

Sun Network 10GbE Switch 72p

Frequently Asked Questions

Overview

10GbE Market

The need for greater networking throughput, the emergence of converged networks (the blending of LAN and SAN traffic over a common backbone), faster access to storage for database, online transactions, backup and recovery, and virtualization are key drivers guiding the adoption of 10GbE in the datacenter. Today's datacenter challenges require accelerated business performance, optimization of asset utilization, lower system acquisition and operational cost, and reduced IT complexity. Enterprise customers want to reduce IT spending while at the same time improving services available to their customers or productivity of their employees. IT managers have the challenge of addressing and solving these seemingly irreconcilable requirements.

Throughput

As the number of sockets and cores have been increasing, so have the required number of Ethernet connections from each server. Typically Ethernet ports are "aggregated" to form bigger pipes, but this consumes too many switch ports and is typically ill-suited to blade architectures. A single 10GbE uplink provides a natural progression by replacing as many as 10 discrete 1GbE connections.

Storage (Bigger pipes, more IOPS)

Storage servers, like Sun's Open Storage 7000 Appliances, benefit from 10GbE connectivity by eliminating traditional bottlenecks and reducing cable sprawl. Storage systems that combine FLASH and SSD technology require significantly larger pipes to balance increased IOPS with network bandwidth.

Virtualization

Server virtualization is gaining traction, from mid-market to high end enterprise. Virtualization effectively brings the networking edge deep into the server kernel where physical

NICs are abstracted and presented to VMs as VNICs. In 1GbE datacenters, a typical server might require 5 NICs in support of virtualization due to the bandwidth limitations of Ethernet pipes. 10GbE NICs provide ample bandwidth that can be effectively shared among multiple VMs, while offering converged services to sharply reduce the number of required cables and drive down operational expenses.

Designed to harness the full capabilities of Oracle's Server Systems and 10GbE NICs, the Sun Network 10GbE Switch 72p brings advanced 10Gb Ethernet Top of Rack switching to datacenter customers. This high-density, non-blocking Top of Rack Switch delivers top performance, low latency, and high throughput to accelerate virtualized applications, and to simplify your datacenter by eliminating the networking bottlenecks and reducing costs.

Customer Benefits

Highest Density Switch on the Market

At 72 ports, the Sun Network 10GbE Switch 72p is the highest 10GbE density Top of Rack (ToR) switch on the market. The Sun ToR switch provides wire rate, low latency, cut through switching on all 72 ports and delivers full, non-blocking bandwidth to each blade or rack server.

Cost effective

This high density 1 RU Top of Rack switch enables a very cost effective scaling for connectivity to the servers, storage and to existing data center networks via 72 10GbE links. The Sun 10GbE ToR switch has enough density to build a sizable fabric in tier 1 that operates at full Bandwidth without the need for spanning tree protocol. A single Sun 10GbE ToR switch can easily displace 3 discrete 24-port or 2 discrete 48-port switches. The Sun 10GbE ToR switch supply of 8 SFP+ uplinks can cost effectively join 64 servers to a tier 2 switch obviating the need for complex tier 1 interconnectivity. This amounts to significant savings in CAPEX. Additionally the Sun 10GbE ToR switch represents a single point of

Sun Network 10GbE Switch 72p Frequently Asked Questions

management for up to 72 servers thus reducing OPEX when compared to meshes of 24-port and 48-port solutions (which constitute upwards of 9 switches necessary to achieve a non-blocking 72 port solution).

Simplification of IT infrastructure and Management

This is the first 10GbE ToR switch that enables multi-rack clusters, simplifies IT infrastructure and reduces network costs by using fewer devices and up to 5x fewer cables, eliminating complex interconnect, and simplifying management.

The Sun ToR is tightly integrated with Sun's standard server management interface, ILOM for initialization and physical management within Oracle Enterprise Manager. This gives customers the ability to seamlessly manage their Network and Server systems with the same interface.

