



LifeKeeper for Linux

7.4

Storage and Adapter Configuration

January 2012

This document and the information herein is the property of SIOS Technology Corp. (previously known as SteelEye® Technology, Inc.) and all unauthorized use and reproduction is prohibited. SIOS Technology Corp. makes no warranties with respect to the contents of this document and reserves the right to revise this publication and make changes to the products described herein without prior notification. It is the policy of SIOS Technology Corp. to improve products as new technology, components and software become available. SIOS Technology Corp., therefore, reserves the right to change specifications without prior notice.

LifeKeeper, SteelEye and SteelEye DataKeeper are registered trademarks of SIOS Technology Corp.

Other brand and product names used herein are for identification purposes only and may be trademarks of their respective companies.

To maintain the quality of our publications, we welcome your comments on the accuracy, clarity, organization, and value of this document.

Address correspondence to:
ip@us.sios.com

Copyright © 2012
By SIOS Technology Corp.
San Mateo, CA U.S.A.
All rights reserved

Table of Contents

Storage and Adapter Configuration Version 7.4.....	1
HP Multipath I/O Configurations.....	14
EMC PowerPath Multipath I/O Configurations.....	16
IBM SDD Multipath I/O Configurations.....	19
Hitachi Dynamic Link Manager Software Multipath I/O Configurations.....	20
Device Mapper Multipath I/O Configurations.....	28

Storage and Adapter Configuration Version 7.4

Item	Description
Multipath I/O and Redundant Controllers	<p>There are several multipath I/O solutions either already available or currently being developed for the Linux environment. SIOS Technology Corp. is actively working with a number of server vendors, storage vendors, adapter vendors and driver maintainers to enable LifeKeeper to work with their multipath I/O solutions.</p> <p>LifeKeeper's use of SCSI reservations to protect data integrity presents some special requirements that frequently are not met by the initial implementation of these solutions.</p> <p>Refer to the technical notes below for supported disk arrays to determine if a given array is supported with multiple paths and with a particular multipath solution. Unless an array is specifically listed as being supported by LifeKeeper with multiple paths and with a particular multipath solution, it must be assumed that it is not.</p>

Item	Description
Heavy I/O in Multipath Configurations	<p>In multipath configurations, performing heavy I/O while paths are being manipulated can cause a system to temporarily appear to be unresponsive. When the multipath software moves the access of a LUN from one path to another, it must also move any outstanding I/Os to the new path. The rerouting of the I/Os can cause a delay in the response times for these I/Os. If additional I/Os continue to be issued during this time, they will be queued in the system and can cause a system to run out of memory available to any process. Under very heavy I/O loads, these delays and low memory conditions can cause the system to be unresponsive such that LifeKeeper may detect a server as down and initiate a failover.</p> <p>There are many factors that will affect the frequency at which this issue may be seen.</p> <ul style="list-style-type: none"> • The speed of the processor will affect how fast I/Os can be queued. A faster processor may cause the failure to be seen more frequently. • The amount of system memory will affect how many I/Os can be queued before the system becomes unresponsive. A system with more memory may cause the failure to be seen less frequently. • The number of LUNs in use will affect the amount of I/O that can be queued. • Characteristics of the I/O activity will affect the volume of I/O queued. In test cases where the problem has been seen, the test was writing an unlimited amount of data to the disk. Most applications will both read and write data. As the reads are blocked waiting on the failover, writes will also be throttled, decreasing the I/O rate such that the failure may be seen less frequently. <p>For example, during testing of the IBM DS4000 multipath configuration with RDAC, when the I/O throughput to the DS4000 was greater than 190 MB per second and path failures were simulated, LifeKeeper would (falsely) detect a failed server approximately one time out of twelve. The servers used in this test were IBM x345 servers with dual Xeon 2.8GHz processors and 2 GB of memory connected to a DS4400 with 8 volumes (LUNs) per server in use. To avoid the failovers, the LifeKeeper parameter LCMNUMHBEATS (in <code>/etc/default/LifeKeeper</code>) was increased to 16. The change to this parameter results in LifeKeeper waiting approximately 80 seconds before determining that an unresponsive system is dead, rather than the default wait time of approximately 15 seconds.</p>
Special Considerations for Switchovers with Large Storage Configurations	With some large storage configurations (for example, multiple logical

Item	Description
	<p>volume groups with 10 or more LUNs in each volume group), LifeKeeper may not be able to complete a sendevent within the default timeout of 300 seconds when a failure is detected. This results in the switchover to the backup system failing. All resources are not brought in-service and the following error message is logged in the LifeKeeper log:</p> <pre data-bbox="695 470 1398 730"> ***ERROR*** recover[51,mes.C]:lcdsendremote: ::receive(300) did not receive message within 300 seconds on incoming_mailbox PM1798.21634 /opt/LifeKeeper/bin/recover: lcdsendremote transfer resource "<resource- name>" to "<resource-name>" on machine "system-name" failed (rt=-1) </pre> <p>The recommendation with large storage configurations is to change SCSIERROR from "event" to "halt" in the /etc/default/LifeKeeper file. This will cause LifeKeeper to perform a "halt" on a SCSI error. LifeKeeper will then perform a successful failover to the backup system.</p>
HP MA8000	Certified by SIOS Technology Corp. with QLogic 2200 adapters. Use the qla2200 driver version 6.04.50 or later.
HP MSA1000 and MSA1500	Certified by SIOS Technology Corp. with HP FCA2214 (QLA 2340) adapters in both single and multiple path configurations. Configuration requirements and notes for support of the MSA1000 in a multipath configuration are provided in the separate HP Multipath I/O Configurations section.
HP 3PAR F200/F400/T400/T800	<p>The HP 3PAR was tested by a SIOS Technology Corp. partner with the following configurations: HP 3PAR T400 (Firmware (InForm OS) version 2.3.1 MU4) using HP 82Q 8Gb Dual Port PCI-e FC HBA AJ764A (Firmware version 5.03.02, driver version 8.03.01.05.05.06-k) with DMMP (device-mapper-1.02.55-2.el5, device-mapper-multipath-0.4.7-42.el5).</p> <p>The test was performed with LifeKeeper for Linux v7.3 using RHEL 5.6 (kernel 2.6.18-238.el5).</p>
HP EVA 3000/5000 and EVA 4X00/6X00/8X00 (XCS 6.x series firmware)	Certified by SIOS Technology Corp. with HP FCA2214 (QLA 2340) adapters in both single and multiple path configurations. Configuration requirements and notes for support of the EVA in a multipath configuration are provided in the separate HP Multipath I/O Configurations section.
HP EVA4400	Certified by Hewlett-Packard Company. Both single path and multipath configuration require the DMMP Recovery Kit and the HP DMMP software. The EVA4400 has been qualified to work with LifeKeeper on Red Hat EL 5 Update 3 and Novell SLES 11. Novell testing was completed by the HP Storage Group.
HP EVA6400/8400	Certified by Hewlett-Packard Company. Both single path and multipath configuration require the DMMP Recovery Kit and the HP DMMP

