



**SIOS Protection Suite for Linux
EC2 Recovery Kit
v9.2.1**

Administration Guide

December 2017

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 © 2017
By SIOS Technology Corp.
San Mateo, CA U.S.A.
All rights reserved

Table of Contents

Chapter 1: Introduction (resource name is EC2)	1
Recovery Kit for EC2	1
SIOS Protection Suite Documentation	1
Principles of Operation	1
Route Table scenario (Backend Cluster):	1
Figure 1. Route Table scenario	2
Elastic IP scenario (Frontend cluster):	3
Figure 2. Elastic IP scenario	4
Chapter 2: Requirements	5
Chapter 3: Configuration	7
Specific Configuration Considerations for Amazon EC2	7
Specific Configuration Considerations for Amazon EC2	7
Adjusting Recovery Kit for EC2 Tunable Values	8
Creating a Resource Hierarchy	9
Deleting a Resource Hierarchy	11
Extending Your Hierarchy	12
Local Recovery and Configuration	13
Local Recovery scenario (Backend Cluster):	13
Elastic IP scenario (Frontend Cluster):	14
Resource Monitoring and Configuration	14
Route Table scenario (Backend Cluster):	15
Elastic IP scenario (Frontend Cluster):	15
Unextending Your Hierarchy	15
User System Setup	16
Route Table scenario (Backend Cluster):	16
Elastic IP scenario (Frontend Cluster):	17

Chapter 1: Introduction (resource name is EC2)

Recovery Kit for EC2

The Recovery Kit for EC2 provides a mechanism to recover an Elastic IP from a failed primary server to a backup server. It also provides a mechanism to enable the IP Recovery Kit to work in multiple availability zones.

Please see the [Principles of Operation](#) for a comparison and additional information about the definition, scenarios, and operation of the Recovery Kit for EC2.

SIOS Protection Suite Documentation

The following is a list of SIOS Protection Suite for Linux related information available from SIOS Technology Corp.

- [SPS for Linux Technical Documentation](#)
- [SPS for Linux Release Notes](#)
- [SIOS Technology Corp. Documentation](#)

Please refer to [Amazon Elastic Compute Cloud \(EC2\) Documentation](#) for more information.

Note: “Amazon Web Services, the “Powered by Amazon Web Services” logo, “AWS”, “Amazon EC2”, “EC2”, “Amazon Elastic Compute Cloud”, “Amazon Virtual Private Cloud”, and “Amazon VPC” are trademarks of Amazon.com, Inc. or its affiliates in the United States and/or other countries.

Principles of Operation

Recovery Kit for EC2 brings an EC2 resource into service by using EC2 API Tools. Therefore each node that Recovery Kit for EC2 is installed on, you must also install EC2 API Tools. Recovery Kit for EC2 must be able to connect to Amazon EC2 web service end point URL (EC2 URL) using https and Amazon EC2 metadata URL (<http://169.254.169.254/>) using http.

Recovery Kit for EC2 provides two functions.

1. The Route Table scenario (Backend Cluster) manages Route Table for LifeKeeper-protected IP resources to be reached from clients within the Amazon VPC™.
2. The Elastic IP scenario (Frontend Cluster) manages Elastic IP available from the Internet.

Route Table scenario (Backend Cluster):

To clarify the administration and operation of Route Table, consider the scenario shown in Figure 1.

Figure 1. Route Table scenario

This example configuration contains one Amazon VPC™, two Availability Zones (AZ).

There are two Subnets in each AZ.

- A first Subnet (hereinafter referred to as "Public Subnet") connects to the Internet via Internet Gateway by Route Table - see Route Table of 10.0.1.0/24 and 10.0.3.0/24.
- A second Subnet (hereinafter referred to as "Private Subnet") connects to the Internet via NAT Instance by Route Table - see Route Table of 10.0.2.0/24 and Route Table of 10.0.4.0/24.

