

White Paper

A High Availability Solution for SAP S/4 HANA Running on IBM Power Systems Implemented by EXPRESSCLUSTER

August 25, 2016
NEC Corporation

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.

Trademark Information

EXPRESSCLUSTER[®] is a registered trademark of NEC Corporation.

IBM, Power Systems, and PowerVM are trademarks of International Business Machines Corporation in the United States and other countries.

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.

Other product names and slogans written in this manual are trademarks or registered trademarks of their respective companies.

Revision History

Revision History

Edition	Revise Date	Description
1	08/25/2016	Newly created

Contents

1.	Introduction.....	1
1-1	Background.....	1
1-2	Purpose of verification.....	1
1-3	Overview	1
1-3-1	Verification procedure	1
1-3-2	Illustration of operation.....	4
2.	Supported scenario and requirements	9
3.	Verification configuration	10
3-1	Configuration diagram.....	10
3-2	Operating environment.....	11
4.	Verification items	12
4-1	Verification scenario	12
5.	Verification results	13
6.	Conclusion	14
7.	Supplement.....	15
7-1	SAP S/4HANA operating procedure	15
7-2	SAP HANA operating procedure.....	17
7-3	Availability of the NFS server	20
8.	Reference URLs	20

1. Introduction

1-1 Background

SAP HANA was released as a commercial in-memory database in December 2010. In-memory computing technology relies on recording and processing all data on memory. This technology achieves a database performance 100,000 or more times higher than that of a general database that processes data recorded on a hard disk. The power of this technology has had a major impact on the entire IT industry. To maximize SAP HANA capabilities, it is important to select and design an appropriate platform on which to use SAP HANA. Constructing a high-speed in-memory database requires, for example, sufficient memory bandwidth. To deploy SAP HANA in a mission-critical system, a robust and highly reliable platform, such as that provided by IBM Power Systems, is required as the execution environment. It is also vital to implement a high availability configuration that can handle both hardware and software failures. To realize non-stop operations, the system must be able to be recovered promptly if a failure occurs without losing data. SAP HANA, an in-memory database, not only enables high-speed processing by processing data on memory, but also guarantees database consistency in the event of a failure by recording and saving all updates and changes of data to logs. If a failure occurs while executing a transaction, therefore, database consistency can be guaranteed by performing log-based recovery. System failures can also be prepared for by implementing a redundant configuration. With clustering, for example, operations can be automatically failed over to a standby server when a failure is detected. EXPRESSCLUSTER, NEC's high availability clustering software, operates as part of this clustering mechanism. By adopting EXPRESSCLUSTER, a high availability configuration that meets today's reliability and fault-tolerance requirements can be implemented.

1-2 Purpose of verification

The purpose of this white paper is to summarize the results of verifying the use of IBM Power Systems and EXPRESSCLUSTER as a way to enhance the fault tolerance of a platform for SAP S/4HANA and make this information widely available as a reference architecture. This white paper introduces a cutting-edge architecture that can realize true high availability; that is, non-stop operability, by combining SAP S/4HANA, IBM Power Systems, and EXPRESSCLUSTER. We hope that this paper will serve as a reference not only to users who are already using SAP HANA but also to users who are now considering using SAP HANA so that they will be able to implement a SAP S/4HANA solution with confidence.

1-3 Overview

1-3-1 Verification procedure

For verification, NEC created an SAP S/4HANA cluster environment on IBM POWER8 by using EXPRESSCLUSTER.

Various types of hardware and software 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.

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

In this configuration, an S/4HANA system is configured as an HA cluster. EXPRESSCLUSTER monitors failures and switches operations, and SAP HANA System Replication synchronizes data.

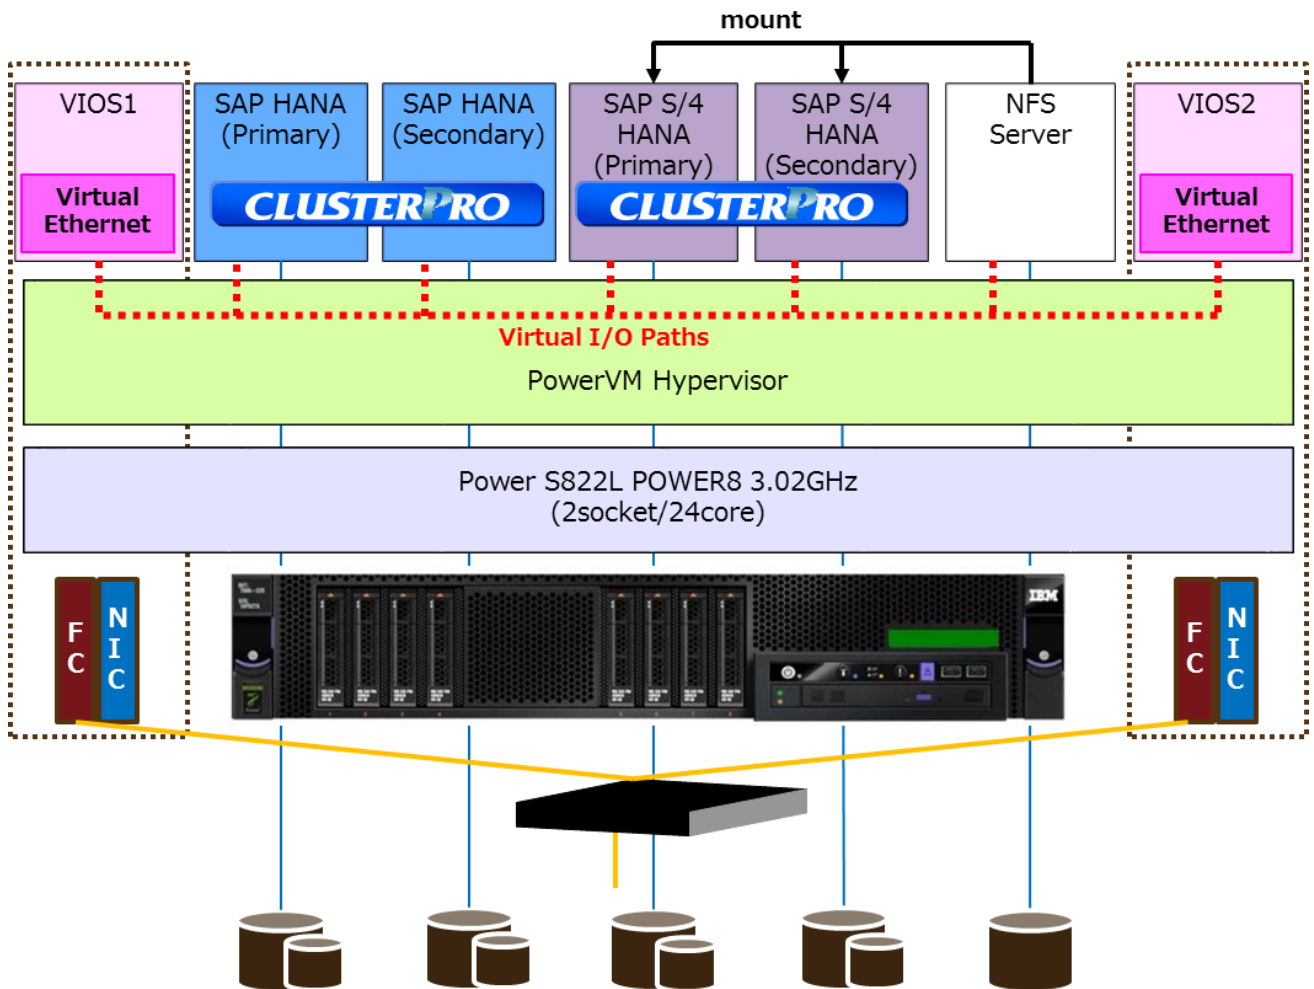


Figure 1-1 System Configuration

- Virtual I/O server (VIOS)

VIOS is dedicated LPAR technology that enables multiple logical partitions (LPAR) to share and use the physical adapter required for disk I/O and communication. Since this technology enable efficient bandwidth usage, servers can be integrated by using only a few physical adapters, leading to lower costs and a smaller server installation footprint.

