



SIOS Protection Suite for Linux v9.2.1

SAP HANA Recovery Kit

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

- SAP HANA Recovery Kit 4
 - Description 4
 - Requirements 4
 - Installation and Configuration..... 5
 - Adjusting SAP HANA Recovery Kit Tunable Values 10
 - Administration of SAP HANA resource..... 10
 - Switch over of SAP HANA resource 10
 - Stop SAP HANA resource..... 10
- Test Scenarios..... 11
 - Test the recovery of the SAP Host Agent 11
 - Test the recovery of sapstartsrv of the SAP HANA instance 11
 - Test the recovery of the secondary SAP HANA DB (replication target) 11
 - Test the recovery of the primary SAP HANA DB 12
 - Test a node failure of the secondary node (reboot -f, power off) 12
 - Test a failover in case of a node failure of the primary node (reboot -f, power off) 12

SAP HANA Recovery Kit

Description

SAP HANA provides three different mechanisms to increase availability.

- **Host Auto-Failover** – At least one standby node is added to the SAP HANA configuration. These nodes are configured to work in standby mode. (SAP HANA scale-out)
- **Storage Replication** – The storage used by the SAP HANA node replicates all data to another SAP HANA node. This replication works without a control process from the SAP HANA system. The Storage Replication is provided by hardware partners.
- **System Replication** – SAP HANA replicates all data by using their own feature to a secondary SAP HANA node. Data is constantly pre-loaded on the secondary SAP HANA node. (SAP HANA scale-up)

With the SAP HANA Recovery Kit, SAP HANA systems can be controlled using system replication in SIOS Protection Suite for Linux (SPS-L).

The Recovery Kit is able to start the SAP HANA system on all nodes and perform the Takeover process for system replication. To ensure the functionality of the SAP HANA system, the following processes and states are continuously tested:

- SAP Host Agent on all nodes
- SAP Start Service for the instance (sapstartsrv) on all nodes
- State of the SAP HANA database on all nodes
- State of the system Replication mode (primary on active node, sync|syncmem|async on secondary nodes)

If the required processes or the database is not active, SPS-L will try to start them again. In the case of an unsuccessful restart of the processes, the status of the SAP HANA resource is marked as faulty (OSF) on the corresponding server. Manual intervention is necessary to correct this condition and to set the state of SAP HANA resource to "in Service" and monitored (ISP).

In the case of an invalid state of the System Replication, the SAP HANA resource is also placed in the state faulty (OSF). It has to be decided with the help of the database administrator if a takeover is to be performed or how the System Replication mode is to be corrected.

In carrying out the "Out of Service" action only the database on the primary node is stopped. The database on the secondary node remains active as does the System Replication mode. If the entire SAP HANA system is to be shut down, the control flag "!volatile!noHANARemove" must first be deleted by using:

```
/opt/LifeKeeper/bin/flg_remove -f "!volatile!noHANARemove"
```

Requirements

Hardware Requirements

- **Servers:** Servers should be configured in accordance with the requirements described in the SAP HANA Master Guide, notes mentioned in this guide, the SPS for Linux Documentation and the SPS for Linux Release Notes.
- **Storage:** For SAP HANA databases in Replication mode, no shared storage is necessary. Special storage requirements may be found in the SAP HANA Master Guide and mentioned in the SAP notes.

Software Requirements

- **SPS Software:** It is required that you install the same version of the SPS software and apply the same versions of the SPS software patches to each server in your cluster.
- **SPS SAP HANA Recovery Kit:** The SAP HANA recovery Kit is available for downloading at the customer's request and must be installed and configured in accordance with the following installation instructions.
- **SAP HANA Software:** SAP HANA Platform Edition 1.0 SPS07 (and later) is recommended for use with the SAP HANA Recovery Kit. The SAP HANA system replication must be enabled and configured.

Installation and Configuration

Installing Scripts for SAP HANA ARK

Download and copy, as the root user, the self-extracting archive HANA-ARK.run in a temporary directory.