Frequently Asked Questions

What is the Sun Network 10GbE Switch 72p?

The Sun Network 10GbE Switch 72p is a 72 port 10 Gigabit Ethernet Top of Rack (TOR) Switch that offers systems and network administrators unprecedented density and functionality wrapped in a single rack unit. The Sun 10GbE ToR switch offers a comprehensive set of layer 2 and layer 3 protocols and has been extensively tested for interoperability with all incumbent brands of datacenter switching.

The Sun 10GbE ToR switch has a unique front panel arrangement specifically designed to accommodate 72 10GbE wire rate ports. Front panel switch ports are composed of 8 conventional SFP+ ports combined with 16 QSFP (Quad Small Form Factor Pluggable) cages. Each QSFP cage supplies 4 individual 10GbE connections (contributing 64 additional ports) that can be carried in a single cable (40GbE of BE) or split into individual 10GbE interconnects using a variety of copper and fiber cabling options. The 8 ports of SFP+ connectivity can be utilized for either 1GbE or 10GbE cable aggregation and legacy interconnect.

Why is this Switch the most cost effective in the market?

As datacenter deployments of 10GbE attached servers accelerate, the demand for 10GbE switching grows in proportion. Choices in 10GbE switching solutions have expanded to meet customer demand. 10GbE Ethernet switch form factors are loosely categorized as: *edge* (or tier 1) solutions, which include embedded blade switching, top of rack (ToR), or end of row configurations, and *aggregation* (or tier 2) which entails some form modular chassis with varying numbers of 10GbE line cards. Modular chassis solutions are extremely costly especially when there is a requirement for redundancy. Individual Ports on modular 10GbE switching platforms often exceed \$2000. Thus it is highly advantageous to interconnect as many as servers as possible at tier 1 to effectively amortize the expense of each tier 2 uplink.

Simply meshing together conventional 24-port ToR switches does not reduce costs because 10GbE ports are rapidly consumed for servicing interconnects. Additionally, layer 2 switching imposes safeguards to avoid “bridging loops” caused by “meshed” interconnects [among individual tier 1 switches]. Consequently most of the necessary interconnects are blocked from passing traffic by spanning tree protocols.

The Sun Network 10GbE Switch 72p fits in the Top of Rack category, but unlike other products in the market place with 24, 26, or even 48 ports, the Sun 10GbE ToR switch has enough density to build a sizable fabric in tier 1 that operates at full bandwidth without the need for spanning tree protocol. A single Sun 10GbE ToR switch can easily displace 3 discrete 24-port or 2 discrete 48-port switches.

What kind of applications is the Sun Network 10GbE Switch 72p best suited for?

Network intensive enterprise clusters requiring 10GbE network:

- Web Services

Sun Network 10GbE Switch 72p Frequently Asked Questions

- Application Services
- Database access and backup
- Network Storage including NAS and iSCSI

Other applications where fabric bandwidth and ultra-low latency are key:

- Virtualized deployments with live migration
- High throughput multicast applications
- Technical computing

How does the Sun 10GbE ToR switch achieve such a high port count while maintaining full bandwidth between all 10GbE ports?

The Sun 10GbE ToR switch has an internal fabric composed of 9 fully meshed 24-port 10GbE switching elements. Of the nine there are 6 24-port switching elements that act as edge or leaf switches. Each leaf switch contributes 12 ports to the front panel (totaling 72 ports) and 12 ports in the direction of the 3 dedicated “spine” switches. The 3 spine switches have no front panel connectivity, instead they each supply their full complement of 24 ports to the 6 leaf switches thus providing constant cross sectional bandwidth among all front panel ports equaling 1.44Tbps of total capacity.

Why is it feasible/advantageous to interconnect 9 24-port switching elements inside the Sun 10GbE ToR switch when doing this with discrete 24-port switches is so impractical and expensive?

