



SteelEye Protection Suite for Linux: PostgreSQL with Shared Storage

Evaluation Guide

This document and the information herein is the property of SIOS Technology Corp. Any unauthorized use and reproduction is prohibited. SIOS Technology Corp. makes no warranties with respect to the contents of this document. SIOS Technology Corp. reserves the right to revise this publication and make changes to the products described herein without prior notification.

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

DB2 Universal Database and IBM Director are trademarks of International Business Machines Corporation.

IBM is a registered trademark of International Business Machines Corporation in the U.S. and other countries.

Microsoft, Windows, Windows Server 2008, Windows Server 2003, Windows 2000, Windows NT, SQL Server and Exchange are trademarks or registered trademarks of Microsoft in the U.S. and other countries.

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

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.

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

Email correspondence to:

evalsupport@us.sios.com

Copyright © 2007-2012

By SIOS Technology Corp.

San Mateo, CA U.S.A.

All Rights Reserved Worldwide

© by SIOS Technology Corp. All rights reserved worldwide.

SteelEye Protection Suite for Linux: Evaluation Guide

Table of Contents

Objective	5
Terms to Know.....	5
The Evaluation Process.....	7
Phase 1 – Prepare to Install.....	8
Hardware Requirements	8
Primary and Secondary Servers.....	8
Client.....	8
Software Requirements.....	8
Primary Server and Secondary Server	8
Network Requirements	10
Primary Server and Secondary Servers.....	11
Client.....	11
Phase 2 – Configure Storage.....	11
Before You Begin	11
Configure iSCSI initiator, discover and login to iSCSI target	11
Phase 3 – Install, Configure, and Start PostgreSQL	13
Primary Server	13
Secondary Server.....	14
Phase 4 – Install SteelEye Protection Suite for Linux.....	15
Download Software	15
Run the SteelEye Protection Suite Installer Script.....	15
Install the Evaluation License Keys	16
Start the SteelEye Protection Suite for Linux	17
Phase 5 - Configure the Cluster	17
Primary Server	17

Access the LifeKeeper GUI	17
Create Communication (Comm) Paths	19
Verify the Communications Paths	24
Create the LifeKeeper Hierarchy	25
Create and Extend an IP Resource.....	25
Create the Shared Filesystem Resource Hierarchy.....	29
Create the PostgreSQLSQL Resource Hierarchy	30
Create the PostgreSQL IP Address Dependency.....	32
Phase 6 –Test Your Environment.....	32
1. Manual Switchover of the PostgreSQL Hierarchy to Secondary Server	33
2. Manual Switchover of the PostgreSQL Hierarchy back to Primary Server	34
3. Simulate a network failure on the Primary Server by failing the IP resource.....	35
4. Hard failover of the resource from the Secondary Server back to the Primary Server	36
5. Bring Failed Server back on line.....	36
6. Verify Local Recovery of PostgreSQL Server	37

Objective

This document is intended to aid you in installing, configuring and using the SteelEye Protection Suite for Linux evaluation product, to make PostgreSQL highly available. If PostgreSQL is not already installed, please allocate some time to install it on your servers. Once this task has been completed, you may install and configure SteelEye Protection Suite for Linux.

There are five phases in this process:

- Phase 1 – Prepare to Install
- Phase 2 – Configure Storage
- Phase 3 – Install and Configure PostgreSQL
- Phase 4 – Install SteelEye Protection Suite for Linux
- Phase 5 – Configure your LifeKeeper Cluster
- Phase 6 – Test Your Environment

Terms to Know

The following terms are used throughout this document and, while some may be familiar to you, it may be helpful to review how SteelEye defines and uses these terms.

Network Communication Terms

Crossover cable – A cable used to directly connect computing devices together, instead of being connected to a network switch, hub or router. This cable creates an isolated, private network to allow cluster-related and data replication traffic to flow between systems.

Types of LifeKeeper Servers

Server – A computer system dedicated to running software application programs.

Active Server – This is the server where the resource hierarchy is currently running (IN SERVICE).

Standby Server – This is the server where the resource hierarchy is defined, but is not currently running. This server is available to bring the resource hierarchy into service should something happen to the resource hierarchy on the Active Server.

Primary Server – This is the server in a LifeKeeper configuration with the highest priority for a given resource hierarchy. It is the server that provides services for the resource hierarchy under normal circumstances.

Secondary Server – This is the server in a LifeKeeper configuration with the 2nd highest priority for a given resource hierarchy.

Source Server - In a LifeKeeper cluster, using data replication, this is the Active Server. It is where the resource hierarchy is currently running and the replicated partition (Source Partition) is accessible for writes.

Target Server – In a LifeKeeper cluster, using data replication, this is the Standby Server. The replicated partition (Target Partition) is updated with writes from the Source Partition by the SteelEye Data Replication system. This partition should not be accessed/modified manually.

SteelEye Data Replication Terms

Replication – Transferring data from one partition to another via a sector-by-sector copy. During replication, the target partition should not be accessed or modified assuring your data integrity.

Synchronous – A replication scheme in which the data is confirmed written and valid on the target before the write operation occurs on the source disk through a series of information exchanges. Synchronous mirrors should only be implemented on high speed (100Mbps+) networks due to the network overhead involved.

Asynchronous – A replication scheme in which the data is released for writing on the source immediately and is sent to the target(s) simultaneously for writing as fast as the data can get there and can be written on them.

Rate of Change – A measure of the amount of data which is changing over a set period of time.

Compression – An algorithm which is optionally implemented to reduce the amount of traffic between source and target nodes. Nine levels of compression are offered. Compression is turned off by default.

Throttling – An optionally implemented mechanism to limit the bandwidth used for replication.

LifeKeeper Product Terms

Communications Path – A mechanism supporting communication between nodes in a LifeKeeper cluster. SIOS highly recommends implementing multiple communication paths between all servers in the cluster to eliminate a single point of failure.

Heartbeat – A periodic message exchanged between nodes in a LifeKeeper cluster that provides server health monitoring. A heartbeat message is one type of inter-node cluster communication sent over a communications path.

Split Brain – A situation in which all communications paths between cluster members fail, but all servers remain up and running. In this situation, both systems believe the other has failed and both believe they should keep or bring resources into service.

Failover – The unplanned migration of a resource hierarchy to the Standby Server because of a system or resource failure on the Active Server.

Switchover – The planned migration of a resource hierarchy from the Active Server to the Standby Server.

Switchback – The setting that governs the recovery behavior of the server where the resource was in service when it failed. If the setting is “intelligent”, the server acts as a possible backup for the given

resource. If the setting is “Automatic”, the server actively attempts to re-acquire the resource without further notice.

Resource – A system asset that can be protected by LifeKeeper. Resources can be used to represent disk partitions, virtual IP addresses, applications, etc.

Extend a Resource – Create or define an already configured LifeKeeper resource onto another server in the cluster and build an equivalency relationship that prevents the resource from coming in service on both systems simultaneously.

Resource Hierarchy – A grouping of resources, in a predetermined order, from high to low. This may also be referred to as simply a Hierarchy.

Shared Storage – One or more logical disk partitions that are physically attached to all nodes in a cluster. LifeKeeper ensures that the volume is only accessible by one server at a time. This is formally called I/O fencing.

Data Replication (Disk Mirroring) – The replication of logical disk partitions to separate physical hard disks in real time to ensure continuous availability, currency and accuracy of data.

Source – The partition on the source server used for replication. The “gold” copy of the data.

Target – The partition on the target server used for replication.

Switchable IP Address – A unique IP address that may be moved between systems in the cluster. Client systems connect to this address and the system where the virtual IP resource is active will respond to requests.

The Evaluation Process

SIOS strongly recommends performing your evaluation of SteelEye Protection Suite for Linux within a test lab environment. SIOS is not responsible and cannot provide support for evaluation software installed in a production environment.

All questions during the evaluation period should be directed to evalsupport@us.sios.com or your local Pre-Sales Engineering contact. Pre-sales support will contact you by the next business day to answer questions. Once you are a licensed customer on software maintenance, you will have access to 24 X 7 post-sales technical support.

Important

Your evaluation license is valid for a limited period of time from the day you receive the SIOS product evaluation package and licenses via email from the SIOS sales team.

Phase 1 – Prepare to Install

Hardware Requirements

Primary and Secondary Servers

- Systems must meet the minimum requirements for the Linux distribution to be used during the evaluation.
- 512MB RAM minimum; 1GB RAM recommended.
- 2GB of available hard disk space recommended.
- Multiple Network Interface Cards (NIC's) are recommended.
- If not using shared storage, configure one or more additional partitions to be used for data replication. On the primary server, these will become the source partitions. On the secondary server(s), these will become the target partitions.
- For replicated partitions, a target partition's size must equal to or larger than the size of its source partition.
- The system (/) and boot (/boot) partitions are not eligible for replication.

Note: You may use more than one partition for replicated data, allowing for separation of multiple database files and/or log directories for performance reasons.

Client

This system is not required but is recommended for testing the cluster environment.

- A Linux terminal using the “psql” command will be used to validate postgresql functionality.