Item	Description
	software. The EVA6400/8400 has been qualified to work with LifeKeeper on Red Hat EL 5 Update 3 and Novell SLES 11. Novell testing was completed by the HP Storage Group.
HP EVA 8100 (XCS 6.x series firmware)	<p>Certified by a SIOS Technology Corp. partner with HP FC 1142SR adapters in DMMP multiple path configurations. Configuration requirements and notes for support of the EVA in a multipath configuration are provided in the separate Device Mapper Multipath I/O Configurations section.</p> <p>EVA 8100 was tested with XCS v6.200 firmware with device-mapper-multipath-0.4.7.-23.el5 with the DMMP Recovery Kit v7.3 with RHEL 5.3.</p>
HP MSA500 (formerly known as Smart Array Cluster Storage or SACS)	<p>Certified by SIOS Technology Corp. with Smart Array 5i and 532 host adapters with the cciss driver. LifeKeeper does support the HP Modular Smart Array 500 system in a redundant controller configuration. This is not a multipath I/O solution, but it does eliminate all potential storage-related single points of failure.</p> <p>Note: LifeKeeper only supports LUN numbers 00 thru 99.</p>
HP MSA500 G2	<p>Certified by SIOS Technology Corp. with the Smart Array 6i and 642 host adapter with the cciss driver. LifeKeeper does support the HP Modular Smart Array 500 G2 system in a redundant controller configuration with both the 2-port and 4-port EMU modules. This is not supported in a multipath I/O configuration, but it does eliminate all potential storage-related single points of failure. The 4-port module provides support for up to a 4-node cluster. LifeKeeper does not support the HP Smart Array Multipath Software for the MSA500 G2.</p> <p>Note: LifeKeeper only supports LUN numbers 00 thru 99.</p>
HP MSA2000fc	<p>Certified by Hewlett-Packard Company with Fibre Channel in both single path and multipath configurations. Models tested were the MSA2012fc and the MSA2212fc with the QLogic QMH2462 HBA using driver version 8.01.07.25 in a single path configuration. The multipath configuration testing was completed using the same models with HP DMMP and the LifeKeeper DMMP Recovery Kit.</p>
HP MSA2000i	<p>Certified by Hewlett-Packard Company with iSCSI in a multipath configuration. The model used for testing was the MSA2012i with HP DMMP. Single path testing was not performed by HP; however, SIOS Technology Corp. supports single path configurations with HP DMMP and the LifeKeeper DMMP Recovery Kit.</p>
HP MSA2000sa	<p>Certified by Hewlett-Packard Company with SA in both single path and multipath configurations. The model used for testing was the MSA2012sa. Both single path and multipath configuration require the DMMP Recovery Kit and the HP DMMP software. HP supports direct connect configurations only at this time.</p>
HP MSA 2300fc	<p>Certified by Hewlett-Packard Company with Fibre Channel in both single and multipath configurations. The model tested was the</p>

Item	Description
	MSA2324fc with the HP AE312A (FC2142SR) HBA using driver version 8.02.09-d0-rhel4.7-04 in a single path configuration. The multipath configuration testing was completed using the same model with HP DMMP and the LifeKeeper DMMP Recovery Kit.
HP MSA 2300i	Certified by Hewlett-Packard Company. Both single path and multipath configuration require the DMMP Recovery Kit and the HP DMMP software.
HP MSA 2300sa	Certified by Hewlett-Packard Company. Both single path and multipath configuration require the DMMP Recovery Kit and the HP DMMP software. Only MSA2300sa rack and tower configurations with DMMP are supported. Blade configurations with LifeKeeper are not supported.
HP P2000 G3 MSA SAS	Certified by SIOS Technology Corp. in multipath configurations using the Device Mapper Multipath Recovery Kit. LifeKeeper for Linux can support up to 11 LUNs in a single cluster with the P2000 G3 SAS array.
HP P4000/P4300 G2	Certified by SIOS Technology Corp. in both a single path and multipath configuration on RHEL 5.5 using the built-in SCSI support in the core of LifeKeeper with iSCSI Software Initiators. Model tested was the HP P4300 G2 7.2TB SAS Starter SAN BK716A. The default kit supports single path storage as well as some multipath storage. In general, the multipath storage is limited to active/passive configurations.
HP P4500 G2	Certified by Hewlett-Packard Company guaranteeing the compatibility of P4500 with P4000 (shown above).
HP P6300 EVA FC	This storage unit was tested by a SIOS Technology Corp. partner in multipath configuration on RHEL 6.1 using the Device Mapper Multipath Recovery Kit.
HP P9500/XP	<p>Certified by Hewlett-Packard Company using SteelEye LifeKeeper for Linux v7.2 or later. Model tested was the HP P9500/XP and has been qualified to work with LifeKeeper on the following:</p> <ul style="list-style-type: none"> • Red Hat Enterprise for 32-bit, x64 (64-bit; Opteron and Intel EMT64) RHEL 5.3, RHEL 5.4, RHEL 5.5 • SuSE Enterprise Server for 32-bit, x64 (64-bit; Opteron and Intel EMT64) SLES 10 SP3, SLES 11, SLES 11 SP1 • Native or Inbox Clustering Solutions RHCS and SLE HA
HP XP20000/XP24000	Certified by SIOS Technology Corp. using LifeKeeper forLinux with DMMP ARK in RHEL 5, SLES10 and SLES 11, configured as multipath by DMMP. The model numbers of tested storage are XP20000 and XP24000. The connection interface is FC. The model number of tested HBA is QLogic QMH2562 and firmware is 4.04.09; driver version is 8.03.00.10.05.04-k. SIOS Technology Corp. recommends that user change setting of path_checker into readsector0 in XP storage.

Item	Description
IBM DS4000 Storage (formerly known as IBM FASTT)	Certified by SIOS Technology Corp. with QLogic 2200 and 2340 adapters in both single and multiple path configurations. Use the qla2200 or qla2300 driver, version 6.03.00 or later, as defined by IBM. When using IBM DS4000 storage arrays systems with Emulex FC adapters, use the lpfc driver versions specified in the Emulex Drivers item below. Single path (i.e. single loop) support: In a single path configuration, a fibre channel switch or hub is required for LifeKeeper to operate properly. Multiple path (i.e. dual loop) support: Multiple paths are supported with the DS4000 storage array models that are released with RDAC support (currently the DS4300, DS4400 and DS4500 models). Fibre channel switches and hubs are not required with multiple path configurations with RDAC. RDAC is a software package that handles path failovers so that an application is not affected by a path failure. The steps to install and setup RDAC are slightly different depending on the version being used. Refer to the IBM RDAC documentation for the instructions to install, build and setup.
IBM DS5000	Certified by partner testing in a multipath configuration using IBM RDAC. Please consult the IBM website to obtain the supported RDAC drivers for your distribution.
IBM DS3500 (FC Model)	Certified by SIOS Technology Corp. in single path and multipath configurations on Red Hat Enterprise Linux Server Release 5.5 (Tikanga), HBA: QLE2560, QLE2460, RDAC: RDAC 09.03.0C05.0331. RDAC is needed for both single path and multipath. Note: SAS and iSCSI connect are not supported.
IBM DS3400 Storage	Certified by SIOS Technology Corp. with QLogic 2300 adapters in both single and multiple path configurations. Use the qla2200 or qla2300 driver, version 6.03.00 or later, as defined by IBM. Please refer to the table entry for IBM DS4000 Storage for more information on single and multiple path support.
IBM System Storage DS3300	Certified by SIOS Technology Corp. with iSCSI Software Initiators. This storage device works in a two node LifeKeeper cluster in both single and multipath configurations. It is required that the IBM RDAC driver be installed on both servers for either single or multipath configurations. If you are using a multipath configuration, you must set SCSIHANGMAX to 50 in the /etc/default/LifeKeeper file. Please consult the IBM website to obtain the supported RDAC drivers for your distribution.
IBM System Storage DS3200	Certified by SIOS Technology Corp. with the IBM SAS HBA (25R8060). This storage device works in a two node LifeKeeper cluster in both single and multipath configurations. It is required that the IBM RDAC driver be installed on both servers for either single or multipath configurations. Please consult the IBM website to obtain the supported SAS and RDAC drivers for your Linux distribution.
IBM DS400	Certified by SIOS Technology Corp. in single path configurations only. Use the firmware version 7.01 build 0838 or later, as defined by IBM.
IBM San Volume	Certified by partner testing in a single path configuration. Certified by

