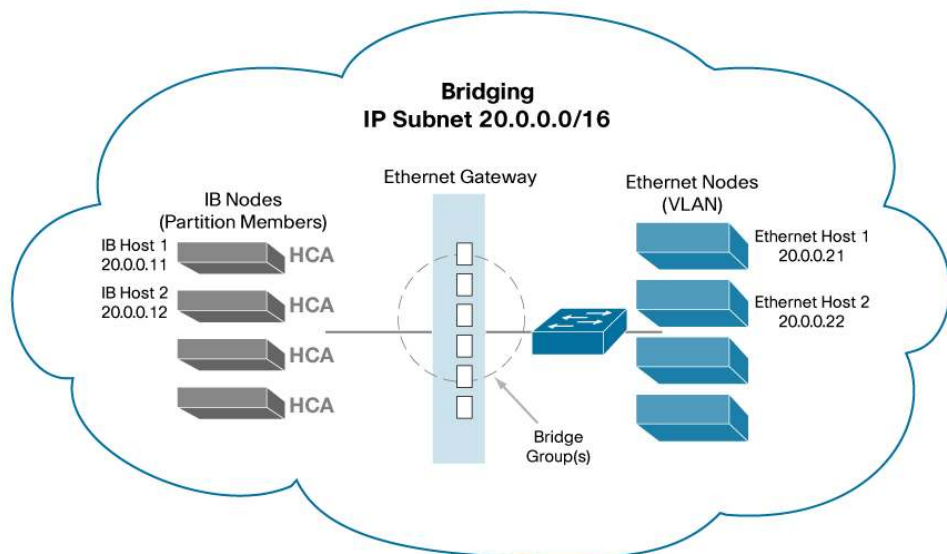
**Figure 5.** Internal Architecture of an Ethernet Gateway



### Ethernet Gateway Bridging

A bridging group is an entity that runs on the Ethernet gateway and allows the bridging of one iPoB broadcast domain to one VLAN (Figure 6). The Ethernet gateway acts as a Layer 2 bridge between InfiniBand and Ethernet and must be configured for Layer 2 bridging (with or without link aggregation and redundancy groups). Outgoing Ethernet packets use the source MAC address of the port or the trunk MAC address, depending on the situation.

**Figure 6.** Bridge Group to Bridge a Single Broadcast Domain



### Ethernet Gateway Addresses

Thirty-two MAC addresses are assigned for each Ethernet gateway card and are allocated as follows:

- 6 Gigabit Ethernet ports +
- Gateway \* 4 (node + 2 ports + 1 reserve) +
- 6 trunk ports + 1 debug Ethernet port +
- 15 reserved

The base number is the GUID taken from the bar code label on the card.

- A GUID refers to an EUI-64 number. The EUI-64 number consists of an EUI-48 number with 2 bytes of 0 inserted in the middle. For example, EUI-48 00:05:ad:01:12:34 becomes EUI-64 00:05:ad:00:00:01:12:34.
- A MAC address is an EUI-48 number.

Gigabit Ethernet port 1: Base number (EUI-48)  
 Gigabit Ethernet port 2: Base number + 1 (EUI-48)  
 Gigabit Ethernet port 3: Base number + 2 (EUI-48)  
 Gigabit Ethernet port 4: Base number + 3 (EUI-48)  
 Gigabit Ethernet port 5: Base number + 4 (EUI-48)  
 Gigabit Ethernet port 6: Base number + 5 (EUI-48)  
 Gateway node: Base number + 6 (EUI-64)  
 Gateway port1: Base number + 7 (EUI-64)  
 Gateway port2: Base number + 8 (EUI-64)  
 Gateway reserved: Base number + 9 (EUI-64)

Debug Ethernet port: EUI-48 MAC address: Use base number (minus middle zeros) + 10 (EUI-48)