- S/4HANA

SAP S/4HANA refers to SAP Business Suite 4 SAP HANA and is a next-generation business suite provided by SAP. Unlike conventional ERP models, the recommended SAP S/4HANA environment consists of the following three servers. In this verification, “Backend Server” and “SAP HANA Database Server” are configured as an HA cluster.

- “Frontend Server” providing UI-related functions.
- “Backend Server” providing application functions.
- “SAP HANA Database Server” providing database functions.

SAP HANA is usually included in the term SAP S/4HANA, but in this white paper, “Backend Server” is described as “SAP S/4HANA” and “SAP HANA Database Server” is described as “SAP HANA.”

- 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 switching virtual IP addresses taking into account the presence or absence of data loss. Before switching a virtual IP address, 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 is the system recommended by NEC.

1-3-2 Illustration of operation

Figure 1-2 shows an illustration of the system when SAP S/4HANA in the application layer and SAP HANA in the database layer are configured in separate HA clusters. Also, in this configuration, Server 1 is running as the primary application server (PAS), Server 2 is running as the additional application server (AAS), Server 3 is running as the primary database server, and Server 4 is running as the secondary database server.

A client can access ABAP SAP Central Services (ASCS) without considering whether ASCS is running on Server 1 or Server 2 by accessing the virtual IP address for ASCS. A client can also access SAP S/4HANA without considering whether SAP HANA is running on Server 3 or Server 4 by accessing the virtual IP address for SAP HANA.

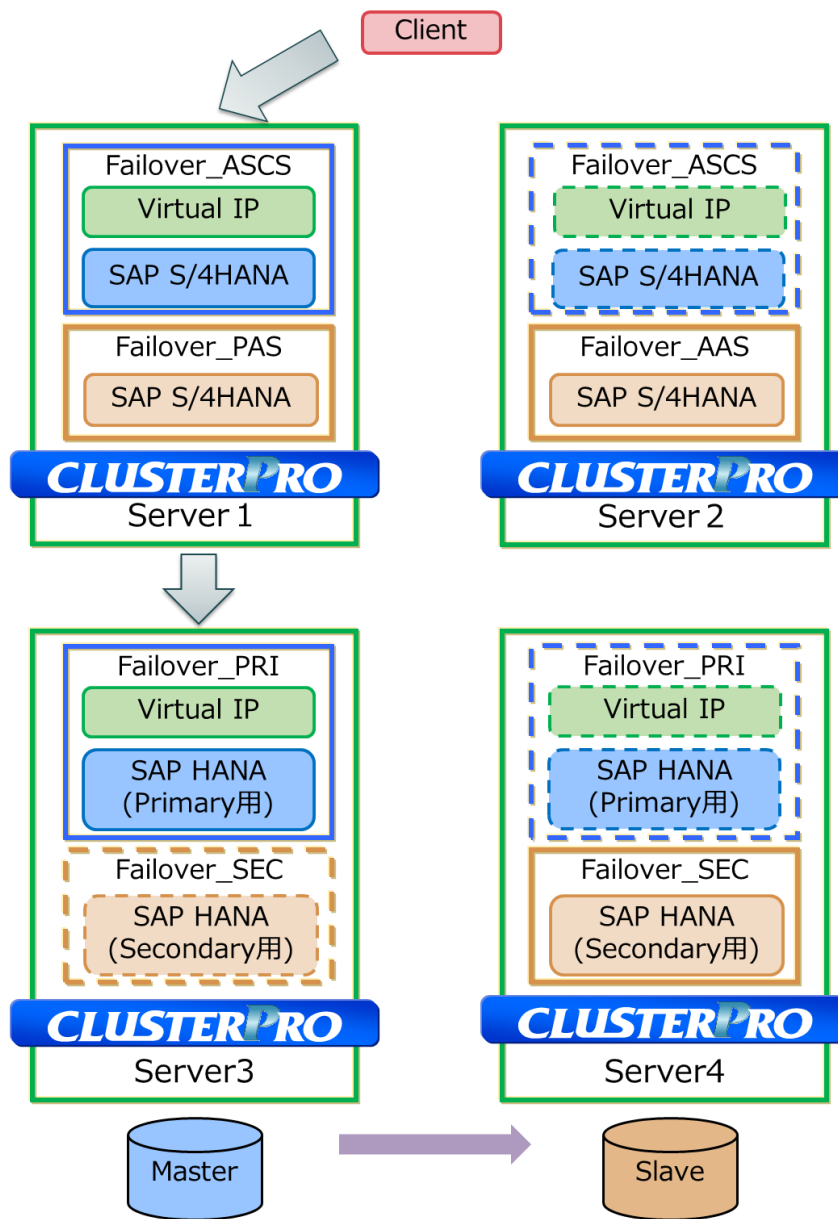


Figure 1-2 Illustration of Normal Operation

- **When a failure occurs on Server 1 on which ASCS is running**

Figure 1-3 shows an illustration of the operation when a failure occurs on the primary server (Server 1). If a failure occurs on Server 1, EXPRESSCLUSTER stops ASCS on Server 1 and then starts ASCS on Server 2, allowing SAP S/4HANA operations to continue. The client can perform operations on SAP S/4HANA, which has become the primary server, by accessing its virtual IP address.