Item	Description
Controller (SVC)	SIOS Technology Corp. in multipath configurations using the SDD Recovery Kit and the Device Mapper Multipath Recovery Kit.
IBM eServer xSeries Storage Solution Server Type445-R / Type445-FR for SANmelody	Certified by partner testing with IBM TotalStorage FC2-133 Host Bus Adapters in multiple path configurations. Use the qla2300 driver, version 7.00.61(non-failover) or later, as defined by IBM. Multiple path support: Multiple paths are supported with the IBM eServer xSeries Storage Solution Server Type445-R / Type445-FR for SANmelody, using the Multipath Linux Driver for IBM SANmelody Solution Server, version 1.0.0 or later.
Dell PowerVault with Dell PERC and LSI Logic MegaRAID controllers	<p>SIOS Technology Corp. has certified the Dell PowerVault storage array for use in a 2-node cluster with the Dell PERC 2/DC, Dell PERC 4/DC, and LSI Logic MegaRAID Elite 1600 storage controllers, as long as the following set of configuration requirements are met. (Note that the Dell PERC 3/DC is the OEM version of the MegaRAID Elite 1600.) These requirements are necessary because these host-based RAID controllers do not provide support for SCSI reservations and unique device IDs, which LifeKeeper normally requires.</p> <ol style="list-style-type: none"> 1. The Dell PowerVault storage should not be mixed with any other types of shared storage to be managed by LifeKeeper within the same cluster. 2. Follow the instructions provided with your hardware for configuring the Dell PowerVault storage and the controllers for use in a cluster. Specifically, this includes getting into the controller firmware setup screens simultaneously on both systems, selecting the adapter properties page, setting "Cluster Mode" to "Enabled", and setting the "Initiator ID" to 6 on one system and to 7 on the other. You should then make sure that both controllers can see the same LUNs, and that the Linuxmegaraid driver is properly configured to be loaded. 3. Because this storage configuration does not support SCSI reservations, you must disable the use of SCSI reservations within LifeKeeper. This is accomplished by adding the option "RESERVATIONS=none" to the LifeKeeper defaults file, /etc/default/LifeKeeper, on both nodes in the cluster. You must manually configure a unique ID for each LUN to be managed by LifeKeeper, using the /opt/LifeKeeper/bin/lkID utility. The assigned ID must be unique within the cluster and should be sufficiently constructed to avoid potential future conflicts. The lkID utility will automatically generate a unique ID for you if desired. Refer to thelkID(8) man page for more details about the use of the utility, the IDs that it generates, where the ID is placed, and any possible restrictions. Also, see the note regarding the use of lkID with LVM in the Known Issues section for the LVM Recovery Kit.

Item	Description
	<p>4. Obtain and configure a STONITH device or devices to provide I/O fencing. This is required due to the lack of SCSI reservation support in this configuration. Note that for this configuration, you should configure your STONITH devices to do a system “poweroff” command rather than a “reboot”. You must also take care to avoid bringing a device hierarchy in-service on both nodes simultaneously via a manual operation when LifeKeeper communications have been disrupted for some reason.</p>
Dell EMC (CLARiiON) CX200	EMC has approved two QLogic driver versions for use with this array and the QLA2340 adapter: the qla2x00-clariion-v6.03.00-1 and the qla2x00-clariion-v4.47.18-1. Both are available from the QLogic website at www.qlogic.com .
DELL MD3000	Certified by Partner testing in both single path and multipath configurations with DELL SAS 5/e adapters. This was specifically tested with RHEL4; however, there are no known issues using other LifeKeeper supported Linux distributions or versions. RDAC is required for both single path and multipath configurations. For single path configurations, use the HBA host type of "Windows MSCS Cluster single path". For multipath configurations, use the HBA host type of "Linux".
Dell PowerVault MD3200/3220	<p>Dell PowerVault MD3200/3220 was tested by a SIOS Technology Corp. partner with the following configuration:</p> <p>DMMP with the DMMP Recovery Kit on RHEL 5.5. Must be used with the combination of Quorum/Witness Server Kit and STONITH. To disable SCSI reservation, set RESERVATIONS=none in <i>/etc/default/LifeKeeper</i>. Server must have interface based on IPMI 2.0.</p>
Dell EqualLogic PS5000	<p>The Dell EqualLogic was tested by a SIOS Technology Corp. partner with the following configurations:</p> <ul style="list-style-type: none"> • Dell EqualLogic PS5000 using SCSI -2 reservations with the iscsi-initiator(Software initiator) using Red Hat Enterprise Linux ES release 4 (Nahant Update 5) with kernel 2.6.9-55.EL. The testing was completed using iscsi-initiator-utils-4.0.3.0-5 and multipath configuration using bonding with active-backup (mode=1). • Dell EqualLogic PS5000 using DMMP with the DMMP Recovery Kit with RHEL 5 with iscsi-initiator-utils-6.2.0.865-0.8.el5. With a large number of luns (over 20), change the REMOTETIMEOUT setting in <i>/etc/default/LifeKeeper</i> to REMOTETIMEOUT=600.
Dell EqualLogic PS4000/4100/6000/6100 6500/6010/6510	<p>The Dell EqualLogic was tested by a SIOS Technology Corp. partner with the following configurations: Dell EqualLogic PS4000/4100/6000/6100/6500/6010/6510 using DMMP with the DMMP Recovery Kit with RHEL 5.3 with iscsi-initiator-utils-6.2.0.868-0.18.el5. With a large number of luns (over 20), change the REMOTETIMEOUT setting in <i>/etc/default/LifeKeeper</i> to REMOTETIMEOUT=600.</p>

Item	Description
FalconStor Network Storage Server (NSS)	<p>Certified by SIOS Technology Corp. The following parameters should be set in <code>/etc/multipath.conf</code>:</p> <pre>polling_interval 5 no_path_retry 36</pre>
Hitachi HDS 9570V, 9970V and 9980V	<p>Certified by SIOS Technology Corp. in a single path configuration with QLogic 23xx adapters. Use the <code>qla2300</code> driver, version 6.04 or later.</p> <p>Note: SIOS Technology Corp. recommends the use of only single controller (i.e. single loop) configurations with these arrays, using a fibre channel switch or hub. However, it is also possible to configure a LifeKeeper cluster in which each server is connected directly to a separate controller or port on the Hitachi array, without the use of a switch or hub, as long as each server has only a single path to the storage. It should be noted that LifeKeeper behaves quite differently from its normal behavior in a split-brain scenario using this configuration. LifeKeeper normally performs a failover of an active hierarchy in a split-brain scenario causing the original primary node to reboot as a result of the stolen SCSI reservation. When the Hitachi arrays are configured with the servers directly connected to multiple controllers or ports, certain timing peculiarities within the Hitachi arrays prevent LifeKeeper from acquiring a SCSI reservation on the backup node and the failover attempt fails, leaving at least part of the hierarchy running on the original primary node. For this reason, it is important that all LifeKeeper resources in such a configuration have a direct line of dependencies to one of the disk resources such that no resources can be brought in-service if the disk resources cannot be transferred. This is particularly true of any IP resources in the hierarchy.</p>
Hitachi HDS 9980V	<p>There are certain specific “host mode” settings required on the Hitachi arrays in order to allow LifeKeeper to work properly in this kind of directly connected configuration. For the 9570V array, the following settings are required:</p> <pre>Host connection mode1 --> Standard mode Host connection mode2 --> Target Reset mode (Bus Device Reset) Third Party Process Logout Spread mode LIP port all reset mode --> LIP port all reset mode</pre> <p>For the 9970V and 9980V arrays, the “host mode” must be set to “SUN”. The HDS 9980V was tested by a SIOS Technology Corp. partner organization in a multipath configuration using DMMP on SLES9 SP3 with LSI Logic Fusion HBAs. Refer to the Device Mapper Multipath I/O Configurations section for details.</p>
nStor NexStor 4320F	<p>This storage was tested by a SIOS Technology Corp. partner organization, in a dual controller configuration with each server in a 2-node cluster directly connected to a separate controller in the array. With this configuration, the LifeKeeper behavior in a split-brain scenario is the same as that described above for the Hitachi HDS storage</p>

