



**EXPRESSCLUSTER X for Linux
SAP HANA System Configuration Guide**

Release 7

NEC Corporation

Apr 08, 2022

TABLE OF CONTENTS:

1	Preface	1
1.1	Who Should Use This Guide	2
1.2	Conventions	3
1.3	EXPRESSCLUSTER X Documentation Set	4
1.4	Related documents	5
2	Overview of SAP HANA Cluster	7
2.1	Overview	7
2.2	Operating Environment	8
2.3	Illustration of operation	9
2.4	Supported scenarios and requirements	12
3	AWS Environment	13
4	Installing and configuring SAP HANA	17
5	Installing and configuring EXPRESSCLUSTER	19
5.1	Installing EXPRESSCLUSTER	20
5.2	Registering the license	21
5.3	Creating a cluster	22
5.4	Creating failover groups	23
5.5	Add group resources	24
5.6	Add monitor resources	25
6	Installing and configuring SAP ERP	27
7	Notes and Restrictions	29
8	Appendix	31
8.1	Detailed Settings	31
8.2	Operating Procedure	36
8.3	Detailed verification results	39
9	Legal Notice	45
9.1	Disclaimer	45
9.2	Trademark Information	46
10	Revision History	47

**CHAPTER
ONE**

PREFACE

This document "EXPRESSCLUSTER X for Linux SAP HANA System Configuration Guide" describes how to create and start a cluster for SAP HANA.

1.1 Who Should Use This Guide

This guide is intended for administrators who want to build a cluster system, system engineers who want to provide user support, and maintenance personnel.

This guide introduces software whose operation in an EXPRESSCLUSTER environment has been verified.

The software and setup examples introduced here are for reference only. They are not meant to guarantee the operation of each software product.

The bundled scripts are for achieving failover.

Please check and (if necessary for their usage environments and their monitoring targets) customize their contents.

1.2 Conventions

In this guide, **Note**, **Important**, **See also** are used as follows:

Note: Used when the information given is important, but not related to the data loss and damage to the system and machine

Important: Used when the information given is necessary to avoid the data loss and damage to the system and machine.

See also:

Used to describe the location of the information given at the reference destination.

The following conventions are used in this guide.

Convention	Usage	Example
Bold	Indicates graphical objects, such as fields, list boxes, menu selections, buttons, labels, icons, etc.	In User Name , type your name. On the File menu, click Open Database .
Angled bracket within the command line	Indicates that the value specified inside of the angled bracket can be omitted.	<code>clpstat -s [-h <i>host_name</i>]</code>
#	Prompt to indicate that a Linux user has logged on as root user.	<code># clpcl -s -a</code>
Monospace	Indicates path names, commands, system output (message, prompt, etc), directory, file names, functions and parameters.	<code>/Linux/server/</code>
bold	Indicates the value that a user actually enters from a command line.	Enter the following: <code># clpcl -s -a</code>
<i>italic</i>	Indicates that users should replace italicized part with values that they are actually working with.	<code>clpstat -s [-h <i>host_name</i>]</code>



In the figures of this guide, this icon represents EXPRESSCLUSTER.

1.3 EXPRESSCLUSTER X Documentation Set

The EXPRESSCLUSTER X manuals consist of the following five guides. The title and purpose of each guide is described below:

EXPRESSCLUSTER X Getting Started Guide

This guide is intended for all users. The guide covers topics such as product overview, system requirements, and known problems.

EXPRESSCLUSTER X Installation and Configuration Guide

This guide is intended for system engineers and administrators who want to build, operate, and maintain a cluster system. Instructions for designing, installing, and configuring a cluster system with EXPRESSCLUSTER are covered in this guide.

EXPRESSCLUSTER X Reference Guide

This guide is intended for system administrators. The guide covers topics such as how to operate EXPRESSCLUSTER, function of each module and troubleshooting. The guide is supplement to the Installation and Configuration Guide.

EXPRESSCLUSTER X Maintenance Guide

This guide is intended for administrators and for system administrators who want to build, operate, and maintain EXPRESSCLUSTER-based cluster systems. The guide describes maintenance-related topics for EXPRESSCLUSTER.

EXPRESSCLUSTER X Hardware Feature Guide

This guide is intended for administrators and for system engineers who want to build EXPRESSCLUSTER-based cluster systems. The guide describes features to work with specific hardware, serving as a supplement to the Installation and Configuration Guide.

1.4 Related documents

1.4.1 SAP HANA documents

For details of SAP HANA please refer to the official SAP documentation available at

https://help.sap.com/viewer/product/SAP_HANA_PLATFORM/1.0/en-US

SAP NOTEs

- #1656099: SAP Applications on AWS: Supported DB/OS and AWS EC2 products
- #1964437: SAP HANA on AWS: Supported AWS EC2 products
- #2063657: HANA System Replication takeover decision guideline
- #2235581 - SAP HANA: Supported Operating Systems

Note: Related documents and URL in this manual are subject to change without notice.

OVERVIEW OF SAP HANA CLUSTER

2.1 Overview

Cloud environments are now being used by the majority of companies, an increasing number of which are deploying SAP HANA on their cloud infrastructure services. Companies are using SAP HANA not only for fast analysis of big data but also for their mission-critical systems. This has led to a growing need to improve the availability of SAP HANA running on cloud infrastructure services.

Although SAP HANA has high availability (HA) functionality, it is still necessary to manually switch servers if a failure occurs. This causes a stoppage in operations from failure detection to completion of server failover, which can potentially lead to lost business opportunities.