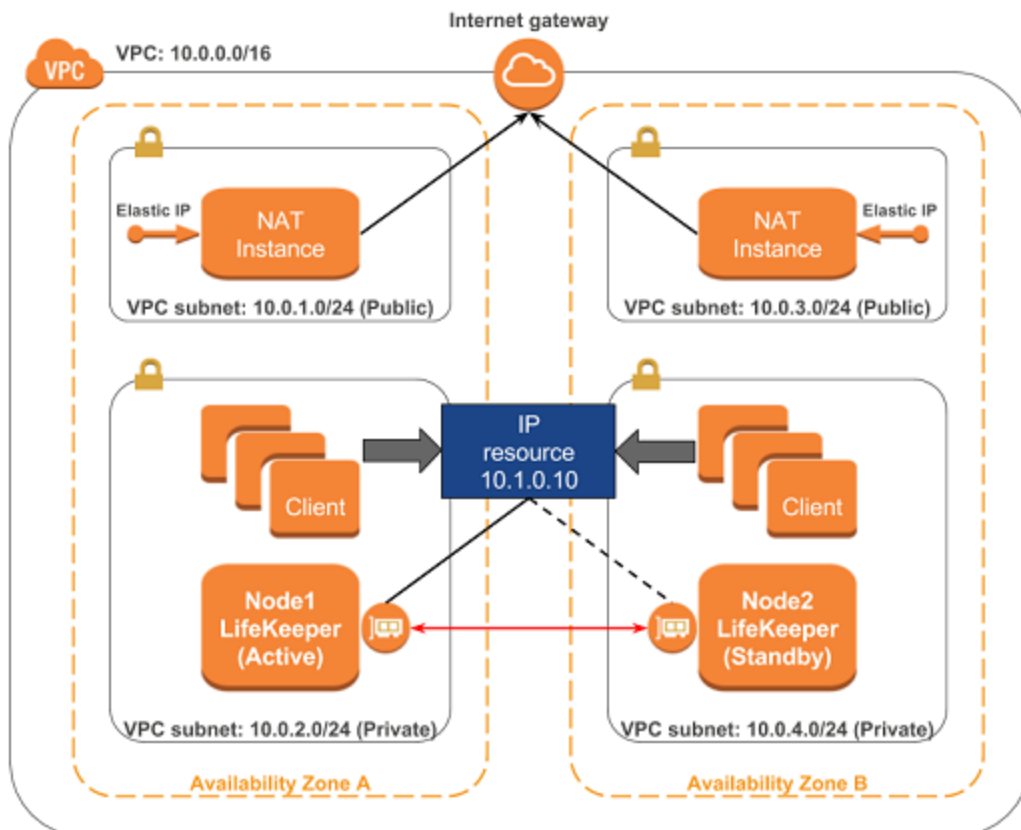
In each Public Subnet, there is an EC2 instance to which you assigned an Elastic IP for NAT (hereinafter referred to as "NAT Instance").

In each Private Subnet, there is an EC2 instance for LifeKeeper Active/Standby (hereinafter referred to as "Node1" and "Node2"), and there are clients that will use the applications protected by Node1/Node2.

Each Node1/Node2 has two Elastic Network Interfaces (ENIs).

Configure the Network ACLs and Security Groups to be able to communicate between each Instance and each Node.

Figure 1. Route Table scenario



Elastic IP scenario (Frontend cluster):

Route Table of 10.0.1.0/24 and 10.0.3.0/24

Destination	Target	Note
10.0.0.0/16	Local	Default
0.0.0.0/0	Internet Gateway	In order to connect to the Internet, requires the allocation of an Elastic IP.

Route Table of 10.0.2.0/24

Destination	Target	Note
10.0.0.0/16	Local	Default
10.1.0.10/32 (IP resource)	Elastic Network Interface (ENI) on LifeKeeper Active Node	This Target is updated by Recovery Kit for EC2 during a switchover.
0.0.0.0/0	NAT instance (10.0.1.0)	Connect to the Internet via NAT

Route Table of 10.0.4.0/24

Destination	Target	Note
10.0.0.0/16	Local	Default
10.1.0.10/32 (IP resource)	Elastic Network Interface (ENI) on LifeKeeper Active Node	This Target is updated by Recovery Kit for EC2 during a switchover.
0.0.0.0/0	NAT instance (10.0.3.0)	Connect to the Internet via NAT

When a resource switchover is performed, LifeKeeper will take the IP resource out of service on Node 1. The Target entry of 10.1.0.10/32 in each Private Subnet will be updated to reflect the ENI of Node2. The IP resource will be brought in-service on Node2. Therefore IP address traffic to 10.1.0.10 is effectively redirected to Node2 by the new Route Table configuration changes in the Private Subnet.

If you need to access the IP address 10.1.0.10 from another subnet containing the public subnet, please add the destination route 10.1.0.10/32 to the route table entry for each subnet. LifeKeeper controls all entries for which the destination is set as "10.1.0.10/32" in the route table within the VPC.