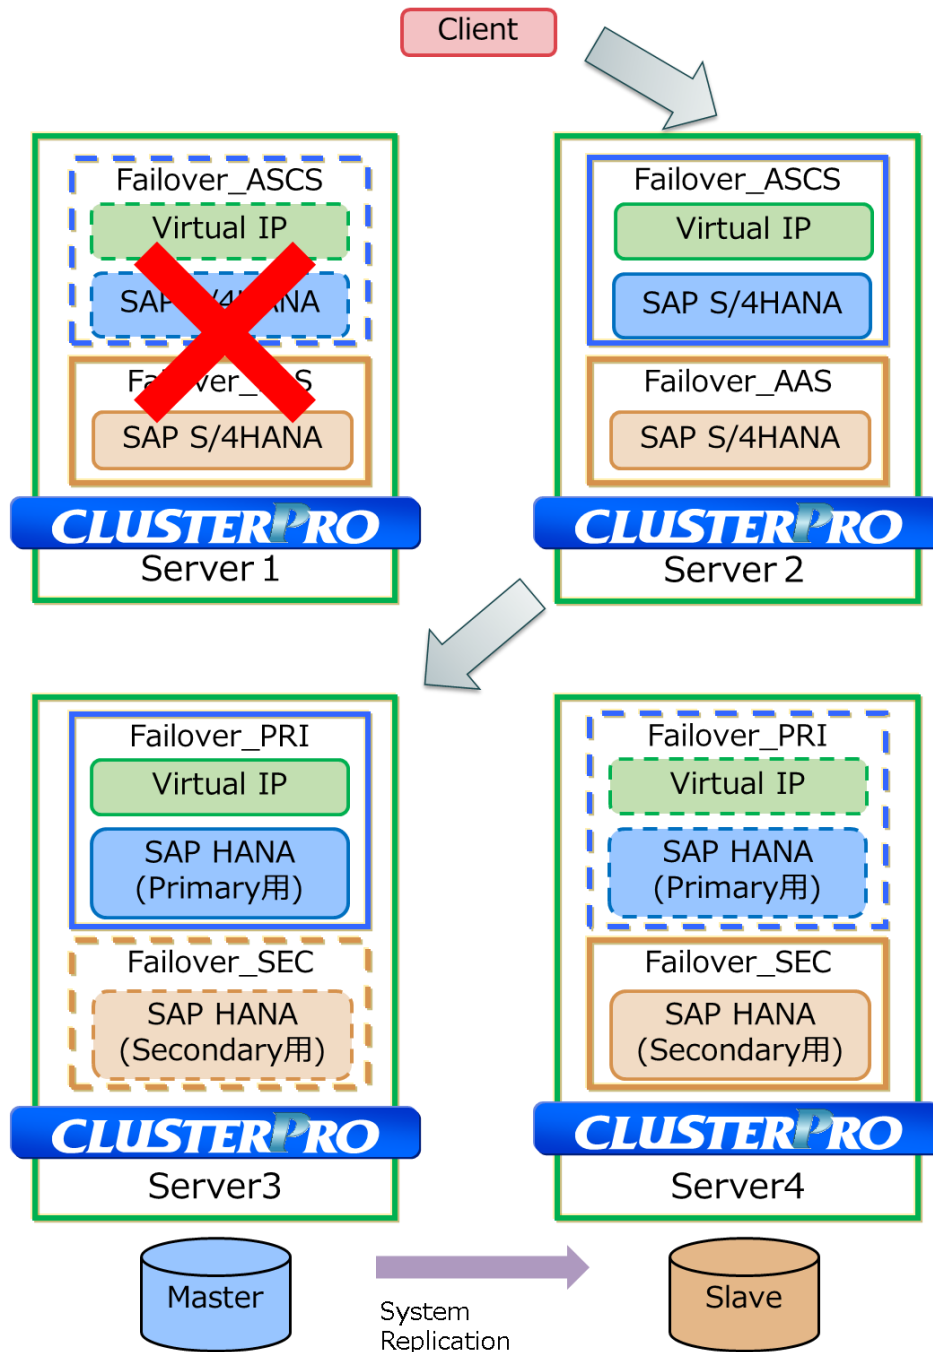


Figure 1-3 Illustration of Operation when a Failure Occurs on Server 1 on Which ASCS Is Running

- **When a failure occurs on Server 2 on which ASCS is standing by**

Figure 1-4 shows an illustration of the operation when a failure occurs on the secondary server (Server 2). If a failure occurs on Server 2, EXPRESSCLUSTER stops SAP S/4HANA on Server 2 and then continues SAP S/4HANA operations on Server 1.

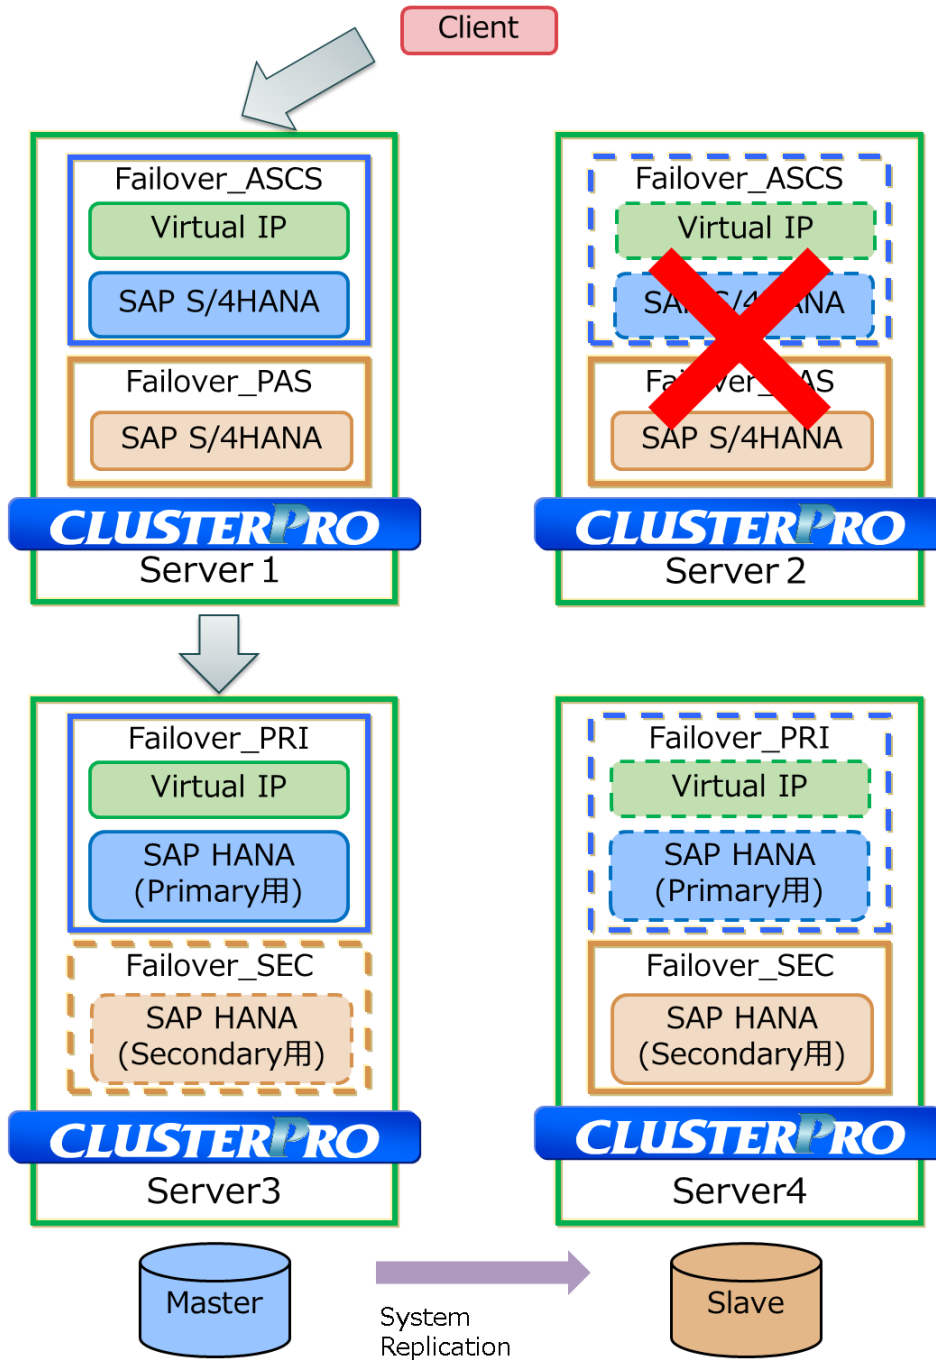


Figure 1-4 Illustration of Operation when a Failure Occurs on Server 2 on Which ASCS Is Standing by

- **When a failure occurs on SAP HANA (primary)**

Figure 1-5 shows an illustration of the operation when a failure occurs on the primary server of SAP HANA. If a failure occurs on the primary server of SAP HANA, EXPRESSCLUSTER stops SAP HANA on Server 3, and changes SAP HANA on Server 4 from the secondary server to the primary server, allowing SAP HANA operations to continue. SAP S/4HANA can perform operations on SAP S/4HANA, which has become the primary server, by accessing its virtual IP address.

