Empowered by Innovation NEC

White Paper

SAP Business Suite Powered by SAP HANA High Availability with EXPRESSCLUSTER

October 12, 2016

EXPRESSCLUSTER X

Disclaimer

Information in this document is subject to change without notice.
NEC Corporation is not liable for technical or editorial mistakes in or omissions from this document.
In addition, whether the customer achieves the desired effectiveness by following the introduction and usage instructions in this document is the responsibility of the customer.
The copyright of the contents in this document belong to 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.
The contents of this document are based on the verification results obtained as of the publication date of this document. The specifications of the related software and infrastructure may change in the future, in which case these verification results will not apply.

Trademark Information

ExpressCluster® X is a registered trademark of NEC Corporation.
Linux is a registered trademark or trademark of Linus Torvalds in the United States and other countries.
SUSE is a registered trademark of SUSE LLC. in the United States and other countries.
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.
Amazon Web Services and all the trademarks related to AWS, other AWS graphics, logs, page headers, button icons, scripts and service names are trademarks, registered trademarks, or trade dress of Amazon Web Services in the United States and/or other countries.
Other product names and slogans written in this manual are trademarks or registered trademarks of their respective companies.

CONTENTS

1 Introduction

1-1 Background

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.

1-2 Purpose of verification

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 standby server (performs failover). NEC wished to verify whether EXPRESSCLUSTER could shorten operational downtime and boost operational efficiency by cooperating with SAP HANA. The verification procedure and results are described in this document.

1-3 Overview

1-3-1 Verification procedure

For verification, NEC created a SAP HANA cluster environment on AWS by using EXPRESSCLUSTER.

Various types of failures were hypothesized on the created environment and it was verified that a cluster system could be restored by data synchronization using the EXPRESSCLUSTER automatic failover function and SAP HANA system replication function, and that operations could be continued without pause (that is, that SAP ERP Application Server automatically connected SAP HANA again and operations continued without stopping).

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.