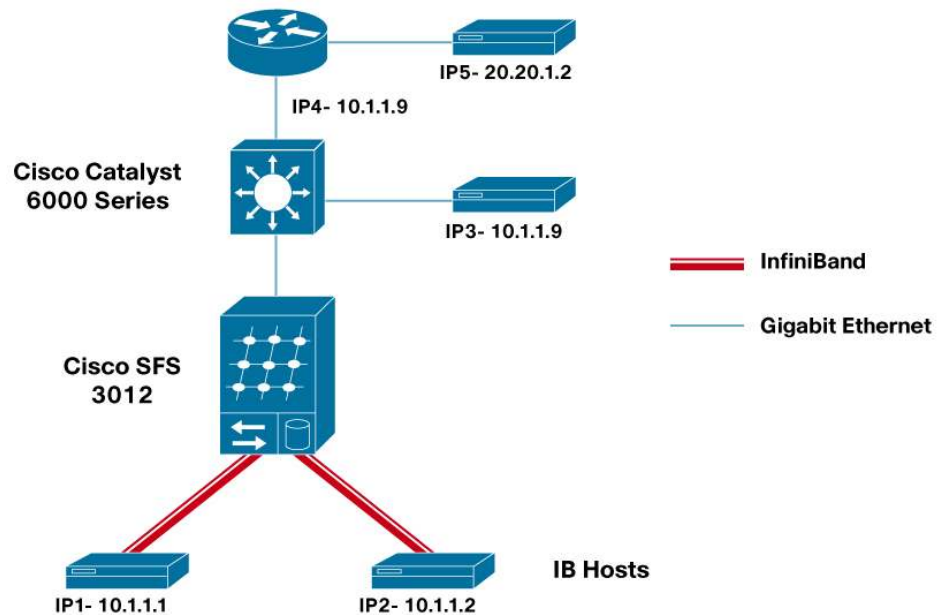
Trunk port 1: Base number + 16 (EUI-48)  
 Trunk port 2: Base number + 17 (EUI-48)

- Trunk port 3: Base number + 18 (EUI-48)
- Trunk port 4: Base number + 19 (EUI-48)
- Trunk port 5: Base number + 20 (EUI-48)
- Trunk port 6: Base number + 21 (EUI-48)

### Packet Flow from InfiniBand Subnet to LAN—Life of a Packet

A packet coming from the InfiniBand fabric can be forwarded to one or multiple subnets. The steps in the packet flow shown in Figure 7 are described in the following sections.

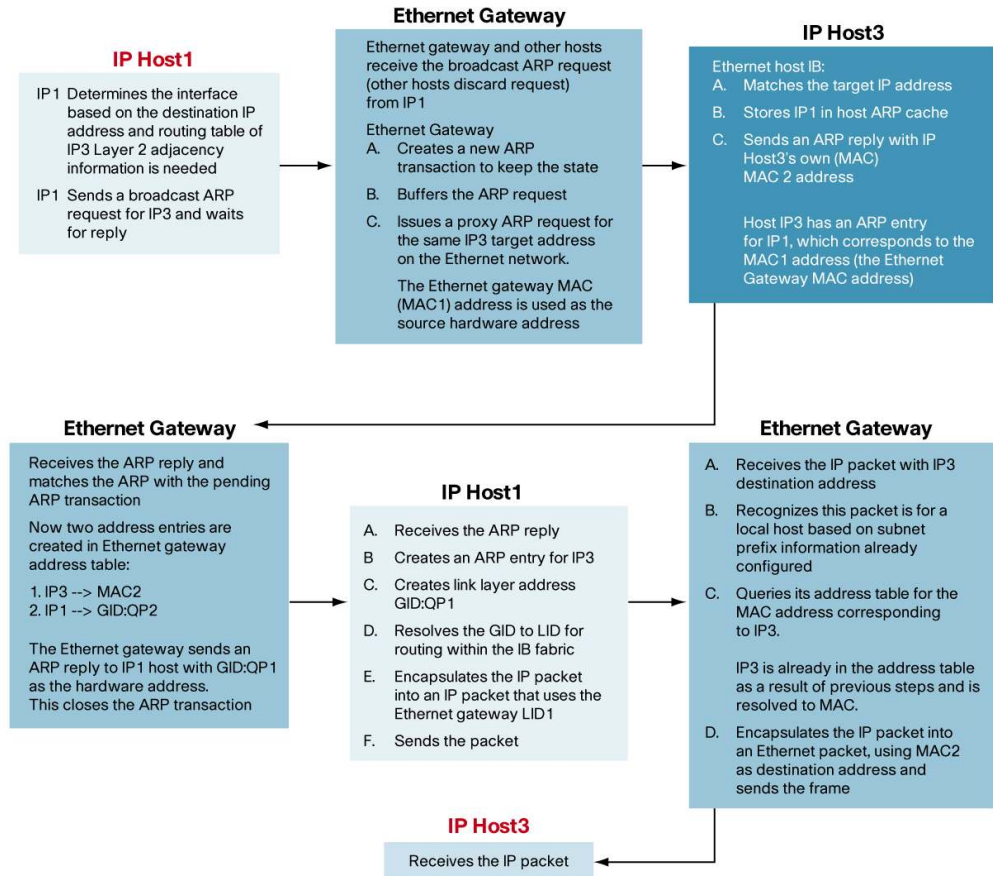
**Figure 7.** Packet Flow—From the InfiniBand Subnet to the LAN