Add execute permission for root to HANA-ARK.run.

```
# chmod 744 <tmp-dir>/HANA-ARK.pm
```

Change to the directory /opt/LifeKeeper and execute the self-extracting archive.

```
# cd /opt/Lifekeeper
```

```
# <tmp-dir>/HANA-ARK.run
```

```
Creating directory HANA-ARK
```

```
Verifying archive integrity... All good.
```

```
Uncompressing SFX archive for SAP HANA Application Recovery Kit Installation 100%
```

```
running /opt/LifeKeeper/HANA-ARK/setup
```

```
Moving HANA.pm to /opt/LifeKeeper/lkadm/subsys/gen/app/bin
```

```
-rwxr-xr-x 1 root root 9096 Sep 26 17:41 /opt/LifeKeeper/HANA-ARK/quickCheck.pl
```

```
-rwxr-xr-x 1 root root 10369 Sep 26 17:41 /opt/LifeKeeper/HANA-ARK/recover.pl
```

```
-rwxr--r-- 1 root root 8793 Sep 26 17:41 /opt/LifeKeeper/HANA-ARK/remove.pl
```

```
-rwxr-xr-x 1 root root 12509 Sep 26 17:41 /opt/LifeKeeper/HANA-ARK/restore.pl
```

```
-rwxr-xr-x 1 root root 11940 Sep 26 17:41
```

```
/opt/LifeKeeper/lkadm/subsys/gen/app/bin/HANA.pm
```

```
Installation of SAP HANA Application Recovery Kit was successful
```

Please repeat these steps on all other cluster nodes.

Creating SAP HANA Resource

1. From the LifeKeeper GUI menu, select **Edit**, then **Server**. From the drop-down menu, select **Create Resource Hierarchy**.

A dialog box will appear with a drop-down list box with all recognized recovery kits installed within the cluster. Select **Generic Application** from the drop-down list.

Please Select Recovery Kit

Click **Next**.

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 should you encounter an error that might require you to correct previously entered information.

If you click **Cancel** at any time during the sequence of creating your resource hierarchy, LifeKeeper will cancel the entire creation process.

2. Select the **Switchback Type**. This dictates how the primary SAP HANA database 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**. **Intelligent switchback** requires administrative intervention to switch the hierarchy back to the primary/original server. **Automatic switchback** means the switchback will occur as soon as the original server comes back on line and re-establishes LifeKeeper communication paths.

We recommend **intelligent** Switchback Type.

Switchback Type

The **Switchback Type** can be changed later, if desired, from the **General** tab of the **Resource properties** dialog box.

Click **Next**.

3. Select the Server where you want to place the primary SAP HANA database. All servers in your cluster are included in the drop-down list box.

Server

Click **Next**.

4. On the next dialog enter the path to the **Restore Script** which will be used to start the SAP HANA database and control the replication mode. All the scripts you need for Restore, Remove, QuickCheck and Recover are located in directory `/opt/LifeKeeper/HANA-ARK`. Enter the complete path in the dialog box.

Restore Script

Click **Next**.

5. Enter the path to the **Remove Script** for the SAP HANA database and click **Next**. This is the script that stops the SAP HANA database.

Remove Script

6. Enter the path to the **QuickCheck Script** for the SAP HANA database. This is the script that monitors the functionality of the SAP HANA database and the Replication Mode.

QuickCheck Script [optional]

Click **Next**.

7. Enter the path to the **Local Recovery Script** for the SAP HANA database. This is the script that attempts to restore a failed component of the SAP HANA

database or the database itself.

Local Recovery Script [optional]

Click **Next**.

8. Enter the necessary **Application Information** for the SAP HANA database on this node.

The following values are required and must be separated by a space:

<SID>_<Instance number>_<Replication mode>_<Site name>

Application Info [optional]

Click **Next**.

9. Select Yes or No for **Bring Resource in Service**, and Click **Next**. Selecting No will cause the resource state to be set to OSU following the create; selecting Yes will cause the previously provided restore script to be executed.

The resource state ISP is required to extend the SAP HANA resource to another cluster member.


Bring Resource In Service

10. Enter the **Resource Tag**, which is a unique name for the SAP HANA resource. This is the label you will see for this resource in the LifeKeeper GUI. The valid characters allowed for the tag are letters, digits and the special characters ` - ` _ ` ` ` . ` ` / ` .

