



## Q&A

# CISCO SFS 3000 SERIES MULTIFABRIC SERVER SWITCHES

## PRODUCT OVERVIEW

**Q.** What are the Cisco® SFS 3000 Series Multifabric Server switches?

**A.** The Cisco Server Switch is a new class of data center infrastructure that provides a platform to interconnect discreet server resources together into a high-performance fabric, to connect that server fabric with shared pools of I/O and storage resources, and to map the resources together to deliver virtual “compute services” based on application or business policy and priority. The Cisco Server Switch enables the delivery of utility computing by consolidating server I/O into a shared pool, significantly simplifying the server I/O architecture and enabling the virtualization of compute resources.

Table 1 shows the complete list of Cisco SFS 3000 Series switches and optional modules.

**Table 1.** Cisco SFS 3000 Series Switches

Part Numbers	Description
<b>SFS-3001-4X012-SK9</b>	Cisco SFS 3001 Multifabric Server Switch. Includes: <ul style="list-style-type: none"><li>• 12 ports 4X InfiniBand</li><li>• (2) System power supplies</li><li>• (1) Fan tray with 3 fans</li><li>• One expansion slot</li></ul>
<b>SFS-3012-4X024-SK9</b>	Cisco SFS 3012 Multifabric Server Switch (Standard). Includes: <ul style="list-style-type: none"><li>• (1) 12 port 4X InfiniBand switch card</li><li>• (1) System Controller module</li><li>• (1) System power supplies</li><li>• (1) Blower module</li><li>• 12 Expansion slots, 2 switch card slots</li></ul>
<b>SFS-3012-4X024-HK9</b>	Cisco SFS 3012 Multifabric Server Switch (HA). Includes: <ul style="list-style-type: none"><li>• (2) 12 port 4X InfiniBand switch card</li><li>• (2) System Controller Modules</li><li>• (2) System power supplies</li><li>• (2) Blower modules</li><li>• 12 Expansion slots, 2 switch card slots</li></ul>
<b>SFS-3012-MFM-K9</b>	Cisco SFS 3012 Management Interface Module
<b>SFS-X3012-04X12K9</b>	Cisco SFS 3012 InfiniBand 4XIB 12-port switch card
<b>SFS-XETH-01C06K9</b>	Cisco SFS 3000 Series InfiniBand-to-Ethernet Gateway Module
<b>SFS-XFCH-02F02K9</b>	Cisco SFS 3000 Series InfiniBand-to-Fibre Channel Gateway Module
<b>PWR-SFS3012</b>	Cisco SFS 3012 Power Supply

Part Numbers	Description
FAN-SFS3012	Cisco SFS 3012 Blower Module
PWR-SFS3001	Cisco SFS 3001 Power Supply (FRU)
FAN-SFS3001	Cisco SFS 3001 FAN (FRU)

**Q.** What business benefits does the Cisco Server Switch provide?

**A.** The Cisco SFS 3000 Multifabric Server Switch family creates a unified fabric that dramatically simplifies the data center architecture by running multiple types of traffic over a single physical connection. With only one interface card in each server, all resources can be managed on one fabric, eliminating the need to install and manage multiple Ethernet, Fibre Channel, and IPC cards. All types of I/O can be aggregated and load balanced on a single 10-Gbps InfiniBand cable, reducing the number of managed host ports and increasing availability. The Cisco Server Switch then connects servers to a pool of shared Fibre Channel and Ethernet ports via line-rate gateways, and creates virtual I/O subsystems on each host, including virtual HBAs and virtual IP interfaces. When used in conjunction with Cisco VFrame Server Virtualization Software, the Cisco Server Switch can dynamically commission or decommission this virtual I/O pool to any physical server on demand.

## CONFIGURABILITY

**Q.** Why are there two base parts listed for the Cisco SFS 3012 Multifabric Server Switch?

**A.** The Cisco SFS3012 Switch can be bought in a standard configuration and a high-availability (HA) configuration. The customer can buy spares to upgrade a standard configuration to an HA configuration in the field.