EXPRESSCLUSTER, NEC's high availability infrastructure software, automatically detects failures in a system that uses SAP HANA running on Amazon Web Services (AWS) and switches to a secondary server (performs failover).

In this guide, NEC verified EXPRESSCLUSTER could shorten operational downtime and boost operational efficiency by cooperating with SAP HANA running on AWS.

To be more precise, various types of failures were hypothesized and it was verified that a cluster system could be normally restored by data synchronization using the EXPRESSCLUSTER automatic failover function and the SAP HANA system replication function, and that operations could be continued without pause.

2.2 Operating Environment

This section describes the OS and SAP HANA versions on which the operation of the Connector for SAP has been verified.

x86_64

HANA Version	EXPRESSCLUSTER Version	OS
SAP HANA 1.0 SPS11 SAP HANA 1.0 SPS12	3.3.0-1~ 4.0.0-1~ 4.1.0-1~	Red Hat Enterprise Linux 7.2 Red Hat Enterprise Linux 7.3 SUSE LINUX Enterprise Server 11 SP4
SAP HANA 2.0 SPS03	4.1.0-1~ 4.2.0-1~	Red Hat Enterprise Linux 7.2

2.3 Illustration of operation

The following figure provides an illustration to show that the system normally operates when Server 1 is running as the primary server and Server 2 is running as the secondary server.

SAP ERP Application server is connected to SAP HANA server by accessing a virtual IP address (VIP).