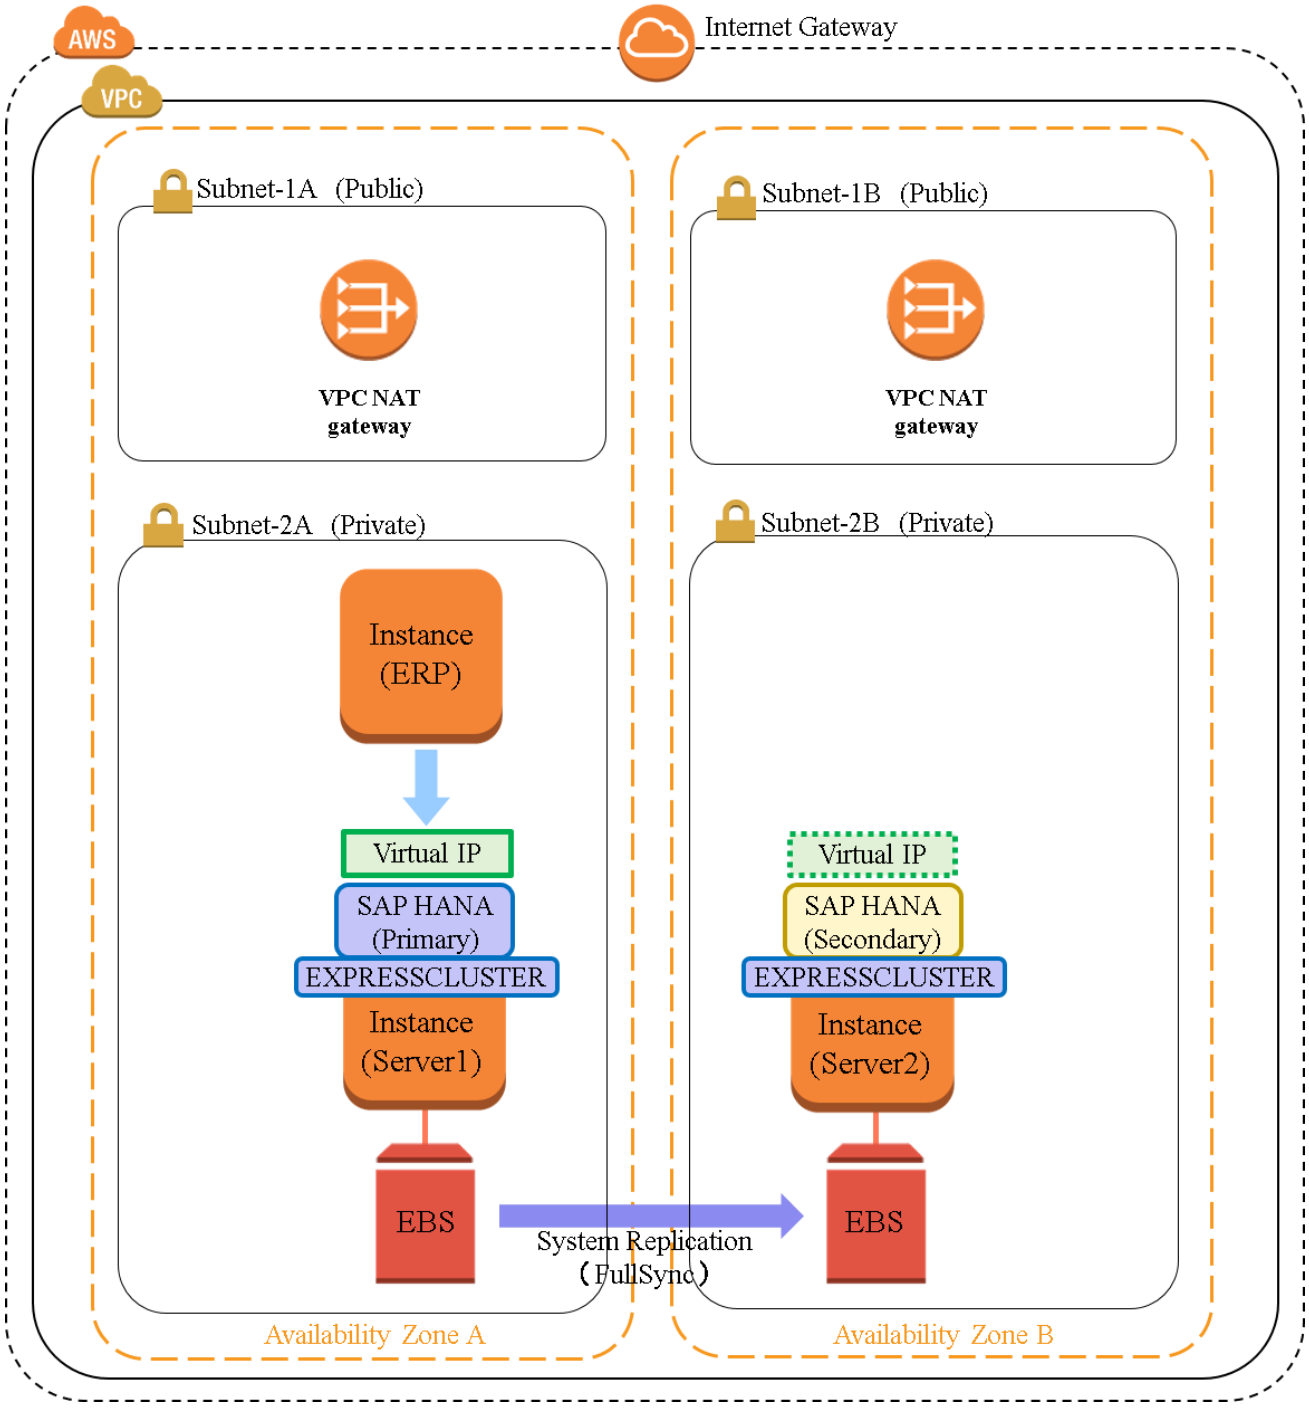

Figure 1-1 System Configuration

- Availability on AWS

  AWS has multiple data centers called Availability Zones in locations such as Tokyo and Singapore. Customers can select the Availability Zone that they want to use and freely determine the Availability Zone in which to allocate an EC2 instance. Availability Zones are connected via high-speed dedicated lines. A system can be created across multiple Availability Zones. To realize the high availability required by mission-critical systems, the two instances composing a cluster must be allocated to different Availability Zones.

- Failover on AWS

  In cluster configuration, the connection destinations of the cluster must be able to be switched transparently. The virtual private cloud (VPC) of AWS can be used to set the network routing (Route Table), and the network routing can be operated by using an application program interface (API). Connection destinations can be switched by using this API and routing a virtual IP address (virtual IP in the above figure) to the elastic network interface (ENI) of the instance.

- Amazon EC2 X1 instance

  The X1 instance is a SAP-certified instance for production workloads. This instance satisfies the performance requirements of both the SAP OLAP and OLTP workloads that are necessary for SAP HANA. The high availability and operational efficiency required by mission critical systems can be easily implemented by leveraging the large-scale and high performance features of the X1 instance.

- 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 (http://service.sap.com/sap/support/notes/2063657)” provides criteria for takeover decision. Before executing the takeover, the operator must check these criteria.

  * To reference SAP Note, you need to register as a user to the SAP Support Portal.

  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.

1-3-2 Illustration of operation

Figure 1-2 shows an illustration of the system 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.

