

Solaris 10 and Emulex Virtual HBA Technology Maximize Server Virtualization Benefits

AT A GLANCE

Solaris™ 10 provides enhanced functionality for virtual server environments. Solaris data centers have long used Fibre Channel to implement high-performance shared storage, with Emulex LightPulse® Fibre Channel host bus adapters (HBAs) being a popular choice for Sun servers. Emulex, a technology leader in Fibre Channel, has been enhancing Fibre Channel to perform well in virtual server deployments. Emulex Virtual HBA technology, based on N_Port ID Virtualization (NPIV) and a new version of HBAnyware®, have been extended to provide virtual port (VPort) management with the Solaris 10 operating system. The result is that Sun users can enhance the security and performance of their Solaris Containers, Logical Domain (LDom) and/or xVM technology implementations in a Fibre Channel storage area network (SAN).

PRODUCTS

- Emulex LightPulse Fibre Channel 4Gb/s and 8Gb/s HBAs
- Emulex HBAnyware Centralized HBA Management
- Solaris StorEdge SAN Foundation Software (SFS) FCA driver (versions 2.31 and above) distributed by Sun Microsystems

Overview

Solaris 10 delivers advanced server virtualization with its support of Solaris Containers, LDom and xVM technologies. Solaris Containers use virtualization with a one-application-per-server deployment model, while at the same time sharing hardware resources. LDom, in conjunction with the CoolThreads technology in Sun Fire UltraSPARC T1 and T2 servers, work at a systems virtualization level, allowing multiple operating systems (OS) called “guests” to run simultaneously. Sun xVM delivers a unique combination of management and virtualization capabilities to Solaris 10 users. Sun environments can run Solaris Containers within an LDom, delivering the flexibility of multiple OS versions, with the scalability and performance of Solaris Containers.

Sun users looking to maximize their data center operations and access to storage have deployed Fibre Channel SANs. Emulex Fibre Channel HBAs are a popular choice for Sun users whether obtained directly from Emulex or through Sun’s reselling of Emulex Fibre Channel HBAs.

Emulex worked closely with Sun Microsystems to deliver seamless support for Emulex LightPulse Fibre Channel HBAs in a Solaris 10 virtual server environment through the SFS FCA driver (Leadville driver) which is fully integrated into Solaris 10. In fact, Emulex’s “emlxs” driver is distributed by Sun as a fully integrated component of its SFS FCA driver for Solaris 10. Fibre Channel management became simpler when the SFS FCA driver was fully integrated into Solaris 10.

A technology leader in Fibre Channel, Emulex has been enhancing Fibre Channel to perform well in virtual server deployments for many years. NPIV is a method for creating VPorts by allowing multiple Fibre Channel fabric addresses to share one physical Fibre Channel port. Emulex played a major role in delivering the first NPIV-enabled HBA in 2005 after developing the technology with IBM and driving it toward American National Standards Institute (ANSI) International Committee T11 for Information Technology Standard (INCITS) standardization.

Emulex and Sun's joint development of NPIV support for the Solaris Operating Environment, and Emulex's incorporation of the feature into its LightPulse Fibre Channel 4Gb/s and 8Gb/s HBAs give Solaris users:

- ▶ Greater security and manageability.
- ▶ Higher availability and increased reliability of each guest.
- ▶ Increased confidence in supporting Service Level Agreements (SLAs).
- ▶ Reduced capital, management and operational costs.

Additionally, Emulex HBAnyware, the industry's most comprehensive HBA management tool, now provides virtual port management for Solaris as well as other operating systems and hypervisors. With HBAnyware, management of virtual ports enabled by NPIV is just a click away.

Emulex LightPulse Virtual HBA technology, based on NPIV and a new version of HBAnyware which provides VPort management, is now available on Solaris 10. This white paper describes how Emulex LightPulse Virtual HBA technology benefits Solaris 10 environments.

Sun Virtualization Technologies

With the largest installed base of any UNIX operating system, Solaris ranges from thin web servers to robust Telco processors and massive enterprise mainframes. Solaris 10 offers a number of virtualization technologies not previously available. They are described below:

