OceanStor SNS2124, SNS2224, SNS2248, SNS3096, SNS5192, and SNS5384

Technical White Paper

Issue 1.0

Date 2014-06-25

HUAWEI TECHNOLOGIES CO., LTD.

Trademarks and Permissions

Notice

Huawei Technologies Co., Ltd.

Address: Huawei Industrial Base
Bantian, Longgang
Shenzhen 518129
People's Republic of China

Website: http://www.huawei.com

Email: support@huawei.com

Tel: 4008837333

Contents

1 Introduction

The OceanStor SNS V1R2 series Fiber Channel switches (the SNS V1R2 series for short) are produced by Huawei Technology Co., Ltd. and are composed of 6 models: OceanStor SNS 2124, OceanStor SNS 2224, OceanStor SNS 2248, OceanStor SNS 3096, OceanStor SNS 5192 and OceanStor SNS5384.

The SNS V1R2 series are adaptive Fiber Channel switches and are compatible with Fiber Channel SAN using the latest SNS single-chip architecture. The SNS V1R2 series are small to high and medium sized business (SMBs) switches, capable of meeting all-sized SAN networking requirements.

The SNS V1R2 series integrate reliability, Availability, and serviceability (RAS) features and comply with scalable requirements of enterprise switches, interoperable and easy to use.

2 Solution

2.1 Enhanced Management Software

1. You can monitor and manage Fabric OS switches using the web-based **Web Tools** without installing it. The **Web Tools** provides a series of functions for switch task management, including switch firmware upgrade, license management, route policy selection, ISL trunking management, and long-distance transmission management.

