



Global Distributed Storage Engineering

EMC Solaris SAN v2.5 Solutions Guide

GDSE-04012-SG

Solution Guide

1.0

Version Control

[illegible]

Legal Notice

Information contained herein is for Internal Use and may be used only for business purposes authorized by Citigroup.

CITI, CITIBANK, CITICORP, CITIGROUP and the Umbrella Device are trademarks and service marks of Citicorp and are used and registered throughout the world. Citibank, Citicorp, Citigroup and their subsidiaries also claim rights in certain other trademarks and service marks contained in these pages.

Contact

This document was created by a joint effort of both Storage Engineering and Unix Engineering. Please address any comments or questions regarding the information contained in this document via email to either the “*GT Global Distributed Storage Engineering” distribution list or alternatively to gdse@citigroup.com and/or “*GT Global Unix Engineering” distribution list or alternatively also gue@citigroup.com.

Latest Versions

It is recommended that the latest version of the document be used from the GDSE website <http://gdse.ny.smb.com>. Newer versions may include expanded scope, or comments on issues discovered after deployment.

Table Of Contents

VERSION CONTROL.....	2
LEGAL NOTICE.....	3
CONTACT	3
LATEST VERSIONS.....	3
TABLE OF CONTENTS.....	4
INTRODUCTION	5
SCOPE.....	6
APPLICABLE ENVIRONMENTS	7
INSTALLATION/CONFIGURATION	8
COMPONENTS OF THE SAN STACK	8
ORDER OF INSTALLATION	8
INSTALLING/CONFIGURING EMULEX HBA DRIVER	9
STEP 1: INSTALLING THE DRIVER	10
STEP 2: UPGRADING FIRMWARE.....	11
STEP 3: CHECKING FIRMWARE VERSION.....	14
STEP 4: EMULEX CONFIGURATION PARAMETERS.....	16
INSTALL POWERPATH	17
VERIFYING THE INSTALLATION	19
INSTALLING VERITAS	21
UPGRADING VxVM TO 4.0 WHEN ROOT DISK IS ENCAPSULATED.....	21
UPGRADING VxVM TO 4.0 WHEN ROOT DISK IS NOT ENCAPSULATED.....	22
CREATING A NON-ROOTDG DATA GROUP AND A VOLUME.....	23
CREATING AN VXFS FILESYSTEM FOR THE NEW VOLUME.....	23
BEST PRACTICES	24
CONSIDERATIONS.....	24
INTEROPERABILITY	24
PLANNING	24
SRDF.....	24
CLUSTERING.....	24
TAPE SAN SUPPORT	24
MONITORING.....	25
REFERENCES	26

Introduction

GDSE is releasing a new guide to enable system administrators and storage administrators to provision CTI SOE SAN Storage to Solaris Server based systems. This guide will cover the upgrade and installation of all components required for provisioning and management SAN storage on Solaris Servers.

Global Unix Engineering has recently released Veritas Volume Manger 4.0 in conjunction with a new patch set for Solaris 8 into production. Given that Veritas Volume Manager is an integral component for management of storage on Solaris, the HBA and the Multi-Pathing components warranted a refresh to maintain complete interoperability between the products. All new Solaris server builds provisioned for SAN storage must adhere to the requirements outlined in this document.

GDSE certifies the following tested configuration of HBA Firmware and HBA Failover Software to work with the new Solaris 8 Server build.

Scope

This scope of this document describes the installation of a GDSE approved stack of components that will allow a Solaris 8 Server to connect to EMC Symmetrix or EMC CLARiiON storage via a SAN connection. In addition, this document will cover best practices and considerations for connecting Solaris servers to SAN storage.

This document is based on testing of standard Intel based servers running a Solaris 8 build with the latest GDSE approved SOE versions of HBA firmware, HBA firmware and path failover software. The versions used will be listed in the certification matrix located in the Release Notes document.

The documentation will cover production server upgrades and new installations.

Topics:

Installation/Configuration

Best Practices

The following areas of focus will be evaluated in this document:

- HBA info (WWN, firmware, driver)
- Provisioning steps/dependencies
- Host I/O configuration and dependencies (I/O stack, VM, etc)
- Clustering implications
- Backup solutions (network based, SAN based)
- Flags / interoperability with other OSs
- Storage Tiers - platform support
- Data Migration Tools and limitations
- Operating System Upgrades

Note: Although SAN based backup is not part of this document, all existing functionality should not be impacted by this Solaris stack.

