



Red Hat Enterprise Linux for SAP Solutions 8

Red Hat HA Solutions for SAP HANA,
S/4HANA and NetWeaver based SAP
Applications

Legal Notice

Copyright © 2023 Red Hat, Inc.

The text of and illustrations in this document are licensed by Red Hat under a Creative Commons Attribution–Share Alike 3.0 Unported license ("CC-BY-SA"). An explanation of CC-BY-SA is available at

<http://creativecommons.org/licenses/by-sa/3.0/>

. In accordance with CC-BY-SA, if you distribute this document or an adaptation of it, you must provide the URL for the original version.

Red Hat, as the licensor of this document, waives the right to enforce, and agrees not to assert, Section 4d of CC-BY-SA to the fullest extent permitted by applicable law.

Red Hat, Red Hat Enterprise Linux, the Shadowman logo, the Red Hat logo, JBoss, OpenShift, Fedora, the Infinity logo, and RHCE are trademarks of Red Hat, Inc., registered in the United States and other countries.

Linux[®] is the registered trademark of Linus Torvalds in the United States and other countries.

Java[®] is a registered trademark of Oracle and/or its affiliates.

XFS[®] is a trademark of Silicon Graphics International Corp. or its subsidiaries in the United States and/or other countries.

MySQL[®] is a registered trademark of MySQL AB in the United States, the European Union and other countries.

Node.js[®] is an official trademark of Joyent. Red Hat is not formally related to or endorsed by the official Joyent Node.js open source or commercial project.

The OpenStack[®] Word Mark and OpenStack logo are either registered trademarks/service marks or trademarks/service marks of the OpenStack Foundation, in the United States and other countries and are used with the OpenStack Foundation's permission. We are not affiliated with, endorsed or sponsored by the OpenStack Foundation, or the OpenStack community.

All other trademarks are the property of their respective owners.

Abstract

This document provides an overview of the available HA solutions for SAP and provides references to corresponding documentation that further describes each solution in detail.

Table of Contents

MAKING OPEN SOURCE MORE INCLUSIVE	3
PROVIDING FEEDBACK ON RED HAT DOCUMENTATION	4
CHAPTER 1. INTRODUCTION	5
CHAPTER 2. HA SOLUTIONS FOR SAP HANA	6
2.1. AUTOMATED SAP HANA SYSTEM REPLICATION	6
2.2. SUPPORTED SCENARIOS FOR AUTOMATED SAP HANA SCALE-UP SYSTEM REPLICATION	6
2.2.1. Support Policies	7
2.2.2. Performance Optimized	7
2.2.2.1. Configuration Guides	7
2.2.3. Cost Optimized	8
2.2.3.1. Configuration Guides	8
2.2.4. Active/Active(Read Enabled)	8
2.2.5. Multitier System Replication	9
2.2.5.1. Configuration Guide	10
2.3. SUPPORTED SCENARIOS FOR AUTOMATED SAP HANA SCALE-OUT SYSTEM REPLICATION	10
2.3.1. Support Policies	11
2.3.2. Configuration Guides for Performance Optimized HANA Scale-Out System Replication HA setups	11
2.3.3. Active/Active (Read Enabled) HANA Scale-Out System Replication	11
2.3.4. Multitarget System Replication (Scale-Out)	11
CHAPTER 3. HA SOLUTION FOR S/4HANA BASED ON ABAP PLATFORM 1809 OR NEWER	12
3.1. STANDALONE ENQUEUE SERVER 2 (ENSA2)	12
3.2. SUPPORTED SCENARIOS	12
3.3. SUPPORT POLICIES	13
3.4. CONFIGURATION GUIDES	13
3.5. COST-OPTIMIZED SAP S/4HANA HA SETUP (HANA SYSTEM REPLICATION AND ENSA2 COMBINED)	13
CHAPTER 4. HA SOLUTION FOR NETWEAVER OR S/4 BASED ON ABAP PLATFORM 1709 OR OLDER	15
4.1. STANDALONE ENQUEUE SERVER 1 (ENSA1)	15
4.2. SUPPORTED SCENARIOS	15
4.2.1. ABAP/Java Dual-Stack	16
4.3. SUPPORT POLICIES	16
4.4. CONFIGURATION GUIDES	16

MAKING OPEN SOURCE MORE INCLUSIVE

Red Hat is committed to replacing problematic language in our code and documentation. We are beginning with these four terms: master, slave, blacklist, and whitelist. Due to the enormity of this endeavor, these changes will be gradually implemented over upcoming releases. For more details on making our language more inclusive, see our [CTO Chris Wright's message](#).