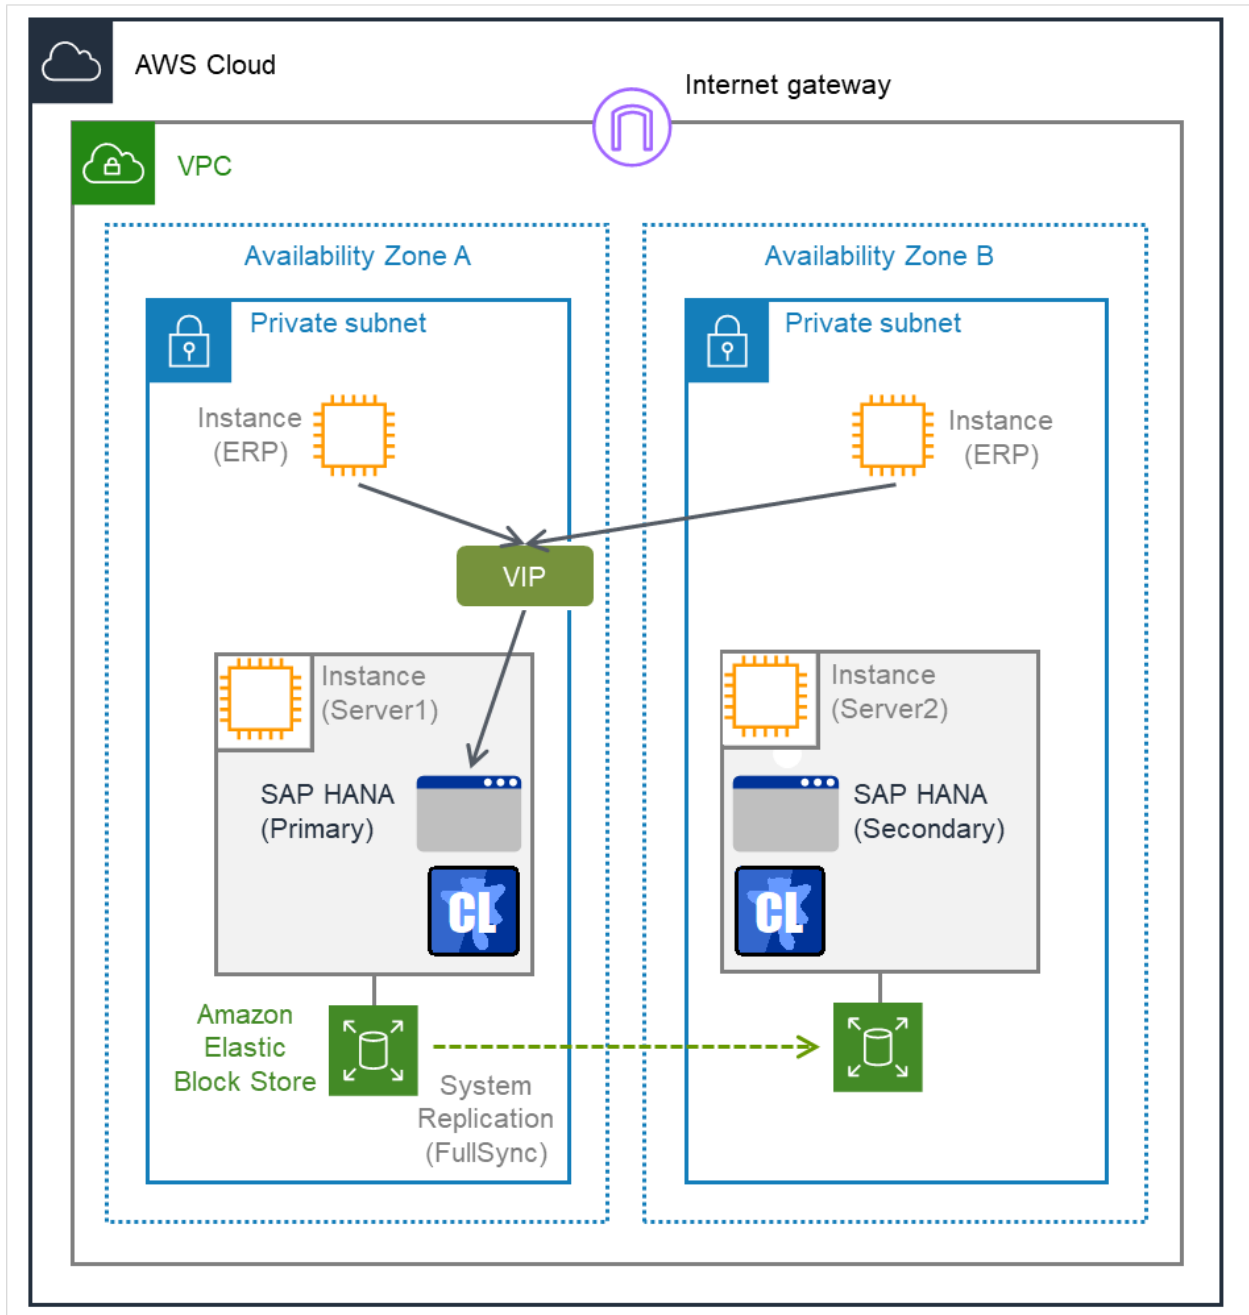


Fig. 2.1: Normal Operation

The following figure is an image of the case where a failure occurs on the primary server.
If a failure occurs on the primary server, EXPRESSCLUSTER stops SAP HANA on Server 1, and changes SAP HANA on Server 2 from the secondary server to the primary server, allowing SAP HANA operations to continue. In addition, EXPRESSCLUSTER switches the virtual IP address of Server 1 to that of Server 2. SAP ERP Application server is connected to the new primary SAP HANA server by accessing its virtual IP address.




Fig. 2.2: Occurrence of Failure on the Primary Server

The following figure is an image of the case where a failure occurs on the secondary.

If a failure occurs on the secondary server, EXPRESSCLUSTER stops SAP HANA on Server 2 and switches the system replication function to Server 1 (that is, disables the full sync option), allowing SAP HANA operations to continue.




Fig. 2.3: Occurrence of Failure on the Secondary Server

2.4 Supported scenarios and requirements

Only the scenarios and parameters indicated below are supported for cooperation between SAP HANA and EXPRESS-CLUSTER. For general system replication requirements, see the guides provided by SAP.

1. Two-node cluster consisting of scale-up (single) configuration x 2
2. It is recommended that both nodes belong to the same network segment in terms of performance. For AWS environment, Single-AZ is recommended.
3. Both nodes must be run as a single instance. No quality assurance or development system is running.
4. The automatic startup attribute of SAP HANA must be set to "off." (SAP HANA startup is managed by EXPRESSCLUSTER.)
5. Multi-tenant database container (MDC) scenario
 - Failover is performed when a failure occurred in a system database or tenant database.
 - Failover is not performed when a tenant database is stopped manually.

AWS ENVIRONMENT

The system configuration used in this verification is shown in the figure below.

In this configuration, EXPRESSCLUSTER monitors failures and switches operations and SAP HANA synchronizes data.




Fig. 3.1: The system configuration in the verification

CIDR	10.0.0/16
Subnet-1A	10.0.1.0/24
Subnet-2A	10.0.2.0/24
Instance (ERP)	10.0.2.100/24
Instance (Server1)	10.0.2.22/24

Continued on next page

Table 3.1 – continued from previous page

Subnet-1B	10.0.11.0/24
Subnet-2B	10.0.12.0/24
Instance (Server2)	10.0.12.22/24
VIP	10.2.0.20/32

In this guide, a cluster environment is configured by allocating SAP HANA instances to different Availability Zones of AWS and installing SAP HANA.

The AWS instance types were determined by referring to the following and selecting a supported environment:

SAP Notes

- #1964437 - SAP HANA on AWS: Supported AWS EC2 products
- #1656099 - SAP Applications on AWS: Supported DB/OS and AWS EC2 products

For SAP HANA, a configuration in which multiple tenant databases are created on SAP Instance is also verified.

SAP HANA (Common)

Region	Asia Pacific (Tokyo)
OS	SUSE Linux Enterprise Server 11 SP4 Red Hat Enterprise Linux 7.4
Instance Type	x1.32xlarge
CPU	128vCPU
Memory	2TB
EBS	/dev/sda1 50GB /dev/sdf 4096 GB /dev/sdb 1024GB /dev/sdc 1024GB /dev/sdd 1024GB /dev/sde 1024GB /dev/sds 50 GB /dev/sdz 50 GB
EIP	-
SAP HANA	SAP HANA SPS12

A NAT Gateway, which is used to control access to the cluster environment, was allocated to each Availability Zone.

An SAP ERP instance was allocated to one of the Availability Zones as SAP ERP Application Server.

SAP ERP

Region	Asia Pacific (Tokyo)
OS	Windows Server 2012 R2
Instance Type	m4.2xlarge
CPU	8vCPU
Memory	32GB
EBS	/dev/sda1 100 GB /dev/sdb 50 GB /dev/sdc 100 GB
EIP	-
SAP ERP	SAP ERP 6.0 EHP7 SR1

INSTALLING AND CONFIGURING SAP HANA

For information of the installation and settings of SAP HANA, refer to the following documents:

https://help.sap.com/viewer/product/SAP_HANA_PLATFORM/1.0/en-US

- SAP HANA Server Installation and Update Guide
- SAP HANA Administration Guide
- SAP HANA Master Guide

https://help.sap.com/viewer/product/SAP_HANA_PLATFORM/2.0.03/en-US

- SAP HANA Server Installation and Update Guide
- SAP HANA Administration Guide
- SAP HANA Master Guide

INSTALLING AND CONFIGURING EXPRESSCLUSTER

For information of the installation and settings of EXPRESSCLUSTER, refer to the "Installation and Configuration Guide".