Applicable Environments

This guide is for servers running Solaris 8 Servers.

Please refer to the GDSE website for appropriate documentation and certifications/release notes.

Supported Configurations:

This Solutions Guide applies to any server certified by Unix Engineer for further details please refer to the GDSE Supported Solaris Configuration Matrix for specific configurations and component versions. The document is located on the GDSE website.

Note: The LP8000 series of the Emulex HBA are only supported as legacy HBAs only. Any environment using LP8000 HBAs is strongly recommended to upgrade to LP9002 HBAs.

INSTALLATION/CONFIGURATION

COMPONENTS OF THE SAN STACK

Component Matrix	
Server Models	Various
Operating Systems	Solaris 8
OS Patch Level	Patch Level 6
HBAs Supported	Emulex LP9002L-E, LP9002DC, LP8000
HBA Firmware	3.92a2
HBA Driver	5.02c
HBA Failover Software	EMC PowerPath 4.1
Volume Manager	Veritas VxVM 4.0
Volume File System	Veritas VxFS 4.0
NaviAgent (CLARiiON Only)	6.6

ORDER OF INSTALLATION

The following is a high level checklist to perform for successful connectivity to the SAN.

CLARiiON

Procedure	Responsibility	Page
Rack Server with HBAs installed and labeled	Infrastructure/Network Integration	NA
Jumpstart Server	SA	
Install and Configure HBA Driver/Firmware	SA	
Submit request to storage management	SA	
Physical connections made to server from SAN	Infrastructure	
Install and Configure Naviagent	SA	
Install and Configure PowerPath	SA	
Submit request for SAN storage	SA	
Zoning while server is up and running	Storage Management	
Register host in Navisphere and assign host to storage group	Storage Management	
Provide persistent binding LUN information to SA	Storage Management	
Install and configure Veritas	SA	

SYMMETRIX

Procedure	Responsibility	Page
Rack Server with HBAs installed and labeled	Infrastructure/Network Integration	NA
Jumpstart Server	SA	
Install and Configure HBA Driver/Firmware	SA	
Submit request to storage management	SA	
Physical connections made to server from SAN	Infrastructure	
Submit request for SAN storage	SA	
Zoning while server is up and running	Storage Management	
Update VCMDB information with server WWPN	Storage Management	
Install and Configure PowerPath	SA	
Install and configure Veritas	SA	

INSTALLING/CONFIGURING EMULEX HBA DRIVER

Emulex HBA Driver package contains both the firmware for both the LP9002 and LP8000 series HBAs. The package is available from GDS under PKG.

Standalone Installation Checklist
Backup any existing files: /kernel/drv/lpfc.conf, sd.conf
Remove old lpfc driver package if necessary;
Install lpfc5.02C
Set the following in /etc/system: set sd:sd_max_throttle = 20 set sd:sd_io_time = 0x3C Other values should be removed and duplicate entries should be removed
Reboot the server
Upgrade firmware and disable boot bios, if necessary (See Emulex release notes for further details and pay close attention to modifying instance numbers)
Verify correct firmware versions and HBA
Remove entries in sd.conf that do NOT map to internal drives or CD/DVD ROM drives. A failure to do so may allow DR (R2) devices to be visible when they should not be.
Add persistent binding information in sd.conf and lpfc.conf using files downloaded from the SAN website.
Reboot the server and verify LUNs are present.

Sun Cluster Installation Checklist
Backup any existing files: /kernel/drv/lpfc.conf, sd.conf
Reboot 1 st Node in non-cluster mode (boot -x)
Remove old/existing lpfc driver package if necessary
Install lpfc5.02C
Set the following in /etc/system
<pre>set sd:sd_max_throttle = 20 set sd:sd_io_time = 0x3C</pre>
Other values should be removed and duplicate entries should be removed
Add persistent binding information in sd.conf and lpfc.conf using files downloaded from the SAN website.
Reboot server once more in non-cluster mode
Upgrade firmware and disable boot bios, if necessary (See Emulex release notes for further details and pay close attention to modifying instance numbers)
Reboot the server in cluster mode and verify LUNs are present.
Follow all the steps above for the second node.