Elastic IP scenario (Frontend cluster):

To clarify the administration and operation of Elastic IP, consider the scenario shown in Figure 2.

This example configuration contains one Amazon VPC™, two Availability Zones (AZ).

There is one Subnet in each AZ.

Each Subnet connects to the Internet via Internet Gateway by Route Table.

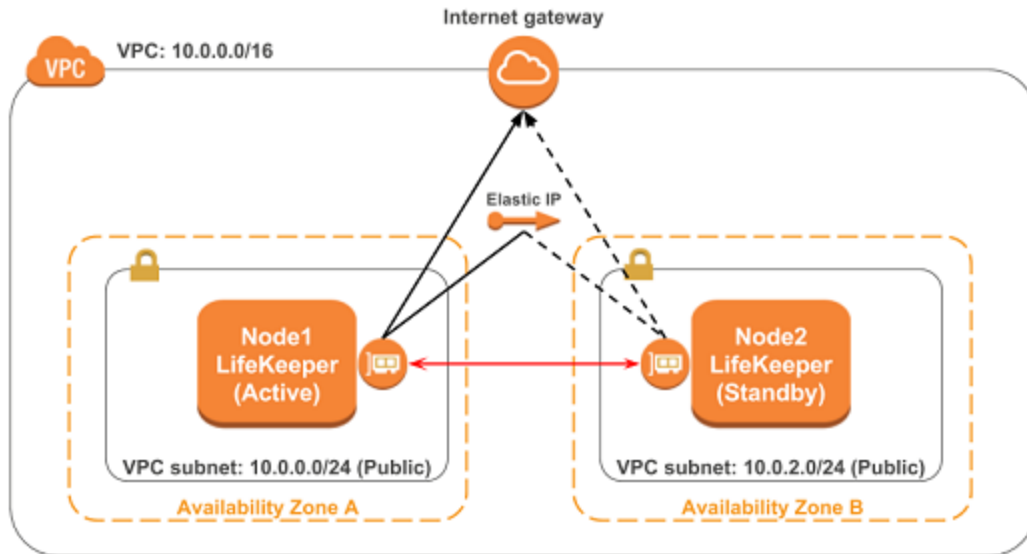
In Subnet, there is an EC2 instance for LifeKeeper Active/Standby (hereinafter referred to as "Node1" and "Node2").

Figure 2. Elastic IP scenario

Each Node1/Node2 has two Elastic Network Interfaces (ENIs).

Configure the Network ACLs and Security Groups to be able to communicate between each Node.

Figure 2. Elastic IP scenario



The system administrator allocates an Elastic IP address of frontend cluster to the ENI.

Assuming that Node1 is the primary server for the resource, the administrator creates the EC2 resource hierarchy on Node1 using the wizard described in the section entitled [Creating a Resource Hierarchy](#).

When resource switchover is performed, Recovery Kit for EC2 disassociates the Elastic IP from the ENI on Node 1. After that Recovery Kit for EC2 determines if the elastic IP is associated with the ENI on Node 2, if not, associates the Elastic IP to the ENI. Therefore client on the Internet can reach Node 2 via the Elastic IP after switchover.

As described at the beginning, Recovery Kit for EC2 must be able to connect to Amazon EC2 web service end point URL (EC2 URL) using https and Amazon EC2 metadata URL (<http://169.254.169.254/>) using http. Please note that this requirement must be satisfied on the standby EC2 instances.

Chapter 2: Requirements

Before attempting to install or remove the Recovery Kit for EC2 you must understand Amazon Web Service software requirements, as well as the installation and removal procedures for the Recovery Kit for EC2 package.

Amazon Web Service and Software Requirements

Before installing and configuring the Recovery Kit for EC2, be sure that your configuration meets the following requirements:

Amazon Web Service (AWS):

You need administrative privileges for the AWS. You must also obtain your AWS Access Key ID and Secret Access Key.

Amazon Virtual Private Cloud (VPC):

- The recovery kit requires a VPC be configured within AWS
- Two or more Subnets created on different Availability Zones (AZ)
- Each Subnet contains associated Route Tables
- If you are configuring a Public (Frontend) Cluster, then one or more Elastic IPs must be allocated

Amazon Elastic Compute Cloud (EC2):

- The recovery kit requires two or more EC2 instances
- The instances are associated on each Subnet
- The instances are attached to an Elastic Network Interface (ENI)
- Amazon EC2 API Tools need to be installed in each of the instances, and each instance must be able to access the Amazon Web Service end point URL. (please refer to <http://awsdocs.s3.amazonaws.com/EC2/ec2-clt.pdf> to install Amazon EC2 API Tools)
- Since the Auto Recovery function may conflict with the recovery function of LifeKeeper, it is not recommended to use these functions together.