The 9 internal 24-port switching elements are interconnected using Sun/Oracle technology that balances load across the entire internal fabric using “Fat Tree” techniques developed for the industry’s largest super computers. These techniques insure that all available data paths are used for switching traffic between all front panel ports. Attempting the same topology with individual 24-port switches cabled together would be hamstrung by conventional loop-detection protocols (Spanning Tree) resulting in 80% less available bandwidth for routing traffic. The total expense of interconnecting 9 COTS ToR

switches would result in per port costs that rival tier-2 connectivity (> \$2K per port)

How many switch hops does an incoming packet traverse across the ToR’s internal fabric and how does this effect total port-to-port latency?

The internal ToR fabric is composed of two “stages” of switching. Ingress packets arriving on a given interface will either experience 1 or 3 switch hops depending the selection of the egress interface. If both ingress and egress interfaces are on the same internal 24-port leaf switch, then a single switch hop is incurred. However if the egress port is on a different leaf switch, the packet will travel to the second stage of switching and then onto the final egress leaf switch resulting in 3 hops. The Sun 10GbE ToR switch utilizes a an innovative forwarding technique called cut-through-switching. Cut-through-switching enables the ingress [internal leaf] switch to begin forwarding arriving bits to the egress port as soon as the destination MAC address is determined. This approach is in stark contrast to many competitive solutions that are required to buffer the entire packet before the forwarding process can be commenced. Each Sun ToR cut-through-switching hop contributes approximately 300 nanoseconds of transit latency. Thus total transit latency ranges from 300 to 900nanoseconds depending the affinity of ingress to egress forwarding.

What is the logic behind the allocation of front panel ports as either discrete 8 SFP+ or 16 QSFP?

The design center for proportioning the Sun 10GbE ToR switch front panel ports is based on its synergy with the Sun Blade 6000 10GbE switched NEM 10GbE 24p. The primary objective was to enable easy build out of high Bandwidth clusters of 160 blade servers (16 chassis, each with one or two 10GbE Switched NEMs). Each SB6000 10GbE switched NEM provides 3 QSFP uplinks. For medium sized clusters of 160 servers, each NEM would be directly attached to one of three Sun ToRs (3 QSFP uplinks from each NEM, each destined to a dedicated ToR). The resulting topology provides [theoretical]

Sun Network 10GbE Switch 72p Frequently Asked Questions

constant cross-sectional Bandwidth across the entire fabric of 160 servers. The initial release of SEFOS (L2/L3 stack) requires the use of Spanning Tree running on external interconnects. Spanning Tree will place a large proportion of individual uplinks into blocking state to prevent L2 loops, thus reducing overall switch to switch Bandwidth, but preserving resilience from link failures. The most important differentiator in SB6000 to ToR interconnection is the resulting 4 to 1 reduction in cabling attributable to QSFP cabling.

What is QSFP?

QSFP (or quad SFP) is widely used in Technical Computing for InfiniBand interconnect. QSFP is the standard connect for 40Gbps IB also known as Quad Data Rate (QDR). QSFP is also the IEEE standard connector for the emerging 40GbE standard. The Sun ToR is using a QSFP as a pre-standard, but fully compliant interconnect for 4 x 10GbE. The Sun ToR leverages the density of QSFP to effectively quadruple available front panel bandwidth in a space only marginally larger than SFP+. Like SFP+, QSFP supports both copper and optical cabling solutions. Copper cables, known as “Direct Attach” or “TwinX”, are terminated with special electrically coupled fittings that snap directly into the QSFP “cage” obviating the need for costly transceivers. For applications that require optical interconnect, Sun provides a 10Gbps QSFP Short Wavelength Transceiver Pluggable, Parallel-Fiber-Optics Transceiver X2124A. Optical cables are available from Sun as well as from a variety of 3rd party suppliers. Compliant Optical cables must be terminated with a MTP/MPO connector. Sun supplies optical Splitter cables with MTP/MPO (for attachment to the QSFP transceiver) and 4 LC connectors applicable to any compliant 10GbE NIC or switch port.