Software Requirements

Primary Server and Secondary Server

- Linux Distribution:
 - RedHat Enterprise Linux 5 (5.4+ recommended) or 6
 - CentOS Linux 5 (5.4+ recommended) or 6
 - Oracle Enterprise Linux 5 (5.4+ recommended)
 - SuSE Linux Enterprise Server 10 or 11 (11 recommended)
 - See <http://docs.us.sios.com/Linux/8.1/LK4L/ReleaseNotes/> for a full list of supported Operating Systems
- Current patches / security updates are recommended.
- On 64-bit Linux systems, ensure that the appropriate 32bit compatibility packages are installed, for example compat-libstdc++

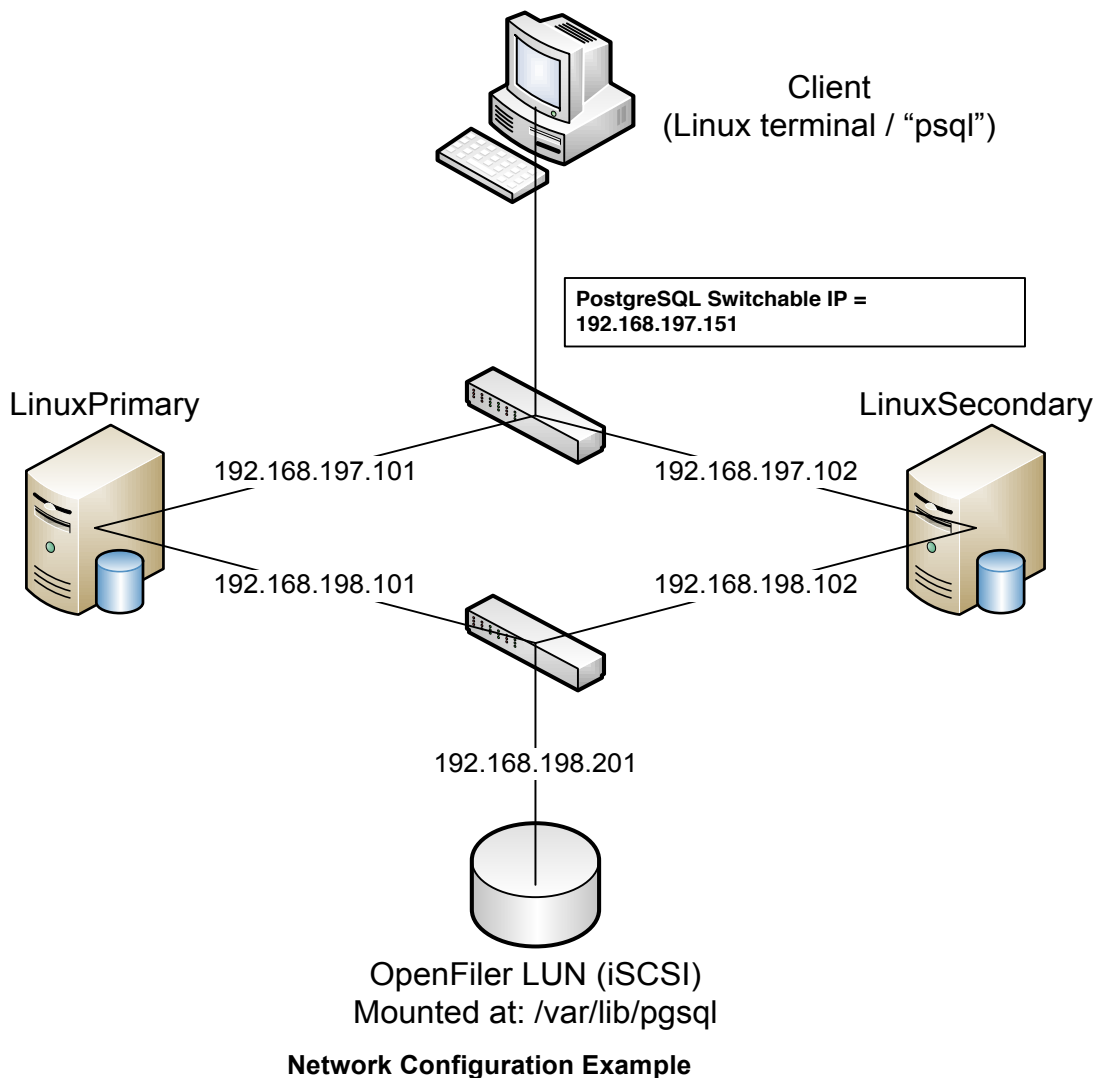
- Its recommended that IPTables is disabled
 - # /etc/init.d/iptables off
 - # chkconfig iptables off
 - See http://docs.us.sios.com/Linux/8.1/LK4L/TechDoc/Content/user_guide/using_lifekeeper_linux/maintenance_tasks/running_lifekeeper_firewall.htm for information regarding the ports SteelEye Protection Suite for Linux uses.
- It's recommended that SELinux is disabled:
 - Edit /etc/selinux/config
 - Set SELINUX=disabled (note: permissive mode is also acceptable)
- Check the configuration of your /etc/hosts file
 - localhost.localdomain and localhost are the only entries that can be on 127.0.0.1
 - create a separate entry for your hostname with a static address

Network Requirements

For your evaluation, we recommend configuring your machines similarly to the following example.

LinuxPrimary and LinuxSecondary are multi-homed, between two LAN segments (the second NIC in each server could even be connected via a cross-over cable if a second physical network is not available). The second NIC is optional in this configuration, but highly recommended in production environments to avoid a single point of failure.

In this example evaluation scenario we will be leveraging Shared (iSCSI) Storage with our PostgreSQL database. OpenFiler is a storage appliance server which will serve an iSCSI target to LinuxPrimary and LinuxSecondary.



Primary Server and Secondary Servers

- Configure the Host file with entries for all LifeKeeper protected servers. This is typically /etc/hosts.

Example:

```
192.168.197.101          LinuxPrimary
192.168.197.102          LinuxSecondary
```

- See your Network Administrator to obtain an unused IP Address to be used as the switchable IP Address. This switchable IP Address will be created later in the configuration process.
- Public Network connection(s) configured with:
 - Static IP address
 - Correct subnet mask
 - Correct gateway address
 - Correct DNS server address(es)
- Private Network connection(s) configured with:
 - Static IP address (on a different subnet from the public network)
 - Correct network mask
 - No gateway IP address
 - No DNS server addresses

Client

Must be able to communicate on the same subnet/network as the servers Public interface addresses. In our example, this is the 192.168.197.0/24 network.

Phase 2 – Configure Storage

Before You Begin

Ensure the following:

- If planning to use replicated storage, have an extra volume/partition on both servers that can be used for data replication. A target volume's size must equal to or larger than the size of its source volume.
- If planning to use shared storage, as in this example, ensure the Shared storage is configured and accessible to your cluster nodes. This can either be Fiber Channel SAN, iSCSI, NAS, etc. In this example we will review configuration of an iSCSI target for use as our PostgreSQL database storage repository.

Configure iSCSI initiator, discover and login to iSCSI target

This Evaluation guide will not cover how to setup an iSCSI Target Server. It is assumed that the shared storage already exists in your environment. If you don't have shared storage and wish to configure it, a

simple solution is to use OpenFiler (<http://www.openfiler.com/>), an Open Source storage management appliance, which can be run on physical hardware or as a virtual machine.

On both Primary and Secondary servers, perform the following functions:

1. If not already installed, ensure that the **iscsi-initiator-utils** rpm package is installed:

```
# yum install iscsi-initiator-utils
```

2. Start the iscsid service and enable it to automatically start when the system boots

```
# service iscsid start  
# chkconfig iscsid on
```

3. Configure the iscsi service to automatically start, which logs into iSCSI targets needed at system start up.

```
# chkconfig iscsi on
```

4. Use the iscsiadm command to discover all available targets on the network storage server (OpenFiler)

```
# iscsiadm -m discovery -t sendtargets -p <name or IP of iSCSI server>
```

```
[root@LinuxPrimary init.d]# iscsiadm -m discovery -t sendtargets -p 192.168.198.201
```

Example **iqn.2006-01.com.openfiler:tsn.postgres**

5. Manually Login to the iSCSI Target

```
# iscsiadm -m node -T iqn.2006-01.com.openfiler:tsn.postgres -p 192.168.198.201 -  
login
```

6. Configure Automatic Login

```
# iscsiadm -m node -T iqn.2006-01.com.openfiler:tsn.postgres -p 192.168.198.201 --op  
update -n node.startup -v automatic
```

7. Use the “fdisk” command to format your iSCSI LUN, if needed

```
# fdisk /dev/sdc
```

8. Create a filesystem on your new iSCSI LUN Partition, sdc1

```
# mkfs.ext3 /dev/sdc1
```

9. Mount your iSCSI LUN at /var/lib/pgsql (assuming a default postgres configuration). If data already exists in this directory, make sure to move it into the shared iSCSI LUN

```
# mount /dev/sdc1 /var/lib/pgsql
```

10. At this point you now have an iSCSI shared LUN, /dev/sdc1, mounted at /var/lib/pgsql. Our disk layout now look as follows (example):

```
[root@LinuxPrimary pgsql]# df
Filesystem      1K-blocks  Used Available Use% Mounted on
/dev/sda2       25967432 3683016 1976400 66% /
Example /dev/sda1     101086   24659   71208 26% /boot
tmpfs           517552    0 517552 0% /dev/shm
/dev/sdc1       966644   38944  878596 5% /var/lib/pgsql
```

Phase 3 – Install, Configure, and Start PostgreSQL

Primary Server

On your Primary server, perform the following actions:

1. Install both the “postgresql-server” and “postgresql” rpm packages if they do not exist on your system. Apply any required dependencies as well

```
# yum install postgresql postgresql-server
```
2. Verify that your Shared iSCSI LUN is still mounted at /var/lib/pgsql via the “df” command
3. If this is a fresh PostgreSQL install, initialize a sample PostgreSQL database:

```
# # su - postgres
# initdb --pgdata=/var/lib/pgsql/data
```
4. Ensure that all files in your PostgreSQL data directory (/var/lib/pgsql) have correct permissions and ownership

```
# chown -R postgres:postgres /var/lib/pgsql
# chmod 755 /var/lib/pgsql
```
5. Finally, manually start the PostgreSQL daemon from the command line. Note: **Do Not** start it via the “service” command, or the /etc/init.d/ scripts

```
# su - postgres  
# pg_ctl start -D /var/lib/pgsql/data -l /var/lib/pgsql/pgstartup.log -o "-p 5432" -w
```

6. Verify PostgreSQL is running by connecting with the psql client (ensure you are still running as the "postgres" linux user):

```
-bash-3.2$ psql  
Welcome to psql 8.1.22, the PostgreSQL interactive terminal.
```

```
Type: \copyright for distribution terms  
      \h for help with SQL commands  
      \? for help with psql commands  
      \g or terminate with semicolon to execute query  
      \q to quit
```

```
postgres=# \q  
-bash-3.2$
```

Secondary Server

On your Secondary Server:

1. Install both the "postgresql" and "postgresql-server" rpm packages if they do not exist on your system. Apply any required dependencies as well
yum install postgresql postgresql-server
2. Ensure that the PostgreSQL data directory (/var/lib/pgsql) has correct permissions and ownership
chown -R postgres:postgres /var/lib/pgsql
chmod 755 /var/lib/pgsql
3. There is no need to perform any of the additional steps taken on the Primary Server

Phase 4 – Install SteelEye Protection Suite for Linux

For the ease of installation, SIOS has provided the SteelEye Protection Suite for Linux with an installation script. After the installation script has been run, the optional Application Recovery Kits (ARKs) should be installed. ARKs are provided as standard RPMs. Instructions for how to obtain the software is described in an email with your evaluation license keys. The software will be installed to the following locations:

SPS for Linux Component	Install Location
LifeKeeper Software	/opt/LifeKeeper
LifeKeeper Config File	/etc/default/LifeKeeper

Perform the following actions on **both** Primary and Secondary server.

Download Software

1. Open the SteelEye Protection Suite evaluation email you received from SIOS.
2. Download the SteelEye Protection Suite Software from the link provided in your email. It is generally easiest to use “wget” to recursively download all files. Example:
 - a. # cd /root
 - b. # wget -r <URL>
 - c. After successful download you will have downloaded contents similar to the follow directory listing:

```
[root@LinuxPrimary ~]# ls -l <directory>
total 63680
-rw-r--r-- 1 root root 23163 May 30 14:03 EULA.pdf
-rw-r--r-- 1 root root 536 May 30 14:03 readme.txt
-rw-r--r-- 1 root root 65179648 May 30 14:03 sps.img
```
3. Download your Evaluation license key from the link specified in your evaluation email. Save the license file to an easy to remember location on both servers.

Run the SteelEye Protection Suite Installer Script

1. Loopback mount the sps.img file previously downloaded, which is an ISO9660 image file. Run the “setup” script inside:

```
# mount -o loop sps.img /mnt
# cd /mnt
# ./setup
```

```
Welcome to the SteelEye Protection Suite for Linux v8.0.0-5104 Setup

A log of this setup will be written to: /var/log/LK_install.log

This script will set up and install SPS for Linux on your system.

Do you wish to continue (answering "no" will abort setup) (y/n) [y] ? █
```

2. During this procedure, you will hit Enter in most cases to accept default values and continue to the next screen. Note the following exceptions:
 - a. On the screen titled "High AvailabilityNFS" you may select "n" as in this particular eval guide we will not be creating a highly available NFS server cluster configuration.
 - b. If you have plans to create a highly available NFS service, adjust your response accordingly.
3. Towards the end of the setup script, you can choose to install a trial license key now, or later. We will install the license key in the next step, so you can safely select "n" at this point
4. In the final screen of the "setup" select the ARKs you wish to install from the list displayed on the screen.
5. The following RPMs should be installed:
 - a. steeleye-lkPGSQL-<version>.noarch.rpm
6. Un-mount the Distribution Enabling disk image:

```
# cd /root
# umount /mnt
```

Install the Evaluation License Keys

The last phase of the setup process installs the licensing keys. You must install the evaluation license key file (".lic") that you downloaded with your evaluation software before starting the SteelEye Protection Suite for Linux.

1. To install your trial license key, run the “lkkeyins” command on both Primary and Secondary Server. This command is located at /opt/LifeKeeper/bin/lkkeyins . Example:

```
# opt/LifeKeeper/bin/lkkeyins <path_to_license/<filename>.lic
```

2. Validate your license keys were installed via the /opt/LifeKeeper/bin/lklicmgr command

```
# /opt/LifeKeeper/bin/lklicmgr
License File: 20101230.lic
Product          Type   Expiry          Other
LifeKeeper for Linux      Eval   27 Mar 2011 (87 days)
Apache Recovery Kit      Eval   27 Mar 2011 (87 days)
SteelEye Data Replication ARK Eval   27 Mar 2011 (87 days)
MySQL Recovery Kit       Eval   27 Mar 2011 (87 days)
...
```

Start the SteelEye Protection Suite for Linux

1. Start:

```
# /opt/LifeKeeper/bin/lkstart
```

Phase 5 - Configure the Cluster

Primary Server

Complete the following steps on the primary server to configure the cluster:

- Create TCP Communication (Comm) Path(s)
- Verify the Communication (Comm) Path(s)

Before you begin, SIOS recommends at least two TCP communications paths between each server within the cluster to each remote server for heartbeat redundancy.

Important


Supported configurations require that you define redundant comm paths, so that the failure of a single communication line will not cause a split brain where resource hierarchies may come in-service on multiple servers simultaneously.

Access the LifeKeeper GUI

The LifeKeeper Graphical User Interface (GUI) is a Java based application that can be run as a native Linux application, or as an applet within your Java-Enabled Web Browser.

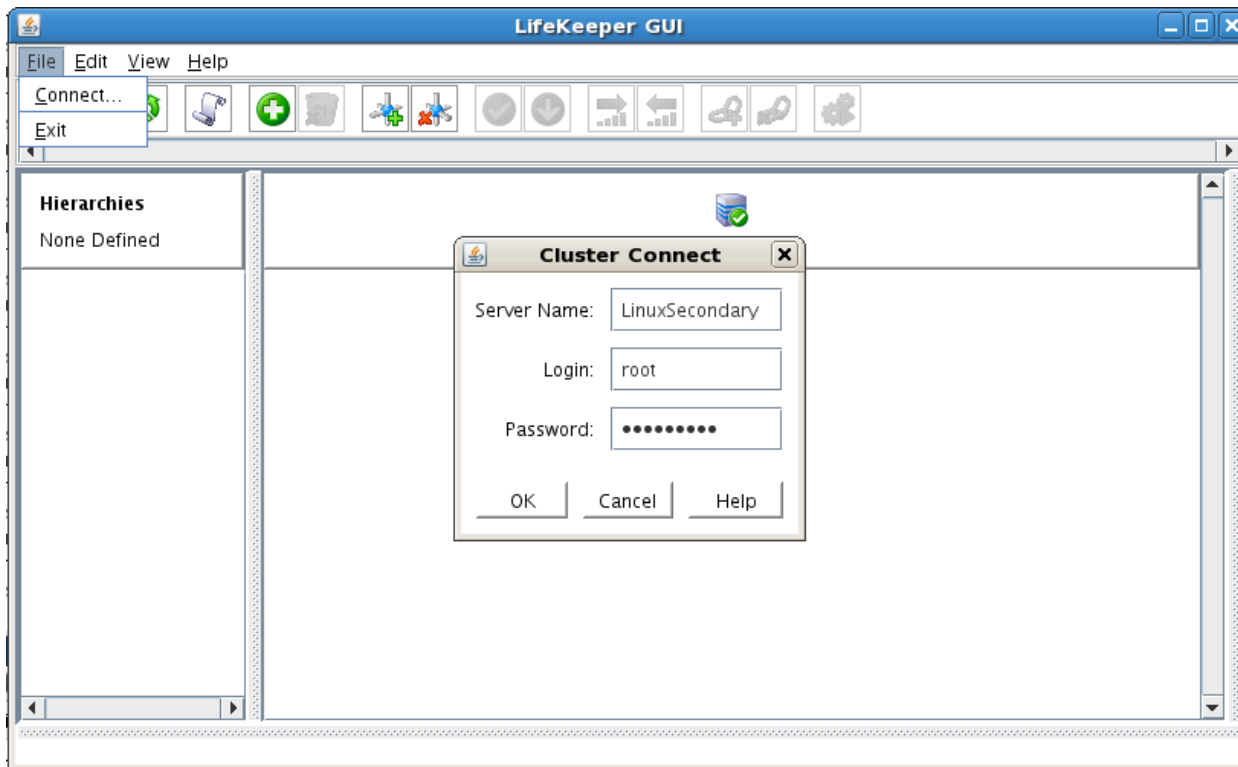
The LifeKeeper GUI is based on Java RMI with callbacks. Hostnames must be resolvable or you may receive a Java 115 or 116 error.

1. To start the LifeKeeper Linux GUI Application:
 - a. `/opt/LifeKeeper/bin/lkGUIapp &`
2. To Connect to the LifeKeeper GUI Applet from a Web Browser, go to:
 - a. <http://<hostname>:81>
3. Enter the name of the server you wish to connect to (this field will be populated with the name of the server you are on, if you are running the GUI from a server with LifeKeeper installed) along with your root credentials and click OK.