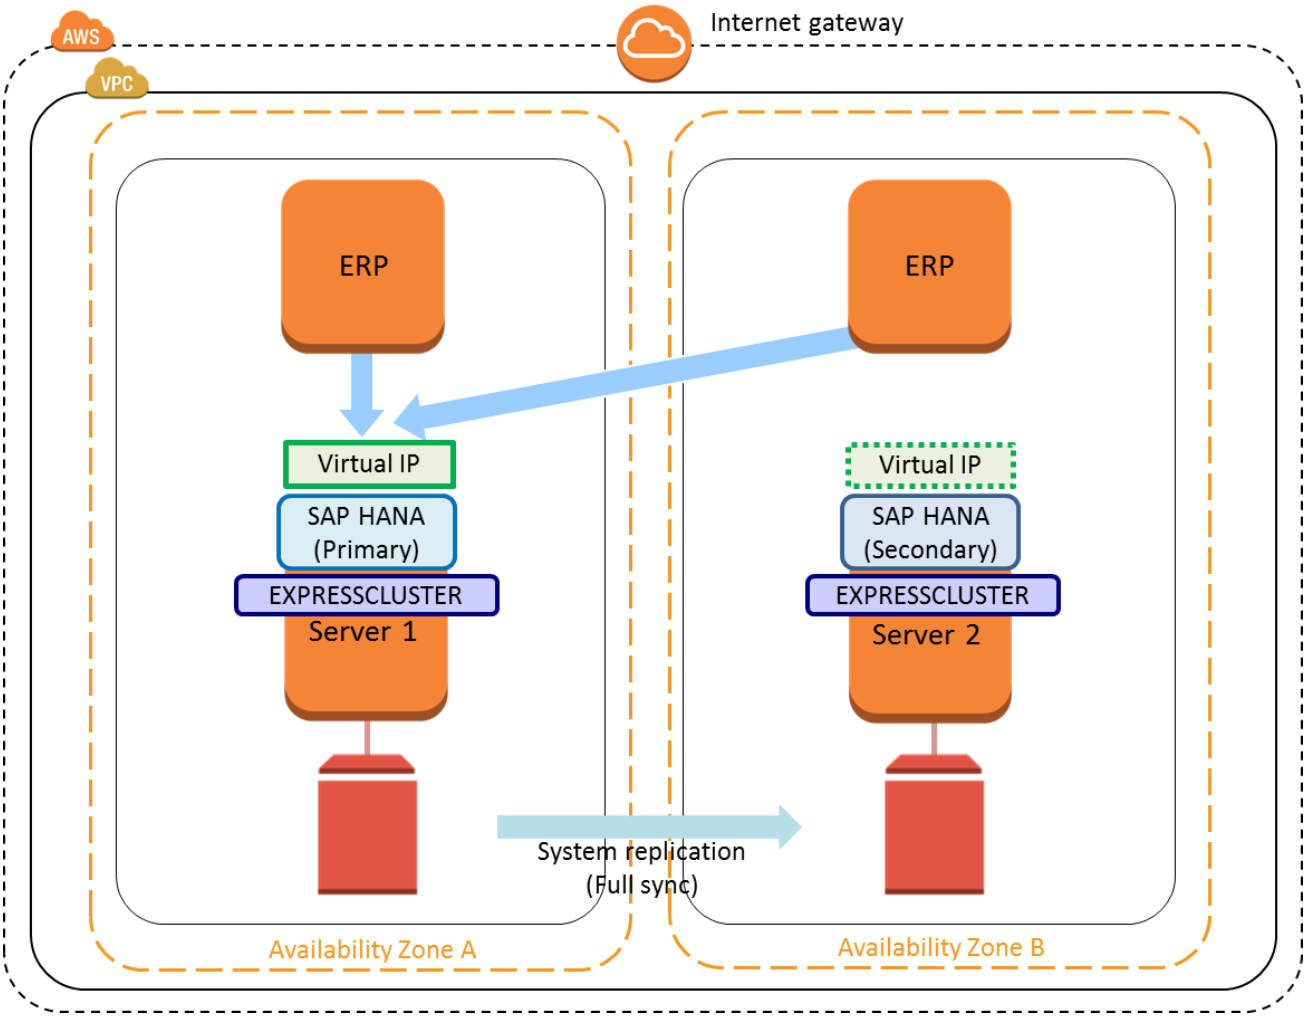

Figure 1-2 Illustration of Normal Operation

Figure 1-3 shows an illustration of the operation when 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.

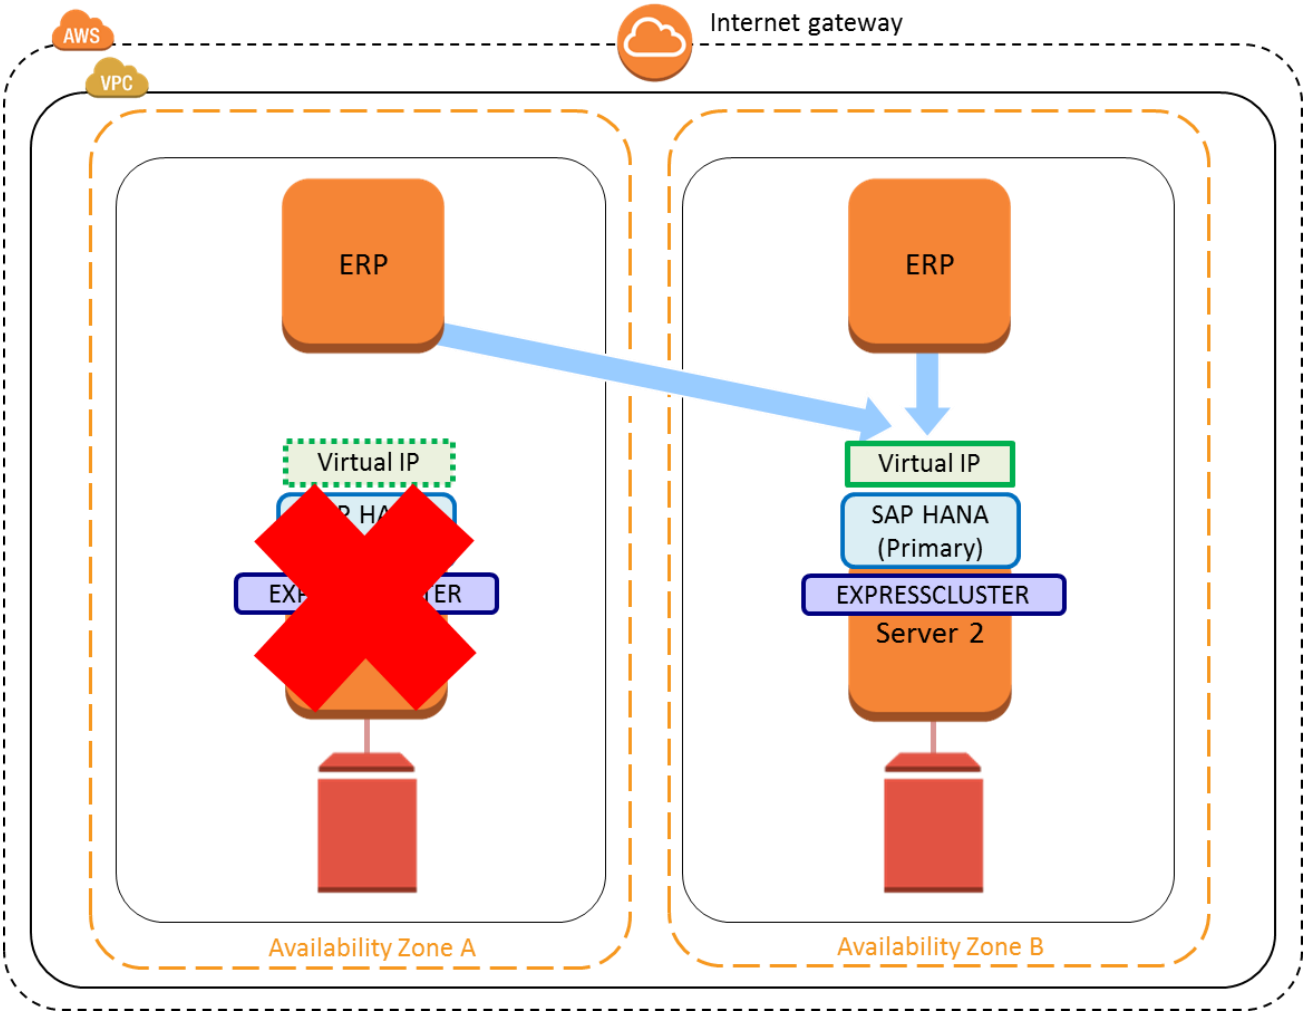

Figure 1-3 Illustration of Operation When a Failure Occurs on the Primary Server

Figure 1-4 shows an illustration of the operation when a failure occurs on the secondary server.