5.1 Installing EXPRESSCLUSTER

Install this product on each node (Node#1 and Node#2).

For details on the installation please refer to the following document:

"Installation and Configuration Guide"
- "Installing EXPRESSCLUSTER"

After installing EXPRESSCLUSTER, execute the following command to install the sample script.

```
# rpm -i expresscls_spnw-<Version of EXPRESSCLUSTER>.x86_64.rpm
```

The sample script bundled with the media is stored under the following directory.

```
media/Linux/<Version of EXPRESSCLUSTER>/common/hana/sample
```

Set the following sample script bundled with the media to EXEC resource and custom monitor resource. The sample script is for SAP HANA control and monitoring.

In addition, it is necessary to add execution environment specific information to each sample script. Add settings according to the comment of the sample script.

File name	Use
start_hana_primary_sample.sh	For group resource exec_primary_hana
stop_hana_primary_sample.sh	For group resource exec_primary_hana
start_hana_secondary_sample.sh	For group resource exec_secondary_hana
stop_hana_secondary_sample.sh	For group resource exec_secondary_hana
genw_hana_primary_status_sample.sh	For monitor resource genw_primary_hana_status
genw_hana_secondary_status_sample.sh	For monitor resource genw_secondary_hana_status

5.2 Registering the license

The license must be registered before you can use EXPRESSCLUSTER.
For details on the installation please refer to the following document:

"Installation and Configuration Guide"
- "Registering the license"

This product contains the following four licenses.

Licensed Product Name
EXPRESSCLUSTER X for Linux
EXPRESSCLUSTER X Database Agent for Linux
EXPRESSCLUSTER X File Server Agent for Linux
EXPRESSCLUSTER X System Resource Agent for Linux

5.3 Creating a cluster

Create a cluster from Cluster WebUI.

For information of creating a cluster, refer to the following document and the online manual of Cluster WebUI.

"Installation and Configuration Guide"
- "Creating the cluster configuration data"

Register the following as network settings.

Application	Paths	Description
Interconnect LAN (doubling as a public LAN)	1	This is used to perform alive monitoring and to exchange cluster information for servers configuring an HA cluster.

5.4 Creating failover groups

Create the failover groups to which each node belongs from Cluster WebUI.

For information of creating a failover group, refer to the following document and the online manual of Cluster WebUI.

"Installation and Configuration Guide"

- "Creating the cluster configuration data"

Register the following failover groups.

Group type	Description
Primary failover group (failover_PRI)	Failover group that starts on the primary server. SAP HANA is started or stopped as the primary server. The virtual IP address used to access SAP HANA is also enabled or disabled.
Secondary failover group (failover_SEC)	Failover group that starts on the secondary server. SAP HANA is started or stopped as the secondary server. If SAP HANA is started on the same server as failover_PRI, SAP HANA is not started or stopped.

5.5 Add group resources

Register an Exec resource in the failover group created in the preceding section.

For information of registering EXEC resources, refer to the following document and the online manual of Cluster WebUI.

"Reference Guide"

- "Understanding EXEC resources"
- "Understanding AWS Virtual IP resources"

Register the following group resources in each failover group.

The SAP HANA services that are controlled by EXPRESSCLUSTER were set to not to start automatically.

Resource type (group resource name)	Failover group	Description
AWS Virtual IP resource (awsvip)	failover_PRI	Assigns a Virtual IP (VIP) address to an active sever instance (primary server), changes the route table of the assigned VIP address, and publishes operations within the VPC.
EXEC resource for primary control (exec_primary_hana)	failover_PRI	The script to start or stop SAP HANA as the primary server is executed. If SAP HANA has already been started as the secondary server, the started SAP HANA is changed to the primary server.
EXEC resource for secondary control (exec_secondary_hana)	failover_SEC	The script to start or stop SAP HANA as the secondary server is executed. If SAP HANA is started on the same node as the failover_PRI group, the full sync option is disabled.

5.6 Add monitor resources

Next, register a custom monitor resource.

For information of registering custom monitor resources, refer to the following document and the online manual of Cluster WebUI.

"Reference Guide"

- "Understanding custom monitor resources"
- "Understanding AWS AZ monitor resources"
- "Understanding AWS Virtual IP monitor resources"

Monitor type (monitor resource name)	Description	Primary	Secondary
Custom monitor resource for monitoring the primary server genw_ACTDB_hoststatus (genw_primary_hana_status)	The state of SAP HANA on the primary server is monitored by running the landscapeHostConfiguration.py command.	✓	
Custom monitor resource for monitoring the primary server genw_STBDB_hoststatus (genw_secondary_hana_status)	The state of SAP HANA on the secondary server is monitored by running the landscapeHostConfiguration.py command.		✓
AWS AZ monitor resource for monitoring Availability Zone (awsazw)	Periodically monitors the health of the AZ in which the local server exists by using Multi-AZ. If Multi-AZ is not used, this monitor resource can be used to monitor whether the AWS CLI is available.	✓	✓
AWS Virtual IP monitor resource for monitoring AWS Virtual IP (awsvipw)	Periodically monitors whether the VIP address assigned by the AWS Virtual IP resource exists in the local server and whether the VPC route table is changed illegally. (This monitor resource is automatically added when the AWS Virtual IP resource is added.)	✓	✓

INSTALLING AND CONFIGURING SAP ERP

Because there are no SAP ERP parameter settings specific to EXPRESSCLUSTER, SAP ERP was installed by using general procedures and parameters.