Figure 2-1 Management interface of the Web Tools

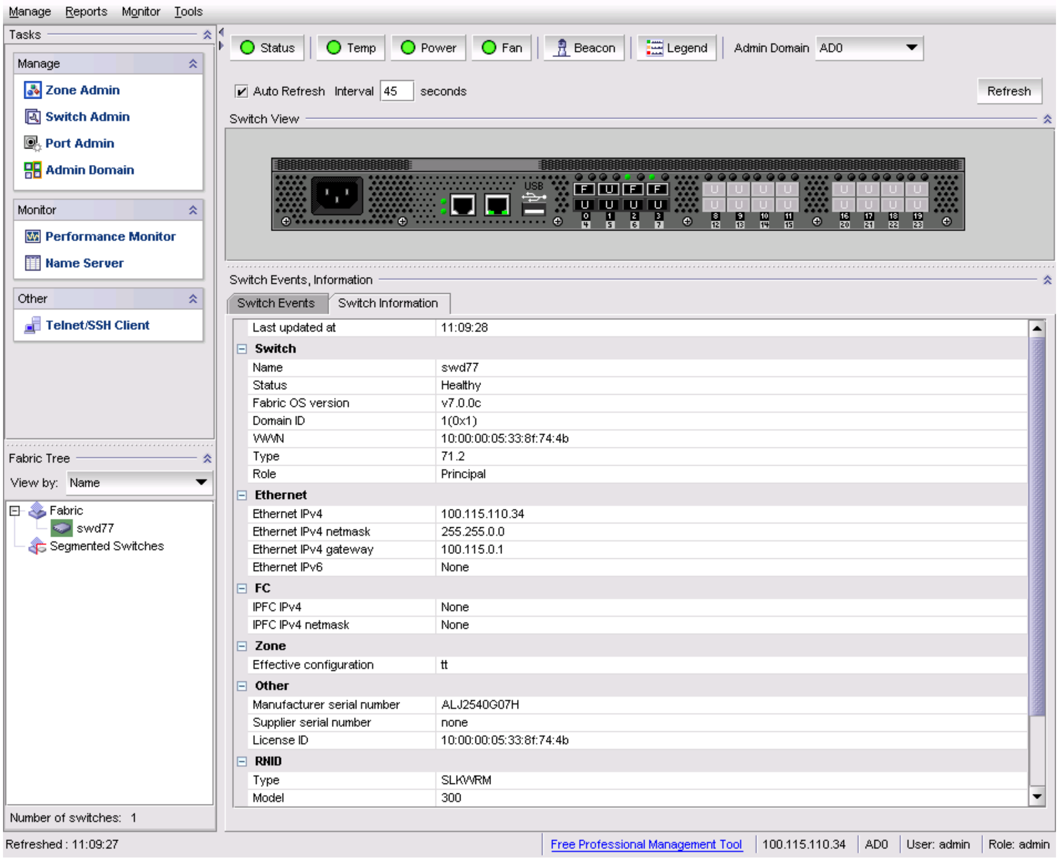

2. You can manage the Fabric using the **CLI**.

3. **BNA** (Brocade Network Advisor)helps organizations simplify network operations by providing a single tool to manage the entire network lifecycle, including monitoring, diagnostics, change management, troubleshooting, and remediation. BNA also provides features that automate time-consuming tasks, freeing up time and resources so organizations can focus on more strategic IT initiatives. Some of these capabilities include: Real-time and historical performance data reporting, Event management, VM-to-storage visibility, Advanced Call Home, etc.

Figure 2-2 Management interface of BNA

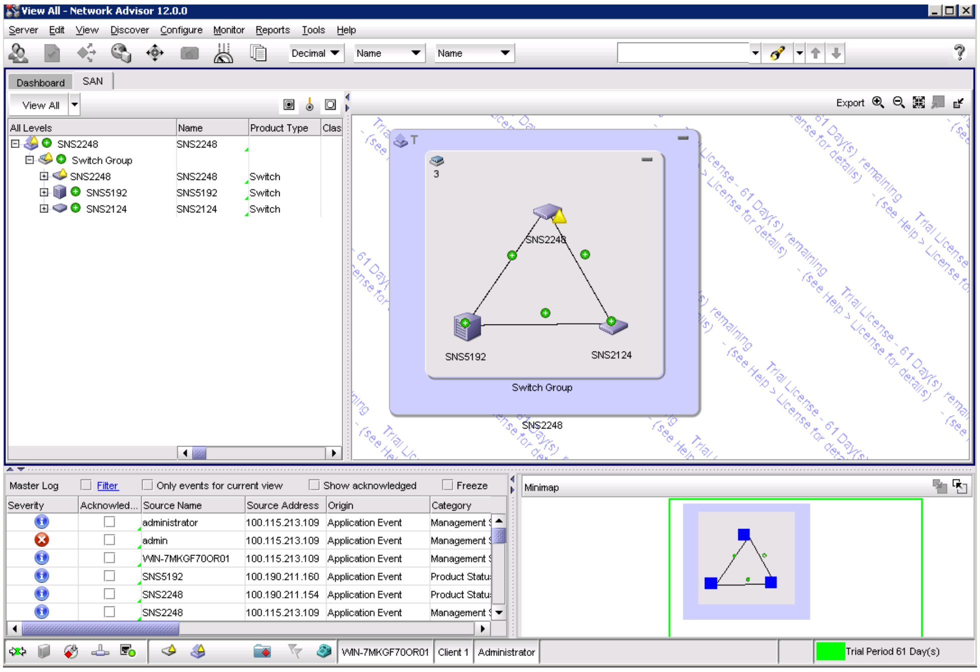

2.2 Access Gateway Mode

The Access Gateway (AG) mode of the SNS V1R2 series is implemented by the N_Port ID Virtualization (NPIV) technology. Switches produced by different vendors cannot be cascaded through the E_Port. This is why the AG mode appears. The SNS V1R2 series can be connected to other switches in AG mode, which technically avoids E_Port cascading. The AG mode ensures connection flexibility of the SAN.

1. After configured to the AG mode, a switch is not connected to other switches in the Fabric through the E_Port, but accesses the Fabric through the N_Port. In AG mode, the switch is connected to a device or a director that supports the NPIV technology. Therefore, the AG mode is not limited to the SAN Fabric of a specified vender.
2. In AG mode, a switch presents an N_Port not an E_Port to the Fabric. Therefore, other switches in the Fabric simply deem the AG-mode switch as a node device not a switch. Additionally, a new switch in AG mode only adds multiple ports. No domain ID assignment needs to be considered. This simplifies Fabric management.