### Same-Subnet Forwarding

The IP subnet prefix information is configured on the InfiniBand switch with the Ethernet gateway. In the example in Figure 8, the subnet prefix is 10.1.1.0/8. The default gateway (called the InfiniBand next hop, or `ib-next-hop`) is also configured for the InfiniBand subnet.

**Figure 8.** Packet Flow for Communication from IP Host 1 to IP Host 3 (Refer to Figure 7)

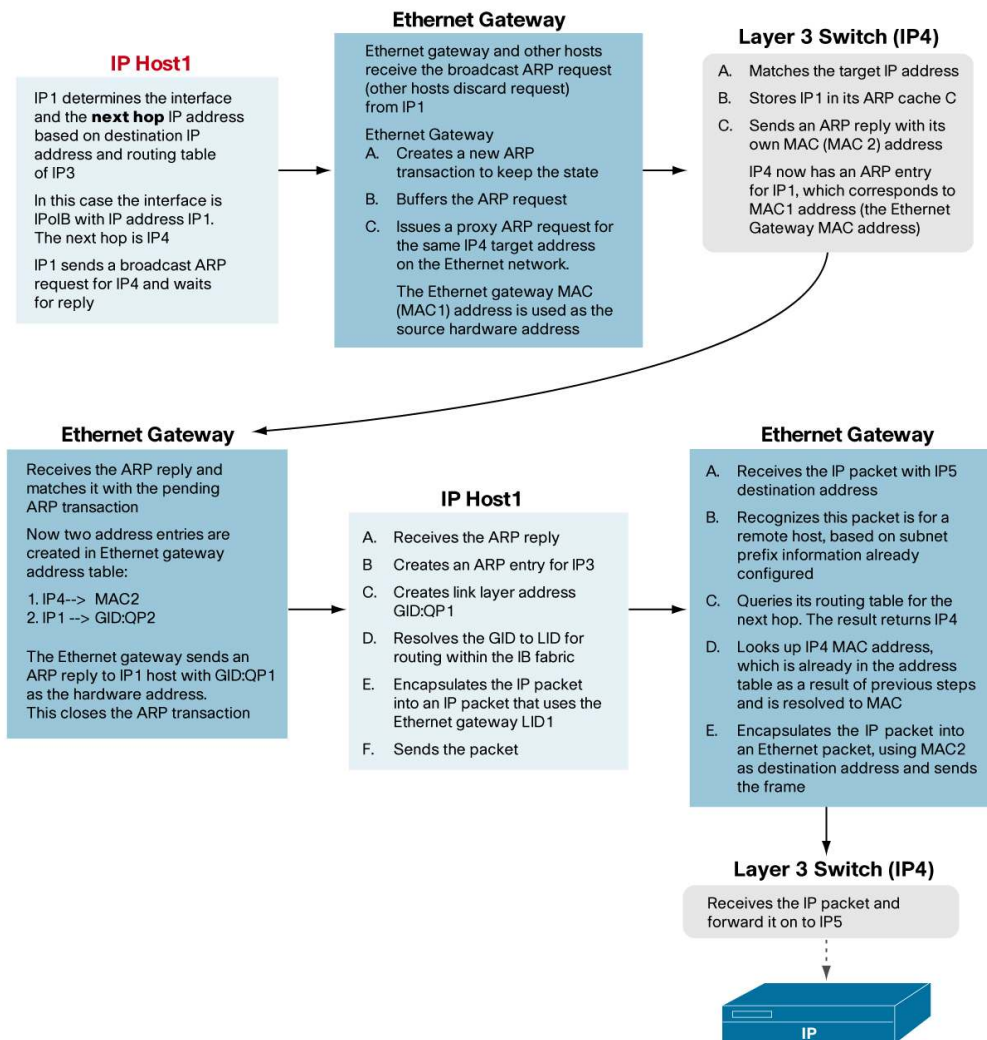


### Remote Subnet Forwarding

Figure 9 shows remote subnet forwarding.

