



Q&A

Cisco SFS 7000 Series Infiniband Server Switches

Q. What are the Cisco® SFS 7000 Series InfiniBand switches?

A. The Cisco SFS 7000 Series InfiniBand switches are a new series of high-performance InfiniBand switches designed with best-in-class reliability, high availability, and serviceability features. The Cisco SFS 7000 Series provides the foundation for interconnecting servers with industry-standard 10-Gbps or 20-Gbps full duplex, low-latency InfiniBand. The Cisco SFS 7000 Series switches are ideal building blocks for high-performance research server clusters or enterprise grids with Cisco VFrame Server Fabric Virtualization Software.

Table 1 shows the complete list of Cisco SFS 7000 Series switches and components.

Table 1. Cisco SFS 7000 Series Products

Part Numbers	Product Name and Description
SFS-7000-4X024-SK9	Cisco SFS 7000 InfiniBand Server Switch <ul style="list-style-type: none">• 24 4X InfiniBand 10-Gbps or 20-Gbps full duplex (read and write) ports• Fully nonblocking design with a 480-Gbps backplane and 200-ns port-to-port cut through latency• Unmatched price-to-performance ratio for building high-performance computing clusters• Excellent building block for enterprise grids with unified I/O and Cisco VFrame Server Fabric Virtualization Software• Intelligent switch with embedded fabric management, capable of running clusters of thousands of nodes• 1 rack unit (1-RU) fixed configuration, with 4-post rack mounting kit included
SFS-7008-4X048-SK9	Cisco SFS 7008 InfiniBand Server Switch (48-port) <ul style="list-style-type: none">• 48 4X InfiniBand 10-Gbps or 20-Gbps full duplex (read and write) ports (field upgradeable to 96 ports)• Fully nonblocking design with a 1.92 terabits per second (Tbps) backplane and less than 600-ns port-to-port cut through latency• Excellent building block for enterprise grids with unified I/O and Cisco VFrame Server Fabric Virtualization software• Director-class high availability and stateful failover of all critical components• SFS Rapid Service Architecture that allows service and replacement of all active electronics without removing cables• Intelligent switch with embedded fabric management, capable of running clusters of thousands of nodes• 6-RU modular chassis, with 4-post rack mounting kit optional
SFS-7008-4X072-SK9	Cisco SFS 7008 InfiniBand Server Switch (72-port) <ul style="list-style-type: none">• 72 4X InfiniBand 10-Gbps or 20-Gbps full duplex (read and write) ports (field upgradeable to 96 ports)• Same specifications as above
SFS-7008-4X096-SK9	Cisco SFS 7008 InfiniBand Server Switch (96-port) <ul style="list-style-type: none">• 96 4X InfiniBand 10-Gbps or 20-Gbps full duplex (read and write) ports• Same specifications as above
SFS-7008-K9	Cisco SFS 7008 Switch Chassis (field-replaceable unit [FRU]) Replacement chassis with no line cards
SFS-7008-HA-K9	Cisco SFS 7008 InfiniBand Server Switch High-Availability Package, includes: <ul style="list-style-type: none">• Cisco SFS 7008 Management Interface Module• Cisco SFS 7008 Power Supply• Cisco SFS 7008 Fan Tray
SFS-7008-SFM-K9	Cisco SFS 7008 Switch Fabric Module, with management Fits in any of the switch fabric module slots in the Cisco SFS 7008
SFS-7008-X7008-04X12K9	Cisco SFS 7008 InfiniBand 4X 12-Port Line Card Line card with 12 ports of 10-Gbps 4X InfiniBand