Figure 2-3 AG mode

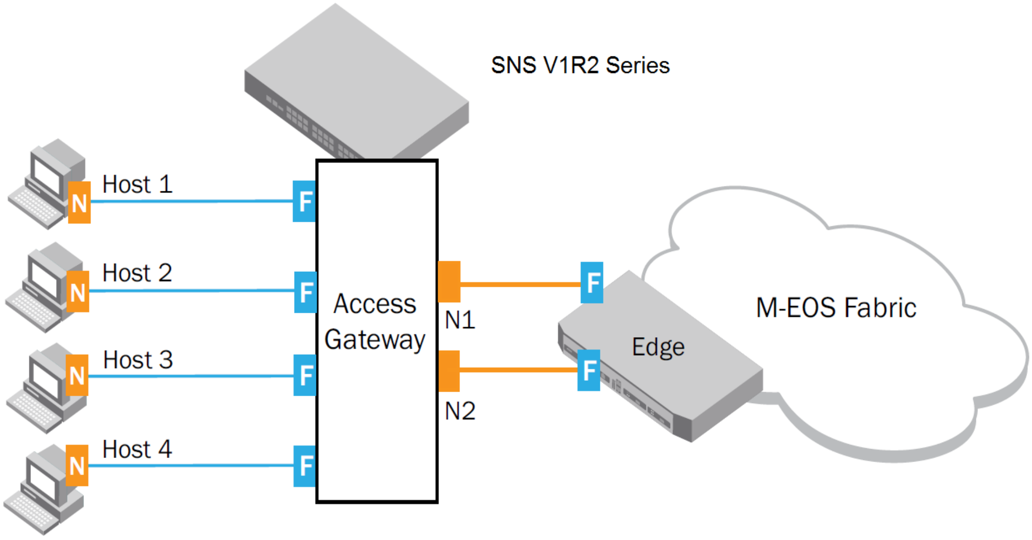

2.3 D_Port (not include SNS2124)

OceanStor SNS V1R2 series Fiber Channel switches supply D_Port which fulfill clearLink Diagnostics. Ensure optical and signal integrity for Gen 5 Fibre Channel optics and cables, simplifying deployment and support of high-performance fabrics.

D_Port test consists of four steps:

(1) Electrical loopback test (E-WRAP)

(2) Optical loopback test (O-WRAP)

(3) Link traffic test

(4) Link latency and distance measurement

Figure 2-4 D_Port

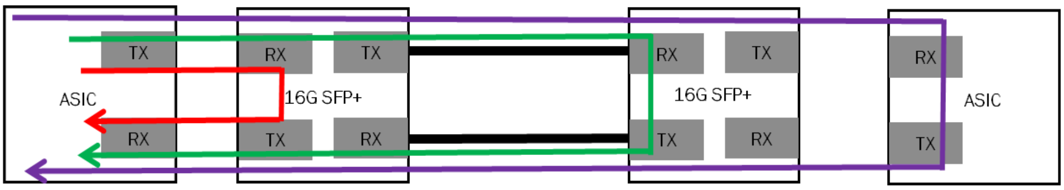

2.4 Virtual Fabric (not include SNS2124/2224)

Virtual Fabric is a hardware virtualized architecture. Conventionally, storage networks are designed and managed on physical switch layer. However, Virtual Fabric allows storage network design and management to be implemented on port layer. The SNS 2248 switch supports the Virtual Fabric feature.

Specifically, Virtual Fabric consists of three functions: logical switch, logical Fabric, and device sharing, which can be customized on demand.

1. Logical switch

   Logical switches are basic components of Virtual Fabric. After enabling the Virtual Fabric function, a switch can be divided into multiple logical switches. Furthermore, ports on the physical switch can be dynamically assigned to any logical switches. Before logical switches are created, the switch with the Virtual Fabric function enabled is deemed as a logical switch by default, as shown in Figure 2-5. A default logical switch cannot be deleted.