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.

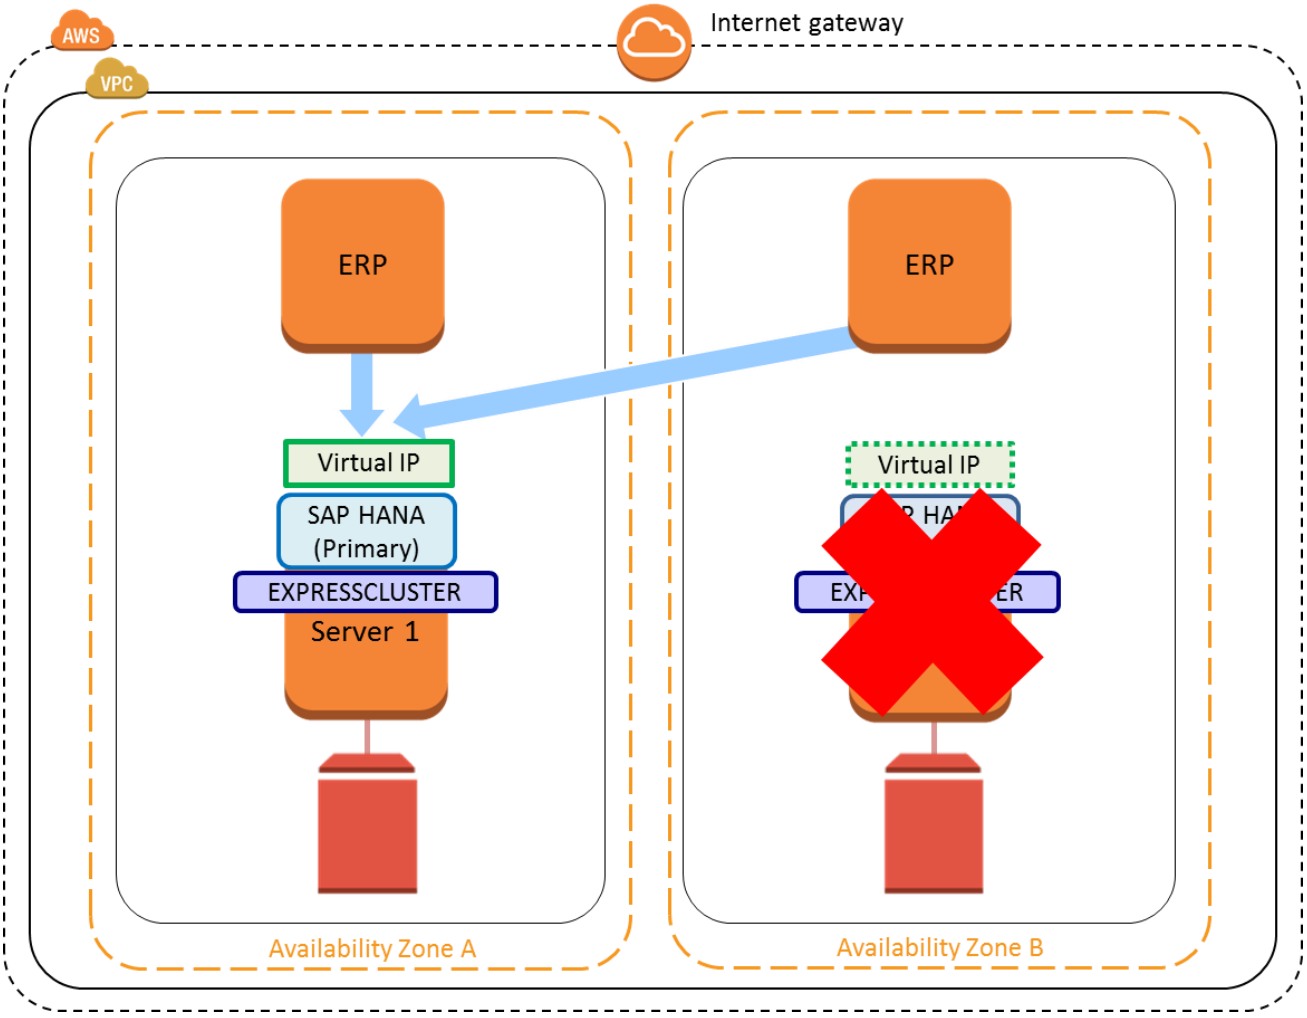

Figure 1-4 Illustration of Operation When a Failure Occurs on the Secondary Server

2 Supported scenarios and requirements

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

1. Two-node cluster consisting of scale-up (single) configuration x 2
2. Both nodes must belong to the same network segment.
3. Both nodes must be run as a single instance. No quality assurance or development system is running.
4. SAP HANA SPS09 (revision 90) or later
5. The automatic startup attribute of SAP HANA must be set to “off.” (SAP HANA startup is managed by EXPRESSCLUSTER.)
6. 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.

3 Verification configuration

3-1 Configuration diagram

This verification uses the following configuration.