PROVIDING FEEDBACK ON RED HAT DOCUMENTATION

We appreciate your feedback on our documentation. Let us know how we can improve it.

Submitting comments on specific passages

1. View the documentation in the **Multi-page HTML** format and ensure that you see the **Feedback** button in the upper right corner after the page fully loads.
2. Use your cursor to highlight the part of the text that you want to comment on.
3. Click the **Add Feedback** button that appears near the highlighted text.
4. Add your feedback and click **Submit**.

CHAPTER 1. INTRODUCTION

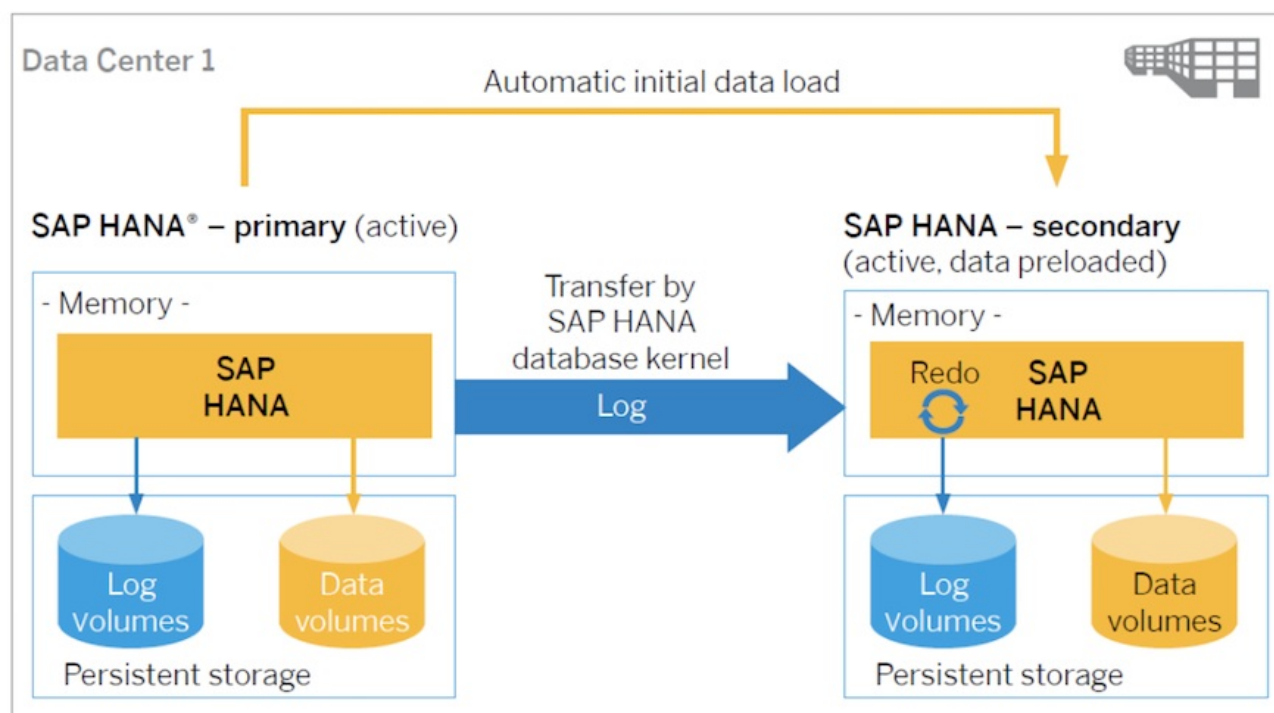
For organisations running SAP production applications, it is essential to ensure the highest possible uptime for their mission critical applications by deploying them in a highly available configuration. With Red Hat HA Solutions for SAP HANA, S/4HANA and NetWeaver based SAP Applications, Red Hat provides such customers with a set of solutions to set up highly available SAP environments on top of the industry leading Red Hat Enterprise Linux High Availability cluster framework.