Figure 2-5 Physical switches and logical switches

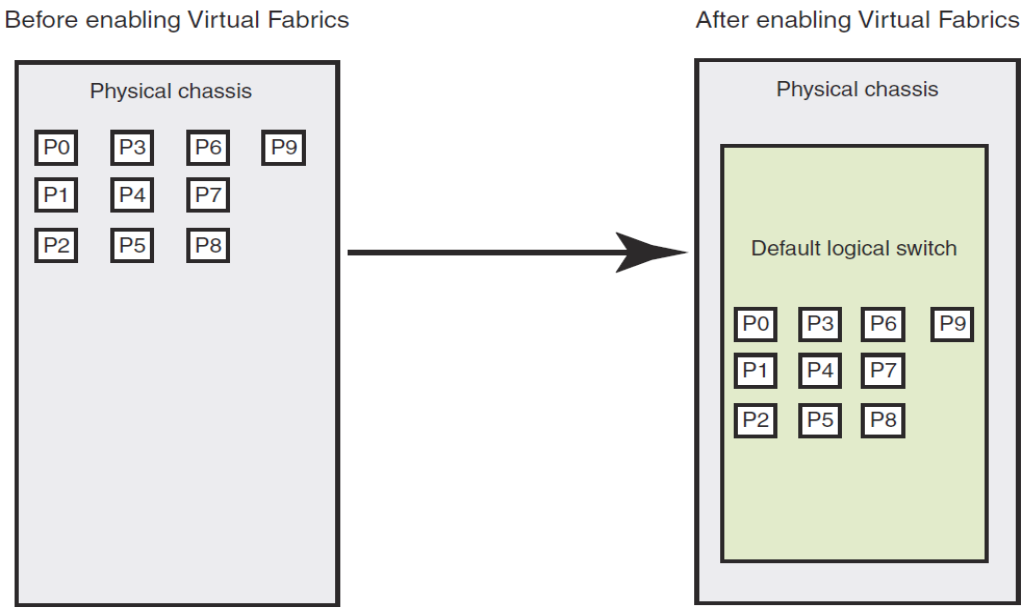

2. Logical Fabric

   The Fabric ID (FID) is the ID for a logical switch in a specified logical Fabric. Logical switches of different physical switches, if owning the same FID, can be assigned to the same logical Fabric. Logical switches in the same logical Fabric support ISL cascading, frame aggregation, and dynamic path selection (DPS).

Virtual Fabric can turn a physical switch into multiple separate logical switches by port-level classification. Meanwhile, Virtual Fabric can improve data transmission, management, and fault isolation. Virtual Fabric also supports Facility Operating System (FOS) features such as frame aggregation, DPS, Fiber Channel routing, access gateway, and Fiber Channel over IP (FCIP).

2.5 Trunking

The trunking function of SNS V1R2 series combines a group of links into a single logical link to optimize bandwidth usage. Traffic is dynamically and sequentially assigned in a trunking group, so that fewer links can implement higher performance, as shown in Figure 2-6. In a trunking group, multiple physical ports serve as a single port to simplify management. The trunking function ensures system reliability by transmitting data in order and avoiding I/O retries after a link breakdown. The frame-based trunking technology of the SNS V1R2 series is more precise and balanced than the switch-based trunking technology.

Figure 2-6 Frame-based trunking

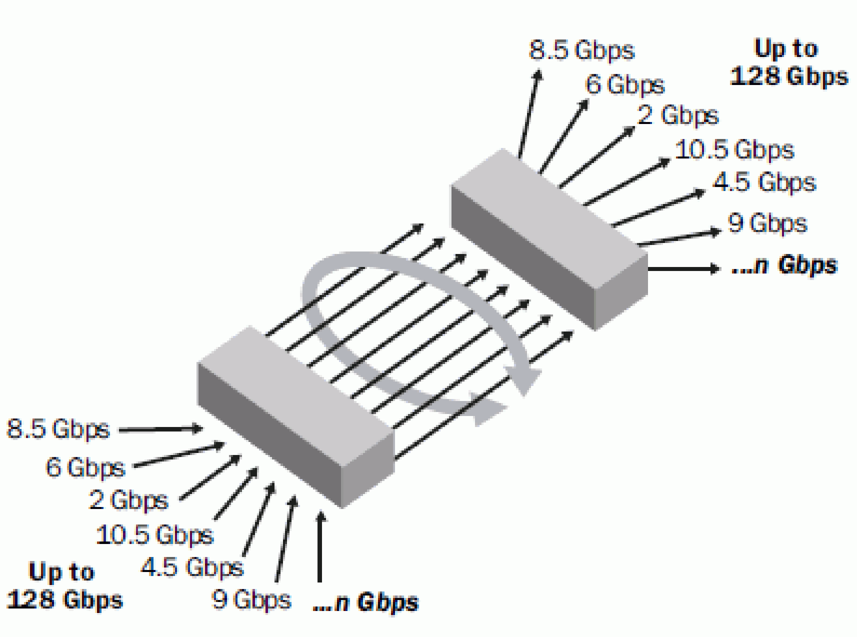

The trunking function of the SNS V1R2 series consists of ISL trunking (E_Port trunking), EX_Port trunking, F_Port trunking, and N_Port trunking, which all require trunking license.