**Note:** For remote subnet packet forwarding, the next hop is determined by the InfiniBand next hop configured on the Cisco SFS. It has no bearing on the default gateway configured on the host itself.

**Figure 9.** Packet Flow for Communication from IP Host 1 to IP Host 5 (Refer to Figure 7)



## Loop Protection

The Ethernet gateway does not participate in the Spanning Tree Protocol and does not exchange or process bridge protocol data units (BPDUs). Upstream Ethernet switches are recommended to run the Spanning Tree Protocol.

InfiniBand hosts appear to the rest of the network as a group of hosts behind an Ethernet link.

To prevent loops in a multiple Ethernet gateway network setup (single- or multiple-chassis), some built-in loop protection mechanisms are provided. All these features can be disabled to suit network needs.

## Disabling Broadcast Forwarding

Broadcast forwarding can be disabled. This option disables forwarding of all IP broadcast packets except services for which the Ethernet gateway has special handling.

Broadcast forwarding is disabled by default and needs to be activated explicitly. For example, Dynamic Host Configuration Protocol (DHCP) packets would require broadcast forwarding to be activated.

## Address Resolution Protocol Packet Painting

The Ethernet gateway provides the option of inserting a signature in every proxy Address Resolution Protocol (ARP) request made on behalf of the IPoIB host. This signature allows the Ethernet gateway to filter the duplicate requests and break the loop. This feature is enabled by default and is recommended when multiple Cisco SFS Ethernet gateways are present.

### Self-Canceling ARP Requests

By default, the self-canceling ARP request feature is active. This feature prevents duplicate ARP requests (which have the same target protocol address) from creating loops. The duplicate ARP request is seen by multiple gateways and discarded.

### Delayed Proxy ARP Transaction

By default, the delayed proxy ARP transaction feature is always active. This feature is an extension of the self-canceling ARP request and is used when a duplicate ARP request is delayed.

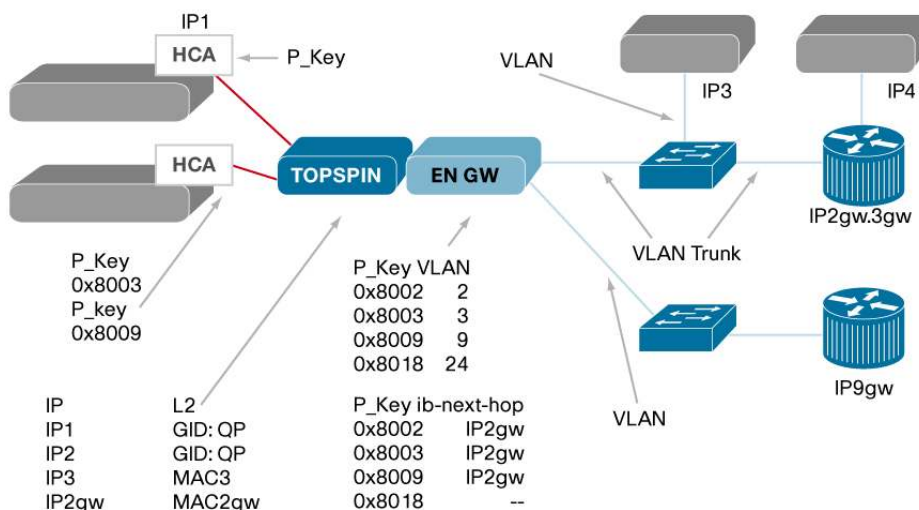
### VLAN and 802.1Q Support

The InfiniBand network supports multiple partitions using the concept of partition keys (p\_keys). The InfiniBand partitions (p\_keys) can be mapped to specific VLANs on the Ethernet network (Figure 10).

