

# **ASM and Multipathing Generic Best Practices and Information Matrix**

March 2008

## INTRODUCTION

In today's competitive business climate, well-planned business continuity strategies must include redundancy to ensure server resources are able to access storage devices even in the event of a failure. Disks and controllers can fail, and data paths can become congested resulting in bottlenecks or can fail entirely cutting off access to key information.

Storage RAID technologies are great solutions for providing high availability for storage devices (disks), however, they do not address the physical path between the server and storage. If there is only a single path from the server to storage and one component fails, no amount of disk redundancy can keep the data available.

Multipathing solutions are designed to provide failover through the use of redundant physical path components—adapters, cables, and switches—between the server and storage. In the event that one or more of these components fail, applications can still access their data, thereby, eliminating single point of failure within the SAN and maintaining continuous access to storage devices and data. Multipathing solutions can also improve system performance by distributing I/O load across all available paths. It provides a higher level of data availability through automatic failover and failback. Active-active (all channels are actively doing I/O) and active-passive (one channel is active I/O while the others are standby) functionalities are optionally available from vendors.

Multipathing software is implemented at the Operation System (OS) device driver level and creates an abstraction (pseudo device) allowing I/O operations to be shared and balanced across all available I/O paths.

Multipathing provides the following key benefits:

- I/O path high availability
- Pseudo device interface for a multi-pathed LUN
- Path failover and failback
- I/O load balancing in active-active configurations

Automatic Storage Management (ASM) is a complimentary solution to multipathing technologies. ASM removes management complexity, improves performance through data distribution and improves data availability but does not have multipathing functionality itself.

ASM should not have any dependencies on host multipathing software and should work transparently in most environments. Oracle suggests following best practices published by Oracle and partners to ensure success.

### Generic Best Practices for ASM and Multipathing Configurations

ASM is designed to produce an error when discovering the same disk via multiple paths. If ASM picked the 1<sup>st</sup> device if found, it might not choose the pseudo device. Returning an error forces the user to be explicit about the path Asm should open. A single disk may have multiple device names in a multipath configuration. Therefore, ASM should be configured so that only one of the pseudo devices is discovered. Here is an example for a disk with 3 paths:

1. The first path to the disk (/dev/rdisk/c3t1d5s4)
2. The second path to the disk (/dev/c4t1d5s4)
3. The pseudo device name(/dev/rdisk/emcpower1)

The ASM initialization parameter, ASM\_DISKSTRING, limits the paths searched by ASM when discovering disks. It should be configured to find only the psudo device names. For example, if you are using EMC PowerPath multipathing software, you may set the ASM\_DISKSTRING to

'/dev/rdisk/emcpower\*'. When I/O is sent to the pseudo device, the multipath driver intercepts it and provides load balancing and failover/failback to the underlying sub-paths.

### **Multipathing When Using ASMLIB Disks**

When using [ASMLIB](#) with ASM on Linux, you can ensure the discovery of the multipath disk by configuring ASMLIB to scan the multipath disk first or to exclude the single path disks when scanning.

ASMLib allows two modifications to the disk scan order. First, it allows exclusion of certain disks. In other words, ASMLib will ignore those disks completely. Second, the system administrator can specify disks that are to be scanned first. Disks in this list are scanned before the rest of the disks in the system.

You can specify the desired disk scan behavior by specifying ASM configuration variables. The ORACLEASM\_SCANORDER variable specifies disks to be scanned first. The ORACLEASM\_SCANEXCLUDE variable specifies the disks that are to be ignored.

Please visit the Oracle Technology Network (OTN) for more information about configuring Oracle ASMLIB with multipath disks: <http://www.oracle.com/technology/tech/linux/asmlib/multipath.html>