The Red Hat Enterprise Linux High Availability Add-On provides all the necessary packages for configuring a pacemaker-based cluster that provides reliability, scalability, and availability to critical production services. On top of this, the Red Hat HA Solutions for SAP HANA, S/4HANA and NetWeaver based SAP Applications also allow for setup and configuration of highly available SAP HANA, S/4HANA and NetWeaver based SAP Applications, providing a standard-based approach to reducing planned and unplanned downtime in the corresponding SAP environment.

CHAPTER 2. HA SOLUTIONS FOR SAP HANA

2.1. AUTOMATED SAP HANA SYSTEM REPLICATION

SAP HANA System Replication (HSR) is a built-in high availability and disaster recovery feature to support business continuity. With HANA System Replication, it's possible to copy and continuously synchronize a SAP HANA database to one or more locations. Data is constantly pre-loaded on the secondary system to minimize recovery time objective (RTO).



However, SAP HANA does not contain any mechanism to automatically trigger a failover in case an issue occurs with any components that are part of a HANA System Replication setup. But 3rd party cluster solutions can be used to monitor the healthiness of the HANA System Replication environment and trigger failover when a failure is detected.

On RHEL the Red Hat Enterprise Linux HA Add-On can be used to automate the failover. Red Hat provides HA solutions for both single system SAP HANA setups (Scale-up) or scaleable multi-system SAP HANA setups (Scale-out).

2.2. SUPPORTED SCENARIOS FOR AUTOMATED SAP HANA SCALE-UP SYSTEM REPLICATION

Supported Scenario	Notes
Performance Optimized	Secondary site is not active for client/application servers
Cost Optimized	Support a QA/Test instance running on the secondary site (Cost-Optimized); QA/Test instance will be shutdown first during the failover of Prod

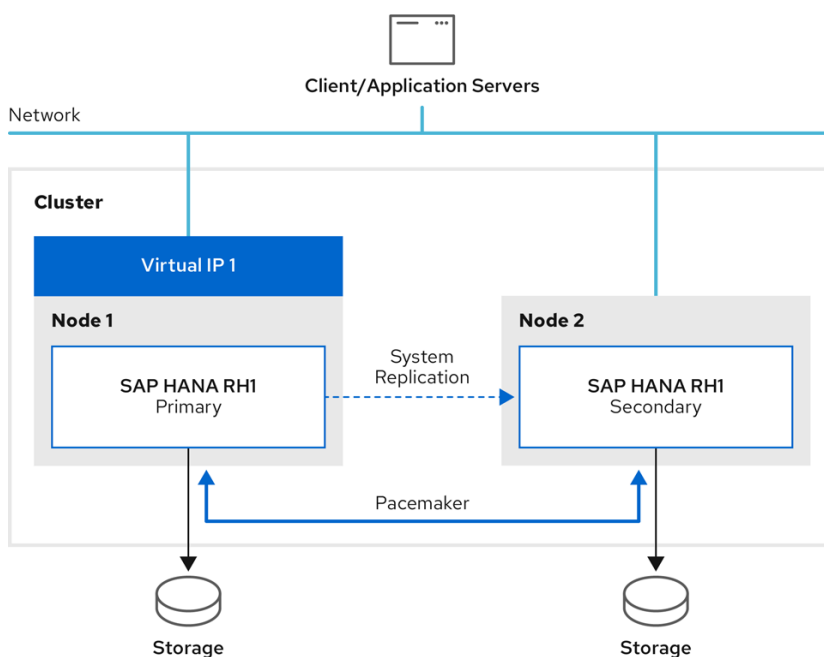
Supported Scenario	Notes
Active/Active (Read Enabled)	The secondary HANA instance can take read-only inquiries
Multitier System Replication	Multitier System Replication is possible, but the tertiary site cannot be managed by the cluster

2.2.1. Support Policies

Please refer to [Support Policies for RHEL High Availability Clusters - Management of SAP HANA in a Cluster](#).