- ▶ **xVM Server** is Sun's implementation of the Open Source Xen technology. This technology provides the ability to support a variety of guests, implemented as "User Domains" or "DomU" on an x86 or x64 server. Each User Domain can run Solaris, Linux, BSD and Windows (the latter on AMD-V or Intel VT processors). The Xen architecture does not set any practical limitation to the number of simultaneous user domains (consult Sun for supported limits on specific platforms). Sun's contribution to Xen, and the uniqueness of xVM, comes from the use of Solaris as the underlying hardware management platform running in a "Control Domain" or "Dom0", enabling use of the rich Solaris environment including self-healing and ZFS, whose scalability enables a large number of very large virtualized file and storage environments. External storage connection is provided by the Control Domain, which presents each virtual machine with abstracted volumes. Sun complements this offer with "xVM Ops Center", covering asset discovery and provisioning, version and patch management, and event monitoring.
- ▶ **Logical Domain, or LDom**, addresses a very different population. This technology leverages the "CoolThreads" multithreading capability in the latest Sun SPARC processors (UltraSPARC T1 and T2), and offers the capability to run separate Solaris "domains" on one or more threads ("strands") for a current maximum of 64 domains (8 strands per core on an eight core T2 server). Domains can either manage their own storage as "direct devices" or share in a pool of storage presented as "virtual devices". Users have the flexibility to run different Solaris 10 updates in each domain, making LDom a good choice for consolidation of small SPARC servers.

► **Solaris Containers** provide a Solaris virtualization technology spanning SPARC and x86 architectures. Unlike LDom or Xen/xVM, which provide virtualized hardware to multiple operating system copies, containers virtualize the operating system to multiple “zones” (which are not the same as “SAN fabric zones”). All Solaris zones (theoretically up to 8,191 of them) use the same Solaris 10 copy, where each update or patch benefits all users, reducing the complexities of version management on large numbers of servers. Each Solaris zone can either own assigned resources including CPU and storage, or share them through the Solaris Resource Manager. Besides consolidation benefits, containers also enable reboot of each Solaris zone in seconds, while the underlying hardware is isolated from most incidents triggering such reboots.

Solaris 10 and Sun’s virtualization technologies (Containers, LDom and xVM) all benefit greatly from the use of Emulex LightPulse Virtual HBA technology, based on NPIV and HBAnyware. Without Emulex LightPulse Virtual HBA technology, every guest would require its own dedicated Fibre Channel HBA. NPIV, implemented through Emulex LightPulse Virtual HBA technology, allows virtual HBAs to be assigned to each guest while HBAnyware manages these virtual HBAs, resulting in isolated workloads, increased asset utilization, and reduced management overhead.

Sun and Fibre Channel

Sun provides a broad product line to deliver Fibre Channel storage solutions to meet the needs of high-end enterprises whose amount of data stored throughout the world has increased significantly over the years. With its strong 10-year relationship with Sun Microsystems, which earned Emulex the 2007 Best-In-Class Supplier award, the Emulex LightPulse family of Fibre Channel HBAs is a key component to Sun Fibre Channel solutions which deliver higher levels of availability, performance, and reliability than IP-based SAN and other storage solutions.

Most operating systems still recognize Fibre Channel devices as parallel SCSI devices through driver emulation, an approach that will not always meet the needs of enterprise SANs with performance-sensitive and large-scale applications. Enterprise SANs are further challenged when presenting Fibre Channel devices and Logical Unit Numbers (LUN) as SCSI devices for administrators have to approach multiple vendors for operational support. Thus, Sun developed the Leadville Fibre Channel stack into Solaris 10 to supersede the traditional SCSI emulation method.

Leadville, the code name for the StorEdge SAN Foundation Software (SFS), is an open standards-based I/O framework and device driver stack to support Fibre Channel. With Solaris 10, Leadville is fully integrated into the OS and is available on SPARC, x64, and x86 platforms. In addition, because the Leadville stack is integrated into the OS, it is part of all the Solaris update releases, enabling continuous innovation and predictable quality.

Emulex and Fibre Channel

Emulex has always been at the forefront of Fibre Channel innovation. Whether it's virtualization, security, new LPe12002 8Gb/s Dual-Channel Fibre Channel HBAs, or simplifying SAN administration, the Emulex strategy is to not only be first, but to engage with other ecosystem members to ensure that solutions are deployable in today's enterprises.

The Emulex LightPulse family of Fibre Channel HBAs supports Virtual HBA technology which provides the scalability necessary to support an increased number of applications running on virtualized servers. Emulex LightPulse Virtual HBA technology allows Solaris 10 users to effectively virtualize the HBA functionality so each guest running on a server can have independent access to its own protected storage. Sun environments can also leverage SAN management tools and best practices, such as fabric zoning and LUN mapping/masking.

Emulex's technological depth in virtualization, SAN management and performance, as well as proven reliability track record, places its Fibre Channel Virtual HBA technology, with NPIV support and HBAnywhere centralized management, as a key component of all server and HBA virtualization deployments.

NPIV Evolution

In 2001, when server virtualization based on non-proprietary technology was still in its infancy, developers from Emulex and IBM began to design a method for virtualizing HBA ports. Their insight into data center requirements deploying virtualization technologies led to NPIV. Figure 1 shows the timeline of NPIV development and delivery from the initial design phase, to the first demonstration given at Storage Network World in 2005, and through the shipment of the first Solaris 10 HBA with Virtual HBA technology support by Emulex in July of 2008.