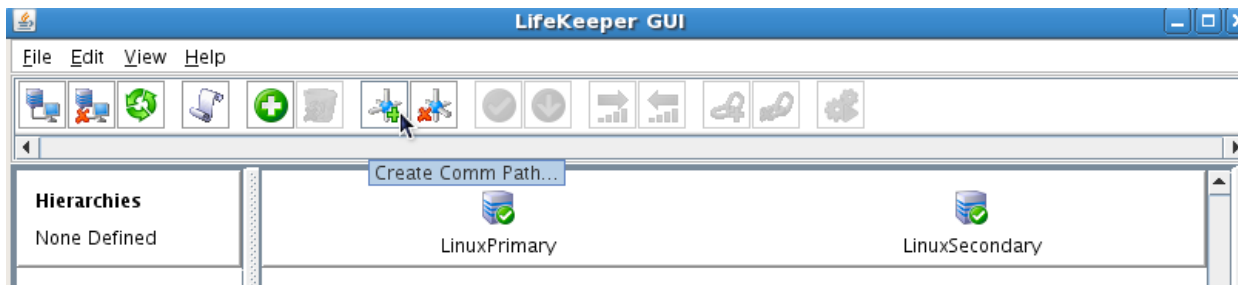


Create Communication (Comm) Paths


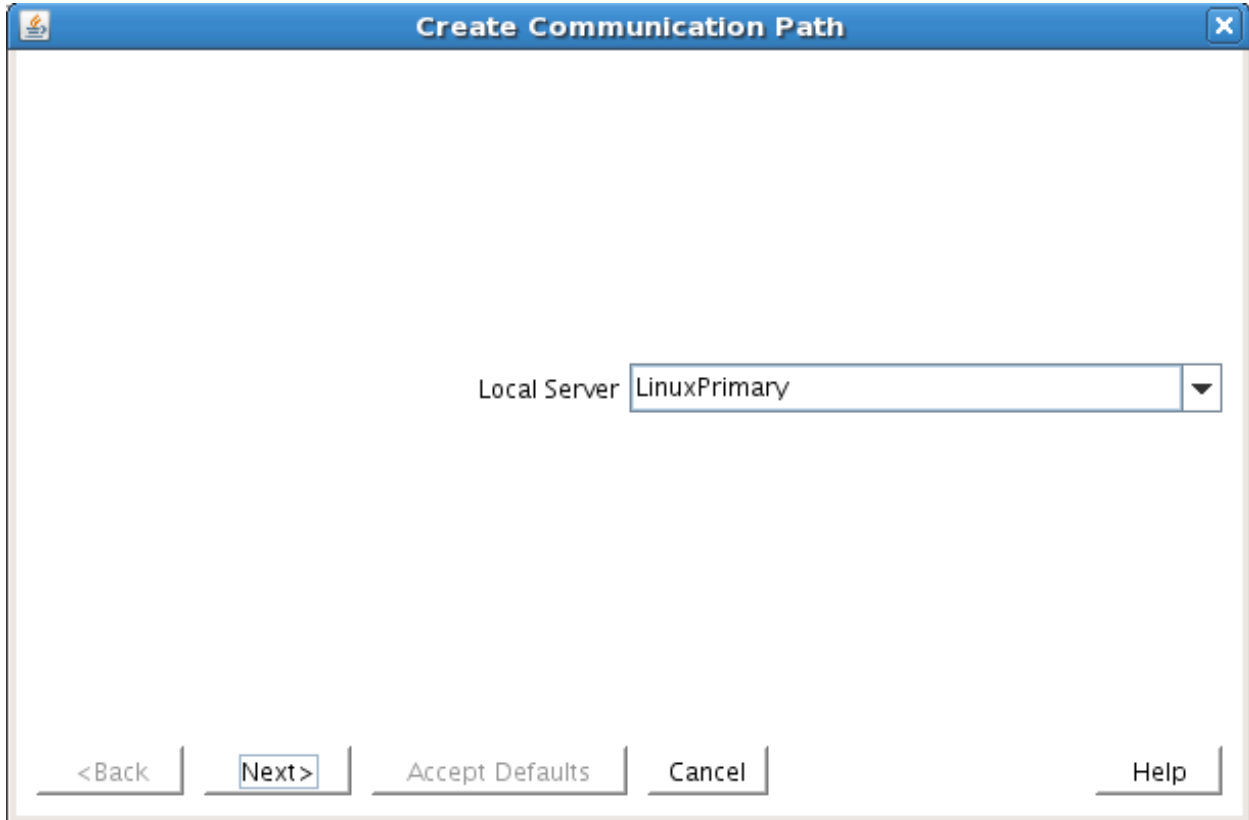
4. Within the LifeKeeper GUI, from the File menu, select Connect. Enter the name of your Secondary server, login and password when the Cluster Connect window displays.



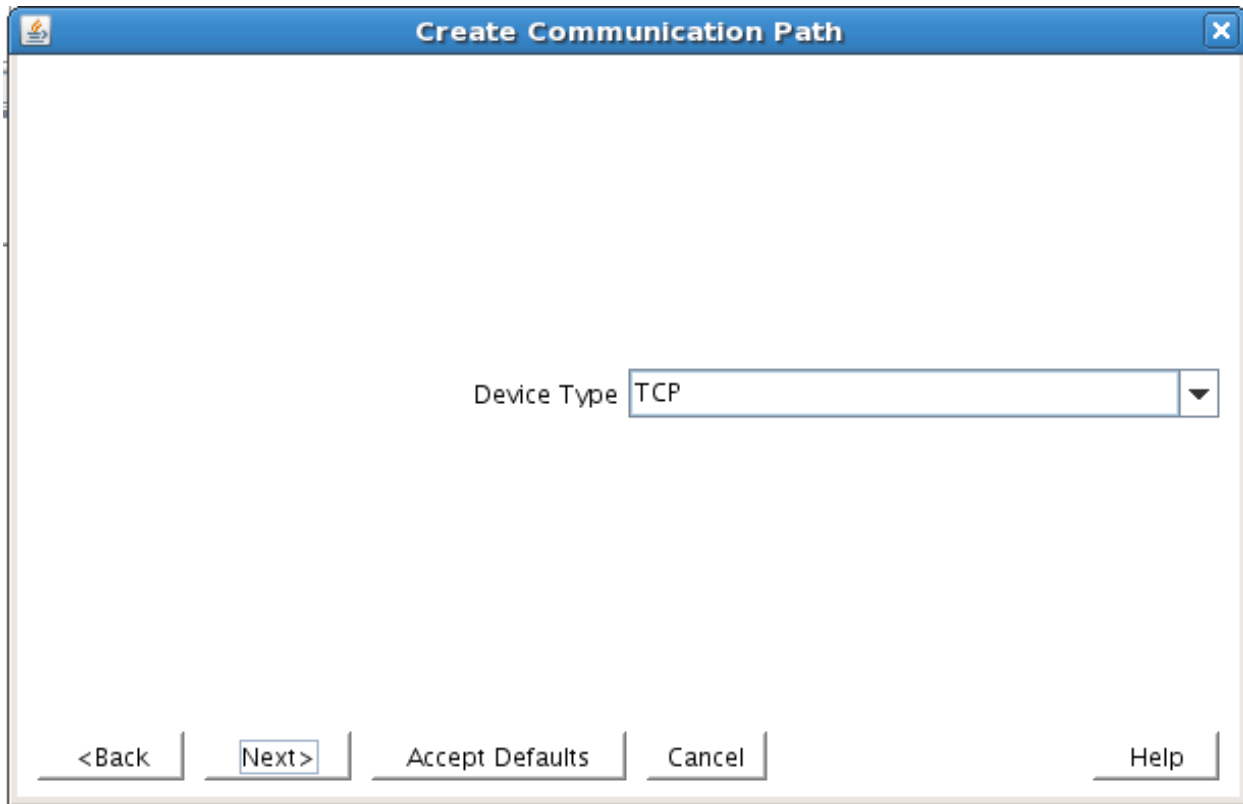
5. Within the LifeKeeper GUI, click the Create Comm Path button on the toolbar. You can also right click one of the servers and click Create Comm Path from the pop-up menu as well.



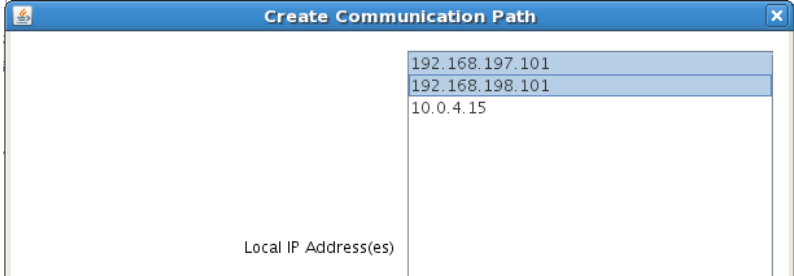
6. Select your Local and Remote Server(s) from the list box. If a server does not appear in the list box, you may enter it by typing its name and clicking the Add Server button. When using the Add Server procedure, you must make sure that the computer names for both network interfaces on the servers respond correctly when you ping them (from all of the partner server(s)) using the **ping -a IP ADDRESS** syntax. If they do not, this must be corrected prior to continuing. Click Next.