What are my choices for QSFP cables?

There are 3 types of QSFP cables to choice from, QSFP to QSFP straight-through, QSFP to SFP+ Copper splitters and QSFP optical splitter. These cables come in different lengths.

Please note that all copper choices DO NOT REQUIRE SEPARATE TRANSCEIVERS.

How does the Sun 10GbE ToR switch mount in the datacenter rack in relation to the hot and cold aisle?

In most ToR deployments it is highly desirable to position the front of the switch (side with data ports) in the same physical plane as the Server NICs (or blade switch ports). The in-server-rack deployment thus places the front of the switch in the hot aisle and requires (X2074-R) which directs the airflow from the rear of the ToR out the front (into the hot aisle)

When the Sun 10GbE ToR switch is deployed as an end-of-row or co-resident with datacenter switching kit, the front of switch will be in the same plane as neighboring switch ports in the cold aisle, thus requiring front to back cooling (X2074A-F).

What are my choices for attaching rack mount servers to the Sun ToR?

The ToR front panel is highly optimized for deployments with the SB6000 10GbE Switched NEM, which takes full advantage of QSFP (reducing NEM to ToR uplink cabling by 75%). However the ToR is also well suited for interconnecting rack mount servers by using a combination of SFP+ and QSFP-splitter cables. When the ToR is collocated or adjacent to 10GbE capable servers, the cabling of choice is copper TwinX for both SFP+ and QSFP. If the ToR is collocated with datacenter switching kit, then optical cabling will be required for both SFP+ and QSFP. Optical cabling requires the use of active transceivers for all connections longer than 5M.

How do I attach the ToR to Oracle Unified Storage?

The current shipping versions of the Oracle Unified Storage use XFP transceivers which restrict all connectivity to the ToR to Optical cabling. As such either the 8 SFP+ or 16 QSFP ports can be used. Future

Sun Network 10GbE Switch 72p

Frequently Asked Questions

generations of unified storage will provide standard SFP+ connectivity allowing for the use of inexpensive copper TwinX cabling, tapping either the ToR's SFP+ or QSFP with the aid of a splitter cable.

What is the management interface for the Sun 10GbE ToR switch and how is the management Ethernet interface on the ToR used?

Industry Standard CLI provides an interface to ToR for configuration and Management (Status, performance, monitoring). Customers familiar with industry standard CLI for network management will require minimal training to use the Sun ToR. See the CLI user guide for more details.

<http://docs.sun.com/app/docs/prod/netwk.10ge.72p#hic>

What is the role of the embedded iLOM agent on the Sun 10GbE ToR switch?

ILOM is a critical component of Sun's advanced distributed management fabric. ILOM is resident on all Sun servers, both SPARC and X64 as well as blade switches and the ToR. ILOM is designed to pro-actively monitor all aspects of the management platform including power, FAN health, HW, and temperature measured at multiple sites inside the system chassis. ILOM also supplies a unified enterprise wide agent that is discovered and managed from Oracle xVM Operations Center. In the case of TOR ILOM serves as the "glue" that combines chassis and systems management with host OS and application management (coincident with the amalgamation Oracle Enterprise Manager with xVM). For a full list of other management agents embedded in the Sun ToR, please refer to the Data Sheet found on the Oracle Product Page located under "Ethernet ToR Switches" at

<http://www.oracle.com/us/products/servers-storage/networking/ethernet/032353.htm>

Where do I find software updates for the Sun 10GbE ToR switch?

Software updates for the Sun ToR can be found on SunSolve in the form of a patch. To access SunSolve see the following URL: <http://www.sunsolve.com>

Where do I find troubleshooting guides for the Sun ToR?