Emulex Driver/firmware: **lpfc5.02C_A3**

STEP 1: INSTALLING THE DRIVER

Obtain the package from the Citigroup GDS package repository. (For the US regions use stealth-nj or other replicated GDS servers). The example below is for a LP9002L HBA.

```
# pkgadd -d lpfc_5.02C
```

The following packages are available:

```
1 lpfc    Emulex LightPulse FC SCSI/IP Host Bus Adapter driver
         (sun4u) 5.02C
```

Select package(s) you wish to process (or 'all' to process all packages). (default: all) [?,??,q]:

Processing package instance <lpfc> from </mnt/lpfc_5.02C>

Emulex LightPulse FC SCSI/IP Host Bus Adapter driver
(sun4u) 5.02C

lpfc (FCP) / lpfn (IP) combo driver

Copyright (c) 2003, Emulex Corporation
3535 Harbor Boulevard, Costa Mesa, CA 92626

All rights reserved. This product and related documentation is protected by copyright and distributed under licenses restricting its use, copying, distribution and decompilation. No part of this product or related documentation may be reproduced in any form by

any means without prior written authorization of Emulex Corporation and its licensors, if any.

Rebuild manual pages database for section 7d [y,n,?] y

Processing package information.

Processing system information.

12 package pathnames are already properly installed.

Verifying package dependencies.

Verifying disk space requirements.

Checking for conflicts with packages already installed.

Checking for setuid/setgid programs.

This package contains scripts which will be executed with super-user permission during the process of installing this package.

Do you want to continue with the installation of <lpfc> [y,n,?] y

Installing Emulex LightPulse FC SCSI/IP Host Bus Adapter driver as <lpfc>

.....
.....

Installation of <lpfc> was successful.

The following packages are available:

1 lpfc Emulex LightPulse FC SCSI/IP Host Bus Adapter driver
(sun4u) 5.02C

Select package(s) you wish to process (or 'all' to process all packages). (default: all) [?,??,q]: q

*** IMPORTANT NOTICE ***

This machine must now be rebooted in order to ensure sane operation. Execute
shutdown -y -i6 -g0
and wait for the "Console Login:" prompt.

reboot

STEP 2: UPGRADING FIRMWARE

The example below describes a step-by-step process of the Emulex firmware upgrade process. The letters in **BOLD** are all user inputs. The firmware can only be loaded after the driver has been installed and server is rebooted.

Note: The firmware loading process should never be interrupted half way through. It should be allowed to complete successfully. If the process is interrupted before finishing, the card maybe damaged.

cd /usr/sbin/lpfc

./lputil

LightPulse Common Utility for Solaris/SPARC. Version 1.5a0 (12/18/2002).
Copyright (c) 2002, Emulex Network Systems, Inc.

Emulex Fibre Channel Host Adapters Detected: 2
Host Adapter 0 (lpfc0) is an LP9K (Ready Mode)

Host Adapter 1 (lpfc1) is an LP9K (Ready Mode)

MAIN MENU

1. List Adapters
2. Adapter Information
3. Firmware Maintenance
4. Reset Adapter
5. Persistent Bindings

0. Exit

Enter choice => **3**

0. lpfc0
1. lpfc1

Select an adapter (0-1) => **0**

FIRMWARE MAINTENANCE MENU

1. Load Firmware Image
2. Display AWC File Contents
3. Display Flash Load List
4. Display RAM Load List
5. Load PCI Configuration File
6. Boot BIOS Maintenance

0. Return to Main Menu

Enter choice => **1**

Enter Image Filename => **cd392a2.awc (USE DD392A2.AWC FOR LP8000)**

Opening File...

Reading Header...

Validating Checksum...

Initializing Host Adapter...

Erasing Flash ROM Sectors...

100% complete

Loading Image...

100% complete

Updating Wakeup Parameters...

Resetting Host Adapter...

Image Successfully Downloaded...

FIRMWARE MAINTENANCE MENU

1. Load Firmware Image
2. Display AWC File Contents
3. Display Flash Load List
4. Display RAM Load List
5. Load PCI Configuration File
6. Boot BIOS Maintenance

0. Return to Main Menu

Enter choice => **0**

MAIN MENU

1. List Adapters
2. Adapter Information
3. Firmware Maintenance
4. Reset Adapter
5. Persistent Bindings

0. Exit

Enter choice => **3**

0. lpfc0
1. lpfc1

Select an adapter (0-1) => **1**

FIRMWARE MAINTENANCE MENU

1. Load Firmware Image
2. Display AWC File Contents
3. Display Flash Load List
4. Display RAM Load List
5. Load PCI Configuration File
6. Boot BIOS Maintenance

0. Return to Main Menu

Enter choice => **1**

Enter Image Filename => **cd392a2.awc (USE DD392A2.AWC FOR LP8000)**

Opening File...