For example, InfiniBand hosts belonging to p\_key 0x8001 could be VLAN 11 and p\_key 0x8002 VLAN 12.

- Standard 802.1Q VLANs are supported.
- Up to 32 VLANs can be supported per gateway.
- Static port-based VLANs are supported.
- One VLAN is mapped to one InfiniBand partition.

Figure 10. P\_Key Mapping



## Link Aggregation

Link aggregation is an optional feature available on the Ethernet gateway and is used with Layer 2 bridging. Link aggregation allows multiple Ethernet gateway ports to merge logically into a single link.

Link aggregation logically combines multiple links to support a larger data stream than a single link can, and it combines related objects or operations into a single one.

Link aggregation offers these benefits:

- Higher aggregate bandwidth to traffic-heavy servers
- Reroute capability in the event of a single port or cable failure

## Link Aggregation Features

One link aggregation group can be assigned to one bridge group or to multiple bridge groups.

- Six link aggregation groups are supported for each Ethernet gateway.
- A link aggregation group cannot span multiple gateways.
- Each link aggregation group can support up to 32 VLANs.
- Static link aggregation (802.3ad) group configuration (Port Aggregation Protocol [PagP] and Link Aggregation Control Protocol [LACP]) is not supported).

On Cisco Ethernet switches, channel mode needs to be turned on.

Several frame distribution types are supported (Table 1). The default frame distribution is src-dst-ip.

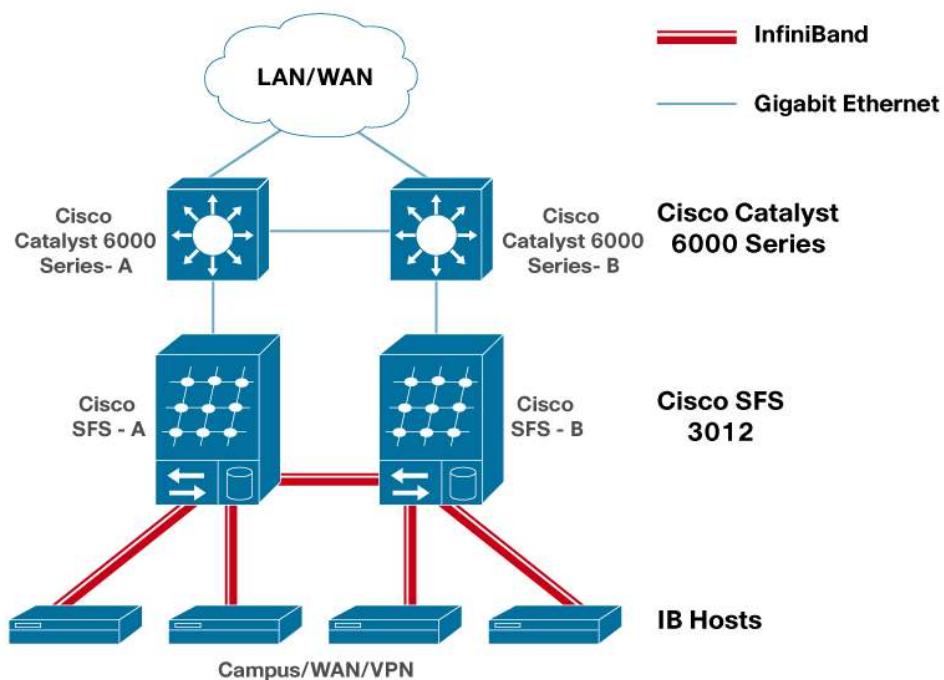
**Note:** This frame distribution applies only to outbound traffic. Distribution for inbound traffic is determined by the frame distribution across Cisco EtherChannel<sup>®</sup> on the Ethernet switch.