Item	Description
	arrays, and the same hierarchy configuration precautions should be observed.
ADTX ArrayMasStor L and FC-II	These storage units were tested by a SIOS Technology Corp. partner organization, both in a single path configuration with a switch and in a dual controller configuration with each server in a 2-node cluster directly connected to a separate controller in the array. In both configurations, the LifeKeeper behavior in a split-brain scenario is the same as that described above for the Hitachi HDS storage arrays, and the same hierarchy configuration precautions should be observed. The ArrayMasStor L was also tested and certified by our partner organization in a multipath configuration using QLogic 2340 and 2310 host bus adapters and the QLogic failover driver, version 6.06.10.
Fujitsu ETERNUS3000	This storage unit was tested by a SIOS Technology Corp. partner organization in a single path configuration only, using the PG-FC105 (Emulex LP9001), PG-FC106 (Emulex LP9802), and PG-PC107 host bus adapters and the lpfc driver v7.1.14-3.
Fujitsu ETERNUS 6000	This storage unit was tested by a SIOS Technology Corp. partner organization in a single path configuration only, using the PG-FC106 (Emulex LP9802) host bus adapter and the lpfc driver v7.1.14-3.
Fujitsu FibreCAT S80	This array requires the addition of the following entry to <code>/etc/default/LifeKeeper</code> : <code>ADD_LUN_TO_DEVICE_ID=TRUE</code>
Fujitsu ETERNUS SX300	This storage unit was tested by a SIOS Technology Corp. partner organization in a multipath configuration only using the PG-FC106 (Emulex LP9802) and PG-FC107 host bus adapters and the lpfc driver v7.1.14. The RDAC driver that is bundled with the SX300 is required.
Fujitsu ETERNUS2000 Model 50	This storage unit was tested by a SIOS Technology Corp. partner organization in a multipath configuration with dual RAID controllers using the PG-FC202 (LPe1150-F) host bus adapter with the EMPD multipath driver. Firmware version WS2.50A6 and driver version EMPD V2.0L12 were used in the testing. Testing was performed with LifeKeeper for Linux v6.2 using RHEL4 (kernel 2.6.9-67.ELsmp) and RHEL5 (kernel 2.6.18-53.el5).
Fujitsu ETERNUS4000 Model 300	This storage unit was tested by a SIOS Technology Corp. partner organization in a multipath configuration with dual RAID controllers using the PG-FC202 (LPe1150-F) host bus adapter with the EMPD multipath driver.
Fujitsu ETERNUS2000 Model 200	This storage unit was tested by Fujitsu Limited in a multipath configuration using PG-FC203L (Emulex LPe1250-F8) host bus adapter (Firmware version 1.11A5, driver version 8.2.0.48.2p) with EMPD multipath driver (driver version V2.0L20, patch version T000973LP-1). The test was performed with LifeKeeper for Linux v7.1 using RHEL5 (kernel 2.6.18-164.el5).

Item	Description
Fujitsu ETERNUS VS850	Certified by vendor support statement in a single path configuration and in multipath configurations using the Device Mapper Multipath Recovery Kit.
NEC iStorage Storage Path Savior Multipath I/O	<p>Protecting Applications and File Systems That Use Multipath Devices: In order for LifeKeeper to configure and protect applications or file systems that use SPS devices, the SPS recovery kit must be installed.</p> <p>Once the SPS kit is installed, simply creating an application hierarchy that uses one or more of the multipath device nodes will automatically incorporate the new resource types provided by the SPS kit.</p> <p>Multipath Device Nodes: To use the SPS kit, any file systems and raw devices must be mounted or configured on the multipath device nodes (/dev/dd*) rather than on the native /dev/sd* device nodes.</p> <p>Use of SCSI-3 Persistent Reservations: The SPS kit uses SCSI-3 persistent reservations with a "Write Exclusive" reservation type. This means that devices reserved by one node in the cluster will remain read-accessible to other nodes in the cluster, but those other nodes will be unable to write to the device. Note that this does not mean that you can expect to be able to mount file systems on those other nodes for ongoing read-only access.</p> <p>LifeKeeper uses the sg_persist utility to issue and monitor persistent reservations. If necessary, LifeKeeper will install the sg_persist(8) utility.</p> <p>Hardware Requirements: The SPS kit has been tested and certified with the NEC iStorage diskarray using Emulex LP952, LP9802, LP1050 HBAs and Emulex lpc driver. This kit is expected to work equally well with other NEC iStorage D and S supported by SPS.</p> <p>Multipath Software Requirements: The SPS kit has been tested with SPS for Linux 3.3.001. There are no known dependencies on the version of the SPS package installed.</p> <p>Installation Requirements: SPS software must be installed prior to installing the SPS recovery kit.</p> <p>Adding or Repairing SPS Paths: When LifeKeeper brings an SPS resource into service, it establishes a persistent reservation registered to each path that was active at that time. If new paths are added after the initial reservation, or if failed paths are repaired and SPS automatically reactivates them, those paths will not be registered as a part of the reservation until the next LifeKeeper quickCheck execution for the SPS resource. If SPS allows any writes to that path prior to that point in time, reservation conflicts that occur will be logged to the system message file. The SPS driver will retry these IOs on the registered path resulting in no observable failures to the application.</p>

Item	Description
	Once quickCheck registers the path, subsequent writes will be successful.
Newtech SweeperStor SATA and SAS	<p>This storage unit was tested by a SIOS Technology Corp. partner organization, in a multipath configuration with dual RAID controllers, using the QLogic PCI to Fibre Channel Host Adapter for QLE2462 (with Firmware version 4.03.01 [IP], Driver version 8.02.08) with storage firmware J200. Testing was performed with LifeKeeper for Linux v6.2 with DMMP Recovery Kit v6.2 using the following distributions and kernels:</p> <p>RHEL4 DMMP</p> <ul style="list-style-type: none"> Emulex LP 11002 8.0.16.32 or later Emulex LPe 11002 8.0.16.32 or later Qlogic QLA 2462 8.01.07 or later Qlogic QLE 2462 8.01.07 or later <p>RHEL5 DMMP</p> <ul style="list-style-type: none"> Emulex LP 11002 8.1.10.9 or later Emulex LPe 11002 8.1.10.9 or later Qlogic QLA 2462 8.01.07.15 or later Qlogic QLE 2462 8.01.07.15 or later <p>SLES10 DMMP</p> <ul style="list-style-type: none"> Emulex LP 11002 8.1.10.9 or later Emulex LPe 11002 8.1.10.9 or later Qlogic QLA 2462 8.01.07.15 or later Qlogic QLE 2462 8.01.07.15 or later <p>Note: DMMP is required for multipath configurations.</p>
TID MassCareRAID	<p>This storage unit was tested by a SIOS Technology Corp. partner with the following single path configuration:</p> <ul style="list-style-type: none"> • Host1 Qlogic QLE2562 (HBA BIOS 2.10, driver version qla2xxx-8.03.01.04.05.05-k*) • Host2 HP AE312A (HBA BIOS 1.26, driver version qla2xxx-8.03.01.04.05.05-k*) • The test was performed with LifeKeeper for Linux v7.3 using Red Hat Enterprise Linux 5.5 (kernel2.6.18-194.el5) <p>LifeKeeper for Linux can support up to 11 LUNs in a single cluster with the TID MassCareRAID array.</p>
TIDMassCareRAID	<p>This storage unit was tested by a SIOS Technology Corp. partner organization in a multipath configuration using the Qlogic driver with SCSI-2 reservations with no Fibre Channel switches. Red Hat Enterprise Linux ES release 4 Update6 was used with the 2.6.9-67.ELsmp kernel. The FAILFASTTIMER setting in /etc/default/LifeKeeper needs to be changed from 5 to 30.</p>

Item	Description
Sun StorageTek 2540	This storage unit was tested by a SIOS Technology Corp. partner organization, in a multipath configuration using RDAC with Dual RAID Controllers, using the StorageTek 4Gb PCI-E Dual FC Host Bus Adapter and the Sun StorageTek 4Gb PCI Dual FC Network Adapter.
QLogic Drivers	For other supported fibre channel arrays with QLogic adapters, use the qla2200 or qla2300 driver, version 6.03.00 or later.
Emulex Drivers	For the supported Emulex fibre channel HBAs, you must use the lpfc driver v8.0.16 or later.
Adaptec 29xx Drivers	For supported SCSI arrays with Adaptec 29xx adapters, use the aic7xxx driver, version 6.2.0 or later, provided with the OS distribution.
DataCore SANsymphony	<p>This storage device was successfully tested with SUSE SLES 9 Service Pack 3, Device Mapper - Multipath and Qlogic 2340 adapters. We expect that it should work with other versions, distributions and adapters; however, this has not been tested. See DataCore for specific support for these and other configurations.</p> <p>One issue was found during failover testing with heavy stress running where multiple server reboots would result in a server only configuring a single path. The test configuration consisted of a 3-node cluster where 2 servers would be killed simultaneously. After the 2 servers rebooted, about 50% of the time one server would only have a single path configured. A reboot of the server would resolve the problem. This issue was never seen when only a single server was rebooted. This issue has been reported to DataCore. This item is not considered a critical issue since at least one path continues to be available.</p>
Xitech Magnitude 3D	<p>This storage device was successfully tested with Red Hat EL 4 Update 3 and Qlogic 2340 adapters. We expect that LifeKeeper would also work with other versions, distributions and adapters; however, this has not been tested. See Xitech for specific support for these and other configurations.</p> <p>The Magnitude 3D was tested in a single path configuration.</p> <p>During setup, one configuration issue was detected where only 8 LUNs were configured in the OS. This is due to the Magnitude 3D specifying in the SCSI inquiry data that it is a SCSI-2 device. The SCSI driver in the 2.6 kernel will not automatically address more than 8 LUNs on a SCSI-2 device unless the device is included in its exception list. The Magnitude 3D is not in that list. Xitech provided a workaround for testing to issue a command to <code>/proc/scsi/scsi</code> to configure each LUN.</p>