* As of October 15, 2014, the following must be observed when installing SAP ERP by using the SAP ERP6.0 EHP7 SR1 media.

If Database Host is set to a virtual host in the SAP System Database parameter, the connection destination of SAP HANA Client is not set on the virtual host after installation, and may be automatically replaced with a master host name of SAP HANA. The existing setting must be deleted and registered again by running the hdbuserstore command on the virtual host to set the virtual host as the connection destination of the Application Server.

NOTES AND RESTRICTIONS

- If a failure occurs before the full sync option is enabled, data might be lost because failover is performed before a full data copy is made.
- Be sure to start the primary failover group on the server that stores the latest data.

When a failover occurs, update differences might occur between the primary server and secondary server. When the primary server is storing the latest data, if the primary failover group is started on the secondary server and the secondary failover group is started on the primary server, data will be synchronized with the primary server, causing data loss.

- Data synchronization (system replication)

The system replication function of SAP HANA can cause data loss when an actual failure occurs, even in Synchronous mode. The "SAP Note 2063657 - HANA System Replication takeover decision guideline" provides criteria for takeover decision. Before executing the takeover, the operator must check these criteria. NEC adopts the full sync option in Synchronous mode. The possibility of data loss can be eliminated by using the full sync option together with EXPRESSCLUSTER. This setting is recommended by NEC.

8.1 Detailed Settings

The following is an example of setting up EXPRESSCLUSTER in this guide.

- **Cluster configuration**

	Parameter	Value
	Cluster Name	cluster
	Number of Servers	2
	Number of Failover Groups	2
Heartbeat resources	Number of LAN Heartbeat Resources	1
Node#1 (master server)	Server Name	hana01
	Public IP Address (Kernel mode, priority 1)	10.0.2.22
Node#2	Server Name	hana02
	Public IP Address (Kernel mode, priority 1)	10.0.12.22

- **1st group**

	Parameter	Value
	Type	Failover
	Group Name	failover_PRI
	Starting Server	Failover available on all servers
	Group Startup Attribute	Manual Startup
	Failover Attribute	Auto Failover Use the startup server settings.

Continued on next page

Table 8.2 – continued from previous page

	Parameter	Value
	Failback Attribute	Manual Failback
	Failover Exclusive Attribute	No Exclusion
	Start Wait Time	-
	Number of Group Resources	2
1st group resource Depth 0	Type	AWS VIP resource
	Group Resource Name	awsvip
	Final Action at Activation Failure	Activation Retry Threshold: 0 Failover Threshold: 1 No operation (Do not activate the next resource.)
	Final Action at Deactivation Failure	Deactivation Retry Threshold: 0 Stop the cluster service and shut down the OS.
	vpc-id	vpc-xxxxxxx
	eni-id(Node#1)	eni-yyyyyyy
	eni-id(Node#2)	eni-zzzzzzz
2nd group resource Depth 1	Type	EXEC resource
	Group Resource Name	exec_primary_hana
	Start Script Timeout	1800seconds ¹
	Stop Script Timeout	1800seconds ¹
	Dependency	awsvip
	Final Action at Activation Failure	Activation Retry Threshold: 0 Failover Threshold: 1 No operation (Do not activate the next resource.)
	Final Action at Deactivation Failure	Deactivation Retry Threshold: 0 Stop the cluster service and shut down the OS.
	Detail	Script list Start script / start.sh Stop script / stop.sh

• 2nd group

¹ Note for Red Hat Enterprise Linux users: For a very large scale system where SAP HANA takeover might take 30 minutes or more, make sure that the system does not time out by setting **Start script timeout** and **Stop script timeout** to 60 minutes (3600 seconds).

	Parameter	Value
	Type	Failover
	Group Name	failover_SEC
	Starting Server	Failover available on all servers
	Group Startup Attribute	Manual Startup
	Failover attribute	Auto Failover Use the startup server settings.
	Failback attribute	Manual Failback
	Start Wait Time	failover_PRI
	Number of Group Resources	1
3rd group resource Depth 0	Type	EXEC resource
	Group Resource Name	exec_secondary_hana
	Start Script Timeout	1800seconds(*)
	Stop Script Timeout	1800seconds(*)
	Final Action at Activation Failure	Activation Retry Threshold: 0 Failover Threshold: 1 No operation (Do not activate the next resource.)
	Final Action at Deactivation Failure	Deactivation Retry Threshold: 0 Stop the cluster service and shut down the OS.
	Detail	Script list Start script / start.sh Stop script / stop.sh

• **1st monitor resource(Default)**

Parameter	Value
Type	user mode monitor
Monitor Resource Name	userw

• **2nd monitor resource**

Parameter	Value
Type	aws vip monitor
Monitor Resource Name	awsvipw
Interval	60 seconds

Continued on next page

Table 8.5 – continued from previous page

Parameter	Value
Timeout	60 seconds
Retry Count	3 times
Recovery Action	Execute failover the recovery target
Recovery Target	awsvip
Final Action	Stop the cluster service and shutdown OS

• **3rd monitor resource**

Parameter	Value
Type	custom monitor
Monitor Resource Name	genw_primary_hana_status
Interval	30 seconds
Timeout	120 seconds
Retry Count	3 times
Wait Time to Start Monitoring	0 seconds
Monitor Target	At activation Target Resource: exec_primary_hana
Script created with this product	genw.sh
Normal Return Value	0
Recovery Action	Execute failover the recovery target
Recovery Target	failover_PRI
Final Action	No operation