LifeKeeper software:

You must install the same version of LifeKeeper software and any patches on each server. Please refer to the [SPS for Linux Technical Documentation](#) and the [SPS for Linux Release Notes](#) for specific LifeKeeper requirements.

LifeKeeper Recovery Kit for EC2:

You must install the same version of Recovery Kit for EC2 software and any patches on each server.

LifeKeeper IP Recovery Kit:

If you are using the Recovery Kit for EC2 to provide protection for the Route Table (Backend Cluster), you must install the same version of LifeKeeper for Linux IP Recovery Kit software and any patches on each server.

Note: Please refer to the [SPS for Linux Release Notes](#) or your sales representative for the latest release compatibility and ordering information. You should refer to the [SIOS Protection Suite Installation Guide](#) for specific instructions on how to install or remove the LifeKeeper Recovery Kit for EC2.

Chapter 3: Configuration

To ensure that your LifeKeeper configuration provides the protection and flexibility you require you'll need to be aware of the configuration requirements. To appropriately plan your configuration you must understand Amazon, Amazon Virtual Private Cloud, (VPC), Amazon Elastic Compute Cloud (EC2), and the user system setup hierarchy options. In addition to planning your configuration, this section also includes the specific tasks required to configure your recovery kit.

Specific Configuration Considerations for Amazon EC2

In order to properly configure your Recovery Kit for EC2 you should review the following topics to ensure that you have the information necessary to complete the configuration tasks:

- [User System Setup](#)

See the following topics for further configuration considerations:

- [EC2 Resource Monitoring and Configuration Considerations](#)
- [EC2 Local Recovery and Configuration Considerations](#)

Specific Configuration Considerations for Amazon EC2

The following configuration tasks for EC2 resources are described in this section. They are unique to an EC2 resource instance and different for each recovery kit.

- [Creating a Resource Hierarchy](#). Creates an application resource hierarchy in your LifeKeeper cluster.
- [Deleting a Resource Hierarchy](#). Deletes a resource hierarchy from all servers in your LifeKeeper cluster.
- [Extending Your Hierarchy](#). Extends a resource hierarchy from the primary server to a backup server.
- [Unextending Your Hierarchy](#). Unextends (removes) a resource hierarchy from a single server in your LifeKeeper cluster.
- Viewing and Editing EC2 Configuration Properties. Displays configuration details for an EC2 resource and allows some of them to be modified.
- [Adjusting Recovery Kit for EC2 Tunable Values](#). Tunes characteristics of the overall behavior of the Recovery Kit for EC2.

The following tasks are described in the [Administration](#) section within the [SPS for Linux Technical Documentation](#). They are common tasks with steps that are identical across all Recovery Kits.

- [Create a Resource Dependency](#). Creates a parent/child dependency between an existing resource hierarchy and another resource instance and propagates the dependency changes to all applicable servers in the cluster

- [Delete a Resource Dependency](#). Deletes a resource dependency and propagates the dependency changes to all applicable servers in the cluster.
- [In Service](#). Brings a resource hierarchy into service on a specific server.
- [Out of Service](#). Takes a resource hierarchy out of service on a specific server.
- [View Properties](#) / [Edit Properties](#). View or edit the properties of a resource hierarchy on a specific server.

The rest of this section explains how to configure your recovery kit by selecting certain tasks from the Edit menu of the LifeKeeper GUI. You may also select each configuration task from the toolbar.

- Right-click on a global resource in the Resource Hierarchy Tree (left-hand pane) of the status display window to display the same drop down menu choices as the Edit menu. This is only an option when a hierarchy already exists.
- Right-click on a resource instance in the Resource Hierarchy Table (right-hand pane) of the status display window to perform all the configuration tasks, except Creating a Resource Hierarchy, depending on the state of the server and the particular resource.

Adjusting Recovery Kit for EC2 Tunable Values

The table below lists and explains the tunable values that are available for modifying the behavior of the Recovery Kit for EC2. These values are set by adding the `/etc/default/LifeKeeper` configuration file. Because none of the components of the Recovery Kit for EC2 are memory resident, changes to these particular values become effective immediately after they are changed in `/etc/default/LifeKeeper` without requiring a LifeKeeper restart.