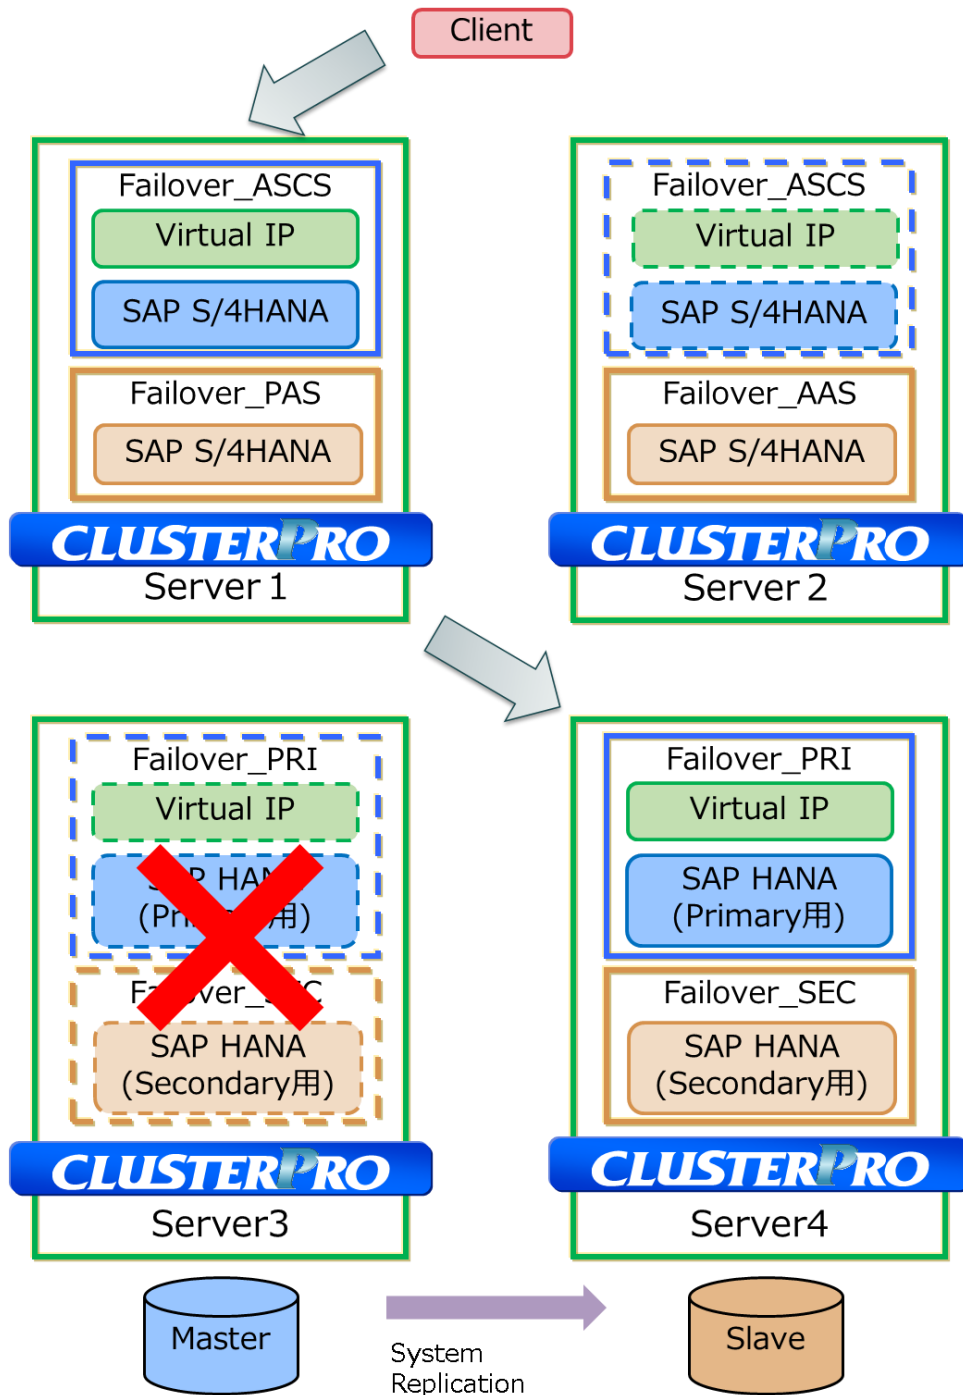


Figure 1-5 Illustration of Operation when a Failure Occurs on the SAP HANA Primary Server

- **When a failure occurs on SAP HANA (secondary)**

Figure 1-6 shows an illustration of the operation when a failure occurs on the secondary server of SAP HANA. If a failure occurs on the SAP HANA secondary server, EXPRESSCLUSTER stops SAP HANA on Server 3 and switches the system replication function to Server 4 (that is, disables the full sync operation), allowing SAP HANA operations to continue.