Troubleshooting guides for the Sun 10GbE ToR switch can be found at the following URL:

<http://docs.sun.com/app/docs/prod/netwk.10ge.72p#hic>

Where do I find a complete list of supported L2 and L3 features?

For up to date information on supported L2 and L3 features, please refer to the Data Sheet found on the Oracle Product Page located under "Ethernet ToR Switches" at

<http://www.oracle.com/us/products/servers-storage/networking/ethernet/032353.htm>

How many Sun Network 10GbE Switch 72p can I install into a Sun Rack?

Up to two Sun Network 10GbE Switch 72p can be installed into a Sun Rack.

When do I need a separate QSFP transceiver for copper cabling?

If copper cabling is acceptable based on available lengths and deployment logistics, then NO additional transceivers are required. The copper direct attach cable includes mechanical connectors that fit directly into the NEM's SFP+ cages.

Does the copper cable plug directly into the NIC?

Yes, the copper cable has an SFP+ connector on both ends of the cable that directly plugs into the NIC. And the SFP+ connector at other end plugs into the SFP+ port of the networking device the NIC is to be connected to.

What is the distance supported by the SFP+ SR Transceiver?

Sun Network 10GbE Switch 72p Frequently Asked Questions

The supported distance is up to 300 meters depending on the quality of the MMF fiber used by the customer.

Does Oracle offer a MMF cable for SFP+ SR Transceiver?

No Oracle does not offer the MMF Cables for the SFP+ SR Transceiver. These are industry standard and widely available or existing in the customer environment.

Where do I go to purchase passive optical cables for SFP+ and QSFP point-to-point applications?

Oracle does not offer passive optical cables for SFP+ and QSFP point-to-point applications. These cables are readily available in the industry from vendors like Tyco and Molex.

What cables are available from Oracle for ToR and which ones require transceivers?

The Sun Network 10GbE Switch 72p (8 SFP+ and 16 QSFP) can be connected to Oracle rackmount servers, blade servers and various external devices via the transceivers and cables. Quad Small Form Factor (QSFP) Pluggable port (4 channels of 10Gb/s) can be connected to each other with QSFP to QSFP passive copper cables. QSFP to 4 SFP+ TwinX Passive Copper Splitter cable provides breakout to 4 SFP+ connectors for connectivity to 4ea 10GbE SFP+ ports. All Passive Copper Cables connect directly into the ToR ports and Transceivers should not be used with these Passive Copper Cables. In general Optical Connectivity requires transceivers and TwinX cables do not require transceivers.

How do I attach QSFP to SFP+?

For direct attach copper Oracle offers the 10Gbps QSFP to 4 SFP+ TwinX Passive Copper Splitter Cable. This Splitter cable provides 4 SFP+ connectors. These are available in different lengths, 1M, 3M and 5M

For optical cabling applications both QSFP and SFP+ transceivers are required. On the NEM side use the x2124A QSFP transceiver. Attach the Optical splitter cable to the

QSFP transceiver and on the remote end use optical SFP+ transceivers that compatible with supplier's interface.

What is the advantage of QSFP to QSFP passive TwinX cables?

The advantages of the QSFP to QSFP passive TwinX Cable are:

- low cost option for shorter distances up to 5 meters to connect to the QSFP ports
- A single cable provides 4 to 1 cable reduction
- QSFP Transceiver is not needed with these cables

What is the advantage of SFP+ TwinX copper cable?

It is a low cost option for shorter distances up to 10meters

What is the advantage of QSFP to 4 SFP+ Splitter passive TwinX cables?

The advantages of the QSFP to 4SFP+ Splitter passive TwinX Cable are:

- low cost option for shorter distances up to 5 meters to connect the QSFP ports to 4 different SFP+ ports with a single cable.
- QSFP Transceiver is not needed with these cables

How do I use the SFP+ ports for 1000BaseT?

You need to purchase the Oracle Adapter SFP+ to 1000BASE-T, Copper SFP+/RJ45, part number x2123A.