Part Numbers	Product Name and Description
SFS-7008-MFM-K9	Cisco SFS 7008 Management Interface Module <ul style="list-style-type: none"> • Management card for the Cisco SFS 7008 chassis • Add an additional management card for redundant management
PWR-SFS7000	Cisco SFS 7000 Power Supply and Cooling Module (FRU) Replacement power supply cooling module for Cisco SFS 7000 chassis
PWR-SFS7008	Cisco SFS 7008 Power Supply and Cooling Module <ul style="list-style-type: none"> • Replacement power supply for Cisco SFS 7008 chassis • Add an extra power supply to a Cisco SFS 7008 chassis for redundant power
SFS-7008-CHID	Cisco SFS 7008 Chassis ID Card (FRU) Replacement chassis ID module for Cisco SFS 7008 chassis
SFS-7008-BLANK	Cisco SFS 7008 Blanking Panels <ul style="list-style-type: none"> • Set of blanking panels for Cisco SFS 7008 line cards, switch modules, and power supplies • Suitable for temporarily removing Cisco SFS 7008 FRUs while the switch stays in operation
SFS-7008-RKIT	Cisco SFS 7008 Rail Kit (Option) <ul style="list-style-type: none"> • Optional rail kit specially designed for mounting Cisco SFS 7008 in a rack • Very solid and suitable for rack-based shipping
SFS-7008-BZEL	Cisco SFS 7008 Bezel (FRU) Replacement Bezel for Cisco SFS 7008 chassis
CAB-SFS7008-SER	Cisco SFS 7000 Series Serial Cable Kit Replacement serial cable kit for Cisco SFS 7000 Series switches

Q. What business benefits does high-performance, low-latency InfiniBand switching provide?

A. The Cisco SFS 7000 Series provides a new class of data center infrastructure that interconnects discreet server resources together into a high-performance, 10-Gbps, low-latency fabric. It brings many of the mainframe “backplane” technologies to commodity servers, including Remote Direct Memory Access (RDMA) for transmitting large amount of data with very little overhead, as well as very low latency (less than 4 microseconds of user-space latency). These technologies enable commodity servers to be combined into effective high-performance clusters and grids, creating the kind of collective performance required for enabling the trend toward commoditization. In addition to high-performance computing, InfiniBand also offers the performance, security, and control required for utility computing. The Cisco SFS 7000 Series interconnects servers through a high-performance, unified fabric, and when combined with Cisco SFS 3000 Series multifabric server switches, allows shared pools of I/O and storage resources that can be mapped together by Cisco VFrame Server Fabric Virtualization Software to deliver virtual “compute services.”

Q. Do the Cisco SFS 7000 Series switches support and work with Cisco VFrame Server Fabric Virtualization Software?

A. Yes. The Cisco SFS 7000 Series switches fully support enterprise and research server grid deployments based on Cisco VFrame Server Fabric Virtualization Software. A node does not have to be connected directly to a Cisco SFS 3000 Series Multifabric Server Switch to participate in a server fabric managed by VFrame software. However, to use VFrame, the network design must include a Cisco SFS 3000 Series Multifabric Server Switch in the overall fabric.

Q. What type of cables does the Cisco SFS 7000 Series support?

A. The Cisco SFS 7000 Series uses copper CX-4 cables. Additionally, it is possible to plug pluggable optical modules (POMs) on to the CX-4 ports to make ribbon fiber cables available. See the question about optical connections in the section about building SFS fabrics.

Q. Is the Cisco SFS 7000 Series compliant with Restriction of Hazardous Substances (RoHS)?

A. Yes. The Cisco SFS 7000 Series is fully RoHS-compliant.

HIGH AVAILABILITY

Q. Do the Cisco SFS 7000 Series switches have redundant power and cooling?

- A.** Yes. The Cisco SFS 7000 Series comes standard with redundant power supply and cooling modules. The Cisco SFS 7008 has an optional high-availability kit, which adds redundant power and cooling as well as redundant management capability. On both switches, both power supplies need to be connected for AC redundancy, and a UPS system is recommended for backup power.