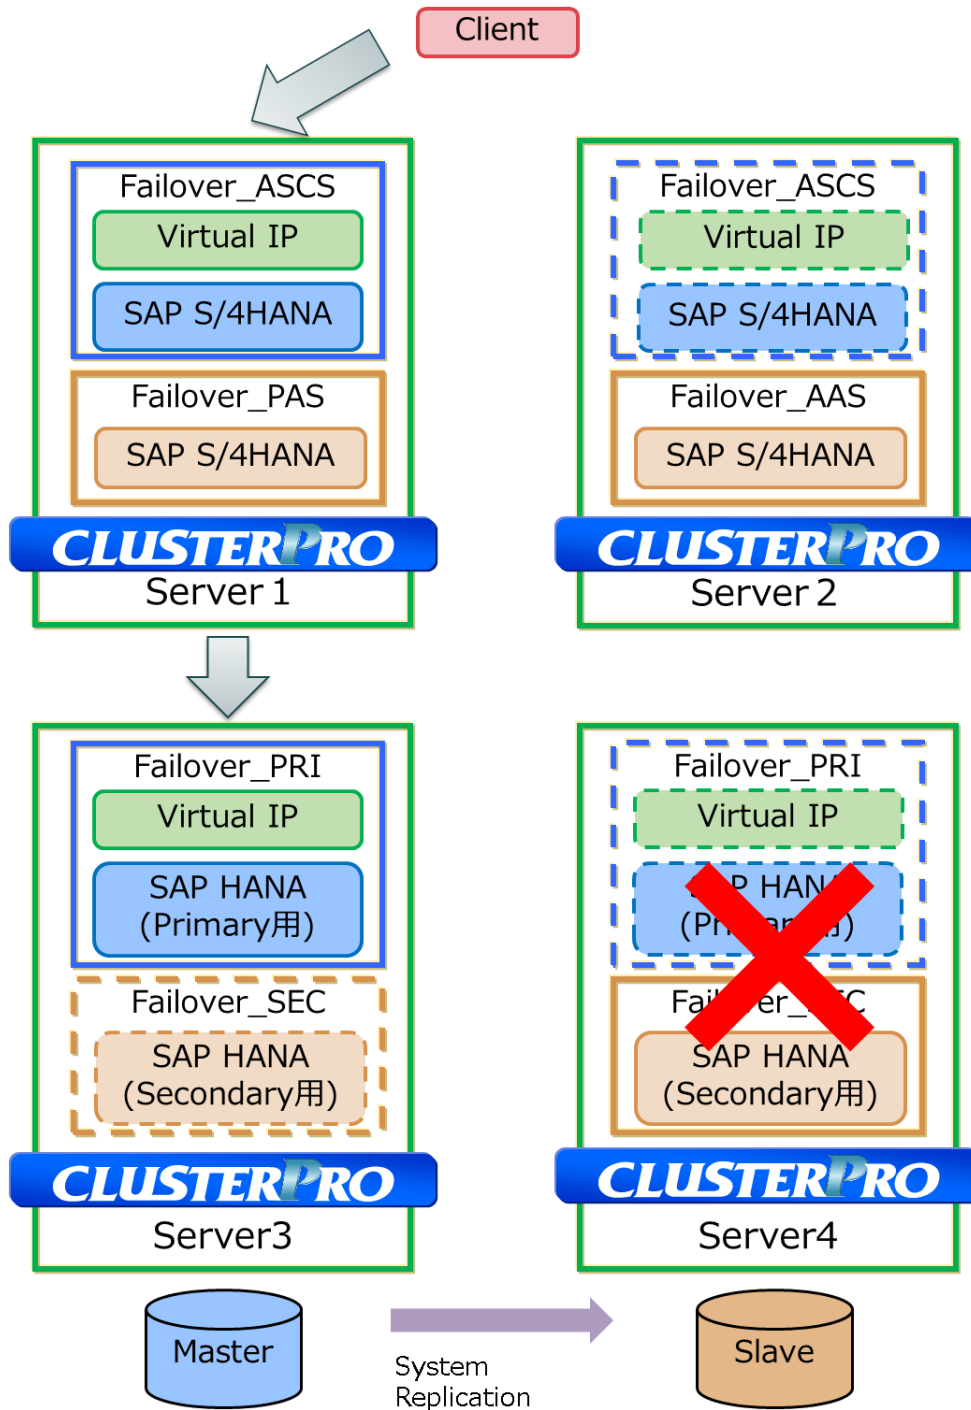


Figure 1-6 Illustration of Operation when a Failure Occurs on the SAP HANA Secondary Server

2. Supported scenario and requirements

The SAP requirements in this verification are as follows. For general system replication requirements, see the guides provided by SAP.

1. A 2-node cluster is configured both for SAP S/4HANA and SAP HANA.
2. The nodes configuring the above mentioned clusters must belong to the same network segment.
3. The version must be SAP HANA SPS11 (revision112) or later.
4. The automatic startup attribute of SAP HANA must be set to “off.”

3. Verification configuration

3-1 Configuration diagram

This verification uses the following configuration.