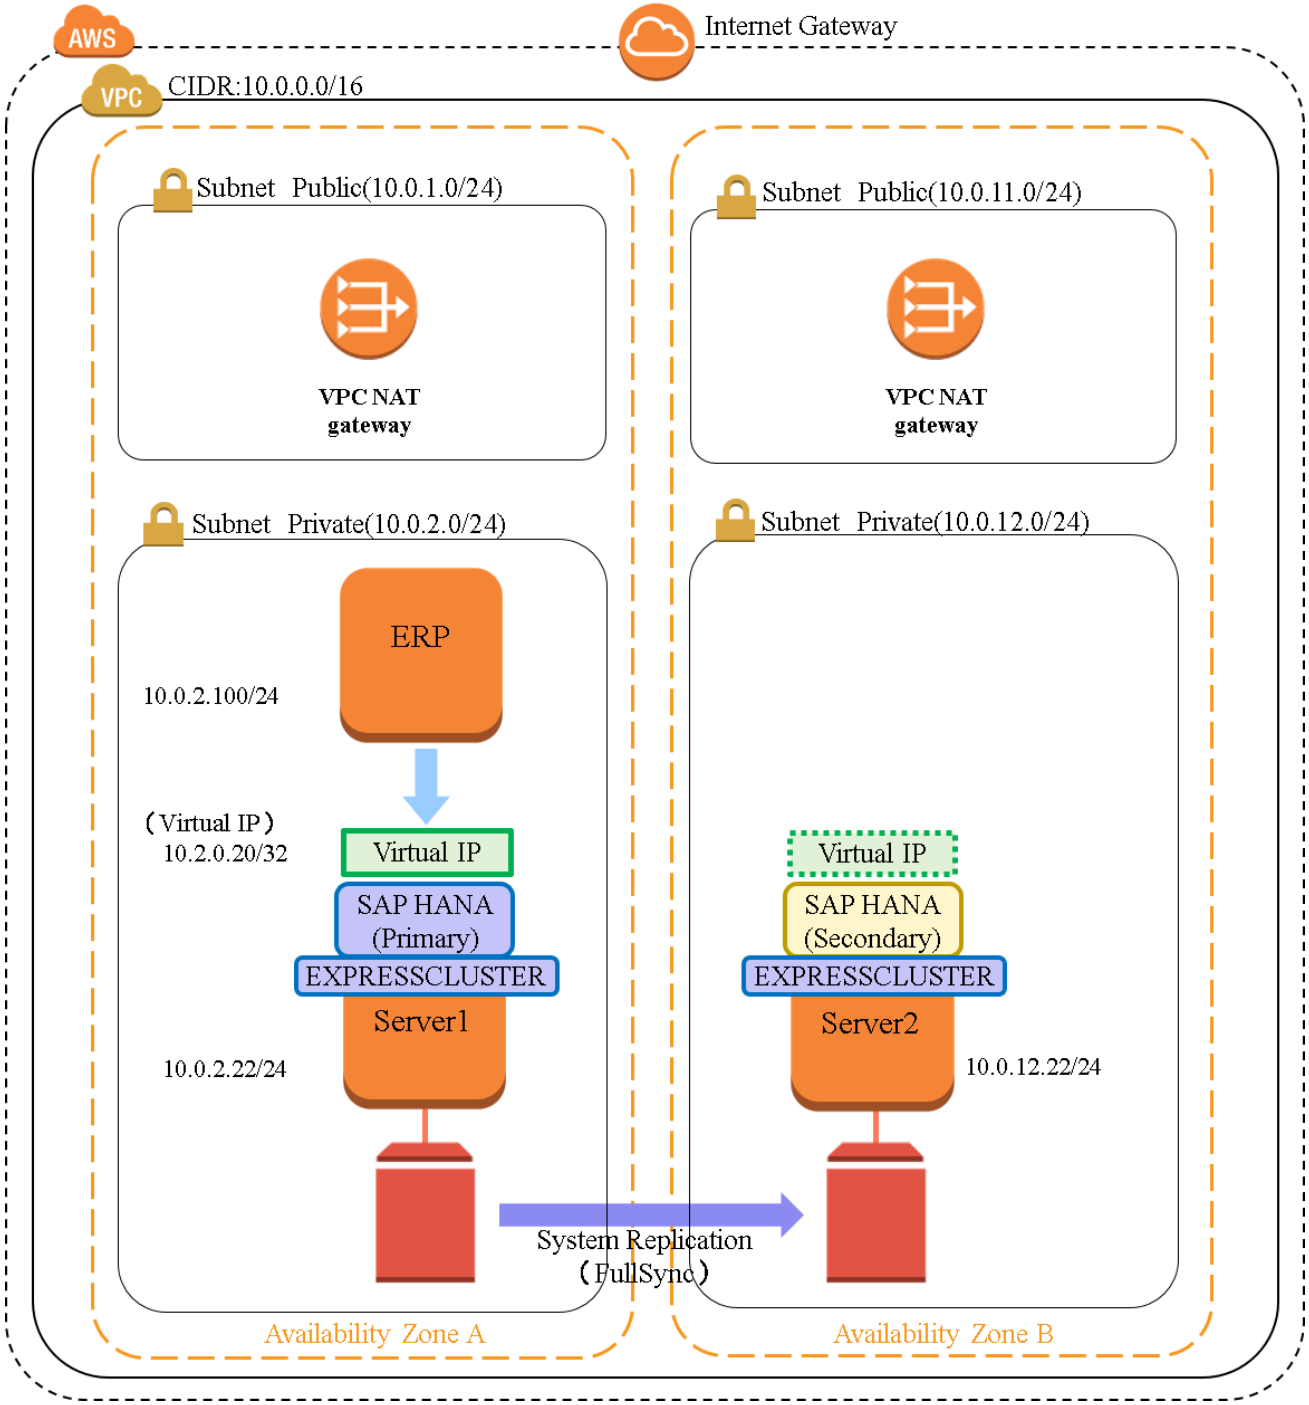

Figure 3-1 System Configuration Diagram

3-2 Operating environment

In this verification, a cluster environment is configured by allocating SAP HANA instances to different Availability Zones of AWS and installing SAP HANA as shown in *Figure 3-1 System Configuration Diagram*.