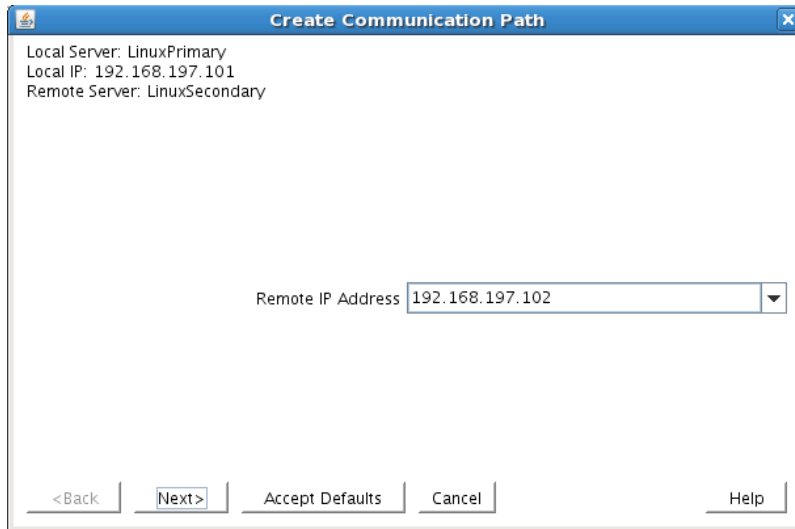


7. Select TCP for Device Type and Click Next.



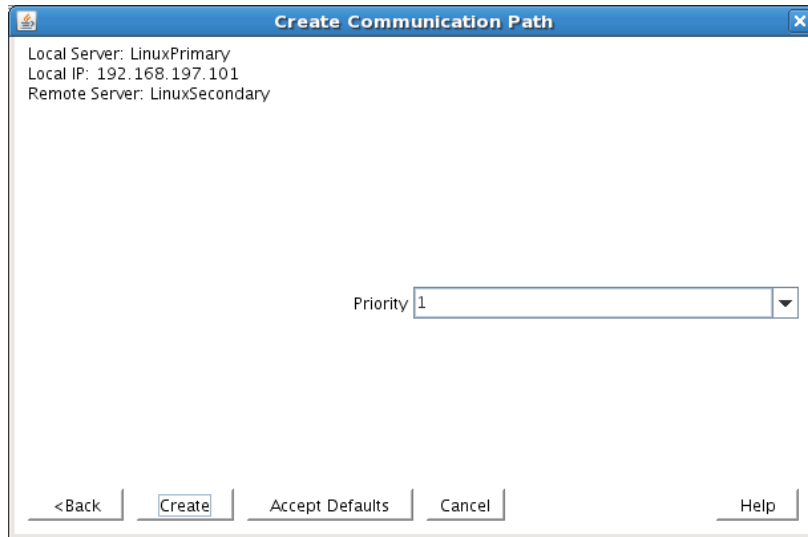
8. Provide all the required information and click Next for the following series of dialog boxes. For each field in the dialog box you can click Help for further information or refer to the table below for an explanation or recommendation.

Field	Tips
<i>For TCP/IP Comm Path...</i>	Choose the IP address to be used by the local server for this comm path. Select both interfaces so that the wizard creates multiple comm. paths during this operation
Local IP Address	
Remote IP Address	Choose the IP address to be used by the remote server for this comm path




Enter the priority for the comm path on the local server. The priority will be used to determine the order that the comm paths between two servers will be used. Priority 1 is the highest; 99 is the lowest.

Priority



9. After entering data in all the required fields, select Create. A message will display indicating the network communication path is successfully created. Click Next.



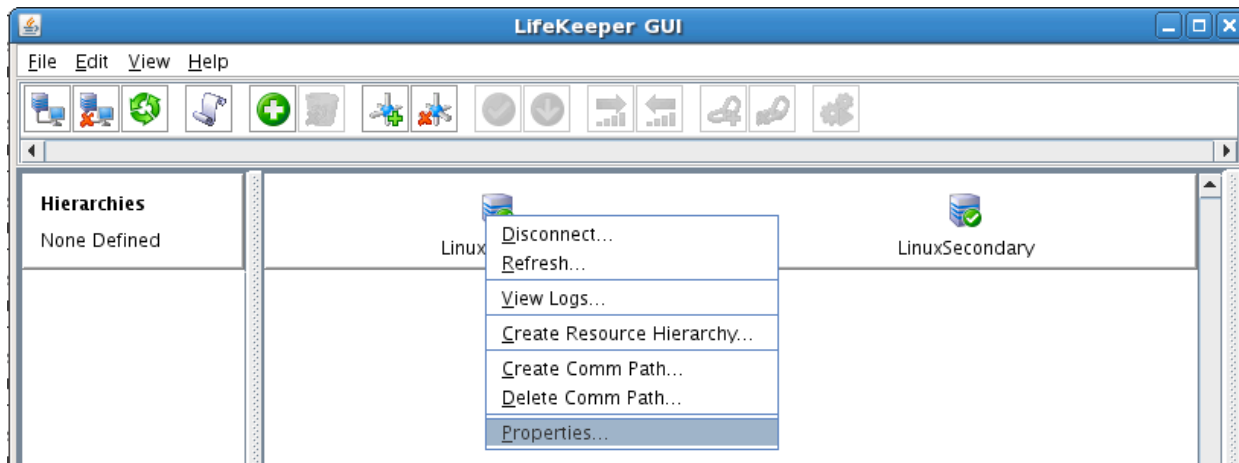
If you selected multiple Local IP Addresses or multiple Remote Servers and the Device Type was set to TCP, then the procedure will return you to the setup wizard the next Comm Path.

10. Select Done in the last dialog box.

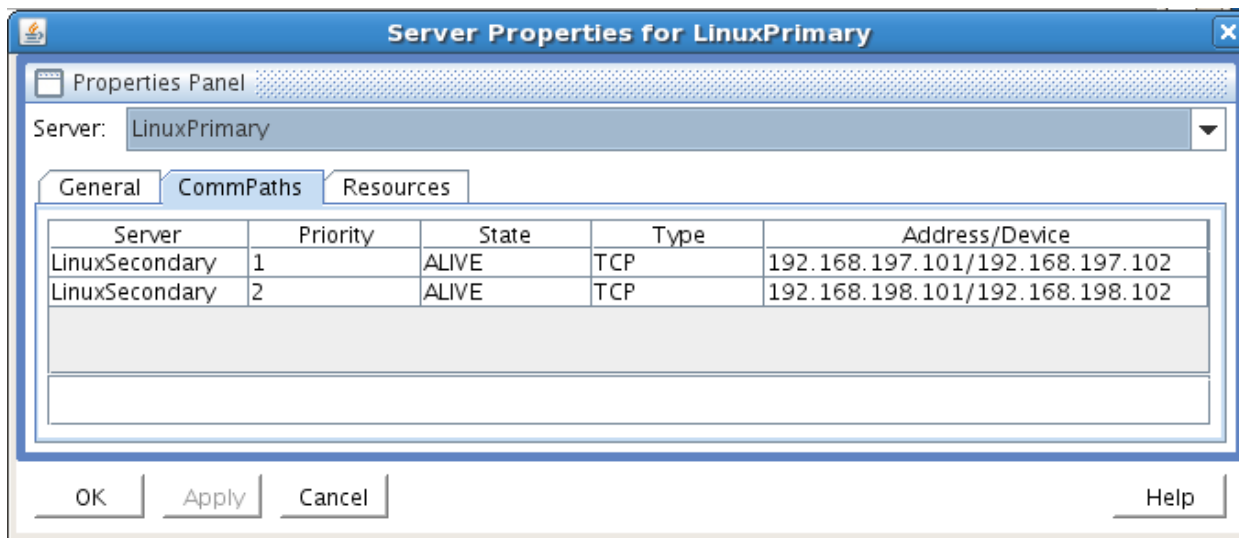
Repeat this process until you have defined all the communication paths you plan to use. SteelEye strongly recommends that you define at least two communication paths for redundancy.

Verify the Communications Paths

1. Verify that the communications paths are configured properly by viewing the Server Properties dialog box. From the LK GUI, select Edit, Server, Properties and then the Comm Paths tab.



2. Note the State displayed is ALIVE. You can also check the server icon in the right, main pane of the GUI. If only one comm path has been created, the server icon shows a yellow warning icon on the server icon, indicating that one comm. path is ALIVE, but there is no redundant comm path. The server icon will display a green heartbeat checkmark when there are at least two comm paths configured and ALIVE.