Reading Header...

Validating Checksum...

Initializing Host Adapter...

Erasing Flash ROM Sectors...

100% complete

Loading Image...

100% complete

Updating Wakeup Parameters...

Resetting Host Adapter...

Image Successfully Downloaded...

#

STEP 3: CHECKING FIRMWARE VERSION

Once the firmware has been upgraded, one may check it by following the process below.
Please pay details to the following info fields:

```
# /usr/sbin/lpfc/lputil
```

LightPulse Common Utility for Solaris/SPARC. Version 1.5a0 (12/18/2002).
Copyright (c) 2002, Emulex Network Systems, Inc.

Emulex Fibre Channel Host Adapters Detected: 2
Host Adapter 0 (lpfc0) is an LP9K (Ready Mode)
Host Adapter 1 (lpfc1) is an LP9K (Ready Mode)

MAIN MENU

1. List Adapters
2. Adapter Information
3. Firmware Maintenance
4. Reset Adapter
5. Persistent Bindings

0. Exit

Enter choice => **2**

ADAPTER INFORMATION MENU

1. PCI Configuration Parameters
2. Adapter Revision Levels
3. Wakeup Parameters
4. IEEE Address
5. Loop Map
6. Status & Counters
7. Link Status
8. Configuration Parameters

0. Return to Main Menu

Enter choice => **2**

0. lpfc0
1. lpfc1

Select an adapter (0-1) => **0**

```
                BIU: 2002606D
Sequence Manager: 00000000
                Endec: 00000000
Operational Firmware: SLI-2 Overlay
                Kernel: 1.30a1
Initial Firmware: Initial Load 3.92a2 (CS3.92A2)
                SLI-1: SLI-1 Overlay 3.92a2 (C1D3.92A2)
                SLI-2: SLI-2 Overlay 3.92a2 (C2D3.92A2)
```

Highest FC-PH Version: 4.3

Lowest FC-PH Version: 4.3

Press any key to continue:

ADAPTER INFORMATION MENU

1. PCI Configuration Parameters
2. Adapter Revision Levels
3. Wakeup Parameters
4. IEEE Address
5. Loop Map
6. Status & Counters
7. Link Status
8. Configuration Parameters

0. Return to Main Menu

Enter choice => **2**

0. lpfc0

1. lpfc1

Select an adapter (0-1) => **1**

BIU: 2002606D

Sequence Manager: 00000000

Endec: 00000000

Operational Firmware: SLI-2 Overlay

Kernel: 1.30a1

Initial Firmware: Initial Load **3.92a2** (CS3.92A2)

SLI-1: SLI-1 Overlay 3.92a2 (C1D3.92A2)

SLI-2: SLI-2 Overlay 3.92a2 (C2D3.92A2)

Highest FC-PH Version: 4.3

Lowest FC-PH Version: 4.3

Press any key to continue:

ADAPTER INFORMATION MENU

1. PCI Configuration Parameters
2. Adapter Revision Levels
3. Wakeup Parameters
4. IEEE Address
5. Loop Map
6. Status & Counters
7. Link Status
8. Configuration Parameters

0. Return to Main Menu

Enter choice => **0**

MAIN MENU

1. List Adapters
2. Adapter Information

3. Firmware Maintenance
4. Reset Adapter
5. Persistent Bindings

0. Exit

Enter choice => 0

STEP 4: EMULEX CONFIGURATION PARAMETERS

The table below illustrates all the parameters settings for the lpfc.conf that EMC recommends for a Fibre Channel Switch implementation. Since both PowerPath and DMP are both enabled in our environment, all variables in the Blue column have to set. The Emulex driver package for Citigroup already contains all the info listed in the table below.

Variable	Default Value	If PowerPath and/or DMP	Comments
log-verbose	0	0x10	(bit mask) Extra activity logging.
Automap	2	0	Only devices with persistent binding will be recognized.
lun-queue-depth	30	20	Default max commands sent to a single logical unit (disk).
tgt-queue-depth	0	512(Symmetrix/Clariion)	Default max commands sent to a single target.
no-device-delay	1	0	Delay to fail back an I/O in seconds.
scan-down	2	0	Select method for ALPA to SCSI ID mapping.
Linkdown-tmo	30	60	(seconds) How long driver waits before deciding FC link is down.
nodev-tmo	0	60	Seconds to hold I/O if device disappears.
Num-iocbs	1024	2048	Adapter command/response structures.
Topology	4	2 (4 for arbitrated loop)	Fibre Channel link topology (defaults to loop only).