Emulex has proven technological depth with the design, development and implementation of the NPIV standard for the virtualized data center. By driving NPIV toward ANSI/INCITS T11 standardization, Emulex ensures that data centers worldwide would have reliable, scalable, and manageable virtual HBA technology. Emulex's partnership with virtualization market leaders, such as Sun Microsystems, results in long-term solutions that meet the needs of ever-growing data centers.

NPIV and Solaris

In a Sun virtualization environment, the administrator uses NPIV to assign a unique Worldwide Name (WWN) to each LDom, Solaris Container, or xVM guest so that they can be recognized as a specific endpoint in the fabric. Fabric zoning enables individual Fibre Channel server ports (initiators) and the corresponding storage ports (targets) to be invisible and inaccessible from other endpoints, thereby restricting interference, adding security, and simplifying management. Fabric zoning uses WWNs to allow access to storage such that a server can see and access only storage LUNs that share a common zone with that server.

In Solaris 10 SANs, NPIV allows the mapping of one virtual VPort to one LDom, Solaris Container or xVM guest in a virtual environment, giving them their own VPort (see Figure 2). Because the VPort and the related port and node worldwide names are not changed in form, storage administrators can use the existing fabric zoning and LUN masking techniques they have used in the past.

Figure 1—NPIV Development Timeline





Figure 2—Server and HBA Virtualization



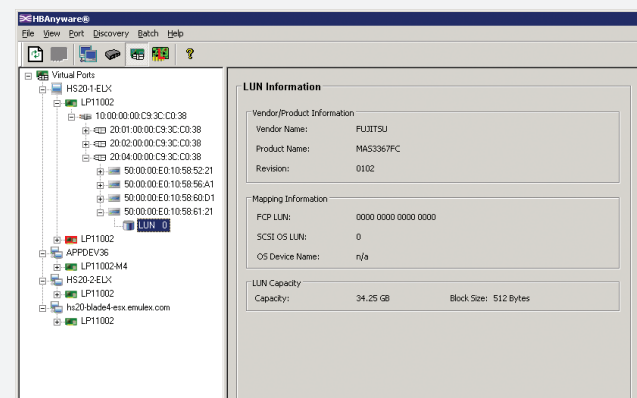
Benefits of Emulex LightPulse Virtual HBA Technology to Server Virtualization

Server virtualization relies on HBA virtual technology to maximize the data center benefits. Emulex LightPulse Virtual HBA technology delivers the following benefits:

- ▶ Maximum security with fabric zoning and LUN masking. Emulex LightPulse Virtual HBA technology allows each LDom, Solaris Container, or xVM guest to be managed separately by allowing each its own address on the Fibre Channel fabric (i.e., allows fabric zoning). Since access to a specific storage LUN is restricted to each OS using the virtual WWN for zoning, access is restricted in the same way as it is restricted to specific physical servers. This results in isolation of all operating systems running on a single physical server, ensuring the highest level of security for the entire system. Furthermore, the system's overall reliability is maximized as management or operations to one OS cannot affect others running on the same physical server.

- ▶ Greater management capabilities by storage administrators. Emulex LightPulse Virtual HBA technology gives each guest its own WWN, allowing the storage attached to it to be managed by storage administrators in the same way physical servers are managed. That is, Solaris administrators do not have to duplicate storage-centered management tasks.

Figure 3—HBAnyware Discovery and Management of Virtual Ports



- ▶ Service Level Agreement (SLA) and Quality of Service (QoS) support. Fabric management tools from vendors such as Cisco and Brocade provide the means to prioritize I/O traffic through the SAN Fibre Channel fabric based on the requester (initiator). When coupled with Emulex LightPulse Virtual HBA technology, these tools provide privileged access to fabric bandwidth, and reduced network latency to selected OS instances and applications. Additionally, Emulex LightPulse Virtual HBA technology enables fabric-based accounting, chargeback and troubleshooting tools to operate for each guest, thus ensuring that SLA and QoS requirements for each guest are met, and corresponding resources are measured and charged to the proper cost center.
- ▶ Flexible provisioning and upgrade. Since zoning and other services are no longer tied to the physical WWN hard-wired to the HBA, it is easier to replace an HBA. You do not have to reconfigure the SAN storage, because the new server can be pre-provisioned independently of the physical HBA WWN.
- ▶ Workload mobility. The virtual WWN associated with each guest follows it when the guest is migrated across physical servers. All settings, including zone membership, LUN masking, and fabric priority, are thereby automatically carried over. No SAN reconfiguration is necessary when the workload is relocated to a new server.