Q. What other redundancy features does the Cisco SFS 7008 have that make it highly available?

- A.** The Cisco SFS 7008 supports fully redundant system and fabric management. If the management subsystem on the switch fails, the management capability is automatically and statefully failed over to the backup system. This failover can happen internally or even to another Cisco SFS 7000 Series switch in the fabric. The Cisco SFS 7008 is designed with a passive chassis architecture. There are no active electronics on the Cisco SFS 7008 chassis, so all repairs are done at the FRU level, and after the chassis is installed, it should never have to be replaced. In addition, the Cisco SFS 7008 has full in-service and out-of-service diagnostics that can proactively notify system administrators of any problems with the chassis through syslog or Simple Network Management Protocol (SNMP) trap messages. If a problem does occur, the Cisco SFS 7008 has a set of LEDs that visually indicate which switch has problems, and which FRU on the switch needs to be replaced.

Q. How does redundant fabric management work and what does this mean?

- A.** An InfiniBand network is a pre-routed switched network, much like the telephone system. This gives InfiniBand many desirable properties including resiliency, predictable performance, and monitoring capabilities. As with any managed, switched network, there is a fabric management entity, called the subnet manager, which oversees the fabric. The subnet manager routes, manages, monitors, and responds to individual node requests in the fabric. Because the subnet manager is critical to operation of the fabric, every Cisco SFS 7000 Series switch comes with an embedded subnet manager capable of running fabrics with thousands of nodes. To make sure there are no single points of failure in the fabric, the Cisco SFS 7000 Series switches synchronize the subnet manager state amongst each other, such that at any given moment, any of the switches is capable of managing the fabric. Thus, if one of the switches fails or has to reboot, the overall fabric is not disturbed because any of the other switches can take over fabric management.

Q. How does the Cisco SFS 7000 Series safely manage system upgrades?

- A.** The Cisco SFS 7000 Series supports separate system image loading and activation. For upgrading multiple switches, it is possible to pre-stage new system images, and activate them simultaneously in your fabric during a maintenance window. Additionally, the Cisco SFS 7000 Series supports a default “recovery image” that will be booted if a switch fails to successfully load the target image. If an upgrade fails, the switch will boot in recovery mode, allowing the administrator to revert back to a known good state. On the Cisco SFS 7008 with redundant management controllers, the upgrades are actually done one controller at a time. If the upgrade fails on the primary controller, the switch will automatically fail over to the redundant controller, which will be still operating in a known good state. If the upgrade is successful, the backup controller can then be optionally upgraded or remain in the known good state while the new system image is tested over time.

BUILDING SFS FABRICS WITH CISCO SFS 7000 SERIES SWITCHES

Q. How do Cisco SFS 7000 Series switches integrate with Cisco SFS 3000 Series switches to provide Ethernet and Fibre Channel connectivity over InfiniBand?

- A.** The base fabric for interconnecting the servers in the cluster is typically built out of one or multiple Cisco SFS 7000 Series switches. After that fabric is built, Cisco SFS 3000 Series switches are added to the cluster by attaching directly to Cisco SFS 7000 Series switches through connecting available InfiniBand ports on each switch. The Cisco SFS 3000 Series provides wire-speed I/O connectivity to Ethernet and Fibre Channel networks using Ethernet-to-InfiniBand and Fibre Channel-to-InfiniBand gateways. Each port-to-port connection between a Cisco SFS 7000 Series switch and a Cisco SFS 3000 Series switch provides 10-Gbps or 20-Gbps full duplex connectivity for line-rate I/O traffic to the entire base InfiniBand server cluster. A server can be connected anywhere in the InfiniBand fabric and still have full access to Ethernet or Fibre Channel I/O connectivity provided by the Cisco SFS 3000 Series.