• **4th monitor resource**

Parameter	Value
Type	custom monitor
Monitor Resource Name	genw_secondary_hana_status
Interval	30 seconds
Timeout	120 seconds
Retry Count	3 times
Wait Time to Start Monitoring	0 seconds
Monitor Timing	At activation Target Resource: exec_secondary_hana
Script created with this product	genw.sh
Normal Return Value	0
Recovery Action	Execute failover the recovery target
Recovery Target	failover_SEC
Final Action	No operation

• **5th monitor resource**

Parameter	Value
Type	AWS AZ monitor

Continued on next page

Table 8.8 – continued from previous page

Parameter	Value
Monitor Resource Name	awsazw
Interval	60 seconds
Timeout	120 seconds
Retry Count	0 times
Wait Time to Start Monitoring	0 seconds
Monitor Timing	Always
Availability Zone	ap-northeast-1a, ap-northeast-1c
Recovery Action	Execute failover the recovery target
Recovery Target	All Groups
Final Action	No operation

8.2 Operating Procedure

This section describes how to start a cluster and how to recover from failure.

- **Starting a cluster**

Server #1 is used as the primary server, and Server #2 is used as the secondary server.

The primary failover group is started on Server #1 and the secondary failover group on Server #2. (SAP HANA starts as the primary database on Server #1 and as the secondary database on Server #2.)

After the failover group has started, a command is run manually on Server #1 to enable the full sync option of SAP HANA.

Note: If a failure occurs before the full sync option is enabled, data might be lost because failover is performed before a full data copy is made.




Fig. 8.1: Normal Operation

- **Recovering from failure that occurred on the primary server**

When a failure occurs on Server #1, the primary failover group fails over to Server #2. SAP HANA on Server #1 stops, and SAP HANA on Server #2 takes over operations.

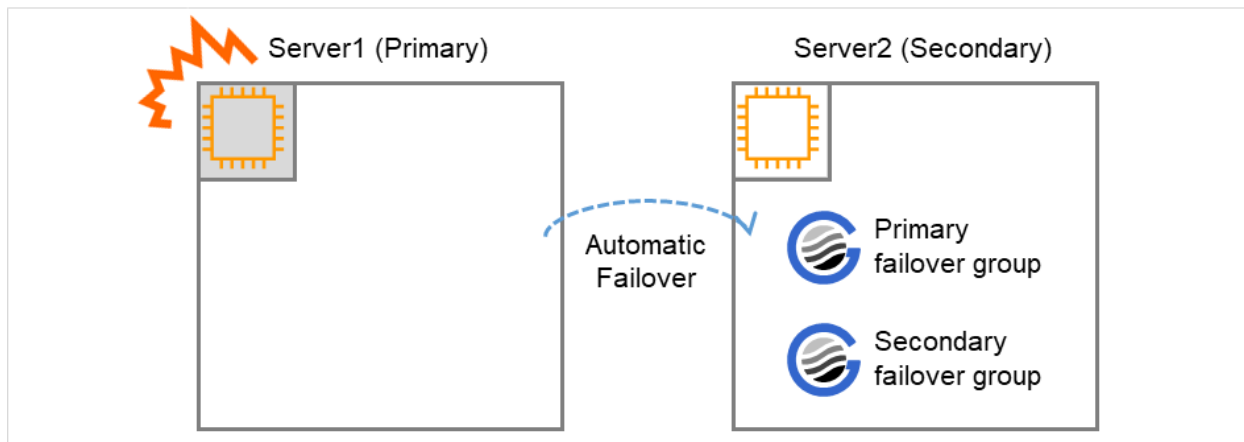


Fig. 8.2: Occurrence of Failure on the Primary Server

- **Recovery procedure**

The secondary failover group is failed over from Server #2 to Server #1 manually.

When the failover is executed, SAP HANA on Server #1 starts as the secondary system.

When the failover is complete, a command is run manually on Server #2 to enable the full sync option of SAP HANA.

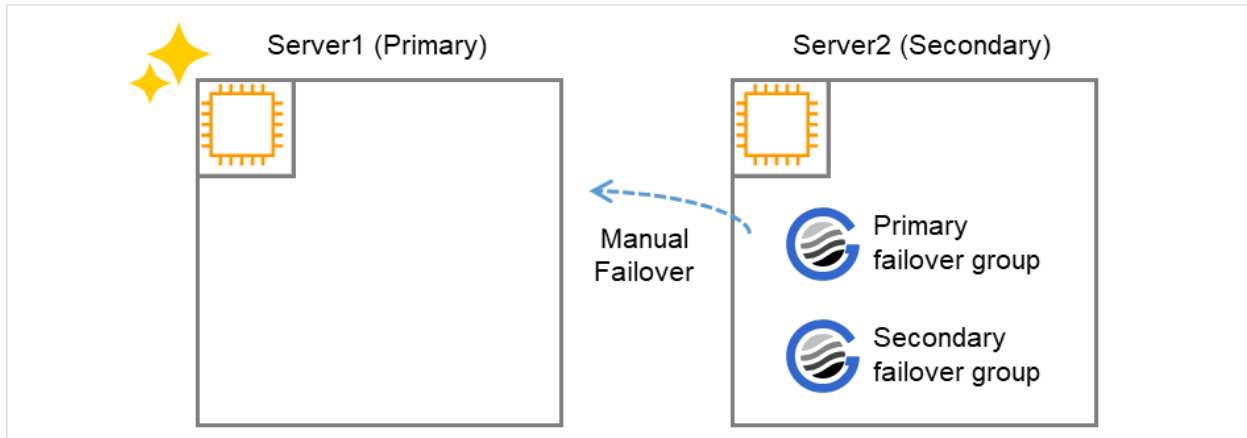


Fig. 8.3: Failure Recovery on the Primary Server

- **When a failure occurs on the secondary server**

When a failure occurs on Server #2, the secondary failover group fails over to Server #1. SAP HANA on Server #2 stops, and operations continue on Server #1 with the full sync option of SAP HANA disabled.