2.6 16 Gbit/s (not include SNS2124)

The 16 Gbit/s Fiber Channel (16 Gbit/s for short) is the latest reformation of SAN for meeting the demand of big data transmission and high performance. On demand of current storage applications, 16 Gbit/s represents higher performance and lower power consumption. Compared with conventional Fiber Channel rates, 16 Gbit/s improves the link efficiency by the 64b/66b encoding technology.

Compared with the 8 Gbit/s Fiber Channel, 16 Gbit/s improves the throughput to 1600 Mbit/s. 16 Gbit/s uses the 64b/66b encoding technology to improve the encoding efficiency to 97%, while the conventional 8b/10b technology only has an encoding efficiency of 80%. Therefore, 16 Gbit/s improves link performance with few increased cost.

16 Gbit/s is also backwards compatible with conventional Fiber Channel rates. The application-specific integrated circuit (ASIC) of 16 Gbit/s supports both the 64b/66b and 8b/10b encoders. The ASIC connecting to the SFP+ optical module has a coupler connected to

each encoder, as shown in Figure 2-7. The switch, according to the selected rate, routes data flows to different encoders. In the process of rate negotiation, two ends of the link determine the maximum rate that the link supports.

Figure 2-7 16 Gbit/s backward compatibility

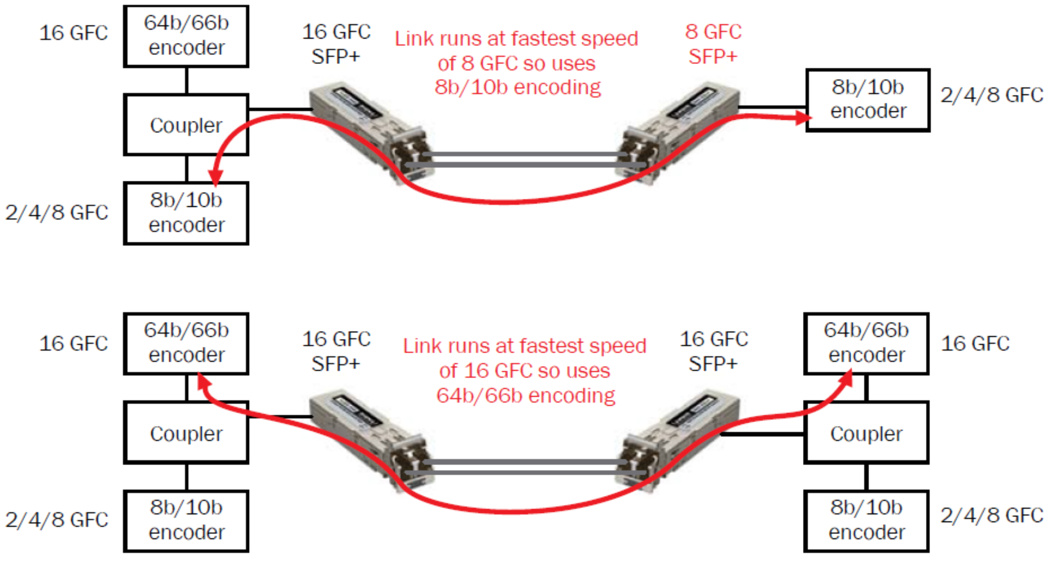

16 Gbit/s has the following advantages:

1. Reduces links, HBAs, and switch ports under the same business load.
2. Lowers power consumption per bit.
3. Simplifies cable management.

2.7 Long-Distance Transmission

According to Fiber Channel protocol standard definitions, Fiber Channel ports on any end of expansion must support high-level buffer credits for long-distance Fiber Channel transmission.

A buffer credit allows the transmit (TX) to send a data frame. The TX does not send the next data frame until receiving an acknowledgement from the receive (RX) indicating that the data frame has arrived. There is an acknowledgement latency during this process. The latency increases as the distance elongates, and the bandwidth available for users decreases. Therefore, the key for long-distance Fiber Channel transmission is to provide a large number of buffer credits for the TX. In this way, the TX can keep sending data frames through the Fiber Channel when waiting for acknowledgements.

The optimal number of buffer credits is determined by distance (frame transmission latency), frame sending latency, link signal frequency, and frame size. When the link rate improves and frame transmission latency reduces, buffer credits must increase to ensure full utilization of the link. However, excessive buffer credits may cause data frame congestion on the RX. Buffer credits are assigned to a group of ports of a switch from a same buffer pool. How buffer credits are assigned varies for different applications and scenarios.

