



LifeKeeper[®] for Linux in VMware ESX Server Virtual Machines Configuration Guidelines

April 2009

SteelEye, SteelEye Technology, and LifeKeeper are registered trademarks of SteelEye Technology, Inc. Other brand and product names used herein are for identification purposes only and may be trademarks of their respective companies.

This document is for informational purposes and is believed to be correct at the time of publication. However, SteelEye Technology does not guarantee the accuracy of the information and reserves the right to change the document at any time. SteelEye Technology makes no warranties, expressed or implied, in this document.

Copyright © 2009
SteelEye Technology, Inc.
Palo Alto, CA U.S.A.
All Rights Reserved

VMware is a virtual machine technology implemented entirely in software, allowing it to be installed and used on nearly any Intel or AMD processor-based system that meets the minimum hardware requirements. LifeKeeper for Linux is supported in VMware ESX Server virtual machines, with ESX Server versions 2.5, 3.0 and 3.5. Additional information about VMware ESX Server can be found on the VMware web site at http://www.vmware.com/products/server/esx_features.html.

LifeKeeper for Linux running in a VMware ESX Server virtual machine is supported with all of the normal LifeKeeper cluster storage options, including shared (SAN) storage, NAS, and data replication. All of these storage options can be used in virtual-to-virtual, virtual-to-physical, and physical-to-virtual failover configurations.

The configuration of NAS and data replication in a VMware ESX Server environment is exactly the same as in a physical machine environment. The configuration of shared (SAN) storage, however, does require some special consideration in a LifeKeeper cluster environment due to the way that VMware ESX Server manages and presents storage controllers and devices to the virtual machines.

In order for LifeKeeper to have the level of access and control that it requires for shared storage devices, all SAN storage that is to be placed under LifeKeeper control in a VMware ESX Server environment should be configured using a technique that VMware calls Raw Device Mapping. Additionally, in virtual-to-virtual configurations with shared storage devices, each VM in the LifeKeeper cluster must reside on a separate ESX Server. Using Raw Device Mapping is an absolute requirement with ESX Server v3.x, as it is the only means of accessing shared storage in a clustered environment that is supported by VMware. With ESX Server v2.5, there is another alternative that will be described later, but raw device mapping is still the recommended approach.

For a VMware ESX Server v2.5 virtual machine, use the following procedure to configure raw device mapping for a shared disk or LUN:

1. The virtual machine should be in the powered off state.
2. In the VMware management interface, click on the virtual machine's Display Name in the main Status Monitor window, to bring up the virtual machine configuration page.
3. Click on the **Hardware** tab.
4. Click **Add Device...** at the bottom of the window.
5. In the Add Device window, click **Hard Disk**.
6. Click **System LUN/Disk** for the type of virtual disk to be added.
7. Choose the appropriate shared storage device from the Target LUN/Disk list.
8. Make sure **Use Metadata** is selected, and choose the desired Metadata File Location.
9. Enter an appropriate Metadata File Name.
10. Choose a virtual SCSI identifier from the Virtual SCSI Node list.
11. Select the **Physical** Compatibility mode.
12. Click **OK**.

For a VMware ESX Server v3.0 virtual machine, use the following procedure to configure raw device mapping for a shared disk or LUN:

1. The virtual machine should be in the powered off state.
2. In the Virtual Infrastructure Client interface, right-click on the virtual machine and select **Edit Settings...** from the menu.
3. On the Hardware tab, click the **Add...** button.
4. Select **Hard Disk** and click the **Next >** button.
5. Select **Mapped SAN LUN** and click the **Next >** button.

6. Select the appropriate shared storage device from the LUN list and click the **Next >** button.
7. For the datastore location, **Store with Virtual Machine** is the recommended selection. If instead you choose **Specify datastore**, select an appropriate datastore location. Click the **Next >** button.
8. Select the **Physical** compatibility mode and click the **Next >** button.
9. Choose a Virtual Device Node from the Node list and click the **Next >** button.
10. Click **Finish**.
11. Click **OK**.

For a VMware ESX Server v3.5 virtual machine, use the following procedure to configure raw device mapping for a shared disk or LUN:

1. The virtual machine should be in the powered off state.
2. In the Virtual Infrastructure Client interface, right-click on the virtual machine and select **Edit Settings...** from the menu.
3. On the Hardware tab, click the **Add...** button.
4. Select **Hard Disk** and click the **Next >** button.
5. Select **Raw Device Mappings** and click the **Next >** button
6. Select the appropriate shared storage device from the **Adapter:Target:LUN** list and click the **Next >** button.
7. Select the datastore location, **Store with Virtual Machine** is the recommended selection. If instead you choose **Specify datastore**, select an appropriate datastore location. Click the **Next >** button.
8. Select the **Physical** compatibility mode and click the **Next >** button.
9. Choose a Virtual Device node from the Node list and click the **Next >** button. Note: select a device node that resides on a different bus from the VM's root disk.
10. Click **Finish**.
11. Click **OK**.
12. Reselect **Edit Settings...** from the menu.
13. Select the **SCSI Controller** for the Virtual Device node selected in step 9 and set the SCSI Bus Sharing policy to either virtual or Physical depending on the configuration (virtual is bus shared between virtual machines, physical if bus is shared virtual-to-physical).
14. Click **OK**.

Following the appropriate procedure above will provide LifeKeeper for Linux running in the virtual machine with direct access to the shared storage device, just as it would appear on a physical machine. This enables LifeKeeper to perform the necessary identification and management of the device in the clustered environment.

With VMware ESX Server v2.5, it is also possible to configure a shared storage device that is to be managed by LifeKeeper as a VMware virtual disk device. This configuration has the limitation, however, that the shared storage device can be used by LifeKeeper only in a virtual-to-virtual failover configuration. A shared storage device configured in this manner cannot be used in a virtual-to-physical or physical-to-virtual configuration. To configure a shared storage device in this manner, the following guidelines and configuration steps must be observed:

1. The SCSI controllers and virtual disks must be configured identically on each virtual machine to ensure that every virtual machine sees the same storage devices and with the same device names.

2. The Bus Sharing mode for the SCSI controller used to access the shared storage must be set to **Physical**.
3. The VMFS-2 file system on the shared storage device must be configured for **Shared** Access Mode.
4. After Linux has been installed and configured in the virtual machines, you should carefully verify that each virtual machine can see the same set of shared storage devices and with the same device names.

We strongly recommend the use of the Linux md software RAID facility with any shared storage devices that are to be protected by LifeKeeper, as further protection against the initial misconfiguration of shared storage devices, as well as to prevent problems due to later configuration changes. By simply configuring each shared storage device into an md linear, raid0, or raid1 device, either separately or in combination with other devices, md will ensure through its own unique ID stamps that a given storage device is the correct one when LifeKeeper attempts to configure it on a different virtual machine during a failover operation.

5. The following line must be added to the LifeKeeper configuration file, `/etc/default/LifeKeeper`, on each virtual machine:

```
LK_NU_EXCEPTIONS=vmware
```

This entry is required in order for LifeKeeper to consider the external virtual disks now seen by the virtual machines as shared.

6. Remember when configuring your LifeKeeper hierarchies that a hierarchy which includes one of these shared virtual disks can only be extended across other VMware virtual machines, and not to a physical machine.