HP Multipath I/O Configurations

Item	Description
MSA1000 and MSA1500 Multipath Requirements with Secure Path	LifeKeeper supports Secure Path for multipath I/O configurations with the MSA1000 and MSA1500. This support requires the use of the Secure Path v3.0C or later.
HP P2000	LifeKeeper supports the use of HP P2000 MSA FC. This storage unit was tested by SIOS Technology Corp. in a multipath configuration on RHEL 5.4.
EVA3000 and EVA5000 Multipath Requirements with Secure Path	<p>LifeKeeper requires the following in order to support the EVA3000 and EVA5000 in a multipath I/O configuration using Secure Path:</p> <ol style="list-style-type: none"> 1. EVA VCS v2.003, or v3.00 or later. For each server, use Command View v3.00 to set the Host OS type to Custom and the Custom Mode Number as hex000000002200282E. See the HP Secure Path Release Notes for detailed instructions. 2. HP Secure Path v3.0C or later.
Multipath Cluster Installation Using Secure Path	<p>For a fresh installation of a multiple path cluster that uses Secure Path, perform these steps:</p> <ol style="list-style-type: none"> 1. Install the OS of choice on each server. 2. Install the clustering hardware: FCA2214 adapters, storage, switches and cables. 3. Install the HP Platform Kit. 4. Install the HP Secure Path software. This will require a reboot of the system. Verify that Secure Path has properly configured both paths to the storage. See Secure Path documentation for further details. 5. Install LifeKeeper.
Multipath Support for MSA1000 and MSA1500 with QLogic Failover Driver	LifeKeeper for Linux supports the use of the QLogic failover driver for multipath I/O configurations with the MSA1000 and MSA1500. This support requires the use of the QLogic driver v7.00.03 or later..
Multipath Support for EVA with QLogic Failover Driver	LifeKeeper supports the EVA 3000/5000 and the EVA 4X00/6X00/8X00 with the QLogic failover driver. The 3000/5000 requires firmware version 4000 or higher. The 4000/6000/8000 requires firmware version 5030 or higher. The latest QLogic driver supplied by HP (v8.01.03 or later) should be used. The host connection must be "Linux". There is no restriction on the path/mode setting by LifeKeeper. Notice that previous restrictions for a special host connection, the setting of the preferred path/mode and the ports that can be used on the EVA do not exist for this version of firmware and driver.

<p>Upgrading a Single Path MSA1000/ MSA1500 or EVA Configuration to Multiple Paths with Secure Path</p>	<p>To upgrade a cluster from single path to multiple paths, perform the following steps (this must be a cluster-wide upgrade):</p> <ol style="list-style-type: none"> 1. Upgrade LifeKeeper to the latest version following the normal upgrade procedures. This step can be accomplished as a rolling upgrade such that the entire cluster does not have to be down. 2. Stop LifeKeeper on all nodes. The cluster will be down until the hardware upgrade is complete and step 5 is finished for all nodes. 3. Install/upgrade the HP Platform Kit on each node. 4. Install the HP Secure Path software on each node. This will require a reboot of the system. Verify that Secure Path has properly configured both paths to the storage. See Secure Path documentation for further details. 5. Start LifeKeeper. All hierarchies should work as they did before the upgrade. <p>Note: This is a change from how the previous version of LifeKeeper supported an upgrade.</p>
<p>Secure Path Persistent Device Nodes</p>	<p>Secure Path supports “persistent” device nodes that are in the form of /dev/spdev/spXX where XX is the device name. These nodes are symbolic links to the specific SCSI device nodes /dev/sdXX. LifeKeeper v4.3.0 or later will recognize these devices as if they were the “normal” SCSI device nodes /dev/sdXX. LifeKeeper maintains its own device name persistence, both across reboots and across cluster nodes, by directly detecting if a device is /dev/sda1 or /dev/sdq1, and then directly using the correct device node.</p> <p>Note: Support for symbolic links to SCSI device nodes was added in LifeKeeper v4.3.0.</p>
<p>Active/Passive Controllers and Controller Switchovers</p>	<p>The MSA1000 implements multipathing by having one controller active with the other controller in standby mode. When there is a problem with either the active controller or the path to the active controller, the standby controller is activated to take over operations. When a controller is activated, it takes some time for the controller to become ready. Depending on the number of LUNs configured on the array, this can take 30 to 90 seconds. During this time, IOs to the storage will be blocked until they can be rerouted to the newly activated controller.</p>
<p>Single Path on Boot Up Does Not Cause Notification</p>	<p>If a server can access only a single path to the storage when the system is loaded, there will be no notification of this problem. This can happen if a system is rebooted where there is a physical path failure as noted above, but transient path failures have also been observed. It is advised that any time a system is loaded, the administrator should check that all paths to the storage are properly configured, and if not, take actions to either repair any hardware problems or reload the system to resolve a transient problem.</p>

EMC PowerPath Multipath I/O Configurations

Protecting Applications and File Systems That Use Multipath Devices	In order for LifeKeeper to configure and protect applications or file systems that use EMC PowerPath devices, the PowerPath recovery kit must be installed. Once the PowerPath kit is installed, simply creating an application hierarchy that uses one or more of the multipath device nodes will automatically incorporate the new resource types provided by the PowerPath kit.
Multipath Device Nodes	To use the PowerPath kit, any file systems and raw devices must be mounted or configured on the multipath device nodes (/dev/emcpower*) rather than on the native /dev/sd* device nodes.
Use of SCSI-3 Persistent Reservations	<p>The PowerPath kit uses SCSI-3 persistent reservations with a "Write Exclusive" reservation type. This means that devices reserved by one node in the cluster will remain read-accessible to other nodes in the cluster, but those other nodes will be unable to write to the device. Note that this does not mean that you can expect to be able to mount file systems on those other nodes for ongoing read-only access.</p> <p>LifeKeeper uses the sg_persist utility to issue and monitor persistent reservations. If necessary, LifeKeeper will install the sg_persist(8) utility.</p> <p>SCSI-3 Persistent Reservations must be enabled on a per LUN basis when using EMC Symmetrix (including VMAX) arrays with multipathing software and LifeKeeper. This applies to both DMMP and PowerPath.</p>
Hardware Requirements	<p>The PowerPath kit has been tested and certified with the EMC CLARiiON CX300 disk array using QLogic QLA2340 HBAs with the EMC-recommendedqla2xxx driver and using Emulex LP10000 HBAs with the EMC-recommended lpfc driver. The PowerPath kit has also been tested and certified with the EMC CLARiX CX3-20 using QLogic QLA2340 HBAs.</p> <p>Note: LifeKeeper on RHEL 6 cannot support reservations connected to an EMC Clariion.</p> <p>This kit is expected to work equally well with other CLARiiON models from EMC or models OEMed from EMC by Dell or other vendors.</p>
Multipath Software Requirements	The PowerPath kit v6.4.0-2 requires PowerPath for Linux v5.3. The PowerPath kit versions prior to v6.4.0-2 requires PowerPath for Linux v4.4.x, v4.5.x, v5.0.x, or v5.1.x
Migrating to the PowerPath v5.3 driver	<p>Option A</p> <ol style="list-style-type: none"> 1. Upgrade to the PowerPath 5.3 driver by doing the following: <ol style="list-style-type: none"> a. Remove the old PowerPath driver b. Install the PowerPath 5.3 driver 2. Upgrade PowerPath 6.4.0-2 kit 3. Reboot the server <p>Note: When the server reboots, the PowerPath 6.4.0-2 kit will be used for the LifeKeeper PowerPath resources. If there is a problem with the Power Path 5.3</p>

driver and the older PowerPath driver needs to be used, this option will require re-installing the version of the PowerPath kit that was used before the installation of the v6.4.0-2 kit.