2.2.2. Performance Optimized

In the **Performance Optimized** scenario, the secondary HANA database is configured to preload the tables into memory, thus the takeover time is normally very fast. However, as the secondary HANA database is dedicated to System Replication and doesn't accept client inquiries, this setup is expensive in terms of hardware cost.



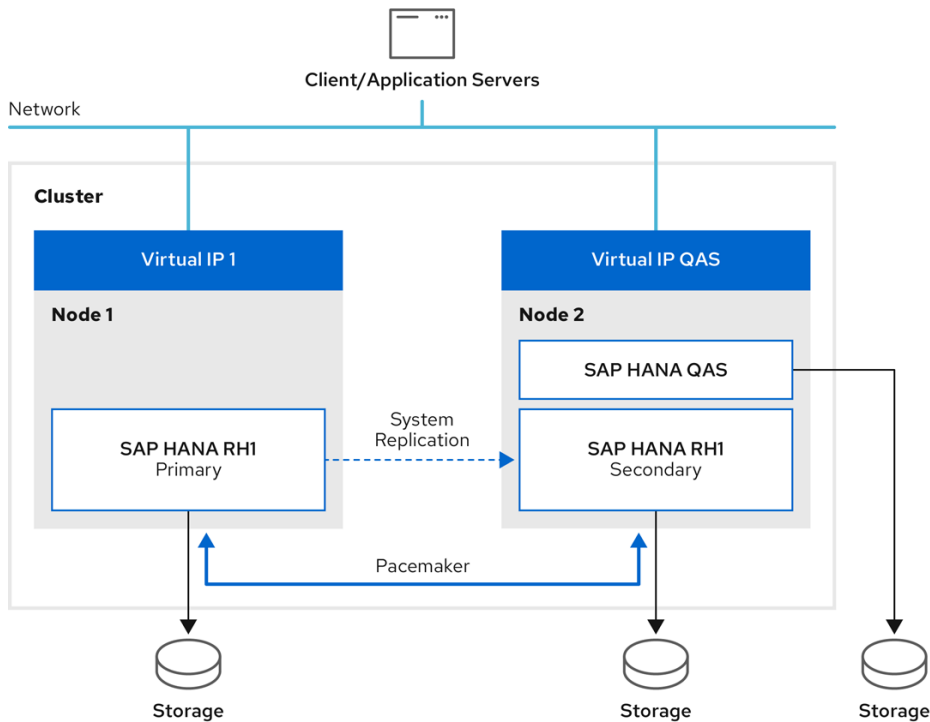
215_RHEL_0222

2.2.2.1. Configuration Guides

- On-Premise: [Automating SAP HANA Scale-Up System Replication using the RHEL HA Add-On](#)
- AWS: [Configuring SAP HANA Scale-Up System Replication with the RHEL HA Add-On on Amazon Web Services \(AWS\)](#)
- Azure: [High availability of SAP HANA on Azure VMs on Red Hat Enterprise Linux](#)
- Google Cloud Platform (GCP): [HA cluster configuration guide for SAP HANA on RHEL](#)
- IBM Power System Virtual Server: [Configure SAP HANA Scale-Up System Replication in a RHEL HA Add-On cluster](#)

2.2.3. Cost Optimized

The **Cost Optimized** scenario supports an additional TEST/QA HANA database on the secondary site, serving client inquiries. Because hardware resources have to be allocated to the TEST/QA instance, the Production HANA database can not be preloaded. Before takeover, the TEST/QA instance has to be shutdown first to free up the hardware resources assigned to it and reassign them to the secondary HANA instance that will be promoted to become the primary instance. The takeover time is therefore longer than for Performance Optimized setups.