**Table 1.** Load Distribution Types and Functions

Distribution Type	Function
<b>dst-ip</b>	Load distribution is based on the destination IP address. Packets to the same destination are sent on the same port, but packets to different destinations are sent on different ports in the channel.
<b>dst-mac</b>	Load distribution is based on the destination-host MAC address of the incoming packet. Packets to the same destination are sent on the same port, but packets to different destinations are sent on different ports in the channel.
<b>src-dst-ip</b>	Load distribution is based on the source logic gate (XOR) destination IP address.
<b>src-dst-mac</b>	Load distribution is based on the source logic gate (XOR) destination MAC address.
<b>src-ip</b>	Load distribution is based on the source IP address. Packets to the same destination are sent on the same port, but packets to different destinations are sent on different ports in the channel.
<b>src-mac</b>	Load distribution is based on the source-MAC address of the incoming packet. Packets from different hosts use different ports in the channel, but packets from the same host use the same port in the channel.
<b>round robin</b>	Round Robin is a load balancing algorithm that distributes load in a circular fashion, thereby creating an evenly distributed load. When using Redundancy Groups and Load Balancing, selecting the Round Robin distribution can increase performance in many cases. Even with a Topology that contains as few as 1 Ethernet host, the performance could benefit from using this distribution type. One exception would be if you have a set restriction on the amount of bandwidth that can be utilized by your Ethernet host. In this case, Round Robin would not provide any further benefit.

## Ethernet Gateway Redundancy

Redundancy can be provided with multiple Cisco SFS switches with Ethernet gateways and multiple upstream Ethernet switches, as illustrated in Figure 11.

**Figure 11.** Ethernet Gateway Redundancy

Load balancing and failover between discrete gateways and switches is achieved by assigning bridge groups to redundancy groups. The redundancy manager process runs on the controller card of the Cisco SFS switch and controls the failover and load-balancing feature.

In redundancy groups configured for failover, traffic is not passed on backup bridge groups. In redundancy groups configured for load balancing, all bridge groups are passing traffic.

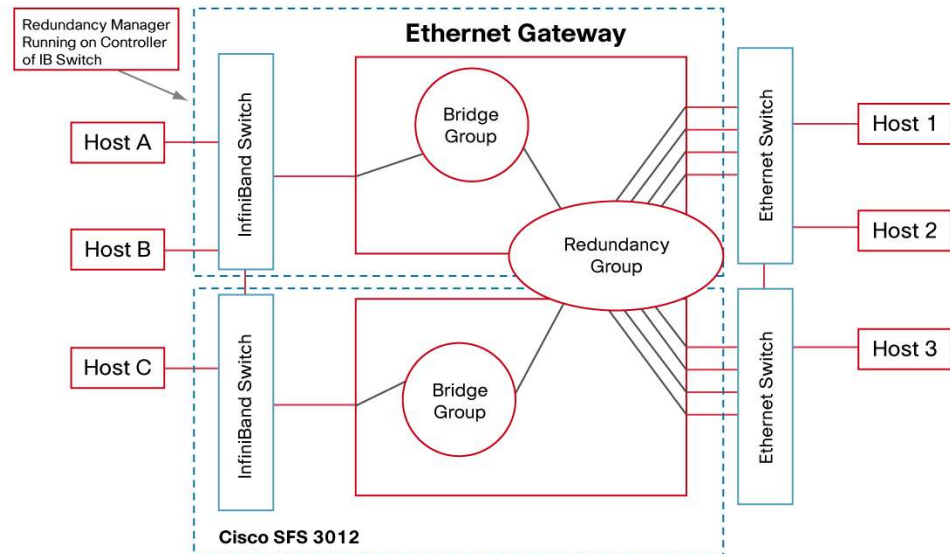
Redundancy groups can be created using multiple gateways across multiple chassis. Load distribution is conversation based and relies on source and destination hardware addresses and source and destination IP addresses.

In the event of an Ethernet gateway failure, the following process occurs:

1. Gratuitous ARP requests are sent on both sides of the fabric.
2. InfiniBand hosts get a gratuitous ARP request that updates the IP address and the GID:QP mapping of the external hosts.
3. External hosts get the updated MAC address of the InfiniBand hosts because they will now use the MAC address of the new gateway.

Figure 12 shows how redundancy works across Ethernet gateways.

**Figure 12.** InfiniBand-to-Ethernet Gateway Failover and Load Balancing



## DHCP

The Ethernet gateway supports the forwarding of DHCP requests when the DHCP server resides on the LAN. When using DHCP, broadcast forwarding must be explicitly activated for the bridge group.

The DHCP server must return the reply with the broadcast flag set for DHCP over IPoIB to work.