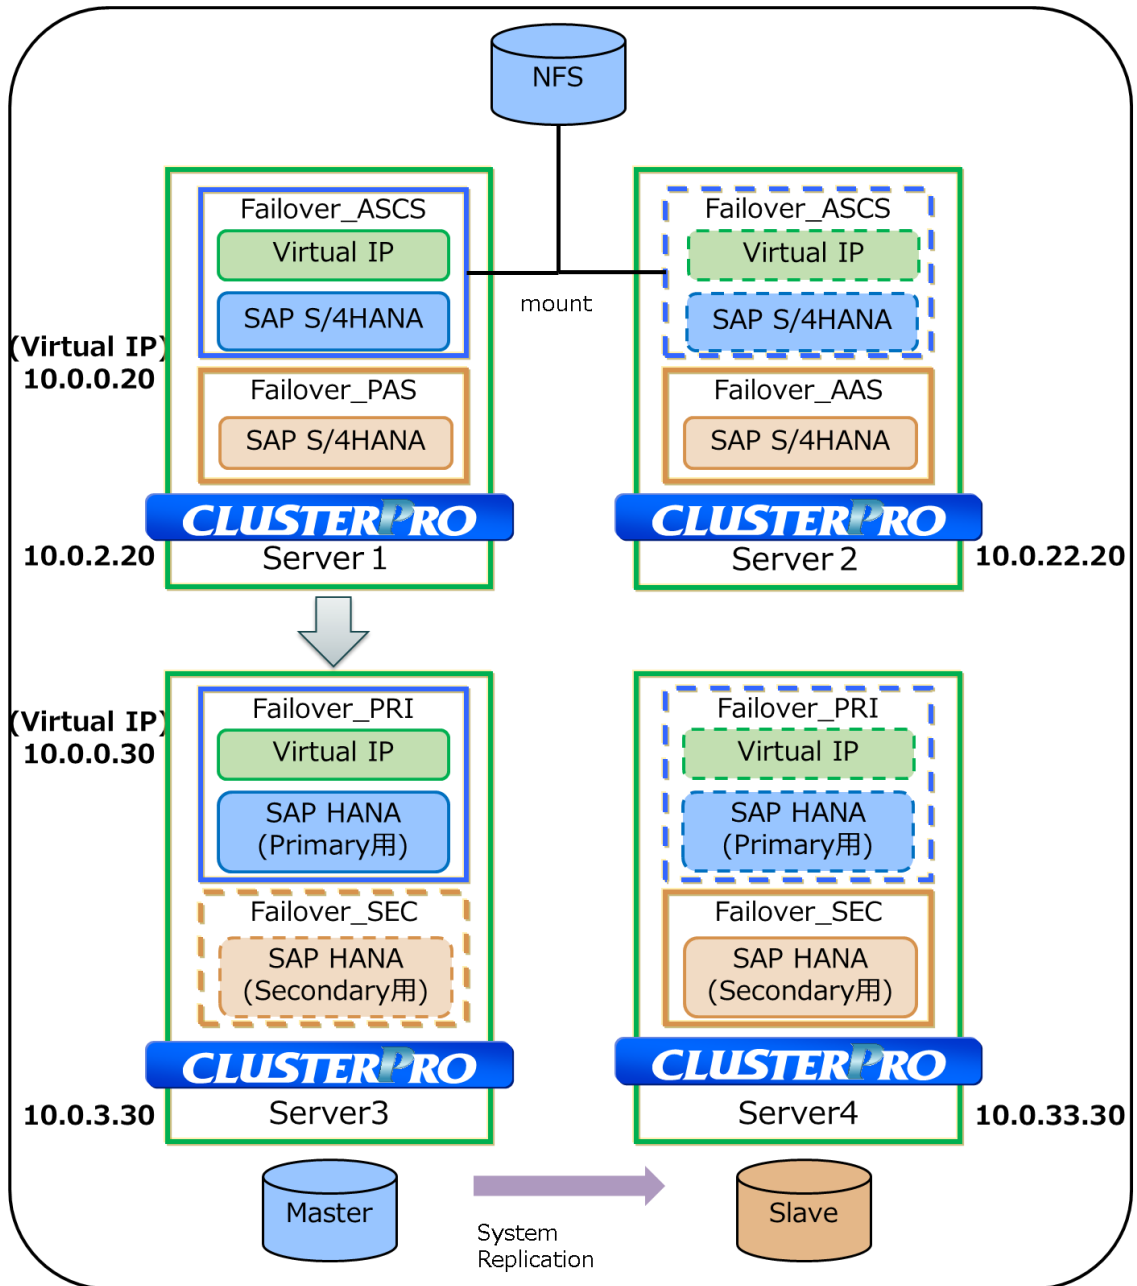


Figure 3-1 System Configuration

3-2 Operating environment

In this verification, a cluster environment is configured by installing SAP S/4HANA and SAP HANA in the virtual machines on IBM PowerVM Hypervisor as shown in *Figure 3-1 System Configuration Diagram*.

SAP S/4HANA (Common)

OS	SUSE Linux Enterprise Server 11 for SAP (SLES11 SP4)
CPU	2vCPU (2 cores)
Memory	24 GB
Disk	200 GB
SAP S/4HANA	SAP S/4HANA 1511
EXPRESSCLUSTER	EXPRESSCLUSTER X 3.3

SAP HANA (Common)

OS	SUSE Linux Enterprise Server 11 for SAP (SLES11 SP4)
CPU	4vCPU (4 cores)
Memory	256 GB
Disk	1024 GB
SAP HANA	SAP HANA SPS11 rev 112
EXPRESSCLUSTER	EXPRESSCLUSTER X 3.3

4. Verification items

4-1 Verification scenario

NEC tested the availability of the SAP S/4HANA and SAP HANA cluster configurations using EXPRESSCLUSTER when the following failures occurred.

The following operations were checked and verified according to each verification scenario below:

<SAP S/4HANA>

- EXPRESSCLUSTER detected the failure and failed over SAP S/4HANA.
- The connection from the client to SAP S/4HANA is available, and operations could continue.

<SAP HANA>

- EXPRESSCLUSTER detected the failure and failed over SAP HANA.
- The connection from SAP S/4HANA to SAP HANA is available, and operations could continue.

- SAP S/4HANA verification scenario

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 hang-up
		SAP S/4HANA	Service down
	Secondary		SAP S/4HANA
		OS	
	SAP S/4HANA		Service down
		SAP S/4HANA	Process down

- SAP HANA verification scenario

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 hang-up
		SAP HANA DB	Service down
	Secondary		SAP HANA DB
		OS	
	SAP HANA DB		Service down
		SAP HANA DB	Process down

5. Verification results

The following table describes the actions that should occur (Desired action) when a SAP S/4HANA failure occurs and whether these actions actually occurred (Result).

Failure type	Server	Component	Failure	Desired action	Result
Hardware failure	Primary	Server	Server down	Failover (to a secondary server)	✓
		Network	Network down	Failover (to a secondary server)	✓
	Secondary	Server	Server down	Failover (to a primary server)	✓
		Network	Network down	Failover (to a primary server)	✓
Software failure	Primary Secondary	OS	OS hang-up	Failover (to a secondary server)	✓
		SAP S/4HANA	Service down	Failover (to a secondary server)	✓
	Process down		Failover (to a secondary server)	✓	
		OS	OS hang-up	Failover (to a primary server)	✓
			Service down	Failover (to a secondary server)	✓
		SAP S/4HANA	Service down	Failover (to a secondary server)	✓
Process down			Failover (to a secondary server)	✓	

The following table describes the actions that should occur (Desired action) when a SAP HANA failure occurs and whether these actions actually occurred (Result).

Failure type	Server	Component	Failure	Desired action	Result
Hardware failure	Primary	Server	Server down	Failover (to a secondary server)	✓
		Network	Network down	Failover (to a secondary server)	✓
	Secondary	Server	Server down	Failover (to a primary server)	✓
		Network	Network down	Failover (to a primary server)	✓
Software failure	Primary Secondary	OS	OS hang-up	Failover (to a secondary server)	✓
		SAP HANA DB	Service down	Failover (to a secondary server)	✓
	Process down		Failover (to a secondary server)	✓	
		OS	OS hang-up	Failover (to a primary server)	✓
			Service down	Failover (to a primary server)	✓
		SAP HANA DB	Service down	Failover (to a primary server)	✓
Process down			Failover (to a primary server)	✓	

In the configuration used in this verification, NEC has verified that EXPRESSCLUSTER executes all operations in coordination with SAP S/4HANA, from quick hardware or software failure detection to automatic failover when a failure occurs.

In the application layer, the connection destination is automatically switched by changing the virtual IP address, allowing operations to continue. In the database layer, with a general HA setting using SAP HANA, servers must be switched manually when a failure occurs. However, by using EXPRESSCLUSTER, switching from the primary server to the secondary server can be executed automatically. 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 cooperation between SAP S/4HANA and EXPRESSCLUSTER can shorten operational downtime and realize the high availability and operational efficiency required for a mission-critical system. NEC has also verified a specific procedure to configure a high availability solution on IBM Power Systems by coordinating with SAP S/4HANA and EXPRESSCLUSTER and SAP HANA and EXPRESSCLUSTER.

7. Supplement

7-1 SAP S/4HANA operating procedure

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

Starting a cluster

Server 1 is used as the PAS server, and Server 2 is used as the AAS server.

The ASCS failover group is started on Server 1. After the ASCS failover group has started, the enqueue replication server instance (ERS) failover group is started manually on Server 2.




Figure 7-1 Normal Operation

Recovering a failure that occurred on the PAS server

When a failure occurs on Server 1, the ASCS failover group fails over to Server 2, allowing operations to continue. When the ASCS failover is complete, the ERS failover group running on Server 2 automatically stops and fails over to Server 1.

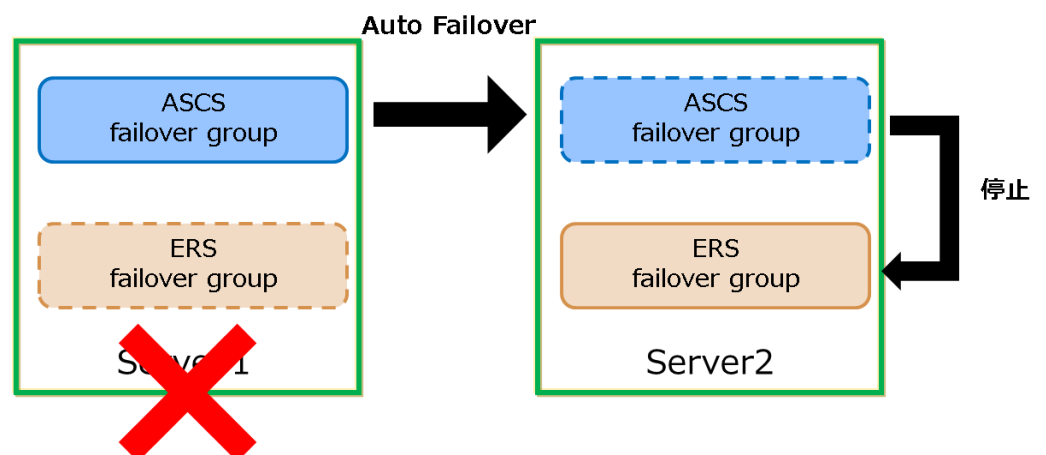


Figure 7-2 Failure Recovery on the PAS Server

Recovery procedure

If the ASCS failover group and ERS failover group are running on a separate node, there is no problem. If the ERS failover group is not running because a failure has occurred on Server 1, the failover group must be started manually.