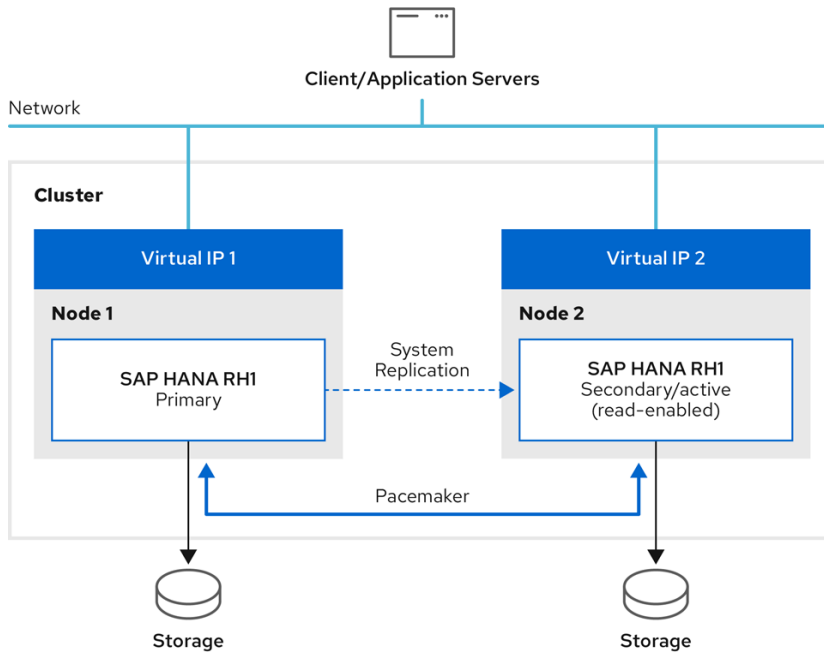
215_RHEL_0222

2.2.3.1. Configuration Guides

- On-Premise: [Automating Cost-Optimized SAP HANA Scale-Up System Replication using the RHEL HA Add-On](#)
- IBM Power System Virtual Server: [Configuring SAP HANA Cost-Optimized Scale-Up System Replication in a RHEL HA Add-On cluster](#)

2.2.4. Active/Active(Read Enabled)

The secondary HANA instance can take read-only inquiries. This setup supports a second virtual IP on the secondary site.