The following RFC lists the specifications to which the Cisco SFS switches and host side drivers conform: <http://www.rfc-archive.org/getrfc.php?rfc=4390>.

## Multicast Support

IP Multicast allows a host to send packets to a specific subset of all hosts as a group transmission. Without the ability to multicast, a host is limited to sending either to a single host or to all hosts.

The InfiBand fabric supports multicast groups, and there is a one-to-one mapping between the InfiBand and IP Multicast groups.

Internet Group Management Protocol (IGMP) membership reports from the InfiBand hosts are sent to the upstream Ethernet device to create the relevant Protocol Independent Multicast (PIM) forwarding entries. The InfiBand hosts interested in receiving the multicast stream are members of this InfiBand multicast group, along with the Ethernet gateway, and receive traffic sent to this group.

When configured in a redundant setup (two Cisco SFS switches connected to the upstream Ethernet switches), only one gateway forwards the IGMP joins and accepts the multicast traffic to prevent loops. In the case of a primary path failure, a gratuitous IGMP report is sent to the upstream switch for the multicast stream to be rerouted.

Similarly, for an IP Multicast sender present on the InfiBand fabric, traffic flows in the other direction.

Multicast support is configured by configuring the bridge group to accept multicast traffic.

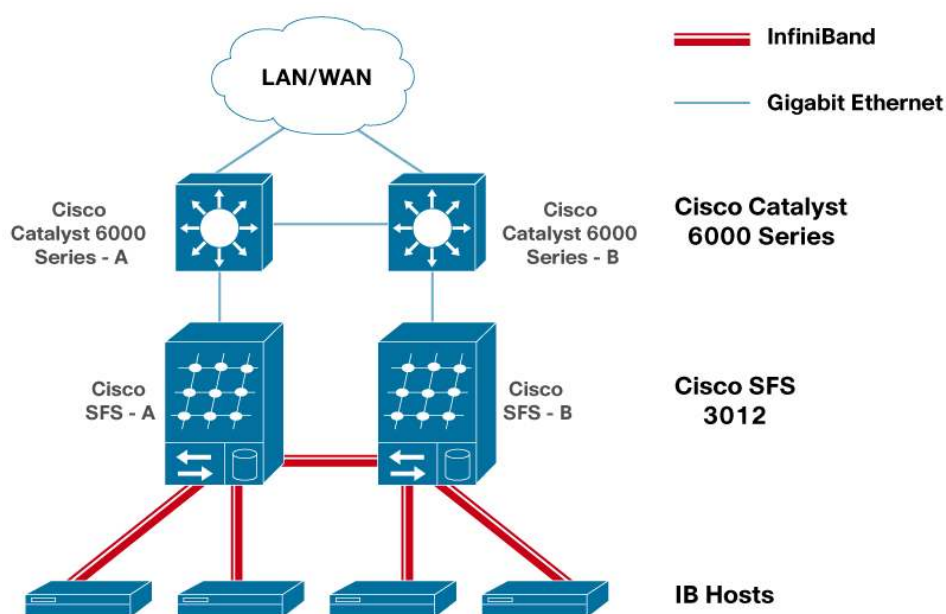
## Load Balancing Across Multigroup Hot Standby Router Protocol Groups

The forwarding of IP traffic to the default gateway is determined by the InfiniBand next hop, which is configured on the Cisco SFS gateway (Figure 13).

To achieve default gateway load balancing across different next hops on the same IP subnet, configuration needs to be performed on the Cisco SFS gateway to segregate the address space and specify a different InfiniBand next hop for each.

**Note:** The InfiniBand hosts and the upstream switches do not need special configuration to achieve load balancing across Multigroup Hot Standby Router Protocol (MHSRP) groups. The Cisco SFS gateway performs the configuration transparently.