**Q.** How many InfiniBand ports are available in the Cisco SFS 3012 Server Switch?

**A.** Each switch card has 12 InfiniBand ports. The standard SKU therefore has 12 ports and the HA SKU has 24. The standard Cisco SFS 3012 can be upgraded to 24 ports by adding a spare switch module (SFS-X3012-04X12K9=) into the chassis. All the ports are 4X InfiniBand, supporting 10 Gbps of total traffic in each direction. (User traffic is 8 Gbps in each direction.)

**Q.** What else does a customer need to access Ethernet and Fibre Channel networks?

**A.** The customer will need to add Cisco SFS 3000 Series InfiniBand-to-Ethernet and InfiniBand-to-Fibre Channel Gateway Modules to access the LAN and SAN. The user will also need 4X IB cables to connect on the IB side, and Ethernet and Fibre-Channel cables to connect to the LAN and SAN respectively.

**Q.** How many gateways can be inserted into the Cisco SFS 3012 Server Switch? Are there any slot requirements?

**A.** The Cisco SFS 3012 Server Switch has 12 slots to accept two types of gateways. InfiniBand-to-Ethernet and InfiniBand-to-Fibre Channel gateways can be deployed in the Cisco SFS 3012, as per the user's needs. There are no slot requirements for the gateways. For example, a Cisco SFS 3012 may have 12 Ethernet Gateways, 12 Fibre Channel gateways, or any combination in-between, in any slots the user prefers.

The Cisco SFS 3001 Server Switch has one expansion slot, and uses the same gateway module as the Cisco SFS 3012.

**Q.** How many ports does the Cisco Ethernet Gateway Module provide?

**A.** The Cisco Ethernet Gateway Module supports six ports that are 100 Mbps / 1GigE, auto negotiated. The ports are copper RJ45 connectors.

**Q.** Does the Cisco Ethernet Gateway Module operate at line rate?

**A.** Yes, all ports perform at line rate, assuming average packet size of 111 bytes or higher.

**Q.** How many ports does the Cisco Fibre Channel Gateway Module provide?

**A.** The Cisco Fibre Channel Gateway supports two FC ports which are auto-negotiated to 1 Gbps or 2 Gbps. The connector is of type "LC".

**Q.** Can you provide sizing guidelines? How many servers can be supported per Ethernet Gateway or Fibre Channel gateway port?

**A.** This depends on the application and the bandwidth loads required for the server fabric. The Cisco Server Switch offers the unique ability to size shared I/O bandwidth independently of the servers themselves. I/O can be dynamically added or removed from the fabric without changing server configurations.

To size bandwidth, we recommend sizing peak and average bandwidth across the server fabric, as well as groups of servers that might share I/O. Then divide that amount of total bandwidth by the bandwidth available per gateway port, and that provides the number of total ports required.

For the Cisco Fibre Channel Gateway, the maximum supported ratio is 32 servers per gateway. So if the cluster has 128 servers, a minimum of four gateways must be used in the Cisco SFS 3012 Server Switch. In some cases, such as heavy media interactions, an inverse consolidation may also make sense. For example, four servers can share six Fibre Channel Gateways. Many customers choose to design their unified fabric with a generic ratio of four servers per Fibre Channel Gateway, and allow available expansion slots to expand later.

For help designing a unified fabric and sizing the number of gateways, please contact SVBU (Server Networking and Virtualization Business Unit ([ask\\_sfs@cisco.com](mailto:ask_sfs@cisco.com))).

**Q.** Do you have recommended designs with the Cisco SFS3012 Server switches?

**A.** Yes, SVBU offers recommended designs for a variety of cluster sizes (16 to 128) in both HA and standard configurations. See the SVBU sales rack site for details.

**Q.** If I want a full HA design, should I implement it with two Cisco SFS 3012 Server switches? Should I use the standard or the HA models?

**A.** For full redundancy, we recommend deploying with two Cisco SFS 3012 switches. In general, the Cisco HA SFS 3012 Server Switch will cause less traffic interruption compared to the standard Cisco SFS 3012. This is because failover of certain components (power, cooling) can take place more quickly without requiring failover of the entire chassis. The final decision will be based on the design and the customer's desire to pay for improved HA performance.