Please make sure the following entries are set in the /etc/system files.

```
* EMC Emulex Settings
set sd:sd_max_throttle = 20
set sd:sd_io_time = 0x3C
```

Step 5: Final Reboot

```
# reboot
```

Note: Please install the *HBainfo* package also located on stealth-nj.

INSTALL POWERPATH

Before upgrading to PowerPath 4.1, you must uninstall the earlier version of PowerPath. However, the uninstall process of PowerPath 3.0.2 leaves 2 configuration files on the system. The files are: **/etc/powermt.custom.saved** and **/kernel/drv/emcp.conf.saved**.

Those files contain previous configuration information that will be used in the upgrade process.

(Note: an upgrade question will appear during the install process – see example below)

In an environment where Veritas is also being used, please make sure that PowerPath is installed after Veritas. This is done to insure that the powerpath entries in /etc/system are after the VxVM entries.

One other very important step is to **backup existing /etc/system** to ensure that existing shared memory settings are set accordingly after the installation.

Obtain the package from the Citigroup GDS package repository. *(For the US regions use **stealth-nj** or other replicated GDS servers)*

```
# cd /net/stealth-nj/export1/home1.localhost/sw/PKG
# pkgadd -d EMCpower4.1.0_b155
```

The following packages are available:

```
1 EMCpower    EMC PowerPathpkg
              (all) 4.1.0_b155
```

Select package(s) you wish to process (or 'all' to process all packages). (default: all) [?,??,q]:

Processing package instance <EMCpower> from </var/tmp>

```
EMC PowerPath
(all) 4.1.0_b155
EMC
```

Enter package base directory (default: **/export/opt/emc**):

Enter directory for logging (default: **/export/opt/emc/EMCpower/log**):

Configuration for PowerPath 3.0.x exists - want to upgrade? [y,n,q,?] (default: y): **y**

Using </export/opt/emc> as the package base directory.

Processing package information.

Processing system information.

7 package pathnames are already properly installed.

Verifying disk space requirements.

Checking for conflicts with packages already installed.

Checking for setuid/setgid programs.

This package contains scripts which will be executed with super-user permission during the process of installing this package.

Do you want to continue with the installation of <EMCpower> [y,n,?] **y**

Installing EMC PowerPath as <EMCpower>

Executing preinstall script.

Installing part 1 of 1.

/etc/S87powervxvm

.....

.....

/usr/man/man1/powermt.1

[verifying class <none>]

Executing postinstall script.

Adding emcp driver...

Adding forceload to /etc/system

setting emcp:bPxEnableInit to 1

setting lwp_default_stksize to 0x6000

setting rpcmod:svc_default_stksize to 0x6000

setting semsys:seminfo_semmns to 0x13d

setting semsys:seminfo_semmni to 0xc

setting semsys:seminfo_semmsl to 0x81

setting semsys:seminfo_semmmap to 0x13d

setting semsys:seminfo_semmnu to 0x13d

Upgrading 30x configuration...

emcpupgrade: configuration files converted to PowerPath 4.1.0 format.

* EMCpower Installation:

*

* Configuration files successfully converted.

* EMCpower Installation:

*

* Installation is successful.

*

* Please REBOOT the host to complete the remaining steps of the installation.

* Refer to the PowerPath for UNIX Installation and Administration Guide for

* more information.

All trademarks used herein are the property of their respective owners.

Installation of <EMCpower> was successful.

The following packages are available:

1 EMCpower EMC PowerPath

(all) 4.1.0_b155

Select package(s) you wish to process (or 'all' to process

all packages). (default: all) [?,??,q]: **q**

reboot

Note: The installation process should always complete successfully. If the installation fails due to the inability to convert older *powermt.custom* files, please follow the directions below:

If files are corrupted

If the PowerPath 3.0.x or 2.1 configuration files are corrupted, your custom configuration is no longer available after you install PowerPath 4.1.

- 1) Remove the /etc/powermt.custom.saved
- 2) Remove the EMCpower package
- 3) Re-install PowerPath 4.1

If any of your applications are configured with PowerPath pseudo devices, you must reconfigure those applications to use native device names. The emcpower pseudo device names were only supported by GDSE under powerpath 1.5.6. All existing configuration should have already been converted to native device names by following the Solaris SAN migration document before upgrading PowerPath.