The resource name should be meaningful.

Resource Tag

11. Click **Create Instance** to start the creation process. A window will display a message indicating the status of the instance creation.



12. Select the **Target Server** on which the SAP HANA resource is to be extended. All remaining cluster members are included in the drop-down list.

Target Server

Click **Next**.

13. Select the same **Switchback Type** that was selected in step 2. Click **Next**.
14. Select or enter the **Template Priority** for the SAP HANA resource. This is the priority for the SAP HANA resource on the server where it is currently in service. Any unused priority from 1 (highest) to 999 (lowest) is valid. The default value is recommended.

Template Priority

Click next.

15. Select or enter the **Target Priority** for the SAP HANA resource. This is the priority for the new extended SAP HANA resource relative to the equivalent resource on other servers. The default value is recommended.

Target Priority

Click **Next**.

After receiving the message that the pre-extend checks were successful, click **Next**.

16. Check or enter the **Resource Tag** for the SAP HANA resource. It is recommended to use the same **Resource Tag** on all cluster members.

Resource Tag

Click **Next**.

17. Enter the necessary **Application Information** for the SAP HANA database on this node.


The following values are required and must be separated by a space:

<SID>_<Instance number>_<Replication mode>_<Site name>

Application Info [optional]

Click **Next**.

18. The creation of the SAP HANA resource on the selected node is started, then the equivalencies between the cluster nodes for the resource are created.



Click **Finish**.

19. When this is completed, the **Done** button will be enabled. Click **Done** to finish.
20. Next, an IP resource should be created and placed in dependency with the SAP HANA resource. Please refer to the administration guide "IP Recovery Kit" for these steps.

Adjusting SAP HANA Recovery Kit Tunable Values

Several of the scripts included with the Recovery Kit have been written with a timeout feature to allow hung scripts to automatically kill themselves.

HANA-DEBUG can be enabled or disabled. The default for HANA_DEBUG is 0 (disabled). To enable, set this parameter to 1 in /etc/default/LifeKeeper.

Variable Name	Description	Default Value
HANA_DEBUG		0
HANA_START_WAIT		2700
HANA_STOP_WAIT		600
HANA_TIMEOUT		value of LKCHECKINTERVAL - 10

The values HANA_START_WAIT and HANA_STOP_WAIT correspond to the timeout parameter of the commands

```
sapcontrol -nr <instance number> -function StartWait <timeout sec> <delay sec>
```

and

```
sapcontrol -nr <instance number> -function Stop <softtimeout sec>
```

Administration of SAP HANA resource

Switch over of SAP HANA resource

If a switch over of the primary database instance is started, the recovery kit performs the following steps:

- The database instance on the previous primary node will be stopped
- A takeover of the system replication will be executed on the previous secondary node
- The database instance on the new secondary node will be re-enabled in system replication to the new primary node
- The database on the new secondary node will be started

Stop SAP HANA resource

When a stop of the SAP HANA resource is executed, only the primary database instance is terminated. The secondary database instance is still running.

In the event that the primary and the secondary database instance are to be stopped, the “!volatile!noHANARemove_<tag>” flag must be removed first. This is done with the following commands:

Example:

```
vmlx-sha1:~ # flg_list
!volatile!noHANARemove_HANA-DB_HN1_00
vmlx-sha1:~ # flg_remove -f '!volatile!noHANARemove_HANA-DB_HN1_00'
vmlx-sha1:~ # flg_list
```

Now it is possible to stop the primary and the secondary database instance.

When the SAP HANA resource is restarted, the “!volatile!noHANARemove_<tag>” flag is created automatically.

Test Scenarios

To understand the behavior of the recovery kit, perform the following tests. The following requirements must be fulfilled.

The LifeKeeper and the SAP HANA database must be installed and configured according to the installation instructions. The SAP HANA database resource is active and running in one of the possible replication modes. The active replication mode must be entered correctly in the properties of the HANA DB resource. See section “Creating SAP HANA Resource”, point 8 and point 17. A dependency to a prior generated protected IP address must be created.