## ASM and Multipathing Best Practices and Information Matrix

Partner	Multi-Pathing Software	Storage Support	OS Support	Configuration Best Practices and Notes
IBM	SDDPCM-MPIO	DS8000 DS6000 ESS (Shark)	AIX version 5.2B, 5.3 and above.	<p>Once SDDPCM is installed and storage devices are configured, a logical device name - /dev/rhdisk is created for per storage lun. These are the devices that are used to build Automatic Storage Management diskgroups.</p> <ul style="list-style-type: none"> <li>Use the /dev/rhdiskn device name when configuring ASM disk groups.</li> <li>The “no reserve” policy needs to be set for each rhdisk device. With DS8000/6000 and ESS storage this setting is automatically handled by the PCM.</li> <li>Change the device characteristics to accommodate reserve policy: chdev -l hdisk1 -a reserve_policy=no_reserve</li> </ul> <p>- Note: Explain the SDD problem with AIX</p> <p><a href="http://www-1.ibm.com/support/docview.wss?uid=ssg1S7000303&amp;aid=1">http://www-1.ibm.com/support/docview.wss?uid=ssg1S7000303&amp;aid=1</a> For information on MPIO see the following: <a href="http://publib16.boulder.ibm.com/doc link/en_US/a doc lib/aixbman/baseadmin/manage_mpio.htm">http://publib16.boulder.ibm.com/doc link/en_US/a doc lib/aixbman/baseadmin/manage_mpio.htm</a></p>
	SDD	DS8000 DS6000 ESS (Shark) SVC	AIX and Linux	<p>- Once configured, SDD will generate /dev/vpath devices. Create a partition on the vpath device and associate the raw device (/dev/raw/rawx) with the vpath partition.</p> <p>- use Linux SSD driver 1.6.0.1-4 or higher.</p> <p>- SDD/AIX does not currently work with ASM.</p> <p>- SDD and the ASMLIB interface option for Linux is currently unsupported.</p>
	RDAC	DS4000	AIX , Windows, Linux	<p>- Use as device: /dev/rhdiskx</p> <p>- Change the device characteristics to accommodate reserve policy: chdev -l hdisk1 -a reserve_policy=no_reserve</p>
Linux	Device Mapper	All	Linux (RedHat and SLES)	<p>- Configure device mapper and use /dev/mapper as the ASM_DISKSTRING.</p> <p>- Note: Explain the device permission change problem upon reboot</p> <p>Refer to the following 2 WPs for details on RH4 and SLES: <a href="http://www.oracle.com/technology/products/database/asm/pdf/device-mapper-udev-asm%20sles9.pdf">http://www.oracle.com/technology/products/database/asm/pdf/device-mapper-udev-asm%20sles9.pdf</a> <a href="http://www.oracle.com/technology/products/database/asm/pdf/device-mapper-udev-crs-asm%20rh4.pdf">http://www.oracle.com/technology/products/database/asm/pdf/device-mapper-udev-crs-asm%20rh4.pdf</a></p>
HP	SecurePath	EVA3000, 4000, 5000, 6000, 8000 and the Xp series	HP-UX	<p>HP Secure Path bundle two drivers, one for Active/Passive model based arrays such as the EVA3000 and EVA5000 and one for Active/Active arrays like the EVA3000/4000/5000/6000/8000 and XP family. The EVA3000 and EVA5000 come in two flavors, Active/Active and Active/Passive. This is purely firmware based.</p> <p>- Active/Passive: SecurePath generates a virtual path "/dev/rdisk/c58t0d0" and associates it with all physical paths, The virtual path is that ASM disk.</p> <p>- Active/Active: Any of the paths can be used as the virtual path.</p> <p>Please read the technical white paper for details: <a href="http://www.oracle.com/technology/products/database/asm/pdf/HP-UX%20-ASM-StgWorks-MP%2002-06.pdf">http://www.oracle.com/technology/products/database/asm/pdf/HP-UX%20-ASM-StgWorks-MP%2002-06.pdf</a></p>
Sun	Traffic			Although Sun Traffic Manager is interoperable with ASM, we currently don't