VERIFYING THE INSTALLATION

Once PowerPath has been installed, one may check it by following the process below:

Checking the package

```
# pkginfo -l EMCpower
PKGINST: EMCpower
NAME: EMC PowerPath
CATEGORY: system
ARCH: all
VERSION: 4.1.0_b155
BASEDIR: /export/opt/emc

VENDOR: EMC

PSTAMP: beavis951018123443
INSTDATE: May 07 2004 14:39
STATUS: completely installed
FILES: 312 installed pathnames
      5 shared pathnames
      17 directories
      151 executables
      116779 blocks use (approx)
```

Checking the /etc/system entries

(Make sure that original application settings for shared memory are included after the upgrade)

```
* EMC Emulex Settings
set sd:sd_max_throttle = 20
set sd:sd_io_time = 0x3C
```

**EMCpower added the following line:*

```
forceload: drv/sd
```

**EMCpower added the following line:*

```
forceload: drv/ssd
```

```
forceload: drv/emcp
```

```
forceload: misc/emcpmp
```

```
forceload: misc/emcpmpc
```

```
forceload: misc/emcpmpaa
```

```
forceload: misc/emcpsapi
```

```
forceload: misc/emcpcg
```

```
forceload: misc/emcphr
```

forceload: misc/emcpsm
forceload: misc/emcpsc
forceload: misc/emcpevm
forceload: misc/emcpdpm
forceload: misc/emcploc
**EMCpower added the following line:*
set emcp:bPxEnableInit=1
**EMCpower added the following line:*
set lwp_default_stksize=0x6000
**EMCpower added the following line:*
set rpcmod:svc_default_stksize=0x6000
**EMCpower added the following line:*
set semsys:seminfo_semmns=0x13d
**EMCpower added the following line:*
set semsys:seminfo_semmni=0xc
**EMCpower added the following line:*
set semsys:seminfo_semmnl=0x81
**EMCpower added the following line:*
set semsys:seminfo_semmmap=0x13d
**EMCpower added the following line:*
set semsys:seminfo_semmnu=0x13d

Checking the registration

Key B9PF-UB4L-VF4T-Q4VX-MZ92-QFYY
Product: PowerPath
Capabilities: All

INSTALLING VERITAS

Veritas is tested and certified by Unix Engineering working closely with Storage Engineering. The complete certification document for Veritas can be found on the Global Unix Engineering website. The high level information listed below is an excerpt from the [Veritas 4.0 Storage Foundation](#) from Global Unix Engineering and provide a guidelines for server configuration. The complete installation process and all relevant materials can be found on the Global Unix Engineering website.

UPGRADING VxVM TO 4.0 WHEN ROOT DISK IS ENCAPSULATED

1. Read the following documents carefully: “Veritas Volume Manager 4.0 Installation/Upgrade guide” - (http://unixeng.smb.com/current/vxvm_ig.pdf) and the “Veritas 4.0 Storage Foundation” - http://unixeng.corp.smb.com/current/RN1.0-VRTS_SF_4.0.pdf

Note: You can not upgrade from versions less than 3.1.1 directly to VxVM 4.0, you must upgrade to VxVM 3.2, then upgrade to VxVM 4.0