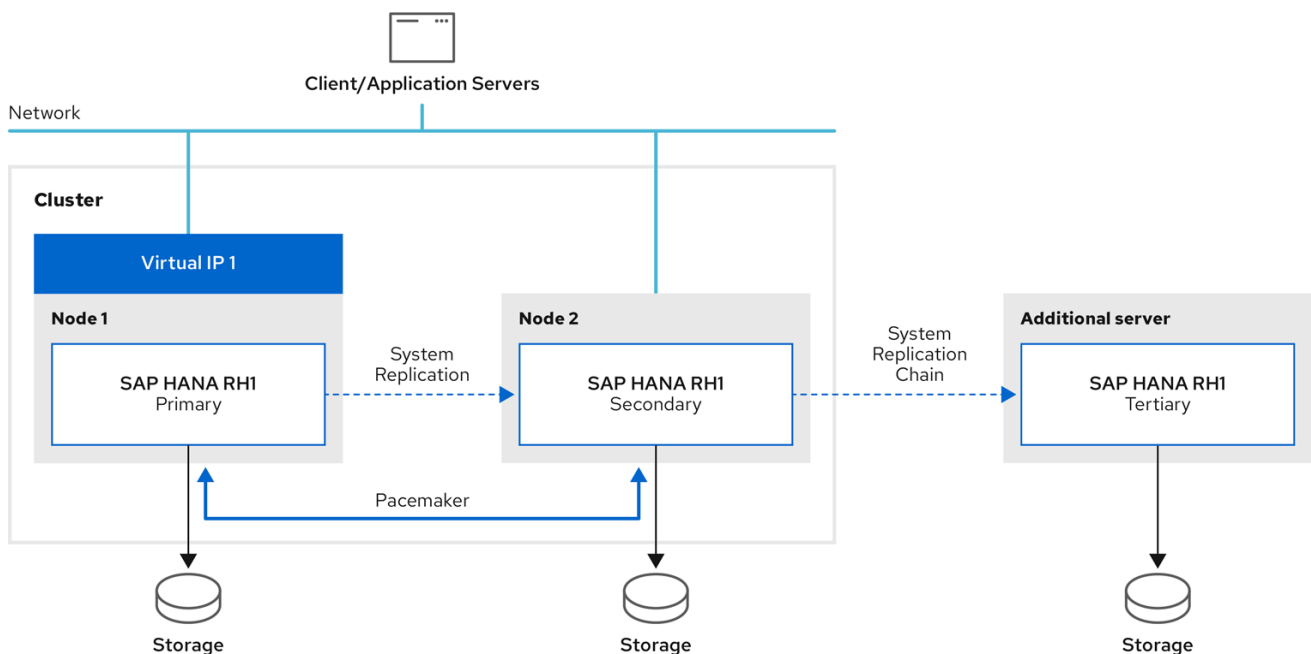


215_RHEL_0222

For more information, please refer to [Adding a secondary virtual IP address for an Active/Active \(Read-Enabled\) HANA System Replication setup](#).

2.2.5. Multitier System Replication

Multitier System Replication is possible, but the tertiary site cannot be managed by the cluster.



215_RHEL_0222

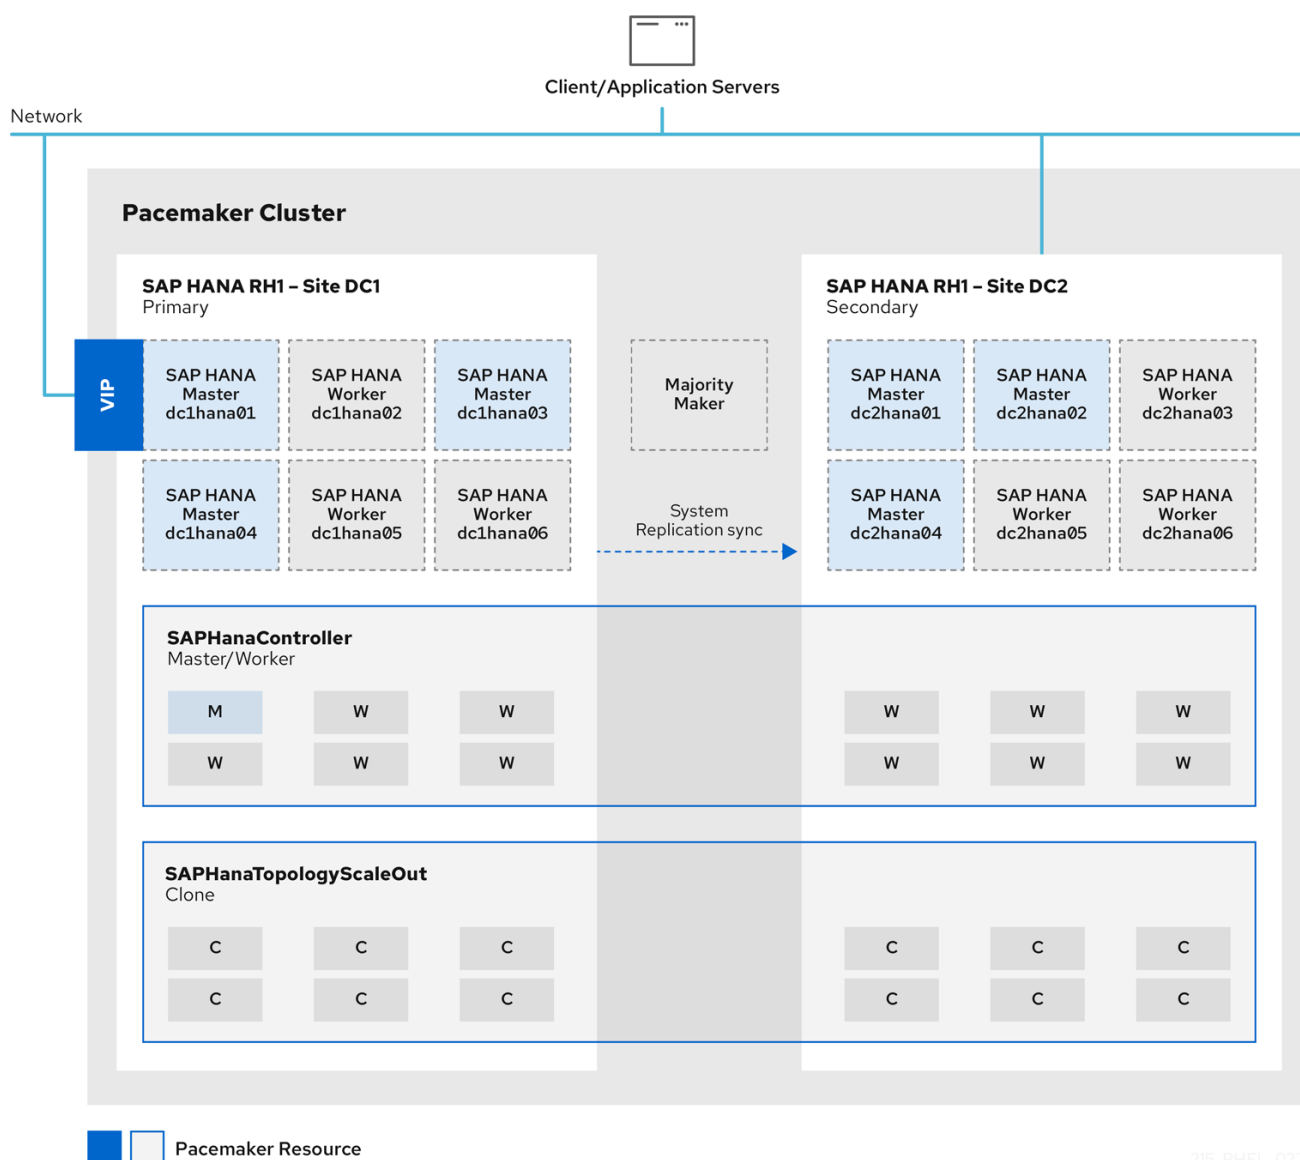
A takeover to the tertiary site will have to be triggered manually, and if the environment should be brought back to the previous state after a manual takeover to the tertiary site, all steps to reconfigure the HANA System Replication setup will have to be carried out manually as well while the cluster is disabled. After it has been verified that the HANA System Replication setup is working correctly again on the HANA instances that should be managed by the cluster, the cluster can be reactivated.