Q. How do I use the Cisco SFS 7000 Series to build InfiniBand fabrics larger than the size of a single switch?

A. InfiniBand fabrics, for both high-performance computing and unified I/O grids, which are larger than a single switch, are built with 2-tier designs by connecting multiple switches together. Switches are connected together using the standard ports on the switch; there are no special ports for connecting Cisco SFS 7000 Series switches together. When building these large fabrics, Cisco SFS 7000 Series switches fulfill two different roles in the big fabric: edge switch and core switch. Any one of the Cisco SFS 7000 Series switches can serve as either an edge switch or a core switch, but not both at the same time. Edge switches use some of their switch ports to connect to servers, and the remaining ports to connect to core switches. Each core switch brings the edge switches together in a web-like fashion by connecting to every edge switch. The number of core switches in the fabric and number of connections between an individual core switch and an edge switch will vary based on fabric design. Core switches do not connect to servers directly. Using this 2-tier model, it is possible to build very large InfiniBand fabrics without sacrificing server-to-server performance.

Q. What is the maximum distance between two ports, either switch ports or HCA ports, in a fabric?

A. Currently, the Cisco SFS 7000 Series supports copper cable distances of up to 15 meters (m) port-to-port in a fabric. It is possible for two servers to be much farther apart than that by stringing switches together, but no distance between two ports can exceed 15 m.

Q. Does the Cisco SFS 7000 Series support optical connections?

A. Yes, the Cisco SFS 7000 Series can use optical connections through external pluggable optical modules (POMs). A POM enables an individual port on a Cisco SFS 7000 Series switch to use 50-micron or 62.5-micron fiber cables. With 50-micron cables, port-to-port fiber distances can exceed 200 m, and with 62.5-micron cables, port-to-port fiber distances go up to 75 m. One POM is needed for each port involved in the connection; so, one switch-to-switch connection would require two POMs. The POMs can be ordered directly from vendors such as Emcore or Fujitsu.

Q. What is a blocking factor in an SFS InfiniBand fabric?

A. InfiniBand fabrics can be designed to be fully nonblocking, which provides full 10-Gbps bandwidth from every server to every other server in the fabric, or partially blocking. Partially blocking designs require fewer switches to build, but provide less average bandwidth to each server. Blocking factors are typically introduced on the edge switch by using more ports to connect servers than to connect the edge switch into the core switches. On the edge switches, if an equal number of ports is used to connect servers and to connect to the core, the fabric is said to be nonblocking. If twice as many ports are used to connect servers as used to connect the edge switch to the core, the blocking factor is said to be two to one or a 50 percent blocking factor. If three times as many ports are used to connect servers as used to connect the edge switch to the core, the blocking factor is said to be three to one or a 67 percent blocking factor. All single switch designs are nonblocking because Cisco SFS 7000 Series switches are internally fully nonblocking.

Q. What effect does the blocking factor have on SFS InfiniBand fabric performance?

A. The blocking factor has absolutely no effect on fabric latency or ability to access I/O resources. However, blocking factor does affect worst-case throughput from server to server. Servers on a 50 percent blocking factor fabric will have half of the maximum bandwidth, 10 Gbps, or 5 Gbps available in the worst case. Servers on a 67 percent blocking factor fabric will on average have one third of the maximum bandwidth or 3.33 Gbps available in the worst case. The reason to consider the worst case is there are many fabric-usage factors that can affect the actual blocking experienced by a server. For example, all communications between servers attached to the same edge switch will be fully nonblocking because, as mentioned in the previous answer, the blocking factor is introduced between the edge switches and the core switches. In addition, blocking will only be experienced if there is contention on the links from the edge switch to the core switch. If other servers on the same edge switch are using the network at less than maximum capacity, it is possible for a server to still have full bandwidth in the fabric because there will not be contention for the uplinks from the edge to the core. Thus, the blocking factor will be the worst-case bandwidth experience by a server node.

Q. How does one calculate the blocking factor?

A. The blocking factor is calculated by dividing the number of ports on an edge switch used for uplinks to core switches by the number of ports used to connect servers. Thus, if on a 24-port Cisco SFS 7000 used as an edge switch, 16 ports are used to connect servers

and 8 ports are used to uplink to the core, the blocking factor is $8 / 16 = 50\%$. All single-switch installations are fully nonblocking because all Cisco SFS 7000 Series switches are fully nonblocking internally.