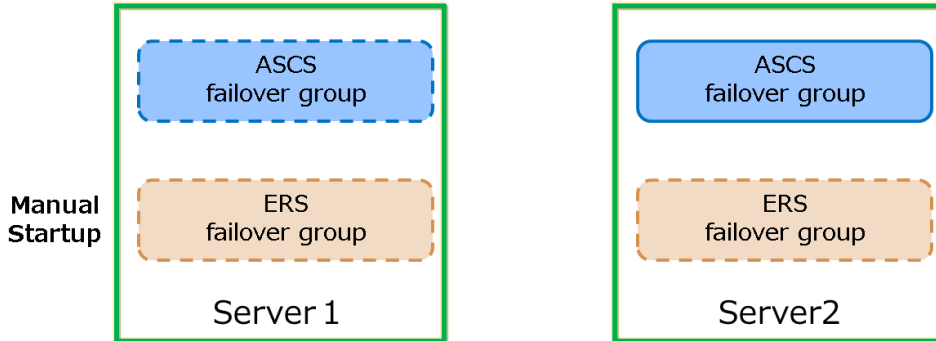


Figure 7-3 Failure Recovery on the PAS Server

When a failure occurs on the AAS server

When a failure occurs on Server 2, the ERS failover group stops. In this case, operations continue because there is no problem in the behavior of the ASCS failover group on Server 1.

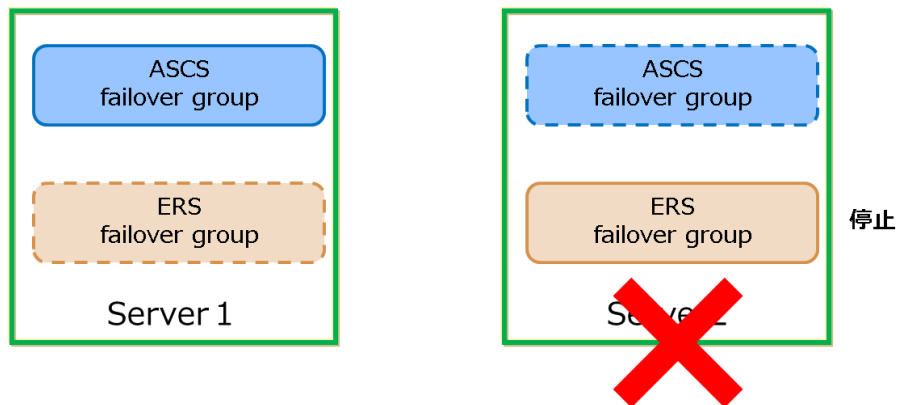


Figure 7-4 Occurrence of Failure on the AAS Server

Recovery procedure

The ERS failover group on Server 2 is started manually.

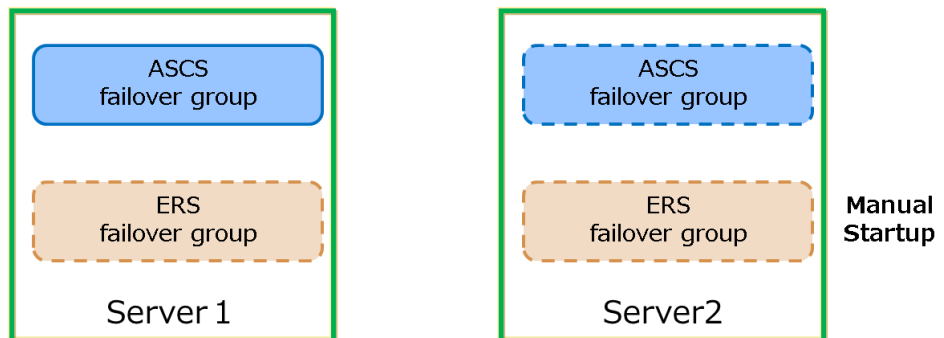


Figure 7-5 Failure Recovery on the AAS Server

7-2 SAP HANA operating procedure

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

Starting a cluster

Server 3 is used as the primary server, and Server 4 is used as the secondary server.

The primary failover group is started on Server 3 and the secondary failover group is started on Server 4. (SAP HANA starts as the primary database on Server 3 and as the secondary database on Server 4.)

After the failover group has started, a command is run manually on Server 3 to enable the full sync option of SAP HANA.

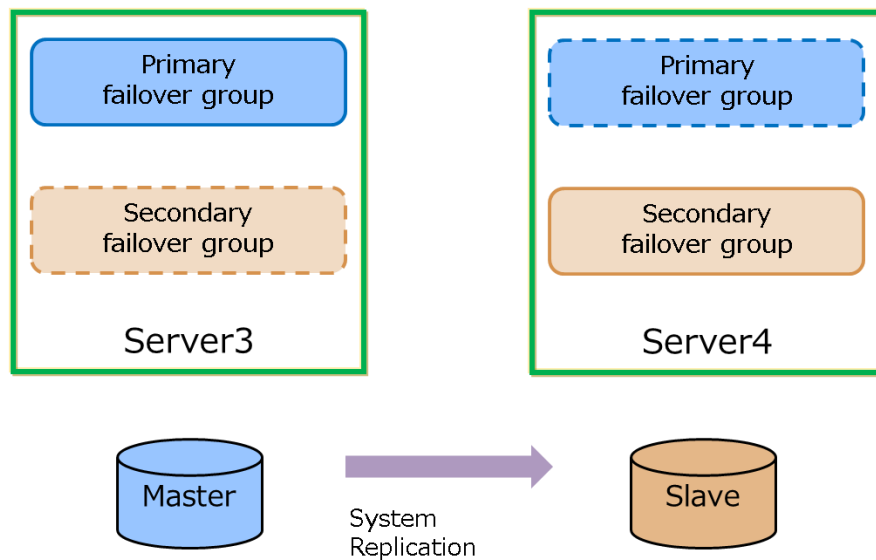


Figure 7-6 Normal Operation

Recovering a failure that occurred on the primary server

When a failure occurs on Server 3, the primary failover group fails over to Server 4. SAP HANA on Server 3 stops, and SAP HANA on Server 4 takes over operations, allowing SAP HANA operations to continue.