2.8 Fabric Vision(not include SNS2124)

Fabric Vision technology, an extension of Gen 5 Fibre Channel, introduces a breakthrough hardware and software solution that maximizes uptime, simplifies SAN management, and provides unprecedented visibility and insight across the storage network. Offering innovative diagnostic, monitoring, and management capabilities, Fabric Vision technology helps administrators avoid problems, maximize application performance, and reduce operational costs.

Fabric Vision technology includes several critical diagnostic, monitoring, and management capabilities that help to increase fabric resiliency, reduce downtime, and optimize application performance. Bottleneck Detection, Forward Error Correction (FEC),Credit Loss Recovery. Flow Vision is a Fibre-Channel SAN network diagnostic tool supported on all platforms supported by Fabric OS 7.2 and later, that provides you with a comprehensive vision of fabric traffic flows and with the ability to non-disruptively create and capture copies of traffic flows for later analysis. Flow Vision also provides a test flow generation capability that you can use to pre-test a SAN infrastructure for robustness. Flow Vision has three features: Flow Monitor, Flow Generator, and Flow Mirror.

2.9 High Availability(only support SNS5192/5384)

High Availability (HA) features provide maximum reliability and nondisruptive replacement of key hardware and software modules. High Availability is available only on SNS5192 and SNS5384. HA supplies redundant active/passive control processor; redundant active/active core switching blades; redundant WWN cards. Designed to provide 99.999 percent uptime capabilities.

Figure 2-8 HA

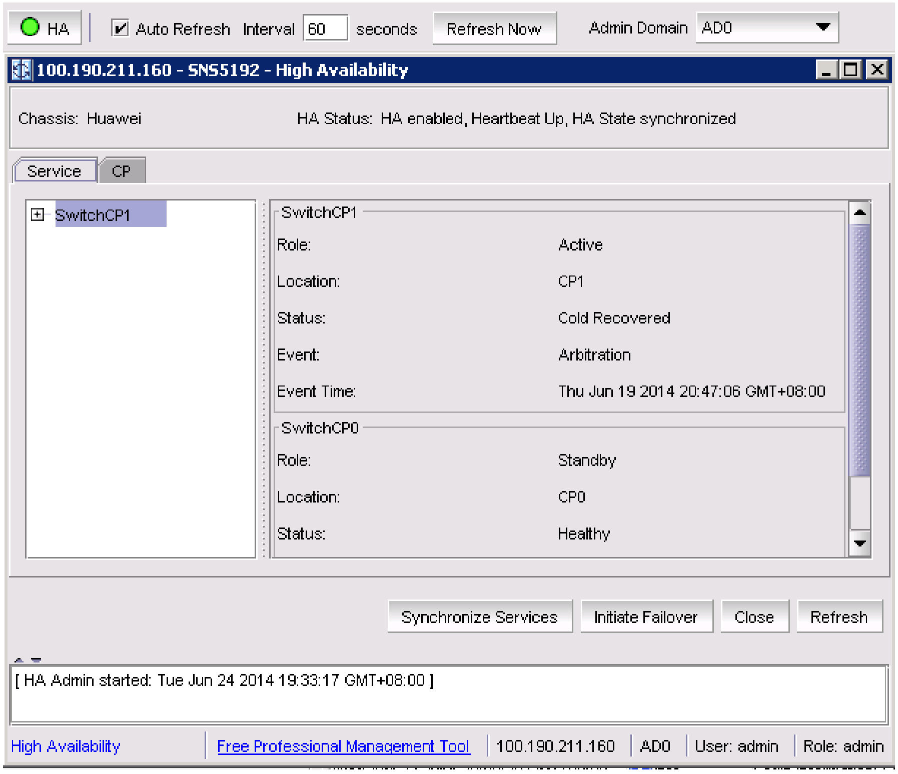

2.10 Simplified Scale-Out Network Design

(only support SNS5192/5384)

Networks are evolving in order to adapt to rapid growth and change in the server and storage infrastructure. UltraScale chassis connectivity leverages optical Inter-Chassis Links (ICLs) of up to 100 meters to connect up to 10 SNS5192/5384s, enabling flatter, faster, and simpler fabrics that increase consolidation while reducing network complexity and costs.

UltraScale ICLs enable scalable core-edge and active-active mesh chassis topologies. These high-density chassis topologies reduce inter-switch cabling by 75 percent and free up to 33 percent of ports for server and storage. This maximizes overall port density in the lowest amount of rack space.