Tunable Value	Explanation
EC2_RESTORE_TIMEOUT	The resource restore timeout. Default value is 300.
EC2_REMOVE_TIMEOUT	The resource remove timeout. Default value is 300.
EC2_RECOVER_TIMEOUT	The local recovery timeout. Default value is 300.
EC2_QUICKCHECK_TIMEOUT	The quickCheck timeout. Default value is 100.
EC2_MAX_RETRY	This value is the number of retries that will be attempted when a resource action or EC2 API command fails. Default value is 3.
IP_NOLINKCHECK	If this variable is set to 1, the link check for the protected network interface will be disabled. Default value is 0 (i.e., The link check is enabled). (This value only applies when protecting an Elastic IP).

Tunable Value	Explanation
IP_WAIT_LINKDOWN	The number of seconds to wait in between taking the protected network interface down and back up. A delay between these two actions is necessary in some environments. Default value is 5. (This value only applies when protecting an Elastic IP).
IP_MAX_LINKCHK	The maximum number of seconds to wait for the link to come back up after it has been repaired. In some environments, it may be necessary to increase this value. Default value is 5. (This value only applies when protecting an Elastic IP).

Creating a Resource Hierarchy

To create a resource instance from the primary server, complete the following steps:

1. From the LifeKeeper GUI menu, select **Edit**, then **Server**. From the drop down menu, select **Create Resource Hierarchy**.
2. A dialog box will appear with a drop down list showing all of the recognized recovery kits installed within the cluster. Select "**Amazon EC2**" from the drop down list and click **Next**.
3. You will be prompted to enter the following information. (When the Back button is active in any of the dialog boxes, you can go back to the previous dialog box. This is especially helpful in the event that you need to correct previously entered information.)

Note: If you click the Cancel button at any time when creating your hierarchy, LifeKeeper will cancel the entire creation process.

Switchback Type	<p>This dictates how the EC2 resource will be switched back to this server when the server comes back up after a failover. You can choose either intelligent or automatic.</p> <ul style="list-style-type: none"> • Intelligent switchback requires administrative intervention to switch the instance back to the primary/original server. • Automatic switchback means the switchback will occur as soon as the primary server comes back on line and reestablishes LifeKeeper communication paths. <p>Note: The switchback type can be changed later from the General tab of the Resource Properties dialog box.</p>
Server	Select the Server for the EC2 resource (typically this is referred to as the primary or template server). All the servers in your cluster are included in the drop down list.
EC2 HOME	Select or enter the EC2_HOME directory PATH. This EC2_HOME is EC2 API tools PATH. Note: Default values is /opt/aws. See if <\$EC2_HOME>/bin/ec2-describe-addresses exists to validate.

EC2 URL	Select or enter the actual EC2_URL. This EC2_URL is the Amazon EC2 web service end point URL.
AWS Access Key	Enter the AWS ACCESS KEY. This AWS ACCESS KEY is the access key ID that identifies you to the AWS Tools. The ACCESS KEY and SECURITY KEY combination is used to make secure REST or Query protocol requests to any AWS service API. Note: The ACCESS KEY can be validated by using the EC2-describe-route-tables API with the ACCESS KEY and SECRET KEY. If the return code is 0, the pair is correct.
AWS Security Key	Enter the AWS SECURITY KEY. This AWS SECURITY KEY is the Secret Access Key. The ACCESS KEY and SECURITY KEY combination is used to make secure REST or Query protocol requests to any AWS service API. Note: The ACCESS KEY can be validated by using the EC2-describe-route-tables API with the ACCESS KEY and SECRET KEY. If return code is 0, the pair is correct.
EC2 Resource type	The EC2 Recovery Kit provides protection for two AWS recovery scenarios. The Route Table and Elastic IP scenario. The Route Table scenario is used in conjunction with a local virtual IP address and is typically used for Backend Clusters. The Elastic IP scenario is used for protection of an Elastic IP and is typically used for Frontend Clusters. Select the EC2 type to be used.
IP resource	This field will only appear and be set in the Route Table scenario. Select the IP resource. This is the virtual IP resource that is protected by LifeKeeper and configured in the Route Table address in the VPC. Note: The list will only show IP resources that are ISP and IPv4 based.
Network Interface	This field will only appear and be set in the Elastic IP scenario. Select the Network Interface to associate with Elastic IPs.
Elastic IP	This field will only appear and be set in the Elastic IP scenario. Obtain the Elastic IPs using the <code>ec2-describe-addresses</code> EC2 API for the selected Network Interface.
EC2 Resource Tag	Select or enter a unique EC2 Resource Tag name for the EC2 resource instance you are creating. This field is populated automatically with a default tag name, <code>ec2-<resource></code> , where <code><resource></code> is the resource name. This tag can be changed.