Create the LifeKeeper Hierarchy

Create and Extend an IP Resource

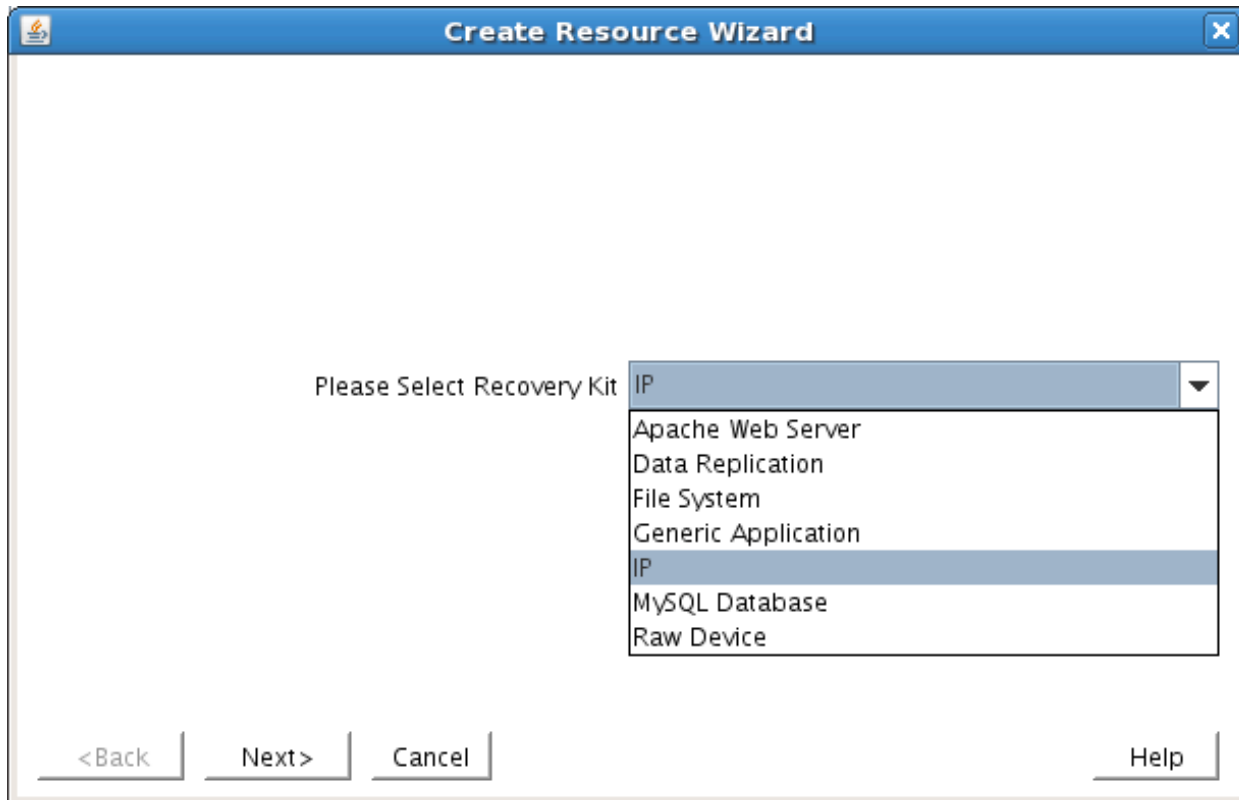
In LifeKeeper, create an IP resource and extend it to the secondary server by completing the following steps. This Virtual IP will have the ability to move between cluster nodes along with the application that depends on it.

1. From the LifeKeeper GUI toolbar, click Create Resource Hierarchy.



The Create Resource Wizard dialog box will appear with a drop down list box displaying all recognized Recovery Kits installed within the cluster.

2. Select IP Address and click Next.



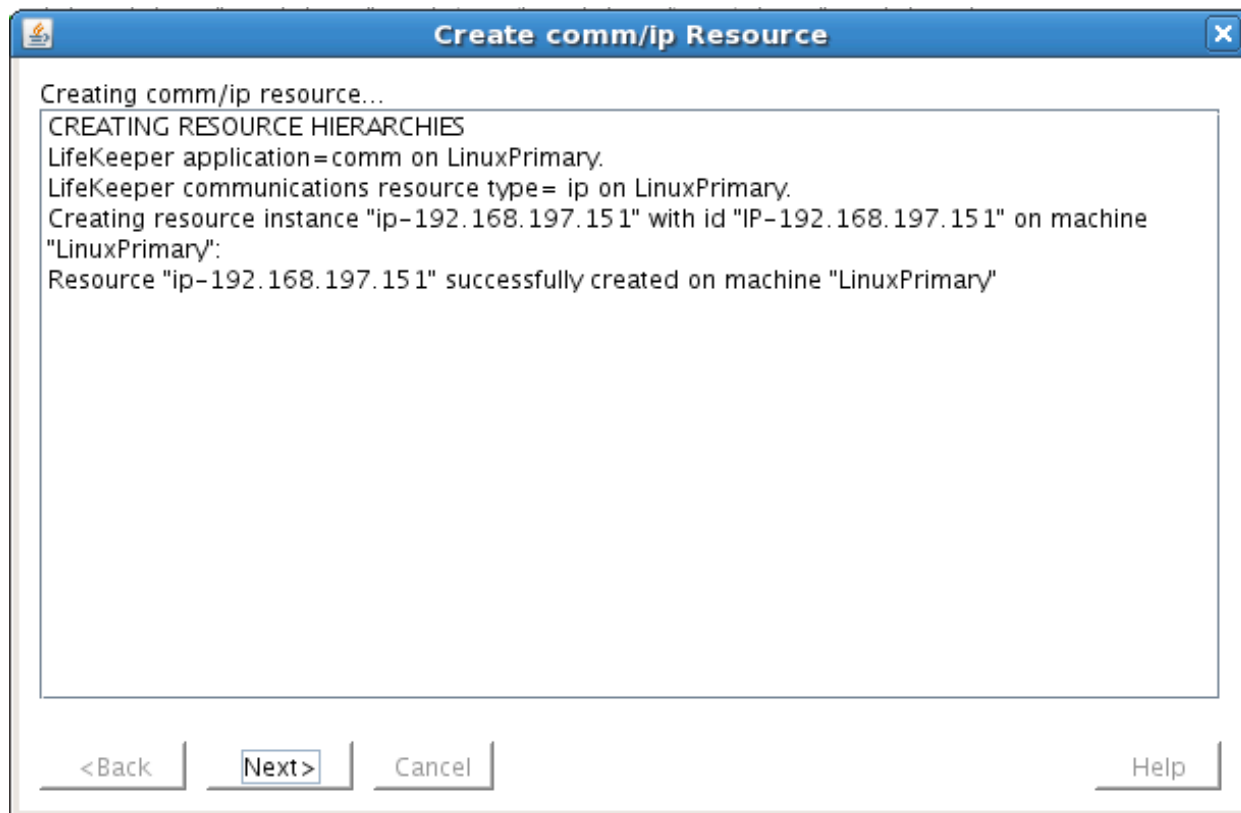
3. Enter the appropriate information for your configuration. The table below contains a list of the fields that display and additional information to assist you as you complete this procedure. Recommended values are also show below. You can also click the Help button for further

information. Press Next to continue after entering the required information.

IP Creation Field Definitions

Field	Tips
Resource Type	Select IP Address as the resource type and click Next.
Switchback Type	Select Intelligent and click Next.
Server	Select the Server where the IP resource will be created. Select your Primary server and click Next.
IP Resource	Enter the virtual IP information and click Next Example: 192.168.197.151 Note: This is an IP address that is not currently in use anywhere on your network. This is the address that all clients will use to connect to the protected resources.
Netmask	The IP subnet mask that your TCP/IP resource will use on the target server. Any standard netmask for the class of the specific TCP/IP resource address is valid. In our sample configuration 255.255.255.0 is used for a subnet mask on both networks. Note: The subnet mask you choose, combined with the IP address, determines the subnet that will be used by the TCP/IP resource and should be consistent with the network configuration.
Network Connection	This is the physical Ethernet card that the IP address is interfacing with. Chose the network connection that will allow your virtual IP address to be routable. Select the correct NIC and click Next.
Backup Interface	Select "none" and click Next
IP Resource Tag	Accept the default value and click Next. This value only affects how the IP is displayed in the GUI. The IP resource will be created on our Primary server.

- LifeKeeper will create and validate your resource. After receiving the message that the resource has been created successfully, click Next when the following dialog box appears so that you can complete the process of Extending the IP Resource to our Secondary server, below.



Extending the IP resource will start automatically after you have finished creating an IP address resource if you clicked Next in the dialog box displayed above. You can also start this from an existing IP address resource by right clicking on the active resource and selecting Extend Resource Hierarchy.

Refer to the table below to complete the Extend IP Resource procedure.

Field	Recommended Entries or Notes
Switchback Type	Leave as "intelligent" and click Next
Template Priority	Leave as default (1)
Target Priority	Leave as default (10)
Network Interface	This is the physical Ethernet card that the IP address is interfacing with. Chose the network connection that will allow your virtual IP address to be routable. The correct physical NIC should be selected by default. Please verify and then click Next.
Backup Interface	Leave as "none"
IP Resource Tag	Leave as default.
Target Restore Mode	Select Enable and click Next.
Target Local Recovery	Select Yes to enable Local Recovery for the SQL resource on the Target server.
Backup Priority	Accept the default value.

- After receiving the message Hierarchy extend operations completed, click Finish and then click Done
- Your first IP resource (192.168.197.150) is now fully protected and has the ability to "float" between cluster nodes as needed. Looking at the LifeKeeper GUI you will notice that the IP resource is Active on the Primary cluster node and Standby on the Secondary cluster node



Create the Shared Filesystem Resource Hierarchy

Create a Filesystem resource to protect the shared iSCSI filesystem and make it high available between cluster nodes. LifeKeeper leverages SCSI Persistent Group Reservations (PGR) to lock the LUN, ensuring that only the active cluster node for the storage resource can access it.

Important

At this point, the shared iSCSI LUN needs to already be mounted on the Primary Server. It should NOT be mounted on the Secondary Server. See section titled “Configure iSCSI initiator, discover and login to iSCSI target” above to review the steps involved.

1. From the LifeKeeper GUI toolbar, click Create Resource Hierarchy.
2. Select File System and click Next.
3. Proceed Through the Resource Creation wizard, providing the following values

Field	Recommended Entries or Notes
Switchback Type	Intelligent
Server	LinuxPrimary (Primary Cluster Node)
Mount Point	Select /var/lib/pgsql . Note that LifeKeeper scans the system for LUNS that are sharable between cluster nodes. The list of possible shared LUNS is presented automatically in this step of the wizard.

4. Select Create Instance to define this resource hierarchy on the Primary Server
5. Click Next to Extend the File System Resource to the Secondary Server
6. In the Extend Wizard, select “Accept Defaults”
7. As a result the File System resource is now protected on both cluster nodes. Click Finish to exit the Extend wizard.
8. Your resource hierarchy should look as follows:



Create the PostgreSQL Resource Hierarchy

Create a PostgreSQL resource to protect the PostgreSQL database and make it high available between cluster nodes.