**Q.** Does user traffic failover and/or load balance across gateways? Does this work even if the gateways are in separate switches?

**A.** Yes, load balancing and failover are supported for both Fibre Channel and Ethernet Gateways, across gateways within a single chassis and across gateways located in separate switches.

For the Ethernet Gateway, link aggregation is supported to merge Ethernet ports into one logical link. All traffic is distributed between the ports of the trunk group based on six different distribution algorithms, including source/destination IP, source/destination MAC, and Round Robin. If one link fails, the bandwidth of the trunk group is reduced but traffic is unaffected, and when a link recovers, the bandwidth is added back to the link aggregation group. Active/passive and active/active load balancing is also supported, based on source/destination hardware addresses and source/destination IP addresses.

The Fibre Channel Gateway similarly supports multipathing and load balancing within the InfiniBand fabric, as well as end-to-end load balancing from storage initiator to target.

## INTEROPERABILITY

**Q.** How do Cisco SFS 3000 Series Multifabric Server switches integrate with Cisco SFS 7000 Series Server switches?

**A.** The base fabric for interconnecting the servers in the cluster is typically built out of one or multiple Cisco SFS 7000 Series switches. Once that fabric is built, Cisco SFS 3000 Series Multifabric Server switches are added to the cluster by attaching directly to Cisco SFS 7000 switches through connecting available InfiniBand ports on each switch. The SFS 3000 Switch Family provides wire-speed I/O connectivity to Ethernet and Fibre channel networks utilizing Ethernet to InfiniBand and Fibre Channel to InfiniBand gateways. Each port-to-port connection between a Cisco SFS 7000 Switch and a Cisco SFS 3000 Switch provides 10-Gbps or 20-Gbps full-duplex connectivity for linerate 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 Switch Family.

**Q.** What Ethernet equipment is the Cisco SFS 3000 Series Switch interoperable with?

**A.** The Cisco SFS 3000 Series Switch is qualified for all Cisco IP-based switches. Please contact SVBU (ask\_sfs@cisco.com), for more detail.

**Q.** What FC storage equipment is the Cisco SFS 3000 Series Switch interoperable with?

**A.** The Cisco SFS 3000 Series is qualified with Cisco MDS switches, Brocade switches and McData Sphereon line of switches. The gateway is also qualified with EMC storage, Hitachi Data Systems, IBM storage, and other storage vendors. Please contact SVBU (ask\_sfs@cisco.com), for the current interoperability matrix and vendor certifications.

**Q.** How do I use the Cisco SFS 3000 and 7000 Series Switch families to build InfiniBand fabrics larger than the size of a single switch?

**A.** InfiniBand fabrics which are larger than a single switch are built with two-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 switches together. When building these large fabrics, Cisco 7000 Series family switches fulfill two different roles in the big fabric: edge switch and core switch. Any one of the Cisco 7000 Series family 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 two-tier model, it is possible to build very large InfiniBand fabrics without sacrificing server-to-server performance.

Once an InfiniBand fabric is built using Cisco 7000 Series switches, that fabric can be connected to a third I/O tier of Cisco 3000 Series switches for shared I/O. There is not a requirement that both a leaf and core layer exist. The Cisco SFS 3012 resides in the I/O layer between the InfiniBand core/leaf and the LAN and SAN fabrics.

## HIGH AVAILABILITY

**Q.** Do the Cisco SFS 3000 Series Multifabric Server Switch have redundant power and cooling?

**A.** Yes. The Cisco SFS 3012 Switch is available in two options: a nonredundant option, and an option with redundant power supply / cooling modules. The Cisco SFS 3012 also has an optional high-availability kit (HA 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.

The Cisco SFS 3001 Switch comes standard with redundant, hot-pluggable power supplies. It also has a pluggable fan tray with redundant fans.

**Q.** Does the Cisco SFS 3012 Switch have the concept of recovery images?

**A.** Yes, the Cisco SFS 3012 can boot a recovery image if the active image is corrupted. The controller card has room for two boot images in addition to the recovery image.