4. Click **Create**. The Create Resource Wizard will then create your EC2 resource.
5. At this point, an information box appears and LifeKeeper will validate that you have provided valid data to create your EC2 resource hierarchy. If LifeKeeper detects a problem an **ERROR** will appear in the information box. If the validation is successful your resource will be created. Click **Next**.

Another information box will appear confirming that you have successfully created an EC2 resource hierarchy. You must extend that hierarchy to another server in your cluster in order to place it under LifeKeeper protection.

When you click **Continue**, LifeKeeper will launch the Pre-Extend configuration task. Refer to [Extending Your Hierarchy](#) for details on how to extend your resource hierarchy to another server.

If you click **Cancel** now, another dialog box will appear alerting you that you will need to manually extend your EC2 resource hierarchy to another server at some other time to put it under LifeKeeper protection.

Deleting a Resource Hierarchy

To delete a resource hierarchy from all of the servers in your LifeKeeper environment, complete the following steps:

1. From the LifeKeeper GUI menu, select **Edit**, then **Resource**. From the dropdown menu, select **Delete Resource Hierarchy**.
2. Select the name of the Target Server that you are deleting from your EC2 resource hierarchy and click **Next**.

Note:This dialog will not appear if you selected the Delete Resource task by right clicking on a resource instance in either pane.

3. Select the Hierarchy to Delete. Identify the resource hierarchy you wish to delete, highlight it then click **Next**.

Note:This dialog will not appear if you selected the Delete Resource task by right clicking on a resource instance in the left or right pane.

4. An information box appears confirming your selection of the target server and the hierarchy you have selected to delete. Click **Delete** to proceed.
5. An information box appears confirming that the EC2 resource was deleted successfully.
6. Click **Done** to exit.

Extending Your Hierarchy

After you have created a hierarchy, you must extend that hierarchy to another server in the cluster. There are three possible scenarios to extend your resource instance from the template server to a target server.

- Continue from creating the resource into extending that resource to another server.
- Enter the Extend Resource Hierarchy task from the edit menu as shown below.
- Right click on an unextended hierarchy in either the left or right hand pane.

Each scenario takes you through the same dialog boxes (with a few exceptions, detailed below).

1. If you are entering the Extend wizard from the LifeKeeper GUI menu, select **Edit**, then **Resource**. From the drop down menu, select **Extend Resource Hierarchy**. This will launch the Extend Resource Hierarchy wizard. If you are unfamiliar with the Extend operation, click **Next**. If you are familiar with the LifeKeeper Extend Resource Hierarchy defaults and want to bypass the prompts for input/confirmation, click **Accept Defaults**.
2. The Pre-Extend Wizard will prompt you to enter the following information.

Note: The first two fields appear only if you initiated the Extend from the Edit menu. It should be noted that if you click Cancel at any time during the sequence of extending your hierarchy, LifeKeeper will cancel the extension process to that particular server. However, if you have already extended the resource to another server, that instance will continue to be in effect until you specifically unextend it.

Field	Tips
Switchback Type	<p>Select the Switchback Type. This dictates how the EC2 instance will be switched back to this server when it comes back into service after a failover to the backup server. You can choose either intelligent or automatic.</p> <ul style="list-style-type: none"> • Intelligent switchback requires administrative intervention to switchback the instance to the primary/original server. • Automatic switchback means the switchback will occur as soon as the primary server comes back on line and reestablishes LifeKeeper communication paths. <p>The switchback type can be changed later, if desired, from the General tab of the Resource Properties dialog box.</p>
Template Priority	<p>Select or enter a Template Priority. This is the priority for the EC2 hierarchy on the server where it is currently in service. Any unused priority value from 1 to 999 is valid, where a lower number means a higher priority (1=highest). The extend process will reject any priority for this hierarchy that is already in use by another system. The default value is recommended.</p> <p>Note: This selection will appear only for the initial extend of the hierarchy.</p>

Field	Tips
Target Priority	Select or enter the Target Priority. This is the priority for the new extended EC2 hierarchy relative to equivalent hierarchies on other servers. Any unused priority value from 1 to 999 is valid, indicating a server's priority in the cascading failover sequence for the resource. A lower number means a higher priority (1=highest). Note: LifeKeeper assigns the number "1" to the server on which the hierarchy is created by default. The priorities do not need to be consecutive and no two servers can have the same priority for a given resource.