**Figure 13.** Load Balancing



### Sample Configuration

This section presents a sample configuration. In this example, the following address space and interfaces are assumed:

Cisco Catalyst® 6000 Series A: VLAN 100—192.168.100.10/24

Cisco Catalyst 6000 Series B: VLAN 100—192.168.100.200/24

InfiniBand hosts: 192.168.100.0/24

192.168.100.1-126—Default gateway 192.168.100.1

192.168.100.129-254—Default gateway 192.168.100.129

**Cisco Catalyst Configuration**Cisco Catalyst 6000 Series A: **show run int vlan 100**

```
interface Vlan100
ip address 192.168.100.10 255.255.255.0
standby 1 ip 192.168.100.1
standby 1 priority 250
standby 1 preempt
standby 2 ip 192.168.100.129
standby 2 priority 50
standby 2 preempt

end
```

Cisco Catalyst 6000 Series B: **show run int vlan 100**

```
interface Vlan100
ip address 192.168.100.200 255.255.255.0
standby preempt
standby 1 ip 192.168.100.1
standby 1 priority 50
standby 1 preempt
standby 2 ip 192.168.100.129
standby 2 priority 250
standby 2 preempt

end
```

**Cisco SFS Configuration**Cisco SFS-A: **interface mgmt-ib** ← mgmt-ib interface must be configured for redundancy groups

```
ip address 192.168.1.1 255.255.255.0
bridge-group 1 subnet-prefix 192.168.100.0 25
bridge-group 1 subnet-prefix 192.168.100.128 25
bridge-group 1 ib-next-hop 192.168.100.1 next-hop-mask 255.255.255.128
bridge-group 1 ib-next-hop 192.168.100.129 next-hop-mask 255.255.255.128
```

!

```
interface gateway 3/2          ← Ethernet gateway is in slot 3
bridge-group 1 pkey ff:ff      ← P_key ff:ff is being bridged
```

!

```
interface ethernet 3/1
bridge-group 1
```

!

```
redundancy-group 1
redundancy-group 1 load-balancing
```

!

**bridge-group 1 redundancy-group 1**

Cisco SFS-B: **interface mgmt-ib**

**ip address 192.168.1.2 255.255.255.0** ← mgmt-ib interface must be configured for redundancy groups

**bridge-group 1 subnet-prefix 192.168.100.0 25**

**bridge-group 1 subnet-prefix 192.168.100.128 25**

**bridge-group 1 ib-next-hop 192.168.100.1 next-hop-mask 255.255.255.128**

**bridge-group 1 ib-next-hop 192.168.100.129 next-hop-mask 255.255.255.128**

!

**interface gateway 3/2** ← Ethernet gateway is in slot 3

**bridge-group 1 pkey ff:ff** ← P\_key ff:ff is being bridged

!

**interface ethernet 3/1**

**bridge-group 1**

!

**redundancy-group 1**

**redundancy-group 1 load-balancing**

!

**bridge-group 1 redundancy-group 1**



**Americas Headquarters**  
Cisco Systems, Inc.  
170 West Tasman Drive  
San Jose, CA 95134-1706  
USA  
www.cisco.com  
Tel: 408 526-4000  
800 553-NETS (6387)  
Fax: 408 527-0883

**Asia Pacific Headquarters**  
Cisco Systems, Inc.  
168 Robinson Road  
#28-01 Capital Tower  
Singapore 068912  
www.cisco.com  
Tel: +65 6317 7777  
Fax: +65 6317 7799

**Europe Headquarters**  
Cisco Systems International BV  
Haarlerbergpark  
Haarlerbergweg 13-19  
1101 CH Amsterdam  
The Netherlands  
www-europe.cisco.com  
Tel: +31 0 800 020 0791  
Fax: +31 0 20 357 1100

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at [www.cisco.com/go/offices](http://www.cisco.com/go/offices).

©2006 Cisco Systems, Inc. All rights reserved. CCVP, the Cisco logo, and the Cisco Square Bridge logo are trademarks of Cisco Systems, Inc.; Changing the Way We Work, Live, Play, and Learn is a service mark of Cisco Systems, Inc.; and Access Registrar, Aironet, BPX, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Enterprise/Solver, EtherChannel, EtherFast, EtherSwitch, Fast Step, Follow Me Browsing, FormShare, GigaDrive, GigaStack, HomeLink, Internet Quotient, IOS, IP/TV, iQ Expertise, the iQ logo, iQ Net Readiness Scorecard, iQuick Study, LightStream, Linksys, MeetingPlace, MGX, Networking Academy, Network Registrar, Packet, PIX, ProConnect, RateMUX, ScriptShare, SlideCast, SMARTnet, StackWise, The Fastest Way to Increase Your Internet Quotient, and TransPath are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or Website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0609R)