Option B

1. Upgrade to the PowerPath 5.3 driver by doing the following:
 - a. Remove the old PowerPath driver
 - b. Install the PowerPath 5.3 driver
 - c. Reboot the server
2. Upgrade to the PowerPath 6.4.0-2 kit and do one of the following to start the PowerPath resources using the upgraded Recovery Kit:

Option 1: Take the PowerPath resources out-of-service and then back in-service.

Note: This will require that all applications using the PowerPath devices be stopped and then restarted. This option allows the actions to be done serially and perhaps at different times to avoid a lot of change.

Option 2: Stop LifeKeeper (lkstop) and start LifeKeeper (lkstart). This will take all resources out-of-service and then back in-service.

Note: As in Option 1, it will stop all applications, but this option requires less user intervention as two commands will ensure that all PowerPath resources are using the new kit.

Option 3: Stop LifeKeeper quickly (lkstop -f) and start LifeKeeper (lkstart).

Note: This will leave the applications running while LifeKeeper reloads how it is accessing the storage. There will be no application downtime.

IBM SDD Multipath I/O Configurations

Protecting Applications and File Systems That Use Multipath Devices	<p>In order for LifeKeeper to configure and protect applications or file systems that use IBM SDD devices, the SDD recovery kit must be installed.</p> <p>Once the SDD kit is installed, simply creating an application hierarchy that uses one or more of the multipath device nodes will automatically incorporate the new resource types provided by the SDD kit.</p>
Multipath Device Nodes	To use the SDD kit, any file systems and raw devices must be mounted or configured on the multipath device nodes (/dev/vpath*) rather than on the native /dev/sd* device nodes.
Use of SCSI-3 Persistent Reservations	<p>The SDD kit uses SCSI-3 persistent reservations with a "Write Exclusive" reservation type. This means that devices reserved by one node in the cluster will remain read-accessible to other nodes in the cluster, but those other nodes will be unable to write to the device. Note that this does not mean that you can expect to be able to mount file systems on those other nodes for ongoing read-only access.</p> <p>LifeKeeper uses the sg_persist utility to issue and monitor persistent reservations. If necessary, LifeKeeper will install the sg_persist(8) utility.</p>
Hardware Requirements	The SDD kit has been tested and certified with IBM ESS, 6800 and 8100 disk arrays and IBM SAN Volume Controller (SVC) using Qlogic QLA2340 HBAs and the IBM-recommended qla2xxx driver. The kit is expected to work equally well with other IBM disk arrays and HBA adapters (Emulex) that are supported by the SDD driver. The IBM-recommended HBA drivers must be used in all cases.
Multipath Software Requirements	The SDD kit requires the IBM SDD driver v1.6.0.1-8 or later.
Adding or Repairing SDD Paths	When LifeKeeper brings an SDD resource into service, it establishes a persistent reservation registered to each path that was active at that time. If new paths are added after the initial reservation, or if failed paths are repaired and SDD automatically reactivates them, those paths will not be registered as a part of the reservation until the next LifeKeeper quickCheck execution for the SDD resource. If SDD allows any writes to that path prior to that point in time, reservation conflicts occur that will be logged in both the SDD log file as well as the system message file. The SDD driver will retry the IOs on the registered path resulting in no observable failures to the application. Once quickCheck registers the path, subsequent writes will be successful.

Hitachi Dynamic Link Manager Software Multipath I/O Configurations

Protecting Applications and File Systems That Use Multipath Devices	<p>In order for LifeKeeper to configure and protect applications or file systems that use HDLM devices, the HDLM recovery kit must be installed.</p> <p>Once the HDLM kit is installed, simply creating an application hierarchy that uses one or more of the multipath device nodes will automatically incorporate the new resource types provided by the HDLM kit.</p>
Multipath Device Nodes	<p>To use the HDLM kit, any file systems and raw devices must be mounted or configured on the multipath device nodes (/dev/sddlm*) rather than on the native /dev/sd* device nodes.</p>
Use of SCSI-3 Persistent Reservations	<p>The HDLM kit uses SCSI-3 persistent reservations with a "Write Exclusive" reservation type. This means that devices reserved by one node in the cluster will remain read-accessible to other nodes in the cluster, but those other nodes will be unable to write to the device. Note that <u>this does not mean</u> that you can expect to be able to mount file systems on those other nodes for ongoing read-only access.</p> <p>LifeKeeper uses the sg_persist utility to issue and monitor persistent reservations. If necessary, LifeKeeper will install the sg_persist(8) utility.</p>
Hardware Requirements	<p>The HDLM kit has been tested and certified with the Hitachi AMS500 disk array using QLogic QLA2340 HBAs and default qla2xxx driver. This kit is expected to work equally well with other Hitachi disk arrays. The HDLM kit has also been tested and certified with the SANRISE AMS series and the SANRISE USP. The HBA and the HBA driver must be supported by HDLM.</p>
Multipath Software Requirements	<p>The HDLM kit has been tested with HDLM for Linux 05-80, 05-81, 05-90, 05-91, 05-92, 05-93, 05-94, 6.0.0, 6.0.1, 6.1.0, 6.1.1, 6.1.2, 6.2.0, 6.2.1, 6.3.0, 6.4.0, 6.4.1 and 6.5.0. There are no known dependencies on the version of the HDLM package installed.</p> <p>Note: The product name changed to "Hitachi Dynamic Link Manager Software (HDLM)" for HDLM 6.0.0 and later. Versions older than 6.0.0 (05-9x) are named "Hitachi HiCommand Dynamic Link Manager (HDLM)".</p> <p>Note: HDLM version 6.2.1 or later is not supported by HDLM Recovery Kit v6.4.0-2. If you need to use this version of HDLM, you can use HDLM Recovery Kit v7.2.0-1 or later with LK Core v7.3.</p>
Linux Distribution Requirements	<p>Linux Distribution Requirements</p> <p>The HDLM kit is supported in the following distributions:</p> <p>RHEL 4 (AS/ES) (x86 or x86_64) Update 1, 2, 3, 4, Update 4 Security Fix (*2), 4.5, 4.5 Security Fix(*4), 4.6, 4.6 Security Fix(*8), 4.7, 4.7 Security Fix(*9), 4.8 (x86/x86_64)(*1)</p> <p>RHEL 5(x86 or x86_64) 5.1, 5.1 Security Fix(*5), 5.2, 5.2 Security Fix(*6), 5.3, 5.3 Security Fix(*10), 5.4, 5.4 Security Fix (*11), 5.5 (x86/x86_64)(*1)</p> <p>(*1) AMD Opteron(Single Core, Dual Core) or Intel EM64T architecture CPU with x86_64 kernel.</p> <p>(*2) The following kernels are supported</p>