Fig. 8.4: Occurrence of Failure on the Secondary Server

- **Recovery procedure**

The secondary failover group is failed over from Server #1 to Server #2 manually.

When the failover is executed, SAP HANA on Server #2 starts as the secondary database.

When the failover is complete, a command is run manually on Server #1 to enable the full sync option of SAP HANA.

Note:

Be sure to start the primary failover group on the server that stores the latest data.
When a failover occurs, update differences might occur between the primary server and secondary server. When the primary server is storing the latest data, if the primary failover group is started on the secondary server and the secondary failover group is started on the primary server, data will be synchronized with the primary server, causing data loss.




Fig. 8.5: Failure Recovery on the Secondary Server

8.3 Detailed verification results

NEC verified that the state transitions of the servers and resource groups were correct by performing the following state transitions.

In the normal system replication setting, servers must be switched manually when a failure occurs. In the configuration with EXPRESSCLUSTER, EXPRESSCLUSTER automatically executes all operations from failure detection to failover when a failure occurs.

NEC has also verified that the potential for data loss can be eliminated by using the full sync option, and that operations can continue without stopping because EXPRESSCLUSTER automatically disables the full sync option when a failure occurs on the secondary server.

Item	Operation	Verification result
Start cluster	The cluster was started from Cluster WebUI. The primary failover group was started on Server #1 and the secondary failover group was started on Server #2 from Cluster WebUI.	The cluster started. The primary failover group started on Server #1, and the secondary failover group started on Server #2. SAP HANA on Server #1 started as the primary database, and SAP HANA on Server #2 started as the secondary database.
Stop cluster	The cluster was stopped from Cluster WebUI.	The cluster stopped. SAP HANA on both Server #1 and Server #2 stopped.
Restart cluster	The primary failover group was started on Server #1 and the secondary failover group was started on Server #2 from Cluster WebUI.	The cluster started. The primary failover group started on Server #1, and the secondary failover group started on Server #2. SAP HANA on Server #1 started as the primary database, and SAP HANA on Server #2 started as the secondary database.
Shut down Server #1	Server #1 was shut down from Cluster WebUI.	Server #1 shut down after SAP HANA stopped. The primary failover group failed over from Server #1 to Server #2. (SAP HANA on Server #1 stopped. SAP HANA on Server #2 took over operations, allowing SAP HANA operations to continue.)
Recover Server #1	Server #1 was started.	Server #1 started and returned to the cluster.

Continued on next page

Table 8.9 – continued from previous page

Item	Operation	Verification result
Move SAP failover group	The secondary failover group was moved from Server #2 to Server #1 from Cluster WebUI.	The secondary failover group moved from Server #2 to Server #1. SAP HANA on Server #1 started as the secondary database.
Shut down Server #1	Server #1 was shut down from Cluster WebUI.	Server #1 shut down after SAP HANA stopped. The secondary failover group failed over from Server #1 to Server #2. (SAP HANA on Server #1 stopped. SAP HANA on Server #2 took over operations, allowing SAP HANA operations to continue.)
Recover Server #1	Server #1 was started.	Server #1 started and returned to the cluster.
Move SAP failover group	The secondary failover group was moved from Server #2 to Server #1 from Cluster WebUI.	The secondary failover group moved from Server #2 to Server #1. SAP HANA on Server #1 started as the secondary database.
Shut down Server #2	Server #2 was shut down from Cluster WebUI.	Server #2 shut down after SAP HANA stopped. The primary failover group failed over from Server #2 to Server #1. (SAP HANA on Server #1 took over operations, allowing SAP HANA operations to continue.)
Recover Server #2	Server #2 was started.	Server #2 started and returned to the cluster.
Move SAP failover group	Move a secondary failover group from Server #1 to Server #2 from Cluster WebUI.	The secondary failover group moved from Server #1 to Server #2. SAP HANA on Server #2 started as the secondary database.

Continued on next page

Table 8.9 – continued from previous page

Item	Operation	Verification result
Shut down Server #2	Server #2 was shut down from Cluster WebUI.	Server #2 shut down after SAP HANA stopped. The primary failover group failed over from Server #2 to Server #1. (SAP HANA on Server #1 took over operations, allowing SAP HANA operations to continue.)
Recover Server #2	Server #2 was started.	Server #2 started and returned to the cluster.
Move SAP failover group	The secondary failover group was moved from Server #1 to Server #2 from Cluster WebUI.	The secondary failover group moved from Server #1 to Server #2. SAP HANA on Server #2 started as the secondary database.
Reboot cluster	The cluster was rebooted from Cluster WebUI. After the cluster was rebooted, the primary failover group was started on Server #1 and the secondary failover group was started on Server #2 from Cluster WebUI.	The cluster rebooted. SAP HANA on both Server #1 and Server #2 stopped. After Server #1 and Server #2 rebooted, the primary failover group started on Server #1, and the secondary failover group started on Server #2. SAP HANA on Server #1 started as the primary database, and SAP HANA on Server #2 started as the secondary database.
Suspend cluster	The cluster was suspended from Cluster WebUI.	The cluster temporarily stopped operations. SAP HANA continued to run.
Resume cluster	The cluster was resumed from Cluster WebUI.	The cluster resumed operations. SAP HANA continued to run.