- An information box will appear explaining that LifeKeeper has successfully checked your environment and that all the requirements for extending this EC2 resource have been met. If there were some requirements that have not been met, LifeKeeper will not allow you to select the **Next** button, and the **Back** button will be enabled. If you click **Back**, you can make changes to your resource extension according to any error messages that may appear in the information box. If you click **Cancel** now, you will need to manually extend your EC2 resource hierarchy to another server to put it under LifeKeeper protection. When you click **Next**, LifeKeeper will launch you into the Extend Resource Hierarchy configuration task.
- The Extend Resource Hierarchy configuration task will prompt you to enter the following information:

Field	Tips
EC2 Resource Tag	Select or enter the EC2 Resource Tag. This is the resource tag name to be used by the EC2 resource being extended to the target server. Note: The field is not editable.

- An information box will appear verifying that the extension is being performed. Click **Next Server** if you want to extend the same EC2 resource instance to another server in your cluster. This will repeat the Extend Resource Hierarchy operation. If you click **Finish**, LifeKeeper will verify that the extension of the EC2 resource was completed successfully.
- Click **Done** to exit from the Extend Resources Hierarchy menu selection.

Note: Be sure to test the functionality of the new instance on all servers.

Local Recovery and Configuration

Local Recovery scenario (Backend Cluster):

When a failure of the protected Route Table is detected by Recovery Kit for EC2, the resulting failure triggers the execution of the EC2 local recovery script. The local recovery gathers specified IP resource entries in all Route Tables and changes the entries' Target to the ENI on the active server. If the local recovery attempt

Elastic IP scenario (Frontend Cluster):

fails, LifeKeeper will perform a failover of the EC2 resource and all dependent resources to a standby server. See the [Principles of Operation](#) section for the configuration of this scenario.

Note: Since the recovery kit will protect the configuration of the route table once the corresponding EC2 resource gets created, the route table should not be modified manually.

The following example shows a typical scenario of the local recovery: When the recovery kit detects a wrong target setting of IP routing in the route table, the local recovery replaces the target to the ENI on the active server. During this process nothing will be changed regarding the entry of -10.1.0.20/32- on the Route Table B.

IP resource	10.1.0.10
ENI on Active Node	eni-01234567

Route Table A - Before

Destination	Target
10.1.0.10/32	eni-89abcdef
10.0.0.0/16	local

Route Table A - After

Destination	Target
10.1.0.10/32	eni-01234567
10.0.0.0/16	local

Route Table B - Before

Destination	Target
10.1.0.10/32	eni-89abcdef
10.1.0.20/32	eni-89abcdef
10.0.0.0/16	local

Route Table B - After

Destination	Target
10.1.0.10/32	eni-01234567
10.1.0.20/32	eni-89abcdef
10.0.0.0/16	local

Elastic IP scenario (Frontend Cluster):

When a failure of the protected Elastic IP is detected by Recovery Kit for EC2, the resulting failure triggers the execution of the EC2 local recovery script. The local recovery allocates the Elastic IP to the ENI on the active node. If the local recovery attempt fails, LifeKeeper will perform a failover of the EC2 resource and all dependent resources to a standby server. See the [Principles of Operation](#) section for the configuration of this scenario.

Resource Monitoring and Configuration

Route Table scenario (Backend Cluster):

The recovery kit uses Amazon EC2 API Tools to perform the monitoring of the Route Table and Route Table settings for the protected IP address. The recovery kit ensures that the target of the protected IP routing in Route Tables is correctly set to the ENI on the active server. Otherwise, the recovery kit performs the EC2 local recovery process.

Elastic IP scenario (Frontend Cluster):

The recovery kit uses Amazon EC2 API Tools to monitor the association of the Elastic IP with the ENI on the active server. The recovery kit ensures that the Elastic IP is correctly associated with the ENI attached on the active server. Otherwise, the recovery kit performs the EC2 local recovery process.

Note: In both scenarios, if the recovery kit gets no reply from EC2 API, the monitoring will reach a timeout at 100 seconds by default. When a timeout occurs, no failover will be performed and the resource will remain in ISP state. Only a timeout related message will be logged in the LifeKeeper log. The recovery kit will execute the monitoring once again after a check interval. See the [Adjusting Recovery Kit for EC2 Tunable Values](#) for more information about how to configure the value for timeout.