2. Obtain new VxVM license key. Save the output of “vxprint -ht
3. Make sure that at least one plex of the volumes on which the OS file systems /, /usr, /var, /opt or /export are defined is created from a single subdisk that begins on a cylinder boundary.
4. Make sure that there is no active I/O on the system. Un-mount all file systems controlled by VxVM excluding all OS file systems.
5. **# cd /net/stealth.nj/export1/home1.localhost/sw/PKG/Veritas_sf_4.0/scripts**
6. **# upgrade_start** It will un-encapsulate and un-mirror the root disk. The script also creates the /VXVM4.0-UPGRADE directory for restoring configuration later
7. Check /etc/vfstab to be sure that all VxVM entries have been removed or commented out by the upgrade_start script.
8. Reboot the system.
9. Remove the following old VxVM packages (do not remove VxFS packages if you do not plan to upgrade it): VRTSvxva (or VRTSvmsa) VRTSvmdoc VRTSvmman VRTSvxvm VRTSlic, if present.
10. Check /etc/system to be sure that all VxVM entries have been removed.
11. If using EMC Power Path 3.0.x or less you must upgrade to PowerPath 4.1. You can find instruction on upgrading PowerPath in this document ([See Above](#))
12. Make sure that your system is on Solaris 8, Patchset 6.
13. Reboot the system.
14. Before installing the new packages, make sure that /opt/VRTS, /opt/VRTSvxvm, /opt/VRTSvxms, /opt/VRTSvlic are auto-mounted or physically exist under /export/opt.
15. **# pkgadd -d /net/stealth.nj/export1/home1.localhost/sw/PKG/Veritas_sf_4.0/pkgs VRTSvlic.**
16. **# pkgadd -d /net/stealth.nj/export1/home1.localhost/sw/PKG/Veritas_sf_4.0/pkgs VRTSvxvm**
17. **# patchadd -d /net/stealth.nj/export1/home1.localhost/sw/PKG/Veritas_sf_4.0/patches/115217-01.**
18. Install new license key: **/opt/VRTS/bin/vxlicinst**
Enter new license key
19. **cd /net/stealth.nj/export1/home1.localhost/sw/PKG/Veritas_sf_4.0/scripts**
20. **# upgrade_finish**
21. Reboot the system. Verify that root disk is encapsulated and mirrored again.
22. Continue to install packages **VRTSvmdoc, VRTSvmman and VRTSvmdev.**
23. Reboot the system.
24. Comparing to the saved output of “vxprint -ht”, import any data volumes that aren’t already imported using vxvol import <diskgroup-name>
25. Start any data volumes that aren’t already started using vxvol -g <diskgroup-name> start <volume-name>

UPGRADING VxVM TO 4.0 WHEN ROOT DISK IS NOT ENCAPSULATED

1. Read the following documents carefully: “Veritas Volume Manager 4.0 Installation/Upgrade guide” - (http://unixeng.smb.com/current/vxvm_ig.pdf) and the “Veritas 4.0 Storage Foundation” - http://unixeng.corp.smb.com/current/RN1.0-VRTS_SF_4.0.pdf

Note: You can not upgrade from versions less than 3.1.1 directly to VxVM 4.0, you must upgrade to VxVM 3.2, then upgrade to VxVM 4.0

2. Obtain new VxVM license key. Save the output of “vxprint -ht”.
3. Make sure that at least one plex of the volumes on which the OS file systems /, /usr, /var, /opt or /export are defined is created from a single subdisk that begins on a cylinder boundary.
4. Make sure that there is no active I/O on the system. Un-mount all file systems controlled by VxVM excluding all OS file systems.
5. **# cd /net/stealth.nj/export1/home1.localhost/sw/PKG/Veritas_sf_4.0/scripts**
6. **# upgrade_start** - The script creates the /VXVM4.0-UPGRADE directory for restoring configuration later.
7. Check /etc/vfstab to be sure that all VxVM entries have been removed or commented out by the **upgrade_start** script.
8. Reboot the system.
9. Remove the following old VxVM packages (do not remove VxFS packages if you do not plan to upgrade it): **VRTSvxva** (or **VRTSvmsa**) **VRTSvmdoc**, **VRTSvmman**, **VRTSvxvm** (remove **VRTSlic** if upgrading both **VxVM** and **VxFS**), if present.
10. Check /etc/system to be sure that all VxVM entries have been removed.
11. If using EMC Power Path 3.0.x or less you must upgrade to PowerPath 4.1. You can find instruction on upgrading PowerPath in this document http://gdse.ny.smb.com/GDSE_docs/GDSE-04002-TM.pdf
12. Make sure that your system is on Solaris 8, Patchset 6.
13. Before installing the new packages, make sure that /opt/VRTS, /opt/VRTSvxvm, /opt/VRTSvxms, /opt/VRTSvlic are auto-mounted or physically exist under /opt.
14. **# pkgadd -d /net/stealth.nj/export1/home1.localhost/sw/PKG/Veritas_sf_4.0/pkgs VRTSvlic.**
15. **# pkgadd -d /net/stealth.nj/export1/home1.localhost/sw/PKG/Veritas_sf_4.0/pkgs VRTSvxvm**
16. **# patchadd -d /net/stealth.nj/export1/home1.localhost/sw/PKG/Veritas_sf_4.0/patches/115217-01**
17. Install new license key.
18. **# /opt/VRTS/bin/vxlicinst**
Enter new license key
19. **# cd /net/stealth.nj/export1/home1.localhost/sw/PKG/Veritas_sf_4.0/scripts**
20. **# upgrade_finish**
21. Reboot the system.
22. Continue to install packages **VRTSvmdoc** and **VRTSvmman**
23. Reboot the system.
24. Comparing to the saved output of “vxprint -ht”, import any data volumes that aren’t already imported using **vxvol import <diskgroup-name>**
25. Start any data volumes that aren’t already started using **vxvol -g <diskgroup-name> start <volume-name>**