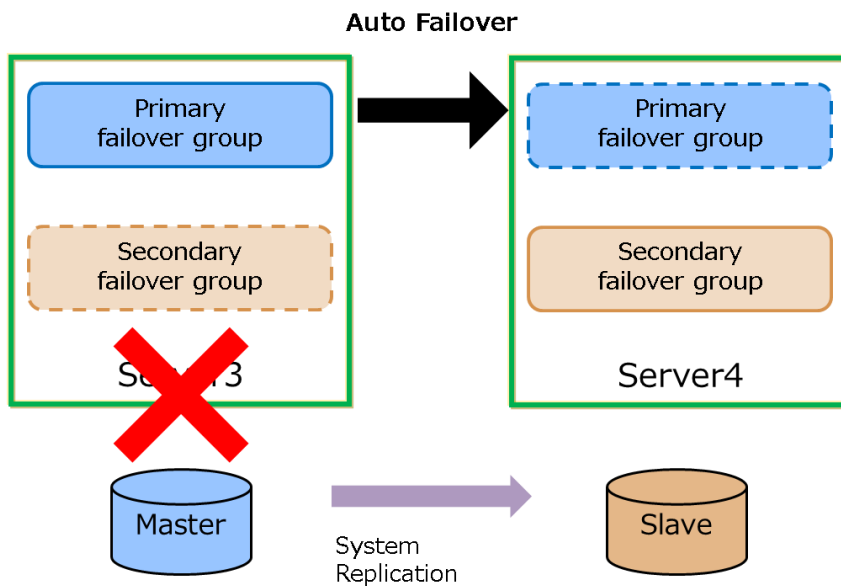


Figure 7-7 Occurrence of Failure on the Primary Server

Recovery procedure

The secondary failover group is failed over from Server 4 to Server 3 manually. When the failover is complete, SAP HANA on Server 3 starts as the secondary database. Also, a command is run manually on Server 3 to enable the full sync option of SAP HANA.

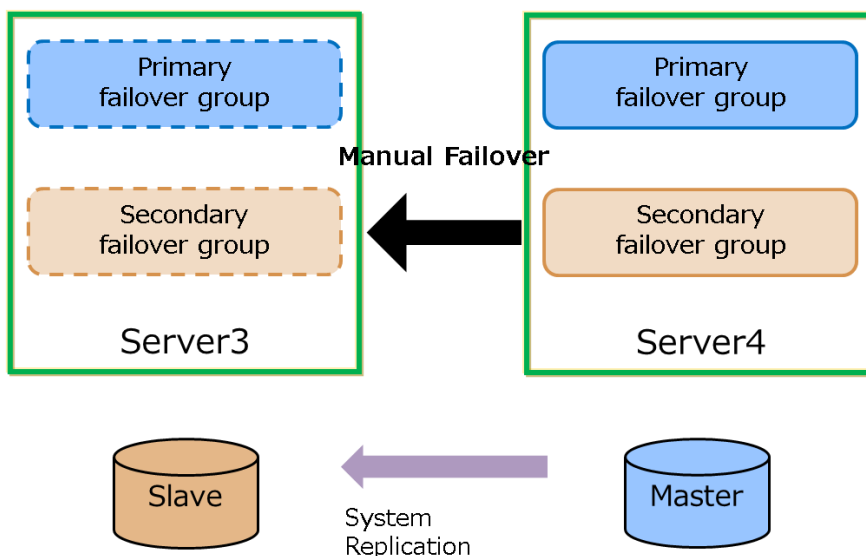


Figure 7-8 Failure Recovery on the Primary Server

When a failure occurs on the secondary server

When a failure occurs on Server 4, the secondary failover group fails over to Server 3. SAP HANA on Server 4 stops, and operations continue on Server 3 with the full sync option of SAP HANA disabled.

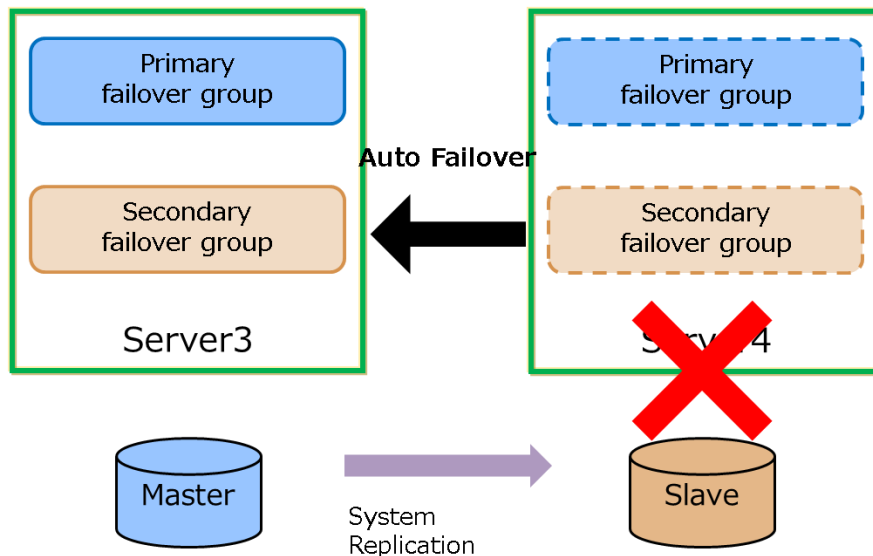


Figure 7-9 Occurrence of Failure on the Secondary Server

Recovery procedure

The secondary failover group is failed over from Server 3 to Server 4 manually.

When the failover is complete, SAP HANA on Server 4 starts as the secondary database.

Also, a command is run manually on Server 1 to enable the full sync option of SAP HANA.

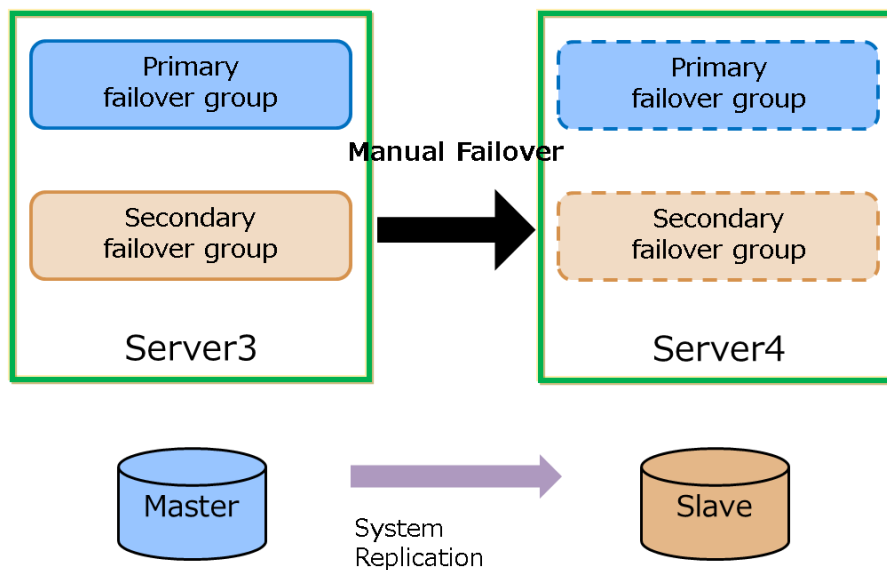


Figure 7-10 Failure Recovery on the Secondary Server

7-3 Availability of the NFS server

In this verification, the server that provides the NFS service was not included in the cluster. Include the NFS server in the cluster if so required by your system.

8. Reference URLs

EXPRESSCLUSTER

<http://jpn.nec.com/clusterpro/index.html>

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 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.