FABRIC MANAGEMENT

Q. How do I manage the Cisco SFS 7000 Series?

A. Management of the Cisco SFS 7000 Series is integrated into the management suite for all Cisco SFS products. Management options include the Java-based Element Manager GUI, the Web-based chassis manager GUI, and a command-line interface (CLI). For management security, the Web-based management supports HTTPS and the CLI supports Secure Shell (SSH) Protocol v2. Additionally, the switches fully support SNMPv2 and v3 MIBs for integration with other network management utilities. The Cisco SFS 7000 Series management features proactive notification of system events through syslog and SNMP traps. The administrator can configure these active messages to many levels, ranging from frequent alerts to alerts for critical failures only.

Q. What fine-grained user management is supported on the Cisco SFS 7000 Series?

A. The Cisco SFS 7000 Series management supports fine-grained user roles with access control lists (ACLs). Each user can be assigned a different role that grants ACL privileges to view and modify different settings on the switch. These ACLs are strictly enforced, making it possible for diverse groups within an organization to safely manage and monitor a Cisco SFS 7000 Series switch.

Q. What management security features does the Cisco SFS 7000 Series support?

A. The Cisco SFS 7000 Series switches are hardened secure devices. The switches support SSHv2 and HTTPS for secure remote management. Additionally, the switches support SNMPv3 encryption. For user control, the switches support RADIUS for external authentication. Each user is assigned a role with a fine-grained ACL. These ACLs are strictly enforced for each user, making it safe for diverse groups within an organization to manage these switches.

Q. Is it possible to observe the entire fabric and its topology through the Cisco SFS 7000 Series management suite?

A. Yes. It is possible to visually view and monitor the entire InfiniBand fabric from a Cisco SFS 7000 Series management connection. Fabric information such as number of switches, servers, error rates across the fabric, and general performance information are all available in the topology viewer in the Cisco SFS 7000 Series management tools. Additionally, information about overall fabric health, and individual switch and server health is available. This information can be viewed in the Cisco SFS management tools, or accessed through SNMP or CLI-based scripts.

Q. What type of monitoring is supported on the Cisco SFS 7000 Series switches?

A. The Cisco SFS 7000 Series has many different monitoring options. It is possible to get pull performance information on a port-by-port basis in real time. This information includes raw performance data as well as error rates and packet statistics. This data can be expressed as updated text in tables or as live moving graphs over time. It is also possible to graph and compare a number of different data sets at the same time. Additionally, this data can be monitored at a high level to answer questions like, "I want to monitor all ports between server 1 and server 2 for performance and error rates." The administrator can either observe these monitored connections manually, or set triggers that will proactively notify the administrator of certain conditions such as transmission rates dropping below a threshold or error rates exceeding a threshold.

Q. Do I need any additional fabric management software to operate my Cisco SFS InfiniBand fabric?

A. To take advantage of advanced server virtualization and network provisioning, unified I/O fabrics, and enterprise grid capability, Cisco VFrame Server Virtualization Software must be added to the SFS fabric. For extremely large SFS fabrics, in excess of 1500 servers, the Cisco SFS High Performance Subnet Manager must be added to the SFS fabric for performance and advanced management capabilities.

SERVICE AND WARRANTY

Q. What is the warranty for the Cisco SFS 7000 Series switches?

A. Cisco SFS 7000 Series switches come with the Cisco Limited Lifetime Hardware Warranty.

This warranty automatically comes with the purchase of eligible Cisco SFS products, free of charge. Ongoing software updates are also available to customers on Cisco.com free of charge. For specific details on the Limited Lifetime Hardware Warranty, visit:

http://www.cisco.com/univercd/cc/td/doc/es_inpk/lh2cen_.htm

Q. What types of service and support are available for the Cisco SFS 7000 Series switches?

A. In addition to worldwide, 24-hour Cisco Technical Assistance Center (TAC) support, a full complement of implementation-based services are available at this time for the Cisco SFS 7000 Series switches.