	Manager (MPXIO)			have specific information or WP describing the best practices.
EMC	PowerPath	EMC Storage	All	<p>Linux example:</p> <ul style="list-style-type: none"> <li>- Create the raw device using PowerPath pseudo name. <ul style="list-style-type: none"> <li>- raw /dev/raw/raw10 /dev/emcpowerg1</li> </ul> </li> <li>- To maintain these raw device bindings upon reboot, they must be entered into the /etc/sysconfig/rawdevices file as follows: <ul style="list-style-type: none"> <li>- /dev/raw/raw10 /dev/emcpowerg1</li> </ul> </li> <li>- To guarantee that the raw device binding occurs during any restart, use the chkconfig utility. <ul style="list-style-type: none"> <li>- /sbin/chkconfig rawdevices</li> </ul> </li> </ul> <p>On AIX: ASM can use concurrent logical raw logical volumes or raw partitions. Must use the pseudo device: /dev/rhdiskpowerx</p> <p>HP-UX: ASM can use the raw partitions. Must use native device path: /dev/rdisk/cxydz</p> <p>Please refer to the best practice guide for details:  <a href="http://www.oracle.com/technology/products/database/asm/pdf/asm-on-emc-5_3.pdf">http://www.oracle.com/technology/products/database/asm/pdf/asm-on-emc-5_3.pdf</a></p>
Hitachi Data System	HDLM	HDS Storage	All	<p>Hitachi Dynamic Link Manager (HDLM) software manages access paths to storage systems. HDLM provides functionality for distributing the load across paths and switching to another path if there is a failure in a path being used, thus improving system availability and reliability.</p> <p>On Linux, each path to a LUN is assigned a device name such as /dev/sda, /dev/sdb, etc. HDLM finds paths to the same actual LUN and assigns a common HDLM name such as /dev/sddlmaa. The HDLM device format on Linux is /dev/sddlm[a-p][a-p][1-15]. Partitions on an HDLM device are represented by the number after the sddlm* name. The whole device is represented without specifying the number. You must use the HDLM device name for HDLM load-balancing and path-failover functionality to be active.</p> <p>For more information, refer to the appropriate HDLM Users Guide for the specific OS platform.</p>
Symantec	DMP	All	All supported platforms	<p>DMP is bundled with the Veritas storage Foundation logical volume manager. ASM disks must be configured as logical volumes to utilize DMP multipathing. ASM must the dmp device which resides over the vxvm logical volumes: /dev/vx/rdmp/cxydzsx</p> <p>Please review the Symantec and Oracle technical white paper (soon to be published).</p>
MicroSoft	MPIO		Windows	<p>The release of Microsoft's multipathing drivers, Microsoft MPIO, allows software and hardware vendors to develop solutions that are not only specific to their storage devices, but also work effectively with the Microsoft Windows 2000/2003 Server platforms—including Microsoft Windows Storage Server 2003.</p>

	Not MPIO (refer to notes section)		Windows	<p>HP, Sun, EMC, HDS and Veritas have multipathing software supporting their specific products on Windows (Secure Path, Traffic Manager, HDML, Powerpath, and VxVM/DMP).</p> <p>The HBA vendor Qlogic has developed a driver and a management application providing multipathing failover and failback when Qlogic HBAs are used. Qlogic's multipathing driver is minimal in nature, but quite adequate for most ASM consumers.</p>
Pillar Data Systems	Axiom Path Manager (APM)	Pillar Storage	All	<a href="http://www.oracle.com/technology/products/database/asm/pdf/Pillar-ASM_TWP%2003-06_0.pdf">http://www.oracle.com/technology/products/database/asm/pdf/Pillar-ASM_TWP%2003-06_0.pdf</a>
Fujitsu	ETERNUS Multipath Driver	Fujitsu Storage		<p>The interfaces between servers and ETERNUS storage systems also deliver a high available mechanism. This mechanism is provided by ETERNUS Multipath Driver or GR Multipath Driver that manages multiple Fibre Channel paths.</p> <p><a href="http://www.oracle.com/technology/products/database/asm/pdf/asm%20best%20practices0907-fujitsu.pdf">http://www.oracle.com/technology/products/database/asm/pdf/asm%20best%20practices0907-fujitsu.pdf</a></p>
NEC	NEC Path Manager	NEC Storage		<a href="http://www.oracle.com/technology/products/database/asm/pdf/NEC_ASM_BestPractice-Aug-06.pdf">http://www.oracle.com/technology/products/database/asm/pdf/NEC_ASM_BestPractice-Aug-06.pdf</a>



ASM and Multipath Generic Best Practices and Information Matrix  
March 2008  
Author: Ara Shakian  
Contributing Author(s): Nitin Vengurlakar

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

Worldwide Inquiries:  
Phone: +1.650.506.7000  
Fax: +1.650.506.7200  
[www.oracle.com](http://www.oracle.com)

Copyright © 2003, Oracle. 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 is a registered trademark of Oracle Corporation and/or its  
affiliates. Other names may be trademarks of their respective owners.