In this verification, SAP HANA is configured by using AWS CloudFormation.

The AWS instance types were determined by referring to the following and selecting a supported environment:
SAP Note 1964437 - SAP HANA on AWS: Supported AWS EC2 products
SAP Note 1656099 - SAP Applications on AWS: Supported DB/OS and AWS EC2 products
In this verification, the X1 instance is used. The X1 instance is the latest memory optimized instance. For SAP HANA, a configuration in which multiple tenant databases are created on SAP Instance is also verified.

SAP HANA (Common)

AMI	suse-sles-11-sp4-v20160301-hvm-ssd-x86_64 (ami-03a0ad6d)
Region	Asia Pacific (Tokyo)
OS	SUSE Linux Enterprise Server 11 SP4
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
EXPRESSCLUSTER	EXPRESSCLUSTER X 3.3

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

AMI	Windows_Server-2012-R2_RTM-English-64Bit-Base-2016.05.11 (ami-447a9d25)
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

3-3 Setup

The verification environment was configured as described below.

1. SAP HANA was installed and set up.
2. EXPRESSCLUSTER was installed and set up.
3. SAP ERP was installed and set up.

3-3-1 SAP HANA

SAP HANA was installed and upgraded to SPS09 or later following the procedures in *SAP HANA Server Installation and Update Guide*.
(http://help.sap.com/hana/SAP_HANA_Server_Installation_Guide_en.pdf)

System replication (Synchronous with full sync option) was set up following the procedures in *SAP HANA Administration Guide*.
(http://help.sap.com/hana/SAP_HANA_Administration_Guide_en.pdf)

3-3-2 EXPRESSCLUSTER

EXPRESSCLUSTER was set up as described below.

Networks

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.

Failover groups

EXPRESSCLUSTER groups the resources required to continue operations as a failover group and performs failover in operation units. In this verification, the following failover groups were registered.

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.

Group type	Description
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.

Group resources

In EXPRESSCLUSTER, the resources required for operations are called *group resources*, and registered in a failover group. In this verification, the following group resources were registered.

Resource type	Failover group	Description
EXEC resource for virtual IP address (exec_IPAddress)	failover_PRI	A virtual IP address is switched by replacing the IP address or by adding an alias by using Amazon EC2 API Tools, an AWS application programming interface (API).
EXEC resource for primary control (exec_HANA_Primary)	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_HANA_Secondary)	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.

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

Monitor resources

In EXPRESSCLUSTER, the resources used for monitoring are called *monitor resources*. In this verification, the following monitor resources were registered.

Monitor type	Description	Primary	Secondary
Custom monitor for monitoring the primary server genw_ACTDB_hoststatus	The state of SAP HANA on the primary server is monitored by running the landscapeHostConfiguration.py command.	✓	
Custom monitor for monitoring the primary server genw_STBDB_hoststatus	The state of SAP HANA on the secondary server is monitored by running the landscapeHostConfiguration.py command.		✓
Custom monitor for monitoring Availability Zone genw_azw	The health of the Availability Zone is checked by running the AWS API, Amazon EC2 API Tools.	✓	✓

Monitor type	Description	Primary	Secondary
IP monitor ipw	Communication with a NAT instance is monitored and the health of communication between subnets is checked.	✓	✓

3-3-3 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.

4 Verification items

4-1 Verification scenario

NEC tested the availability of the SAP HANA cluster configuration running on AWS using EXPRESSCLUSTER when the following failures occurred.

Failure type	Server	Component	Failure
Hardware failure	Primary	Server	Server down
		Network	Network down
	Secondary	Server	Server down
		Network	Network down
Software failure	Primary	OS	OS hung-up
		SAP HANA DB	Service down
			Process down
	Secondary	OS	OS hung-up
		SAP HANA DB	Service down
			Process down
Cloud failure	Primary	Availability Zone	Zone down
	Secondary	Availability Zone	Zone down

The following operations when the above mentioned failures occurred were checked and verified:

- EXPRESSCLUSTER detected the failure and failed over SAP HANA.
- The connection from SAP ERP remained available, and operations could continue. (Data could be updated and referenced.)

5 Verification results

This section describes the actions that should occur when a failure occurs.

Failure type	Server	Component	Failure	Desired action	Result
Hardware failure	Primary	Server	Server down	Failover (to a standby server)	✓
		Network	Network down	Failover (to a standby server)	✓
	Secondary	Server	Server down	Failover (to an active server)	✓
		Network	Network down	Failover (to an active server)	✓
Software failure	Primary	OS	OS hung-up	Failover (to a standby server)	✓
		SAP HANA DB	Service down	Failover (to a standby server)	✓
			Process down	Failover (to a standby server)	✓
	Secondary	OS	OS hung-up	Failover (to an active server)	✓
		SAP HANA DB	Service down	Failover (to an active server)	✓
			Process down	Failover (to an active server)	✓
Cloud failure	Primary	Availability Zone	Zone down	Failover (to a standby server)	✓
	Secondary	Availability Zone	Zone down	Failover (to an active server)	✓

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.

6 Conclusion

NEC has verified that the SAP environment can be configured on the Amazon EC2 X1 instance, enabling monitoring of a wide range of failures, from failures in the OS layer to failures in SAP, thereby allowing failures to be detected quickly. The SAP environment also provides business continuity by performing automatic failover when a failure is detected. NEC has also verified that cooperation between SAP HANA and EXPRESSCLUSTER can shorten operational downtime and realize the high availability and operational efficiency required for mission-critical systems.

7 Supplement

7-1 Detailed settings

This section describes an example of the EXPRESSCLUSTER settings used for the configuration in this document.

For how to install and set up EXPRESSCLUSTER, see the relevant manual.

* The required EXPERSSCLUSTER resources vary depending on the OS used. This section describes setting examples for SUSE Linux Enterprise Server and Red Hat Enterprise Linux.