TECHNICAL SUPPORT SERVICE

Technical Support Service is available through Cisco SMARTnet[®] support and Cisco SMARTnet Onsite. Cisco SMARTnet support augments the resources of the operations staff by providing access to expertise, both online and on the telephone, and a range of hardware Advance Replacement options. Cisco SMARTnet Onsite complements the hardware Advance Replacement feature by adding the services of a field engineer—services that can be critical for those locations where staffing is insufficient or unavailable to perform parts replacement activities. For more information about Cisco SMARTnet support, visit:

http://www.cisco.com/en/US/products/svcs/ps3034/ps2827/ps2978/serv_group_home.html

ADVANCED SERVICES

Total Implementation Solutions (TIS) offers a full range of implementation solutions including project management, project engineering, configuration, staging, and rollout coordination, to help ensure correct installation and deployment. Configuration services now include development and verification of configuration for intelligent services such as Cisco VFrame Server Virtualization Software. For more information about Advanced Services, please visit <http://www.cisco.com/go/services>

PRODUCT AND CONTACT INFORMATION

Q. Where can I find technical and product specifications and other additional information about the Cisco SFS 7000 Series?

A. For product literature including data sheets and product specifications, see the Server Networking and Virtualization Website at <http://www.cisco.com/en/US/products/ps6418/index.html>.

Q. What are the part numbers for the Cisco SFS Series products?

A. The part numbers and ordering information for the new Cisco SFS Series products is available at <http://www.cisco.com/en/US/products/ps6418/index.html#products>



Corporate 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: 408 526-4100

European Headquarters
Cisco Systems International BV
Haarlerbergpark
Haarlerbergweg 13-19
1101 CH Amsterdam
The Netherlands
www.europe.cisco.com
Tel: 31 0 20 357 1000
Fax: 31 0 20 357 1100

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-7660
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

Cisco Systems has more than 200 offices in the following countries and regions. Addresses, phone numbers, and fax numbers are listed on the **Cisco.com Website at www.cisco.com/go/offices.**

Argentina • Australia • Austria • Belgium • Brazil • Bulgaria • Canada • Chile • China PRC • Colombia • Costa Rica • Croatia • Cyprus • Czech Republic
Denmark • Dubai, UAE • Finland • France • Germany • Greece • Hong Kong SAR • Hungary • India • Indonesia • Ireland • Israel • Italy
Japan • Korea • Luxembourg • Malaysia • Mexico • The Netherlands • New Zealand • Norway • Peru • Philippines • Poland • Portugal
Puerto Rico • Romania • Russia • Saudi Arabia • Scotland • Singapore • Slovakia • Slovenia • South Africa • Spain • Sweden
Switzerland • Taiwan • Thailand • Turkey • Ukraine • United Kingdom • United States • Venezuela • Vietnam • Zimbabwe

Copyright © 2006 Cisco Systems, Inc. All rights reserved. CCSP, CCVP, the Cisco Square Bridge logo, Follow Me Browsing, and StackWise are trademarks of Cisco Systems, Inc.; Changing the Way We Work, Live, Play, and Learn, and iQuick Study are service marks of Cisco Systems, Inc.; and Access Registrar, Aironet, BPX, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, 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, FormShare, GigaDrive, GigaStack, HomeLink, Internet Quotient, IOS, IP/TV, iQ Expertise, the iQ logo, iQ Net Readiness Scorecard, LightStream, Linksys, MeetingPlace, MGX, the Networkers logo, Networking Academy, Network Registrar, Packet, PIX, Post-Routing, Pre-Routing, ProConnect, RateMUX, ScriptShare, ScriptShare, SlideCast, SMARTnet, 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. (0601R)

Printed in USA

C78-376498-00 10/06