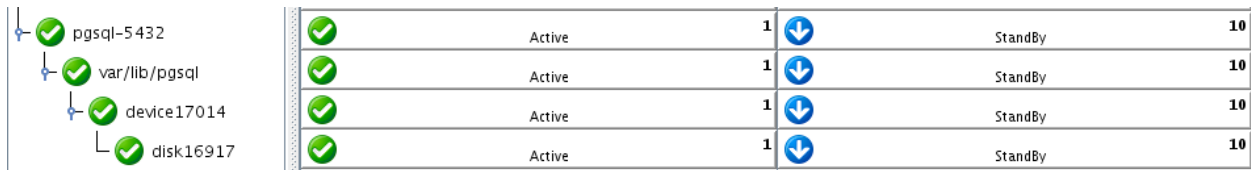
Important

At this point, PostgreSQL needs to be already running on the Primary Server. It should NOT be running on the Secondary Server. See section titled “Install, Configure, and Start PostgreSQL” above to review the process to configure and start PostgreSQL as needed.

1. From the LifeKeeper GUI toolbar, click Create Resource Hierarchy.
2. Select **PostgreSQL Database** and click Next.
3. Proceed Through the Resource Creation wizard, providing the following values

Field	Recommended Entries or Notes
Switchback Type	Intelligent
Server	LinuxPrimary (Primary Cluster Node)
PostgreSQL Executable Location	Leave as default (/usr/bin) since we are using a standard PostgreSQL install/configuration in this example. This field is used to specify the directory path containing the PostgreSQL executables.
PostgreSQL Client Executable Location	Leave as default (/usr/bin/psql) . This field is used to specify the directory path containing the PostgreSQL executable psql.
PostgreSQL Administration Executable Location	Leave as default (/usr/bin/pg_ctl). This field is used to specify the directory path containing the PostgreSQL executable pg_ctl.
PostgreSQL Data Directory	/var/lib/pgsql/data . This field is used to specify the location of the PostgreSQL data directory (datadir) that will be placed under LifeKeeper protection. The specified directory must exist and reside on a shared or replicated file system.
PostgreSQL Port	5432 . This field is used to specify the TCP/IP port number on which the postmaster daemon is listening for connections from client applications.
PostgreSQL Socket Path	Leave as default (/tmp/.s.PGSQL.5432) . This field is used to specify the full path to the Unix- domain socket on which the postmaster daemon is listening for connections from client applications.
Enter Database Administrator User	Enter "postgres" . This field is used to specify a PostgreSQL Database Administrator User name for the specified database instance with connection and administrator privileges for the instance.
PostgreSQL Logfile	Leave as default (/tmp/pgsql-5432.lk.log) . This field is used to specify the log file path that will be used for the PostgreSQL log file.
Database tag	Leave as default

4. Select Create to define the PostgreSQL resource hierarchy on the Primary Server
5. Click Next to Extend the File System Resource to the Secondary Server
6. In the Extend Wizard, select "Accept Defaults"
7. As a result the PostgreSQL resource is now protected on both cluster nodes. Click Finish to exit the Extend wizard.
8. Note: LifeKeeper will automatically identify that the PostgreSQL resource has a dependency on the FileSystem resource (/var/lib/pgsql). The Filesystem Resource will appear underneath the PostgreSQL resource in the GUI
9. Your resource hierarchy should look as follows:



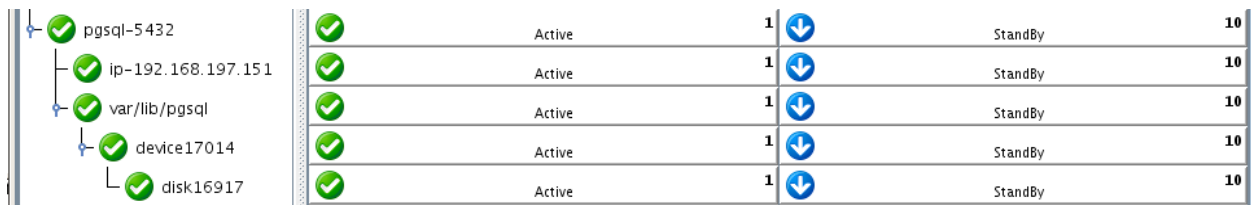
The screenshot shows a resource hierarchy on the left with a tree view: pgsql-5432 (root), var/lib/pgsql (child), device17014 (child of var/lib/pgsql), and disk16917 (child of device17014). To the right is a table with 6 columns: a green checkmark icon, a text label, a status 'Active', a count '1', a blue arrow icon pointing down, a state 'StandBy', and a priority '10'.

✓		Active	1	↓	StandBy	10
✓		Active	1	↓	StandBy	10
✓		Active	1	↓	StandBy	10
✓		Active	1	↓	StandBy	10

Create the PostgreSQL IP Address Dependency

In this step will define an additional dependency: that PostgreSQL depends on a Virtual IP (192.168.197.151) so that the IP address follows the PostgreSQL database should it move.

1. From the LifeKeeper GUI toolbar, right-click on the “pgsql-5432” resource
2. Select “Create Dependency” from the right-click context menu
3. In the Child Resource Tag dropdown menu, select “ip-192.168.197.151”
4. Click Next
5. Click Create Dependency
6. Click Done
7. The Virtual IP address resource (192.168.197.151) will now appear underneath the PostgreSQL resource in the LifeKeeper user interface. This ensures that resources move together, and are started/stopped in the proper order.
8. Your resource hierarchy should look as follows



The screenshot shows an updated resource hierarchy on the left: pgsql-5432 (root), ip-192.168.197.151 (child), var/lib/pgsql (child), device17014 (child of var/lib/pgsql), and disk16917 (child of device17014). The table to the right now has 7 rows, with the new IP resource added as the second row.

✓		Active	1	↓	StandBy	10
✓		Active	1	↓	StandBy	10
✓		Active	1	↓	StandBy	10
✓		Active	1	↓	StandBy	10
✓		Active	1	↓	StandBy	10

At this point in the Evaluation, we have fully protected PostgreSQL, and its dependent resources: IP addresses, and Shared Storage.

Phase 6 –Test Your Environment

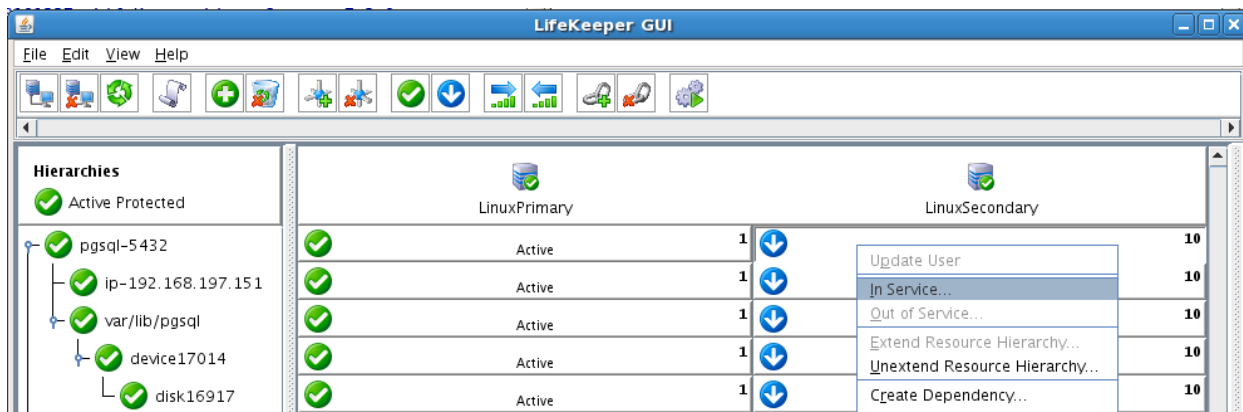
The following test scenarios have been included to guide you as you get started evaluating SteelEye Protection Suite for Linux. **Before beginning these tests, make sure the data replication resources are in the mirroring state.**

Note: For these test examples, the Primary Server is referred to as LINUXPRIMARY. The Backup or Secondary Server is referred to as LINUXSECONDARY.

1. Manual Switchover of the PostgreSQL Hierarchy to Secondary Server

Procedure:

- From the LifeKeeper GUI, right click on the PostgreSQL resource on the Secondary Server (LINUXSECONDARY) and choose IN SERVICE.
- Click “In Service” in the window that pops up



Expected Result:

- Beginning with the PostgreSQL resource, all resources will be removed from service on the Active Server (LINUXPRIMARY).
- Beginning with the dependent resources (IP and Shared Volume), all resources will be brought in service on LINUXSECONDARY

Tests/Verification:

- Using the LifeKeeper GUI, verify that the PostgreSQL and dependent resources are active on LINUXSECONDARY.
- Run “ifconfig -a” on LINUXSECONDARY to validate that the IP Address 192.168.197.151 is active on LINUXSECONDARY
- Run “df -h” to verify that the /var/lib/postgresql shared iSCSI filesystem is mounted on LINUXSECONDARY
- Verify the PostgreSQL services are running on LINUXSECONDARY by running “ps -ef | grep -i postgres”
- On LINUXSECONDARY run the following command to verify client connectivity to the PostgreSQL database:
 - # su - postgres
 - # psql
- Verify that the SCSI reservation has properly locked the iSCSI LUN to the currently active cluster node. On LINUXPRIMARY, run “mount /dev/sdc1 /var/lib/postgresql”. This should FAIL because LINUXPRIMARY does not own the SCSI reservation on this LUN.