3 Acronyms and Abbreviations

Table 3-1 OceanStor SNS acronyms and abbreviations

Acronym and Abbreviation	Definition
Activity LED	A port LED that indicates when frames are entering or leaving the port
AL_PA	Arbitrated Loop Physical Address.A unique one-byte value assigned during loop initialization to each NL_Port on a loop.
Alarm	A message generated by the switch that specifically requests attention. Alarms are generated by several switch processes. Some alarms can be configured
Alias	A named set of ports or devices. An alias is not a zone, and can not have a zone or another alias as a member
Arbitrated Loop	A Fibre Channel topology where ports use arbitration to establish a point-to-point circuit
Arbitrated Loop Physical Address (AL_PA)	A unique one-byte value assigned during loop initialization to each NL_Port on a loop
ASIC	Application Specific Integrated Circuit. A chip designed for a specific applications, such as a transmission protocol or a computer
Auto Save	Zoning parameter that determines whether changes to the active zone set that a switch receives from other switches in the fabric will be saved to permanent memory on that switch
BootP	A type of network server

Buffer Credit	In buffer-credit flow control, the source and destination set the number of unacknowledged frames (buffer credits) allowed to accumulate before the source stops sending data
Cascade	A serial connection
Cascade Topology	A fabric in which the switches are connected in series. If you connect the last switch back to the first switch, you create a cascade-with-a-loop topology
Class 2 Service	A service which multiplexes frames at frame boundaries to or from one or more N_Ports with acknowledgment provided
Class 3 Service	A service which multiplexes frames at frame boundaries to or from one or more N_Ports without acknowledgment
CLI	Command Line Interface
Default Visibility	Zoning parameter that determines the level of communication among ports/devices when there is no active zone set
Domain ID	User defined number that identifies theswitch in the fabric
Effective configuration	The zone set that defines the current zoning for the fabric
E_Port	Expansion Port
Event Log	Log of messages describing events that occur in the fabric
F_Port	Fabric Port
Fabric Management Switch	The switch through which the fabric is managed
Fabric Name	User defined name associated with the file that contains user list data for the fabric
Failover	Automatically switching control from one CPU to another due to an error condition
FDMI	Fabric Device Management Interface

FL_Port	Fabric Loop Port
Flash Memory	Memory on the switch that contains the chassis control firmware
Frame	Data unit consisting of a start-of-frame (SOF) delimiter, header, data payload, CRC, and an end-of-frame (EOF) delimiter
FRU	Field Replaceable Unit
FT feature	Fault Tolerant feature,which implement the fail-over feature between two CPU blade in one 5120
GUI	graphical user interface
Hard Zone	Hard zoning divides the fabric for purposes of controlling discovery and inbound traffic
Hop	A measure of fabric latency represented by the ISL that any frame crosses when travelling from one switch to another. A frame that travels from one switch to another over an ISL
Inactive Firmware	The firmware image on the switch that is not in use
In-band Management	The ability to manage a switch through another switch over an inter-switch link
Initiator	The device that initiates a data exchange with a target device
In-Order-Delivery	A feature that requires that frames be received in the same order in which they were sent
Input Power LED	A chassis LED that indicates that the switch logic circuitry is receiving proper DC voltages
Inter-Switch Link	The connection between two switches using E_Ports
LIP	Loop Initialization Primitive sequence

Logged-in LED	A port LED that indicates device login or loop initialization status
Management Information Base	A set of guidelines and definitions for SNMP functions
N_Port	Node Port
NPIV	N_port ID virtualization
Pending Firmware	The firmware image that will be activated upon the next switch reset
Port Binding	An authorization method that defines a list of device WWNs that can login to a switch port.
POST	Power-On Self Test Diagnostics
Principal Switch	The switch in the fabric that manages domain ID assignments
RSCN	Registered State Change Notification
SFP	Small Form-Factor Pluggable
SNMP	Simple Network Management Protocol,An application protocol that manages and monitors network communications and functions
Soft Zone	Soft zoning divides the fabric for purposes of controlling discovery. Members of the same soft zone automatically discover and communicate freely with all other members of the same zone
Target	A storage device that responds to an initiator device
trunk	A batch of port which can be transmit concurrently
WWN	Worldwide Name,A unique 64-bit address assigned to a device by the device manufacturer