CREATING A NON-ROOTDG DATA GROUP AND A VOLUME

The steps below simply represents the steps that will aid in creating a non-rootdg data group. SAs should use whatever method they prefer (i.e. vxdiskadm or the GUI).

```
# vxdisk list
```

DEVICE	TYPE	DISK	GROUP	STATUS
c1t0d0s2	auto:sliced	disk01	rootdg	online
c1t1d0s2	auto:sliced	-	-	online
c2t16d0s2	auto	-	-	error
c2t16d7s2	auto:cdsdisk	-	-	online
c2t16d8s2	auto:cdsdisk	-	-	online

```
# /etc/vx/bin/vxdisksetup -i c2t16d7
```

```
# /etc/vx/bin/vxdisksetup -i c2t16d8
```

```
# vxdg init datadg c2t16d7=c2t16d7s2
```

```
# vxdg -g datadg adddisk c2t16d8=c2t16d8s2
```

```
# vxdisk list
```

DEVICE	TYPE	DISK	GROUP	STATUS
c1t0d0s2	auto:sliced	disk01	rootdg	online
c1t1d0s2	auto:sliced	-	-	online
c2t16d0s2	auto	-	-	error
c2t16d7s2	auto:cdsdisk	c2t16d7	datadg	online
c2t16d8s2	auto:cdsdisk	c2t16d8	datadg	online

```
# vxdg -g datadg free
```

DISK	DEVICE	TAG	OFFSET	LENGTH	FLAGS
c2t16d7	c2t16d7s2	c2t16d7s2	0	104853546	-
c2t16d8	c2t16d8s2	c2t16d8s2	0	104853546	-

```
# vxassist -g datadg make vol01 104853546 c2t16d7
```

```
# vxassist -g datadg make vol02 104853546 c2t16d8
```

CREATING AN VXFS FILESYSTEM FOR THE NEW VOLUME

```
# mkfs -F vxfs /dev/vx/dsk/datadg/vol01
```

```
# mkdir -p /export/vol01
```

```
# mount /dev/vx/dsk/datadg/vol01 /export/vol01/
```

BEST PRACTICES

CONSIDERATIONS

GDSE recommends a full backup before any major upgrade of a current environment.

Adding additional luns may require a server reboot to recognize new available disks. Once a signature is written they should be available for use.

Always use native device names under Solaris. EMCpower device names are not supported in this stack. Please check the upgrade document for further details: [GDSE-02017-TM](#) "Solaris SAN Upgrade v2.4"

INTEROPERABILITY

The solution compiled above has been verified by EMC and is therefore fully supported by their support staff. Any deviations from the stack may cause some unexpected result in the environment.

No known issues were discovered during testing.

PLANNING

Please consult your local storage integration team for available storage sizes and timeframes for production rollout.

SRDF

Switching to an R2 disk after an RDF split may require a server reboot to recognize the disks. The same procedure to configure a production host should be applied to an R2 host build. It is presently not recommended for a DR host to be configured for both RDF and BCV devices at the same time.

CLUSTERING

No clustering solution has been tested using VxVM 4.0 and therefore not supported at the present.

TAPE SAN SUPPORT

Although this stack was not tested with Tape San functionality, all existing functionality should not be affected. Please refer to: http://gdse.ny.smb.com/GDSE_docs/GDSE-03017-B.pdf for further details.

MONITORING

The EMC PowerPath Monitor will show a visible alert when a path failure occurs. SANmon presently handles this monitoring process. SANmon is a tool written to periodically check HBA and path status on the Solaris environment.

REFERENCES

Document	Location
Emulex 5.02c-3.92a2 Release Notes	http://gdse.ny.smb.com/GDSE_docs/GDSE-03052-RN.pdf
Veritas 4.0 Certification for Solaris	http://unixeng.corp.smb.com/current/RN1.0-VRTS_SF_4.0.pdf
PowerPath 4.1 Release notes	http://gdse.ny.smb.com/GDSE_docs/GDSE-04002-TM.pdf
GDSE Website	http://gdse.ny.smb.com/
GUE Website	http://unixeng.smb.com/