Do the SFP+ optical transceivers support 1GbE operation?

Yes, they Support 1GE and 10GE dual rates.

Will the SFP+ optical transceivers auto negotiate between 1GbE and 10GE?

Auto-negotiation is not supported between the 10GE and 1GE speed. The transceiver must be manually configured to operate at 1GE speed.

Sun Network 10GbE Switch 72p Frequently Asked Questions

Does the SFP+ TwinX copper cable plug directly into the NIC?

Yes it includes SFP+ connector on both ends of the cable that directly plugs into the NIC. The SFP+ connector at the other end plugs into the SFP+ port of the networking device that the NIC is to be connected to.

What is the distance supported by the SFP+ SR Transceiver?

The SFP+ SR Transceiver supports up to 300 meters depending on the quality of the MMF fiber used by the customer.

What is cut-through switching?

Cut-through switching begins forwarding arriving bits as soon as the destination is determined (Ethernet Destination MAC address received). The entire packet is rarely buffered.

Do I need to configure VLANs on the Switch?

VLANs are configurable on the Switch per application need.

What are the systems management and security options for the Sun Network 10GbE Switch 72p?

- 1) Industry Standard CLI provides an interface to ToR for configuration and Management - Customers familiar with industry standard CLI for network management will require minimal training to use the ToR
- 2) Integration with Sun ILOM for initialization and physical management - The ToR is managed as a chassis component in a manner consistent with blades and power supplies
- 3) Simple Network Management Protocol or SNMPv3 (v1, v2, v3) for management with standard network management solutions - ToR can be managed with standard off-the-shelf network management solutions. The ToR ships with dozens of enterprise and proprietary MIBs (Management Information Base)

4) Support for Access Control Lists for creating fine grain policies for network security - ACLs are the fundamental building blocks to network security and allow customers to create fine grain policies to protect servers and the network from aberrant traffic which is intentionally or anecdotally malicious

5) Port based authentication - Management simplification

6) Port Mirroring for troubleshooting and network management - Port mirroring is an essential troubleshooting tool for server and network administrators. Port mirroring is considered a best practice for network management

7) The remote network monitoring (RMON) runs entirely on the NEM and is controlled either by the Common Line Interface (CLI) or remotely by a RMON management system - RMON like port mirroring is an essential troubleshooting tool used by network administrators to troubleshoot network issues. RMON is designed to offload system monitoring from a centralized management host by performing polling of critical functions as an embedded function.

What Servers does the Sun Network 10GbE Switch 72p support?

The Sun Network 10GbE Switch 72p supports all current and future Sun Blade servers, Sun rackmount servers and Sun Storage devices that support SFP+ and QSFP ports.

What kind of service and warranty is offered on the Sun Network 10GbE Switch 72p?

The Sun Network 10GbE Switch 72p warranty is 1 year, 2nd business day.

Where can I find details on the Sun 10GbE ToR switch?

A link to the Sun ToR web page is located under the Ethernet ToR Switches section on the Ethernet Networking Products landing page at

<http://www.oracle.com/goto/networking/ethernet>

Sun Network 10GbE Switch 72p Frequently Asked Questions

Where can I find documentation for Sun Network 10GbE Switch 72p?

Documentation for the Sun Network 10GbE Switch 72p can be found at:

<http://docs.sun.com/app/docs/prod/netwk.10ge.72p#hic>



Sun Network 10GbE Switch 72p Frequently Asked Questions



Oracle Corporation

Worldwide Headquarters

500 Oracle Parkway
Redwood Shores, CA
94065
U.S.A.

Worldwide Inquiries

Phone
+1.650.506.7000
+1.800.ORACLE1

Fax
+1.650.506.7200

oracle.com

Copyright © 2010, Oracle and/or its affiliates. All rights reserved. This document is provided for information purposes only and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. UNIX is a registered trademark licensed through X/Open Company, Ltd. 0110