Test the recovery of the SAP Host Agent

Determine the status and the process numbers of the SAP Host Agent processes by using:

```
vlmx-sha1:~ # /usr/sap/hostctrl/exe/saphostexec -status
saphostexec running (pid = 3818)
sapstartsrv running (pid = 3867)
11:30:49 17.11.2016 LOG: Using PerfDir (DIR_PERF) = /usr/sap/tmp
saposcol running (pid = 3965)
```

Kill one of the processes or execute

```
/usr/sap/hostctrl/exe/saphostexec -stop
```

to influence the functionality of the SAP Host Agent.

The SAP HANA recovery kit will recognize that the SAP Host Agent is not working properly and restart the SAP Host Agent on that node.

The behavior can be observed by the following command:

```
tail -f /var/log/lifekeeper.log
```

The SAP HANA resource does not change the state. After a successful recovery, the SAP Host Agent is fully functional again.

If the recovery kit is unable to restart the SAP Host agent, the HANA database and the resource remains in the previous state. The SAP Host Agent will be checked again and if possible restarted later.

Test the recovery of `sapstartsrv` of the SAP HANA instance

To test the recovery of `sapstartsrv`, the service must be stopped. One method is to terminate the `sapstartsrv` with the command

```
sapcontrol -nr <ID> -function Stop
```

as user `<sid>adm`. The other method is to kill the `sapstartsrv` process.

In this case the `sapstartsrv` will be restarted by SAP HANA recovery kit.

The resource does not change its state.

Test the recovery of the secondary SAP HANA DB (replication target)

In event of a failure of the secondary database instance (replication target), the replication mode of the database instance on that node is determined. When the replication mode is set correctly, the database instance will be started again. If the mode is set to primary, the recovery kit will enable the correct replication mode and start the database instance.

The behavior can be observed by looking at the log file `/var/log/lifekeeper.log`.

After the recovery, the state of the database and the replication can be tested as user `<sid>adm` by using:

```
sapcontrol -nr 0 -function GetProcessList
```

```
hdbnsutil -sr_state
```

In the event that the secondary database instance cannot be started by the recovery kit, the SAP HANA resource is flagged as faulty (OSF) on the corresponding node.

Hierarchies		vmlx-sha1		vmlx-sha2	
! Unprotected					
! HANA-DB_HN1_00		Failed	1	Active	10
└─ ip-10.0.0.176		StandBy	1	Active	10

If the cause of unsuccessful start is eliminated, the SAP HANA recovery kit starts the database instance in the subsequent check cycle. The resource will be flagged as standby at the corresponding node.

Test the recovery of the primary SAP HANA DB

In the event of a failure of the primary database instance (replication source), the replication mode of the database instance on the primary node is determined. When the replication mode is set to primary, the database instance will be started again. If the mode is not set to primary, the recovery kit will take over the replication mode and start the database instance.

The behavior can be observed by looking at the log file `/var/log/lifekeeper.log`.

After the recovery, the state of the database and the replication can be tested as user `<sid>adm` by using:

```
sapcontrol -nr 0 -function GetProcessList
```

```
hdbnsutil -sr_state
```

In the event that the primary database instance cannot be started by the recovery kit at that node, LifeKeeper will initiate a failover of the entire hierarchy to the former secondary node.

On this node the HANA recovery kit performs a takeover of the replication mode. This node becomes the primary node of replication. For the faulty node, a re-enable to the primary replication node is started. If this is successful, the secondary database is also restarted.

In event of a failed re-enable the replication mode or a failed restart of the database instance on secondary node, the SAP HANA resource will be flagged as failed (OSF) on the corresponding node. A manual intervention is necessary to eliminate the cause.

If the failover of the primary database instance failed, the resource is flagged as faulty (OSF) and remains in this state.

Test a node failure of the secondary node (reboot -f, power off)

If an error causes the secondary node to fail, the resource remains on the active node. The replication is disrupted. If the secondary is restarted and LifeKeeper is active, the database instance is restarted as a replication target.

Test a failover in case of a node failure of the primary node (reboot -f, power off)

If an error causes the primary node to fail, the remaining node is starting a takeover of resources. The HANA recovery kit also takes the role of primary source of replication. If the faulty node restarted and LifeKeeper is active, the database instance is restarted as a replication target on that node.