Example of EXPRESSCLUSTER settings for SUSE Linux Enterprise Server

	Parameter	Value
Cluster configuration	Cluster Name	cluster
	Number of Servers	2
	Number of Failover Groups	2
	Number of Monitor Resources	4
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	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.
	Failback Attribute	Manual Failback
	Failover Exclusive Attribute	No Exclusion
	Start Wait Time	----------
	Number of Group Resources	2
1st group resource Depth 0	Type	EXEC resource
	Group Resource Name	exec_IPAddress
	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.

		Parameter	Value
		Detail	Script list Start script / start.sh Stop script / stop.sh
	2nd group resource Depth 1	Type	EXEC resource
		Group Resource Name	exec_primary_hana
		Dependency	exec_IPAddress
		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		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
	1st group resource Depth 0	Type	EXEC resource
		Group Resource Name	exec_secondary_hana
		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)		Type	user mode monitor
		Monitor Resource Name	userw
		Recovery Target	Local Server
		Final Action	Stop the cluster service and shut down the OS.
2nd monitor resource		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

	Parameter	Value
	Script created with this product	genw.sh
	Nice Value	0 Target Resource:
	Recovery Action	Execute failover the recovery target
	Recovery Target	failover_PRI
	Final Action	No operation
3rd monitor resource	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
	Nice Value	0
	Recovery Action	Execute failover the recovery target
	Recovery Target	failover_SEC
	Final Action	No operation
4th monitor resource	Type	Custom monitor
	Monitor Resource Name	genw_azw
	Interval	60 seconds
	Timeout	120 seconds
	Retry Count	0 times
	Wait Time to Start Monitoring	0 seconds
	Monitor Timing	Always
	Script created with this product	genw.sh
	Nice Value	0
	Recovery Action	Custom setting
	Recovery Target	Local Server
	Recovery Script Execution Count	0 times
	Maximum Reactivation Count	3 times
	Maximum Failover Count	Once
	Final Action	No operation

Example of EXPRESSCLUSTER settings for Red Hat Enterprise Linux

	Parameter	Value
Cluster configuration	Cluster Name	cluster
	Number of Servers	2
	Number of Failover Groups	2
	Number of Monitor Resources	4
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	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.
	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-xxxxxxxx
	eni-id(Node#1)	eni-yyyyyyyy
	eni-id(Node#2)	eni-zzzzzzzz
2nd group resource Depth 1	Type	EXEC resource
	Group Resource Name	exec_primary_hana
	Start Script Timeout	1800seconds(*)
	Stop Script Timeout	1800seconds(*)
	Dependency	awsvip
	Final Action at Activation Failure	Activation Retry Threshold: 0
		Failover Threshold: 1
		No operation (Do not activate the next resource.)

	Parameter	Value
	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	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
1st 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)	Type	user mode monitor
	Monitor Resource Name	userw
	Recovery Target	Local Server
	Final Action	Stop the cluster service and shut down the OS.
2nd monitor resource	Type	aws vip monitor
	Monitor Resource Name	awsvipw
	Interval	60 seconds
	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	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

	Parameter	Value
	Nice Value	0 Target Resource:
	Recovery Action	Execute failover the recovery target
	Recovery Target	failover_PRI
	Final Action	No operation
4th monitor resource	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
	Nice Value	0
	Recovery Action	Execute failover the recovery target
	Recovery Target	failover_SEC
	Final Action	No operation
5th monitor resource	Type	Custom monitor
	Monitor Resource Name	genw_azw
	Interval	60 seconds
	Timeout	120 seconds
	Retry Count	0 times
	Wait Time to Start Monitoring	0 seconds
	Monitor Timing	Always
	Script created with this product	genw.sh
	Nice Value	0
	Recovery Action	Custom setting
	Recovery Target	Local Server
	Recovery Script Execution Count	0 times
	Maximum Reactivation Count	3 times
	Maximum Failover Count	Once
	Final Action	No operation

Caution for Users of Red Hat Enterprise Linux

For a very large scale system in which SAP HANA takeover might take 30 minutes or more, make sure that the system does not time out by setting the timeout time to 60 minutes.

7-2 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.

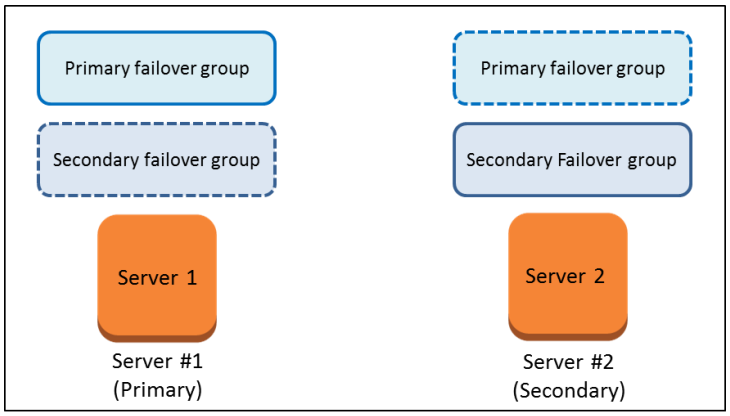

Figure 7-1 Normal Operation

Caution:
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.

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.

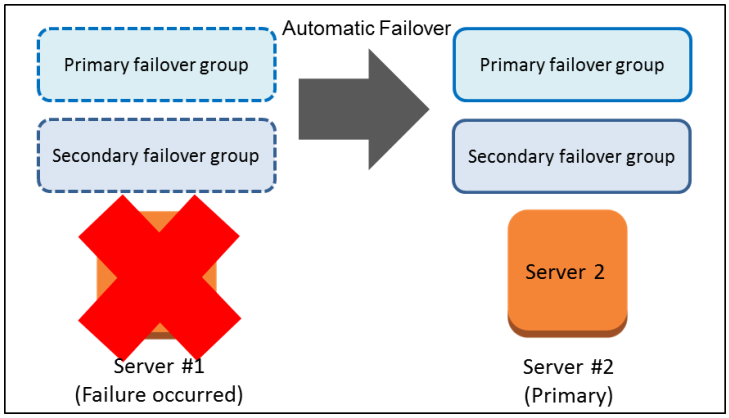

Figure 7-2 Occurrence of Failure on the Primary Server

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.