NEC verified that no problems occurred in any of the above operations by hypothesizing hardware and software failure and generating pseudo failures on the following components.

- **AWS infrastructure**

Item	Operation	Verification result
AWS AZ monitor resource for Availability Zone failure (awsazw)	A pseudo failure (verification mode) was generated on Server #1 while Server #1 was the primary server and Server #2 was the secondary server.	The failure was detected and the primary failover group was failed over. (SAP HANA on Server #1 stopped. SAP HANA on Server #2 took over operations, allowing SAP HANA operations to continue.)
AWS AZ monitor resource for Availability Zone failure (awsazw)	A pseudo failure (verification mode) was generated on Server #2 while Server #1 was the primary server and Server #2 was the secondary server.	The failure was detected and the secondary failover group failed over. (SAP HANA on Server #2 stopped. Operations continued on Server #1, with the SAP HANA full sync option disabled.)

• **Network**

Item	Operation	Verification result
Network failure (Primary)	A network failure was generated on Server #1 while Server #1 was the primary server and Server #2 was the secondary server. (The network access control list (ACL) of the Server #1 subnet was changed on the AWS console and all communications were blocked.)	The IP monitor detected the failure and Server #1 shut down. The primary failover group failed over. (SAP HANA on Server #2 took over operations, allowing SAP HANA operations to continue.)
Network failure (Secondary)	A network failure was generated on Server #2 while Server #1 was the primary server and Server #2 was the secondary server. (The network ACL of the Server #2 subnet was changed on the AWS console and all communications were blocked.)	The IP monitor detected the failure and Server #2 shut down. The secondary failover group failed over. (Operations continued on Server #1, with the SAP HANA full sync option disabled.)

• **OS**

Item	Operation	Verification result
Server alive monitoring (Primary)	Server #1 was stopped while Server #1 was the primary server and Server #2 was the secondary server. (The shutdown -n -r now command was run.)	The primary failover group failed over. (SAP HANA on Server #2 took over operations, allowing SAP HANA operations to continue.)
Server alive monitoring (Secondary)	Server #2 was stopped while Server #1 was the primary server and Server #2 was the secondary server. (The shutdown -n -r now command was run.)	The secondary failover group failed over. (SAP HANA on Server #1 took over operations, allowing SAP HANA operations to continue.)

• **SAP HANA**

Item	Operation	Verification result
custom monitor (genw_primary_hana_status)	The SAP HANA process (Indexserver) was stopped on Server #1 while Server #1 was the primary server and Server #2 was the secondary server. (kill -9 was run.)	The failure was detected and the primary failover group failed over. (SAP HANA on Server #1 stopped. SAP HANA on Server #2 took over operations, allowing SAP HANA operations to continue.)
custom monitor (genw_secondary_hana_status)	The SAP HANA process (Indexserver) was stopped on Server #2 while Server #1 was the primary server and Server #2 was the secondary server. (kill -9 was run.)	The failure was detected and the secondary failover group failed over. (SAP HANA on Server #2 stopped. SAP HANA on Server #1 took over operations, allowing SAP HANA operations to continue.)

LEGAL NOTICE

9.1 Disclaimer

- Information in this document is subject to change without notice.
- NEC Corporation is not liable for technical or editorial errors or omissions in the information in this document.
- You are completely liable for all risks associated with installing or using the product as described in this manual to obtain expected results and the effects of such usage.
- The information in this document is copyrighted by NEC Corporation.
- No part of this document may be reproduced or transmitted in any form by any means, electronic or mechanical, for any purpose, without the express written permission of NEC Corporation.

9.2 Trademark Information

- EXPRESSCLUSTER® is a registered trademark of NEC Corporation.
- SAP, SAP HANA and other SAP products and services mentioned in this document as well as their respective logos are trademarks or registered trademarks of SAP SE in Germany and other countries.
- Linux is a registered trademark of Linus Torvalds in the United States and other countries.
- RPM is a registered trademark of Red Hat, Inc. or its subsidiaries in the United States and other countries.
- SUSE is a registered trademark of SUSE LLC in the United States and other countries.
- Amazon Web Services and all AWS-related trademarks, as well as other AWS graphics, logos, page headers, button icons, scripts, and service names are trademarks, registered trademarks or trade dress of AWS in the United States and/or other countries.
- Microsoft, Windows, Windows Server, Internet Explorer, Azure, and Hyper-V are registered trademarks of Microsoft Corporation in the United States and other countries.
- Other product names and slogans written in this manual are trademarks or registered trademarks of their respective companies.

REVISION HISTORY

Edition	Revised Date	Description
1st	Apr 10, 2019	New guide
2nd	June 5, 2019	Added SAP HANA 2.0 SPS03 on Operating Environment <i>2.2. Operating Environment</i>
3rd	Mar 13, 2020	Added link to SAP HANA Platform 2.0 SPS 03 guide <i>4. Installing and configuring SAP HANA</i>
4th	Apr 10, 2020	Added the information of internal version 4.2.0-1 <i>2.2. Operating Environment</i>
5th	May 13, 2020	Corrected Edition Number and Revision History
6th	Apr 09, 2021	Corrected the appearance.
7th	Apr 08, 2022	Corrected typographical errors.

© Copyright NEC Corporation 2019. All rights reserved.