2.2.5.1. Configuration Guide

- IBM Power System Virtual Server: [Configuring SAP HANA multitier system replication in a RHEL HA Add-On cluster](#)

2.3. SUPPORTED SCENARIOS FOR AUTOMATED SAP HANA SCALE-OUT SYSTEM REPLICATION

Supported Scenario	Description
Performance Optimized	Secondary site is not active to client/application servers
Active/Active (Read Enabled)	The secondary HANA instance can take read-only inquiries
Multitarget System Replication	Primary HANA instance is replicated to multiple secondary HANA instances



215_RHEL_0222

2.3.1. Support Policies

Please refer to [Support Policies for RHEL High Availability Clusters - Management of SAP HANA in a Cluster](#).

2.3.2. Configuration Guides for Performance Optimized HANA Scale-Out System Replication HA setups

- On-Premise: [Red Hat Enterprise Linux HA Solution for SAP HANA Scale Out and System Replication](#)
- AWS: [Configuring SAP HANA Scale-Out System Replication with the RHEL HA Add-On on Amazon Web Services \(AWS\)](#)
- Azure: [High availability of SAP HANA scale-out system on Red Hat Enterprise Linux](#)

2.3.3. Active/Active (Read Enabled) HANA Scale-Out System Replication

In HANA 2.0, the secondary instance can take Read-Only inquiries. This setup supports a second virtual IP on the secondary site. For more details please check chapter "Adding a secondary virtual IP address resource for Active/Active (Read-Enabled) setup" in [Red Hat Enterprise Linux HA Solution for SAP HANA Scale Out and System Replication](#).

For more information see [Active/Active\(Read-Enabled\)](#).

2.3.4. Multitarget System Replication (Scale-Out)

Starting with HANA 2.0 SPS 04 Multi Target Replication is supported in a cluster environment. The primary site is replicated to a secondary site and also replicated to an additional secondary site to meet additional availability requirements. In terms of failure this additional third site will be automatically registered to the new primary site, which was the former secondary.