	<p>x86: 2.6.9-42.0.3.EL, 2.6.9-42.0.3.ELsmp, 2.6.9-42.0.3.ELhugemem</p> <p>x86_64: 2.6.9-42.0.3.EL, 2.6.9-42.0.3.ELsmp, 2.6.9-42.0.3.ELLargesmp</p> <p>(*3) Hitachi does not support RHEL4 U2 environment</p> <p>(*4) The following kernels are supported</p> <p>x86: 2.6.9-55.0.12.EL, 2.6.9-55.0.12.ELsmp, 2.6.9-55.0.12.ELhugememx</p> <p>x86_64: 2.6.9-55.0.12.EL, 2.6.9-55.0.12.ELsmp, 2.6.9-55.0.12.ELLargesmp</p> <p>(*5) The following kernels are supported</p> <p>x86: 2.6.18-53.1.13.el5, 2.6.18-53.1.13.el5PAE, 2.6.18-53.1.21.el5, 2.6.18-53.1.21.el5PAE</p> <p>x86_64: 2.6.18-53.1.13.el5, 2.6.18-53.1.21.el5</p> <p>(*6) The following kernels are supported</p> <p>x86: 2.6.18-92.1.6.el5, 2.6.18-92.1.6.el5PAE, 2.6.18-92.1.13.el5, 2.6.18-92.1.13.el5PAE, 2.6.18-92.1.22.el5, 2.6.18-92.1.22.el5PAE</p> <p>x86_64: 2.6.18-92.1.6.el5, 2.6.18-92.1.13.el5, 2.6.18-92.1.22.el5</p> <p>(*7) The following kernels are supported</p> <p>x86: 2.6.9-34.0.2.EL, 2.6.9-34.0.2.ELsmp, 2.6.9-34.0.2.ELhugemem</p> <p>x86_64: 2.6.9-34.0.2.EL, 2.6.9-34.0.2.ELsmp, 2.6.9-34.0.2.ELLargesmp</p> <p>(*8) The following kernels are supported</p> <p>x86: 2.6.9-67.0.7.EL, 2.6.9-67.0.7.ELsmp, 2.6.9-67.0.7.ELhugemem, 2.6.9-67.0.22.EL, 2.6.9-67.0.22.ELsmp, 2.6.9-67.0.22.ELhugemem</p> <p>x86_64: 2.6.9-67.0.7.EL, 2.6.9-67.0.7.ELsmp, 2.6.9-67.0.7.ELLargesmp, 2.6.9-67.0.22.EL, 2.6.9-67.0.22.ELsmp, 2.6.9-67.0.22.ELLargesmp</p> <p>(*9) The following kernels are supported</p> <p>x86: 2.6.9-78.0.1.EL, 2.6.9-78.0.1.ELsmp, 2.6.9-78.0.1.ELhugemem, 2.6.9-78.0.5.EL, 2.6.9-78.0.5.ELsmp, 2.6.9-78.0.5.ELhugemem, 2.6.9-78.0.8.EL, 2.6.9-78.0.8.ELsmp, 2.6.9-78.0.8.ELhugemem, 2.6.9-78.0.17.EL, 2.6.9-78.0.17.ELsmp, 2.6.9-78.0.17.ELhugemem, 2.6.9-78.0.22.EL, 2.6.9-78.0.22.ELsmp, 2.6.9-78.0.22.ELhugemem</p> <p>x86_64: 2.6.9-78.0.1.EL, 2.6.9-78.0.1.ELsmp, 2.6.9-78.0.1.ELLargesmp, 2.6.9-78.0.5.EL, 2.6.9-78.0.5.ELsmp, 2.6.9-78.0.5.ELLargesmp, 2.6.9-78.0.8.EL, 2.6.9-78.0.8.ELsmp, 2.6.9-78.0.8.ELLargesmp, 2.6.9-78.0.17.EL, 2.6.9-78.0.17.ELsmp, 2.6.9-78.0.17.ELLargesmp, 2.6.9-78.0.22.EL, 2.6.9-78.0.22.ELsmp, 2.6.9-78.0.22.ELLargesmp</p> <p>(*10) The following kernels are supported</p>
--	--

	<p>x86: 2.6.18-128.1.10.el5, 2.6.18-128.1.10.el5PAE, 2.6.18-128.1.14.el5, 2.6.18-128.1.14.el5PAE, 2.6.18-128.7.1.el5, 2.6.18-128.7.1.el5PAE</p> <p>x86_64: 2.6.18-128.1.10.el5, 2.6.18-128.1.14.el5</p> <p>(*11) The following kernels are supported</p> <p>x86: 2.6.18-164.9.1.el5, 2.6.18-164.9.1.el5PAE, 2.6.18-164.11.1.el5, 2.6.18-164.11.1.el5PAE</p> <p>x86_64: 2.6.18-164.9.1.el5, 2.6.18-164.11.1.el5</p> <p>(*12) The following kernels are supported</p> <p>x86: 2.6.9-89.0.20.EL, 2.6.9-89.0.20.ELsmp, 2.6.9-89.0.20.Elhugemem</p> <p>x86_64: 2.6.9-89.0.20.EL, 2.6.9-89.0.20.ELsmp, 2.6.9-89.0.20.Ellargesmp</p> <p>(*13) The following kernels are supported</p> <p>x86: 2.6.18-194.11.1.el5, 2.6.18-194.11.1.el5PAE</p> <p>x86_64: 2.6.18-194.11.1.el5</p>
Installation Requirements	<p>HDLM software must be installed prior to installing the HDLM recovery kit. Also, customers wanting to transfer their environment from SCSI devices to HDLM devices must run the Installation Support setup script after configuring the HDLM environment. Otherwise, sg3_utils will not be installed.</p>
Adding or Repairing HDLM Paths	<p>When LifeKeeper brings an HDLM resource into service, it establishes a persistent reservation registered to each path that was active at that time. If new paths are added after the initial reservation, or if failed paths are repaired and HDLM automatically reactivates them, those paths will not be registered as a part of the reservation until the next LifeKeeper quickCheck execution for the HDLM resource. If HDLM allows any writes to that path prior to that point in time, reservation conflicts that occur will be logged to the system message file. The HDLM driver will retry these IOs on the registered path resulting in no observable failures to the application. Once quickCheck registers the path, subsequent writes will be successful. The status will be changed to "Offline(E)" if quickCheck detects a reservation conflict. If the status is "Offline(E)", customers will need to manually change the status to "Online" using the online HDLM command.</p>

		OS version / Architecture									
		RHEL4									
		U1- U4	U3 Security Fix(*7)	U4 Security Fix(*2)	4.5	4.5 Security Fix(*4)	4.6	4.6 Security Fix(*8)	4.7	4.7 Security Fix(*9)	4.8
		x86/x86_64									
HDLM	05-80 05-81 05-90	X									
	05-91 05-92	X		X							
	05-93	X(*3)		X	X						
	05-94	X(*3)		X	X	X	X	X			
	6.0.0	X(*3)		X	X	X	X	X	X	X	X
	6.0.1	X(*3)		X	X	X	X	X	X	X	X
	6.1.0	X(*3)		X	X	X	X	X	X	X	X
	6.1.1	X(*3)	X	X	X	X	X	X	X	X	X
	6.1.2	X(*3)	X	X	X	X	X	X	X	X	X
	6.2.0	X(*3)	X	X	X	X	X	X	X	X	X
	6.2.1	X(*3)	X	X	X	X	X	X	X	X	X
	6.3.0	X(*3)	X	X	X	X	X	X	X	X	X
	6.4.0	X(*3)	X	X	X	X	X	X	X	X	X

LifeKeeper	v6.0											
	v6.0(v6.0.1-2 or later)	X	X	X	X	X	X	X				
	v6.1											
	(v6.1.0-5 or later)	X	X	X	X	X	X	X				
	v6.2											
	(v6.2.0-5 or later)	X	X	X	X	X	X	X				
	v6.2											
	(v6.2.2-1or later)	X	X	X	X	X	X	X				
	v6.3											
	(v6.3.2-1or later)	X	X	X	X	X	X	X				
	v6.4											
	(v6.4.0-10 or later)	X	X	X	X	X	X	X	X	X	X	X
	v7.0											
	(v7.0.0-5 or later)	X	X	X	X	X	X	X	X	X	X	X
V7.1												
(v7.1.0-8 or later)	X	X	X	X	X	X	X	X	X	X	X	
HDLM ARK	6.0.1-2	X	X	X	X	X	X	X				
	6.1.0-4	X	X	X	X	X	X	X				
	6.2.2-3	X	X	X	X	X	X	X				
	6.2.3-1	X	X	X	X	X	X	X				
	6.4.0-2	X	X	X	X	X	X	X	X	X	X	X
	7.0.0-1	X	X	X	X	X	X	X	X	X	X	X
X = supported blank = not supported												