Unextending Your Hierarchy

To unextend a hierarchy complete the following steps:

1. From the **LifeKeeper GUI menu**, select **Edit**, then **Resource**. From the dropdown menu, select **Unextend Resource Hierarchy**.
2. Select the Target Server that you are unextending from the EC2 resource. It cannot be the server that the EC2 resource is currently in service on. Click **Next**.
Note: If you selected the **Unextend** task by right-clicking from the right pane on an individual resource instance, the dialog box will not appear.
3. Select the EC2 Hierarchy to unextend. Click **Next**.
Note: If you selected the **Unextend** task by right-clicking from either the left pane on a global resource or the right pane on an individual resource instance, the dialog will not appear.
4. An information box will appear confirming the target server and the EC2 resource hierarchy you have chosen to unextend. Click **Unextend**.
5. An information box will appear confirming the EC2 resource was unextended successfully.
6. Click **Done** to exit.

User System Setup

Route Table scenario (Backend Cluster):

The Route Table protection option in the Recovery Kit for EC2 provides the ability to automatically update the routing in the VPC. During a failover the recovery kit will update the route table to reflect the new Elastic Network Interface (ENI) location of the virtual IP address on the target server. In order for LifeKeeper to protect, monitor and update the Route Table in the VPC, the following configuration steps must be performed:

- The virtual IP address to be protected by the LifeKeeper for Linux IP Recovery Kit must be out of range of the allocated CIDR in the VPC.
- The virtual IP address must be protected by LifeKeeper prior to creating the Recovery Kit for EC2 resource.
- The Source/Dest Checking of the ENI must be disabled. This is required in order for the instance to accept network packets for the virtual IP address.
- Broadcast PING checking of the LifeKeeper IP resources must be disabled. LifeKeeper monitors IP resources by executing the Broadcast PING test of the IP address on the local subnet. In multiple availability zone environments this feature would not be useable because of the different subnets that exist between multiple availability zones. To disable this feature you must set the NOBCASTPING entry in the `/etc/default/LifeKeeper` configuration file as follows:

```
NOBCASTPING=1
```

- The Route Table should have a route entry for the virtual IP address and the ENI of the active server.

Note: Since the EC2 recovery kit will protect the configuration of the Route Table once the corresponding EC2 resource has been created, the Route Table should not be modified manually after hierarchy creation.

Example:

Destination: VIP 10.1.0.10/32

Target: eni-a2cc76e8

Elastic IP scenario (Frontend Cluster):

The screenshot shows the AWS Management Console interface for configuring a Route Table. The left sidebar contains navigation options like VPC Dashboard, Virtual Private Cloud, and Security. The main content area shows the 'Routes' tab for the route table 'rtb-bbcf11df'. A table lists the routes, with the last row highlighted in red:

Destination	Target	Status	Propagated
10.0.0.0/16	local	Active	No
0.0.0.0/0	igw-2e3f674b	Active	No
10.1.0.10/32	eni-a2cc76e8 / i-04c88e3eca15d3493	Active	No

Elastic IP scenario (Frontend Cluster):

The Elastic IP (EIP) protection option in the Recovery Kit for EC2 provides the ability to automatically re-associate an EIP with a specific ENI (the ENI used by EC2 resource on the active or backup server).

In order for LifeKeeper to protect, monitor and update the association of an EIP with the ENI on the active or backup server, the following configuration steps must be performed:

- One ENI can be associated with only one Elastic IP. No other EIPs (any EIPs other than the one used by EC2 resource) should be associated with the specific ENIs. Otherwise the recovery kit will dis-associate any other EIPs that are already associated with the specific ENIs.

Notes:

- Since an Elastic Block Store (EBS) of AWS can only be attached to one EC2 instance, DataKeeper for Linux is recommended when creating an HA cluster configuration using EBS.
- We recommend increasing RESRVRECTIMEOUT in /etc/default/LifeKeeper to 300 from 60 as the default. RESRVRECTIMEOUT is the number of seconds that a LifeKeeper process will sleep when waiting to reserve a resource for "recovery", while another process already has the resource reserved.

The [Message Catalog](#) provides a listing of all messages that may be encountered while using SIOS Protection Suite for Linux and, where appropriate, provides additional explanation of the cause of the errors and necessary action to resolve the error condition. This full listing may be searched for any error code received, or you may go directly to the [Recovery Kit for EC2 Message Catalog](#) which contain listings of all messages that may be encountered while utilizing the Recovery Kit for EC2.