HBAnyware Centralized Management

Emulex developed HBAnyware knowing that data center administrators deploying SANs would need a tool to centralize the management of their HBAs. Independent studies have found that HBAnyware is more intuitive, more scalable and offers greater functionality than the competitive management solution, making HBAnyware the logical choice for large Sun data centers.

“...Our test concluded that Emulex’s management tool was an easier, more intuitive management application, lending itself to large SAN environments.”

—Dennis Martin

President of Demartek Testing laboratory

HBAnyware for Solaris 10 delivers all the functionality needed to manage HBAs across the server virtualization environment, regardless of the number or type of OS running on each physical server. HBAnyware delivers:

- ▶ State-of-the-art, powerful management capabilities, no matter where the HBAs are located, what platform they are running on, or how they can be accessed—whether in-band (over the Fibre Channel link, an exclusive feature of Emulex) or out-of-band (over the local area network).
- ▶ Automation of key management tasks for unparalleled HBA configuration and management efficiency. Driver profiles, which contain the driver’s parameter settings, can be replicated and efficiently applied to other HBAs in a single step.
- ▶ Batch mode propagation of driver profiles and firmware versions effortlessly across the virtual SAN.
- ▶ Scriptable command-line interface giving storage administrators the flexibility to tailor or automate their HBA management actions using customized scripts.

- ▶ Global driver parameter setting, either globally (for all Emulex LightPulse HBAs) or locally (for a given HBA), thereby making configuration easier to accomplish and also preventing errors by consolidating all common settings and only requiring changes (exceptions) to be itemized.
- ▶ Ability to easily capture configuration information about enterprise SAN components and export the data into an open, flexible XML format, which IT administrators can in-turn import into off-the-shelf data analysis tools such as Excel in order to create custom reports.
- ▶ Ability to fully configure “Boot-from-SAN” settings while the server is online, thereby pre-staging required changes ahead of the next maintenance window.
- ▶ Failover process whereby HBAnyware allows the HBA WWN to be changed online to match the ones from a failed server when a standby server is used to replace a failed server on the spot.

HBAnyware supports HBAs running on Solaris, Windows, Linux or VMware operating environments, as well as a graphical user interface running locally on the managed server, or remotely from a central console.

Enabling the creation and management of virtual ports, HBAnyware allows Solaris users to create the VPorts on a particular Solaris server and monitor traffic volume, error rates and trends, or even capture and log fabric incidents related to a specific virtual machine environment, enabling fine-grained performance optimization and preventive maintenance (Figure 3).

Summary

Emulex Virtual HBA technology is based on a strong foundation of virtualization expertise, proven with leading edge NPIV-based capabilities supported on multiple operating systems, now including Sun Solaris 10. Table 1 summarizes the many benefits to Solaris Containers, LDom, and xVM technologies from Emulex LightPulse Fibre Channel Virtual HBA technology, based on NPIV-enabled HBAs and HBAnyware.

Table 1—Server Virtualization Benefits from Emulex Virtual HBA Technology

Total cost of ownership (TCO)	Reduced with lower capital, management and operational costs
System overhead	Reduced through full FCP protocol processing at HBA level, flexible provisioning and workload mobility
System reliability	Enhanced with advanced troubleshooting support and security through fabric zoning and LUN masking
Vport management and troubleshooting	Accomplished with HBAnyware, portable connections, and guest-level zoning
SLA and QoS support	Enhanced with fabric prioritization for LDom, Solaris Containers and xVM guests, as well as HBAnyware-based performance and event logging

With Emulex LightPulse Virtual HBA technology, Sun Solaris users can take advantage of the latest HBA product features and tools from Emulex that reduce TCO, while delivering the system performance they have come to expect in the SAN. With a key role in the design of the NPIV standard and delivery of a robust HBA management tool, Emulex continues to be an industry leader in anticipating and addressing the performance needs of the next generation data center. Look to Emulex as a leader in creating and driving standards that deliver enterprise class functionality, performance, and reliability to data centers around the world.

This document refers to various companies and products by their trade names. In most, if not all cases, their respective companies claim these designations as trademarks or registered trademarks. This information is provided for reference only. Although this information is believed to be accurate and reliable at the time of publication, Emulex assumes no responsibility for errors or omissions. Emulex reserves the right to make changes or corrections without notice. This report is the property of Emulex and may not be duplicated without permission from the Company.

08-1095 - 7/08



www.emulex.com

World Headquarters 3333 Susan Street, Costa Mesa, CA 92626 +1 714 662 5600

Wokingham, UK +44 (0) 118 977 2929 | **Munich, Germany** +49 (0) 89 97007 177

Paris, France +33 (0) 158 580 022 | **Beijing, China** +86 10 68499547