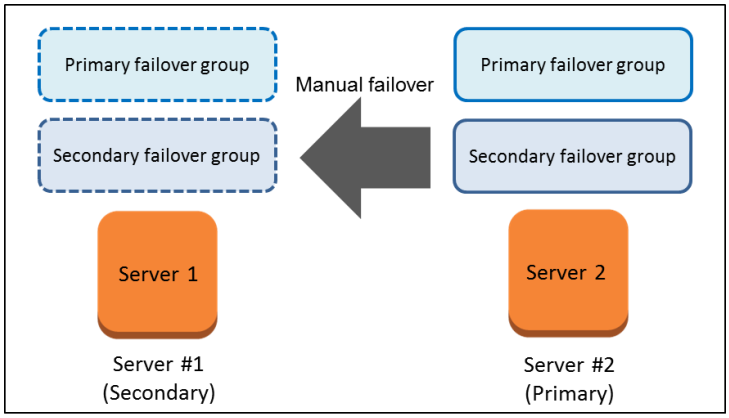

Figure 7-3 Failure Recovery on the Primary Server

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.

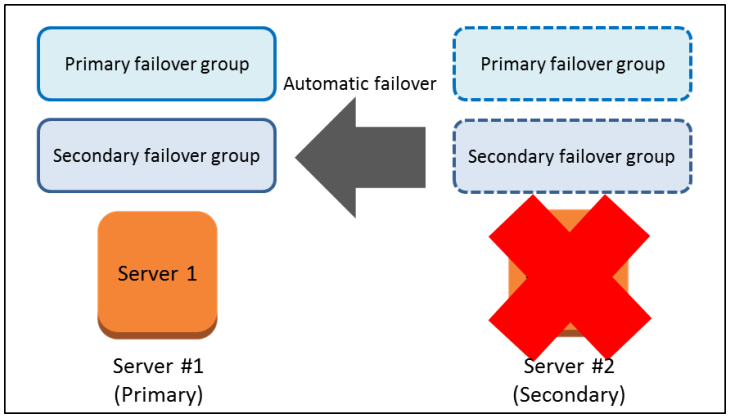

Figure 7-4 Occurrence of Failure on the Secondary Server

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.

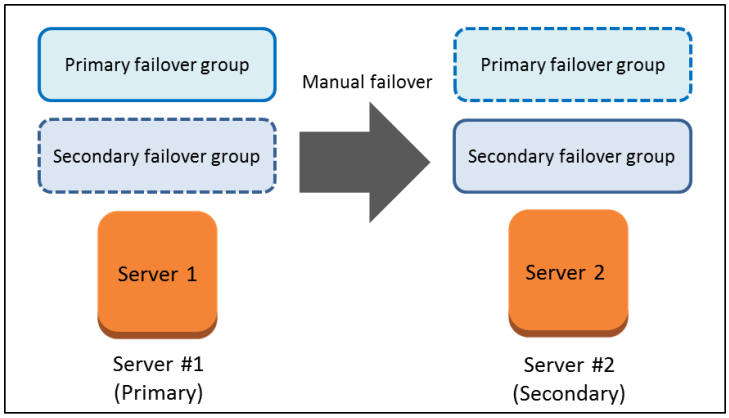

Figure 7-5 Failure Recovery on a Secondary Server

Caution

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.

7-3 Detailed verification results

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

Item	Operation	Verification result
Start cluster	The cluster was started from WebManager. The primary failover group was started on Server #1 and the secondary failover group was started on Server #2 from WebManager.	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 WebManager.	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 WebManager.	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 WebManager.	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.
Move SAP failover group	The secondary failover group was moved from Server #2 to Server #1 from WebManager.	The secondary failover group moved from Server #2 to Server #1. SAP HANA on Server #1 started as the secondary database.

Item	Operation	Verification result
Shut down Server #1	Server #1 was shut down from WebManager.	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 WebManager.	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 WebManager.	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 WebManager.	The secondary failover group moved from Server #1 to Server #2. SAP HANA on Server #2 started as the secondary database.
Shut down Server #2	Server #2 was shut down from WebManager.	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 WebManager.	The secondary failover group moved from Server #1 to Server #2. SAP HANA on Server #2 started as the secondary database.

Item	Operation	Verification result
Reboot cluster	The cluster was rebooted from WebManager. 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 WebManager.	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 WebManager.	The cluster temporarily stopped operations. SAP HANA continued to run.
Resume cluster	The cluster was resumed from WebManager.	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
Custom monitoring for Availability Zone failure (`genw_azw`)	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.)
Custom monitoring for Availability Zone failure (`genw_azw`)	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.)v

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.)

8 Reference URLs

EXPRESSCLUSTER

http://www.nec.com/en/global/prod/expresscluster/

SAP HANA Server Installation and Update Guide

http://help.sap.com/hana/SAP_HANA_Server_Installation_Guide_en.pdf

SAP HANA Administrator Guide

http://help.sap.com/hana/SAP_HANA_Administration_Guide_en.pdf

SAP Note 1656099 - SAP Applications on AWS: Supported DB/OS and AWS EC2 products

http://service.sap.com/sap/support/notes/1656099

SAP Note 1964437 - SAP HANA on AWS: Supported AWS EC2 products

http://service.sap.com/sap/support/notes/1964437

SAP Note 2063657 - HANA System Replication takeover decision guideline

http://service.sap.com/sap/support/notes/2063657

* To reference SAP Note, you need to register as a user to the SAP Support Portal.