2. Manual Switchover of the PostgreSQL Hierarchy back to Primary Server

Procedure:

- From the LifeKeeper GUI, right click on the PostgreSQL resource on the Primary Server (LINUXPRIMARY) and choose IN SERVICE.
- Click “In Service” in the window that pops up

Expected Result:

- Beginning with the PostgreSQL resource, all resources will be removed from service on the Active Server (LINUXSECONDARY).
- Beginning with the dependent resources (IP and Shared Volume), all resources will be brought in service on LINUXPRIMARY

pgsq-5432	Active	1	StandBy	10
ip-192.168.197.151	Active	1	StandBy	10
var/lib/pgsql	Active	1	StandBy	10
device17014	Active	1	StandBy	10
disk16917	Active	1	StandBy	10

Tests/Verification:

- Using the LifeKeeper GUI, verify that the PostgreSQL and dependent resources are active on LINUXPRIMARY.
- Run “ifconfig -a” on LINUXPRIMARY to validate that the IP Address 192.168.197.151 is active on LINUXPRIMARY
- Run “df -h” to verify that the /var/lib/pgsql shared iSCSI filesystem is mounted on LINUXPRIMARY
- Verify the PostgreSQL services are running on LINUXPRIMARY by running “ps -ef | grep -i postgres”
- On LINUXPRIMARY run the following command to verify client connectivity to the PostgreSQL database:
 - # su - postgres
 - # psql
- Verify that the SCSI reservation has properly locked the iSCSI LUN to the currently active cluster node. On LINUXSECONDARY, run “mount /dev/sdc1 /var/lib/pgsql”. This should FAIL because LINUXSECONDARY does not own the SCSI reservation on this LUN.

3. Simulate a network failure on the Primary Server by failing the IP resource

IMPORTANT NOTE: Only perform this test if you have more than one communications path configured.

If you perform this test and have only one communications path configured, your system will enter a split-brain scenario as described in the LifeKeeper Administration Guide found [here](#). Refer to this document for more information or contact SteelEye pre-sales technical support for assistance in resolving this condition.

Procedure:

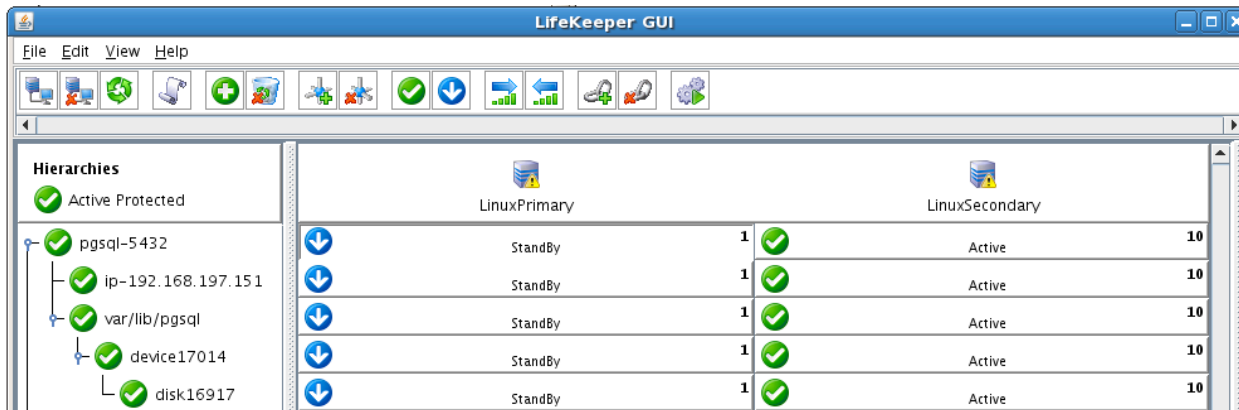
- On LINUXPRIMARY, pull the network cable attached to the NIC that the virtual IP address is configured on

Expected Result:

- The IP Resource should fail first.
- The entire hierarchy should failover to LINUXSECONDARY

Tests/Verification:

- Check the LifeKeeper Log to verify the IP resource failed – “/opt/LifeKeeper/bin/lk_log log”
- Using the LifeKeeper GUI, verify the PostgreSQL resource hierarchy fails over successfully to LINUXSECONDARY
- After this test has been completed, re-connect the network cable on LINUXPRIMARY



4. Hard failover of the resource from the Secondary Server back to the Primary Server

Procedure:

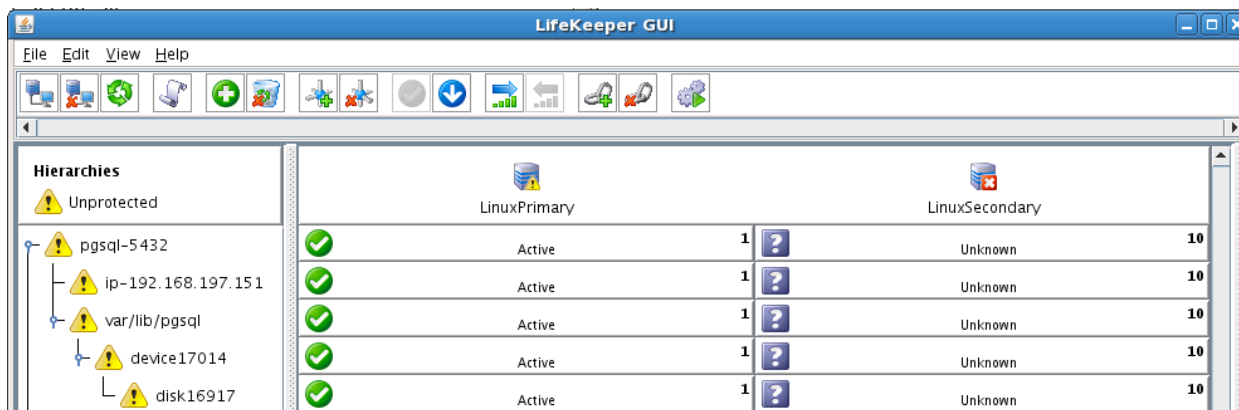
- Pull the power cord on LINUXSECONDARY, as this is the server with all resources currently In Service.

Expected Result:

- After failure has been detected, beginning with the dependent resources (IP and Volume), all resources will be brought in service on LINUXPRIMARY.

Tests/Verification:

- Using the LifeKeeper GUI, verify the mirror has reversed and is in a Resync Pending state waiting for LINUXSECONDARY to come back on line.
- Verify the PostgreSQL Server services are running on LINUXPRIMARY.
-



5. Bring Failed Server back on line

Procedure:

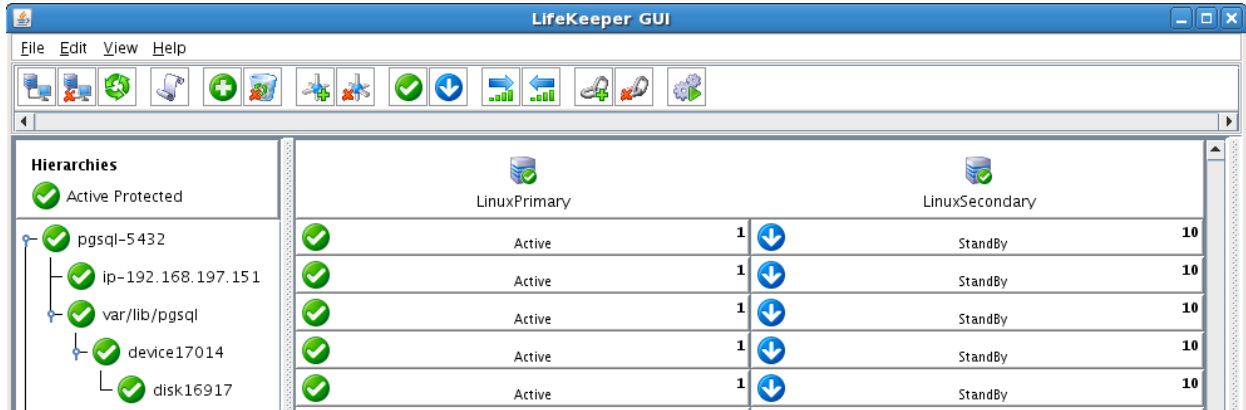
- Plug the power cord back into LINUXSECONDARY and boot it up.

Expected Result:

- Using the LifeKeeper GUI, verify that LINUXSECONDARY is coming back up and has become the Standby Server.

Tests/Verification:

- Verify the PostgreSQL Hierarchy is in service on LINUXPRIMARY and standby on LINUXSECONDARY.



6. Verify Local Recovery of PostgreSQL Server

Procedure:

- Kill the PostgreSQL processes via the command line:
 - # ps -ef | grep postgres
 - # (kill -9 the PIDs returned)
 - run "ps -ef | grep postgres once again to verify that the processes no longer exist

Expected Result: (Assumes Local Recovery for SQL resource is set to YES)

- The PostgreSQL Server service should stop.
- The PostgreSQL quickcheck process will automatically restart the PostgreSQL Server Service when it runs periodically.
- No failure of PostgreSQL should occur.

Tests/Verification:

- Execute "ps -ef | grep postgres" once again to verify that the postgresql processes have been restored locally on LINUXPRIMARY.
- Verify connectivity to the postgresQL database by running:
 - # su - postgres
 - # psql
- If you inspect the LifeKeeper logs, you will see information indicating that LifeKeeper detected the failure of the PostgreSQL service and recovered it locally. Run /opt/LifeKeeper/bin/lk_log log for more information.