

Oracle ASM and Multi-Pathing Technologies

Nitin Vengurlekar

November 2009

This document will briefly discuss Multipath I/O and how Oracle Automatic Storage Management leverages multi-pathing.

An I/O path generally consists of an initiator port, fabric port, target port, and LUN. Each permutation of this I/O path is considered an independent path. Dynamic Multi-pathing/failover tools aggregate these independent paths into a single logical path. This path virtualization provides I/O load-balancing across the HBAs, as well as non-disruptive failover on HBA failures. Multi-pathing software requires all the required disks to be visible on each available and eligible HBA. These multi-paths will be detected by probing the World Wide Device Identifiers (WWD ID).

Multi-pathing software also provides multi-path software drivers. Most multi-pathing drivers are SCSI compliant pseudo nexus drivers, that support multipath services for fibre channel attached SCSI-3 devices. These drivers receive naming and transport services from one or more physical HBA devices. To support multi-pathing, a physical HBA driver must comply with the multi-pathing services provided by this driver.

Multipaths provides the following benefits:

- Provide a single block device interface for a multi-pathed LUN
- Detect any component failures in the I/O path; e.g., fabric port, channel adapter, or HBA.
- When a loss of path occurs, ensure that I/Os are re-routed to the available paths, with no process disruption.
- Reconfigure the multipaths automatically when events occur.
- Ensure that failed paths get revalidated as soon as possible.
- Provide auto-failback capabilities.
- Configure the multipaths to maximize performance using various load balancing methods; e.g., round-robin, least I/Os queued, or least service time.

Examples of multi-pathing software include EMC PowerPath, Sun Traffic Manager and IBM MPIO.

Although most of these tools are storage array agnostic, some require additional software components and drivers specific for the storage

Note, Oracle Corporation does not certify or qualify these tools. Oracle10g includes a new feature called Automatic Storage Management (ASM). ASM provides filesystem and volume manager capabilities built into the Oracle database kernel. With this capability, ASM simplifies storage management tasks, such as creating/laying out databases and disk management. Although ASM does not provide multi-pathing capabilities, ASM does leverage multi-pathing tools, as long the path or device produces a successful return code from an fstat system call.

The table below describes the attributes of the utilities and how it can be used by ASM. Note, that Oracle Corporation does not certify ASM against these multipathing utilities. The MP utilities listed below are ones that known working solutions. As we do more testing, additional MP utilities will be listed here, thus, this document is an active document.

<u>OS Platform</u>	<u>Multi-pathing tool</u>	<u>ASM Device Usage</u>	<u>Notes</u>
AIX	EMC PowerPath	Use raw partitions thru the pseudo device /dev/rhdiskpowerx	
	IBM SDD (Vpath)		As of this writing, SDD-AIX is known to cause discovery and device handling problems for ASM, and thus is not viable solution. See SDDPCM section below for an alternative solution to SDD for AIX
	IBM SDDPCM (MPIO)	Use /dev/rhdiskx device:	SDDPCM works with the following IBM storage arrays: DS8000,DS6000,DS3000, DS4x00, DS5x00 and Enterprise Storage Server (ESS), HDS, HP, and EMC storage. Note, Starting with AIX 6.1, the DS4000 and its successor, the DS5000, support MPIO driver.
	IBM RDAC	Use /dev/rhdiskx devices	AIX RDAC supports only the DS4x00 storage units. Check IBM RDAC for latest array support matrix. RDAC also supports specific Sun StorageEdge models as well. Check Sun RDAC support Matrix.
	Hitachi Dynamic Link Manager - HDLM	Use /dev/rdsk/dlmdrvx that's generated by HDLM	HDLM generates a scsi (cxydzx) address where the controller is highest unused controller number. HDLM no longer requires HACMP.
	Veritas DMP	Use /dev/vx/rdsk/<dg-name>/<rlvolname>	Note that access to DMP requires VxVM; i.e., the appropriate package license will be needed
	Fujitsu ETERNUS GR Multipath	Use /dev/rhdisk device	