		RHEL5									
		No Updates	5.1	5.1 Security Fix (*5)	5.2	5.2 Security Fix (*6)	5.3	5.3 Security Fix(*10)	5.4	5.4 Security Fix(*11)	5.5
		x86/x86_64									
HDLM	05-80 05-81 05-90	X									
	05-91 05-92										
	05-93	X									
	05-94	X	X								
	6.0.0	X	X	X	X	X					
	6.0.1	X	X	X	X	X					
	6.1.0	X	X	X	X	X					
	6.1.1	X	X	X	X	X					
	6.1.2	X	X	X	X	X	X	X	X	X	X
	6.2.0	X	X	X	X	X	X	X	X	X	X
	6.2.1	X	X	X	X	X	X	X	X	X	X
	6.3.0	X	X	X	X	X	X	X	X	X	X
6.4.0	X	X	X	X	X	X	X	X	X	X	

LifeKeeper	v6.0 (v6.0.1-2 or later)										
	v6.1 (v6.1.0-5 or later)	X	X								
	v6.2 (v6.2.0-5 or later)	X	X								
	v6.2 (v6.2.2-1 or later)	X	X	X							
	v6.3 (v6.3.2-1 or later)	X	X	X	X	X					
	v6.4 (v6.4.0-10 or later)	X	X	X	X	X	X	X	X	X	X
	v7.0 (v7.0.0-5 or later)	X	X	X	X	X	X	X	X	X	X
	v7.1 (v7.1.0-8 or later)	X	X	X	X	X	X	X	X	X	X
	v7.2 (v7.2.0-10 or later)						X	X	X	X	X
HDLM ARK	6.0.1-2										
	6.1.0-4	X	X								
	6.2.2-3	X	X	X							
	6.2.3-1	X	X	X	X	X					
	6.4.0-2	X	X	X	X	X	X	X	X	X	X
	7.0.0-1	X	X	X	X	X	X	X	X	X	X
	7.2.0-1						X	X	X	X	X
X = supported blank = not supported											

Device Mapper Multipath I/O Configurations

Protecting Applications and File Systems That Use Device Mapper Multipath Devices	<p>In order for LifeKeeper to operate with and protect applications or file systems that use Device Mapper Multipath devices, the Device Mapper Multipath (DMMP) Recovery Kit must be installed.</p> <p>Once the DMMP Kit is installed, simply creating an application hierarchy that uses one or more of the multipath device nodes will automatically incorporate the new resource types provided by the DMMP Kit.</p>
Multipath Device Nodes	<p>To use the DMMP Kit, any file systems and raw devices must be mounted or configured on the multipath device nodes rather than on the native <code>/dev/sd*</code> device nodes. The supported multipath device nodes to address the full disk are <code>/dev/dm-#</code>, <code>/dev/mapper/<uuid></code>, <code>/dev/mapper/<user_friendly_name></code> and <code>/dev/mpath/<uuid></code>. To address the partitions of a disk, use the device nodes for each partition created in the <code>/dev/mapper</code> directory.</p>
Use of SCSI-3 Persistent Reservations	<p>The Device Mapper Multipath Recovery Kit uses SCSI-3 persistent reservations with a "Write Exclusive" reservation type. This means that devices reserved by one node in the cluster will remain read-accessible to other nodes in the cluster, but those other nodes will be unable to write to the device. Note that this does not mean that you can expect to be able to mount file systems on those other nodes for ongoing read-only access.</p> <p>LifeKeeper uses the <code>sg_persist</code> utility to issue and monitor persistent reservations. If necessary, LifeKeeper will install the <code>sg_persist(8)</code> utility.</p> <p>SCSI-3 Persistent Reservations must be enabled on a per LUN basis when using EMC Symmetrix (including VMAX) arrays with multipathing software and LifeKeeper. This applies to both DMMP and PowerPath.</p>
Hardware Requirements	<p>The Device Mapper Multipath Kit has been tested by SIOS Technology Corp. with the EMC CLARiiON CX300, the HP EVA 8000, HP MSA1500, HP P2000, the IBM SAN Volume Controller (SVC), the IBM DS8100, the IBM DS6800, the IBM ESS, the DataCore SANsymphony, and the HDS 9980V. Check with your storage vendor to determine their support for Device Mapper Multipath.</p> <p>Enabling support for the use of reservations on the CX300 requires that the hardware handler be notified to honor reservations. Set the following parameter in <code>/etc/multipath.conf</code> for this array:</p> <pre>hardware_handler "3 emc 0 1"</pre> <p>The HP MSA1500 returns a reservation conflict with the default path checker setting (<code>tur</code>). This will cause the standby node to mark all paths as failed. To avoid this condition, set the following parameter in <code>/etc/multipath.conf</code> for this array:</p> <pre>path_checker readsector0</pre> <p>For the HDS 9980V the following settings are required:</p> <ul style="list-style-type: none"> • Host mode: 00

	<ul style="list-style-type: none"> • System option: 254 (must be enabled; global HDS setting affecting all servers) • Device emulation: OPEN-V <p>Refer to the HDS documentation "Suse Linux Device Mapper Multipath for HDS Storage" or "Red Hat Linux Device Mapper Multipath for HDS Storage" v1.15 or later for details on configuring DMMP for HDS. This documentation also provides a compatible multipath.conf file.</p> <p>For the EVA storage with firmware version 6 or higher, DMMP Recovery Kit v6.1.2-3 or later is required. Earlier versions of the DMMP Recovery Kit are supported with the EVA storage with firmware versions prior to version 6.</p>
<p>Multipath Software Requirements</p>	<p>For SUSE, multipath-tools-0.4.5-0.14 or later is required.</p> <p>For Red Hat, device-mapper-multipath-0.4.5-12.0.RHEL4 or later is required.</p> <p>It is advised to run the latest set of multipath tools available from the vendor. The feature content and the stability of this multipath product are improving at a very fast rate.</p>
<p>Linux Distribution Requirements</p>	<p>Some storage vendors such as IBM have not certified DMMP with SLES 11 at this time.</p> <p>SIOS Technology Corp. is currently investigating reported issues with DMMP, SLES 11, and EMCs CLARiiON and Symmetrix arrays.</p>
<p>Transient path failures</p>	<p>While running IO tests on Device Mapper Multipath devices, it is not uncommon for actions on the SAN, for example, a server rebooting, to cause paths to temporarily be reported as failed. In most cases, this will simply cause one path to fail leaving other paths to send IOs down resulting in no observable failures other than a small performance impact. In some cases, multiple paths can be reported as failed leaving no paths working. This can cause an application, such as a file system or database, to see IO errors. There has been much improvement in Device Mapper Multipath and the vendor support to eliminate these failures. However, at times, these can still be seen. To avoid these situations, consider these actions:</p> <ol style="list-style-type: none"> 1. Verify that the multipath configuration is set correctly per the instructions of the disk array vendor. 2. Check the setting of the "failback" feature. This feature determines how quickly a path is reactivated after failing and being repaired. A setting of "immediate" indicates to resume use of a path as soon as it comes back online. A setting of an integer indicates the number of seconds after a path comes back online to resume using it. A setting of 10 to 15 generally provides sufficient settle time to avoid thrashing on the SAN. 3. Check the setting of the "no_path_retry" feature. This feature determines what Device Mapper Multipath should do if all paths fail. We recommend a setting of 10 to 15. This allows some ability to "ride out" temporary events where all paths fail while still providing a reasonable recovery time. The

	<p>LifeKeeper DMMP kit will monitor IOs to the storage and if they are not responded to within four minutes LifeKeeper will switch the resources to the standby server. NOTE: LifeKeeper does not recommend setting "no_path_retry" to "queue" since this will result in IOs that are not easily killed. The only mechanism found to kill them is on newer versions of DM, the settings of the device can be changed:</p> <pre data-bbox="565 451 1344 478">/sbin/dmsetup message -u 'DMid' 0 fail_if_no_path</pre> <p>This will temporarily change the setting for no_path_retry to fail causing any outstanding IOs to fail. However, multipathd can reset no_path_retry to the default at times. When the setting is changed to fail_if_no_path to flush failed IOs, it should then be reset to its default prior to accessing the device (manually or via LifeKeeper).</p> <p>If "no_path_retry" is set to "queue" and a failure occurs, LifeKeeper will switch the resources over to the standby server. However, LifeKeeper will not kill these IOs. The recommended method to clear these IOs is through a reboot but can also be done by an administrator using the dmsetup command above. If the IOs are not cleared, then data corruption can occur if/when the resources are taken out of service on the other server thereby releasing the locks and allowing the "old" IOs to be issued.</p>
--	---