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

**A.** InfiniBand is a prerouted 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 (SM), which oversees the fabric. The SM routes, manages, monitors, and responds to individual node requests in the fabric. Because the SM is critical to operation of the fabric, every Cisco SFS 7000 family switch comes with an embedded SM capable of running fabrics thousands of nodes in size. To make sure there are no single points of failure in the fabric, the Cisco SFS 7000 Series switches synchronize the SM state among each other, such that at any given moment, any of the Cisco SFS 7000 Series switches is capable of managing the fabric. Thus, if any one of the switches dies or has to reboot, the overall fabric is not disturbed because any of the other switches are able to take over fabric management.

## FABRIC MANAGEMENT

**Q.** How do I manage the Cisco SFS 3000 Series switches?

**A.** Management of the Cisco SFS 3000 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 SSH v2. Additionally, the switches fully support SNMP v2 and v3 MIBs for integration with other network management utilities. The Cisco SFS 7000 family management features proactive notification of system events through syslog and SNMP traps. These active messages can be configured by the administrator to many levels, from very verbose to only for critical failures.

**Q.** What fine-grained user management is supported on the Cisco SFS 3000 Series switches?

**A.** The Cisco SFS 7000 Series Switch management supports fine-grained user roles with access control lists (ACLs). Each user can be assigned a different role which grants ACL privileges to view and/or modify different settings on the switch. These ACLs are strictly enforced, making it possible for diverse groups within an organization to safely management and/or monitor a Cisco SFS 7000 Series Switch. For example, SAN administrators can be granted access rights to manage the Cisco Fibre Channel gateway modules, and LAN administrators can be granted rights to manage the Cisco Ethernet gateway modules. These privileges are enforced across all management interfaces, including GUI, CLI, and SNMP.

**Q.** What management security features do the Cisco SFS 3000 family of switches support?

**A.** The Cisco SFS 7000 series are hardened secure devices. The switches support SSH v2 and HTTPS for secure remote management capability. Additionally, the switches support SNMP v3 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 safely manage these switches.

**Q.** Is it possible to observe the entire fabric and its topology through the Cisco SFS 3000 Family Management Suite?

**A.** Yes. It is possible to visually view and monitor the entire InfiniBand fabric from a Cisco SFS 3000 family 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 3000 management tools. Additionally, information about overall fabric health, individual switch, and server health is available. This information can be viewed in the Cisco SFS management tools, or accessed via SNMP or CLI-based scripts.

**Q.** What type of monitoring is supported on the Cisco SFS 3000 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 admin if certain conditions are met 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 advance server virtualization and network provisioning, unified I/O fabrics, and enterprise grid functionality, Cisco VFrame Server Virtualization Software must be added to the Cisco SFS fabric.

## SERVICE AND WARRANTY

**Q.** What is the warranty for the Cisco SFS 3000 Series Switch?

**A.** Cisco SFS 3000 Series switches come with the Cisco Limited Lifetime Hardware Warranty. Ongoing software updates are available to customers on the Cisco Website free of charge.

## Limited Lifetime Warranty

The hardware warranty available on the Cisco SFS 3000 Series is the Limited Lifetime Hardware Warranty. This warranty automatically comes with the purchase of eligible Cisco SFS products, 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](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 3000 Series switches?

**A.** In addition to World Wide 7 x 24 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<sup>®</sup> and SMARTnet Onsite. SMARTnet augments the resources of the operations staff by providing them with 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, visit [http://www.cisco.com/en/US/products/svcs/ps3034/ps2827/ps2978/serv\\_group\\_home.html](http://www.cisco.com/en/US/products/svcs/ps3034/ps2827/ps2978/serv_group_home.html).

## Advanced Services

Cisco Total Implementation Solutions (TIS) offers a full range of implementation solutions, including project management, project engineering, configuration, staging, and rollout coordination, ensuring 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 Total Implementation Solutions, 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 3000 family of switches?

**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 switches and family of products?

**A.** The part numbers and orderability information for the new Cisco SFS switches and SFS family of 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 Website at [www.cisco.com/go/offices](http://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 © 2005 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, ASIST, 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, Empowering the Internet Generation, 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, SlideCast, SMARTnet, StrataView Plus, TeleRouter, 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. (0502R) 205410.Z\_ETMG\_JH\_9.05