<i>HP</i>	<i>EMC PowerPath</i>	<i>ASM can use the raw partitions. Must use native device path: /dev/rdisk/cxydz</i>	
	<i>HP Native MP</i>	<i>Use /dev/rdisk/disk*. This paths are based on WWID.</i>	<i>Available only in HP-UX 11i v3. With Native MP, you must use these DSF paths (/dev/rdisk).</i>
	<i>HP SecurePath</i>	<i>ASM can use the raw partitions. /dev/rdisk/cxydz. Since Securepath doesn't support pseudo-paths, use a single path for any given disk</i>	
	<i>Veritas DMP</i>	<i>Use /dev/vx/rdsk/<dg-name>/<rlvolname></i>	<i>Note that access to DMP requires VxVM; i.e., the appropriate package license will be needed</i>
	<i>Hitachi Dynamic Link Manager – HDLM</i>	<i>Use /dev/rdsk/cHtydz that's generated by HDLM</i>	<i>HDLM generates a scsi address like cHtydz.</i>
	<i>Sun Traffic Manager</i>	<i>ASM can use the pseudo device.</i>	<i>The pseudo device will have the Global Unique Identifier included its name.</i>

<i>Sun</i>	<i>PowerPath</i>	<i>ASM will use raw partitions, associated with the pseudo device: /dev/rdsk/emcpowerx</i>	
	<i>Veritas DMP</i>	<i>Use /dev/vx/rdsk/<dg-name>/<rlvolname></i>	<i>Note that access to DMP requires VxVM; i.e., the appropriate package license will be needed</i>
	<i>Sun Traffic Manager</i>	<i>ASM can use the pseudo device.</i>	<i>The pseudo device will have the Global Unique Identifier included its name.</i>
	<i>Hitachi Dynamic Link Manager - HDLM</i>	<i>Use /dev/rdsk/cxydz that's generated by HDLM</i>	<i>HDLM generates a scsi (cxydzx) address where the controller is highest unused controller number.</i>
	<i>Fujitsu ETERNUS GR Multipath Driver</i>	<i>Use /dev/FJSVmplb/[r]rdsk/mplb*1s{0-7}. *1 "I" is MPLB Instance number</i>	

Please note that ASM can leverage ASMLIB over any of the multipath devices produced by the MP driver. See Metalink Note 309815.1 for details

<i>Linux</i>	<i>IBM-SDD</i>	<i>ASM can use the raw device or ASMLIB devices associated with the Vpath.device</i>	<i>You must use SDD-Linux driver 1.6.05+ . Go to IBM SDD site to get latest driver SDD works with the following IBM storage components: DS8000 DS6000 Enterprise Storage Server (ESS) SAN Volume Controller (SVC)</i>
	<i>PowerPath</i>	<i>ASM can use raw partitions associated with the /dev/emcpowerx device.</i>	<i>Associate the raw devices with the PP device.</i>
	<i>HP SecurePath</i>	<i>ASM should use raw devices associated with /dev/spx</i>	<i>Currently SecurePath is only available on x86 and Itanium, all other 64bit platforms are not supported</i>
	<i>Veritas DMP</i>	<i>Use /dev/vx/rdisk/<dg-name>/<rlvolname></i>	<i>Note that access to DMP requires VxVM; i.e., the appropriate package license will be needed</i>
	<i>IBM RDAC</i>	<i>Use /dev/sdx devices</i>	<i>Supports 2.4 and 2.6 kernels. AIX RDAC supports the DS4x00 storage units. Check IBM RDAC for latest array support matrix.</i>
	<i>Device Mapper (DM)</i>	<i>ASM can use the name generated by DM: /dev/disk/by-name/WWID</i>	<i>Avail in 2.6 kernel. Refer to following note for specifics</i>
	<i>Hitachi Dynamic Link Manager - HDLM</i>	<i>ASM can use /dev/sddlm devices</i>	<i>HDLM generates a device called /dev/sddlm</i>
	<i>Fujitsu ETERNUS GR Multipath Driver</i>	<i>Associate raw device to /dev/FJSVmplb/[r]dsk/mplbI*1s{0-7}</i>	
	<i>MD</i>	<i>ASM can use raw device associated with /dev/mdx device.</i>	<i>Though not a real multipath solution, it does provide rudimentary active/passive support</i>