For more details please refer to:

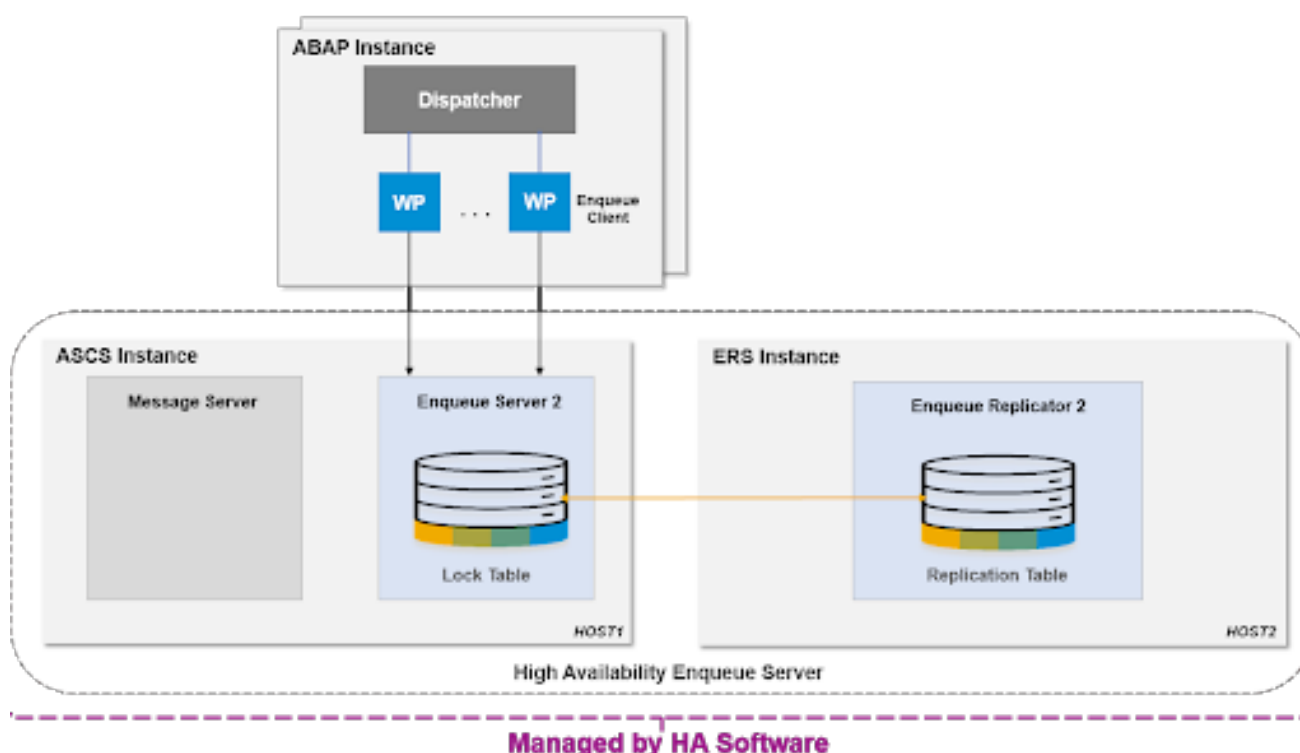
- [Automating SAP HANA Multi Target System Replication in a Pacemaker-based cluster on Red Hat Enterprise Linux \(RHEL\)](#)
- [Multitarget System Replication](#)

CHAPTER 3. HA SOLUTION FOR S/4HANA BASED ON ABAP PLATFORM 1809 OR NEWER

3.1. STANDALONE ENQUEUE SERVER 2 (ENSA2)

Standalone Enqueue Server, a component of Application Server ABAP, is a mechanism to ensure the high availability of the lock table and its entries.

Standalone Enqueue Server has evolved into generation 2 since NetWeaver 7.51, known as Standalone Enqueue Server 2, or ENSA2. In ENSA2, if the ASCS fails, it can start on a separate node in the cluster, and copy the lock entries from the enqueue replicator 2.

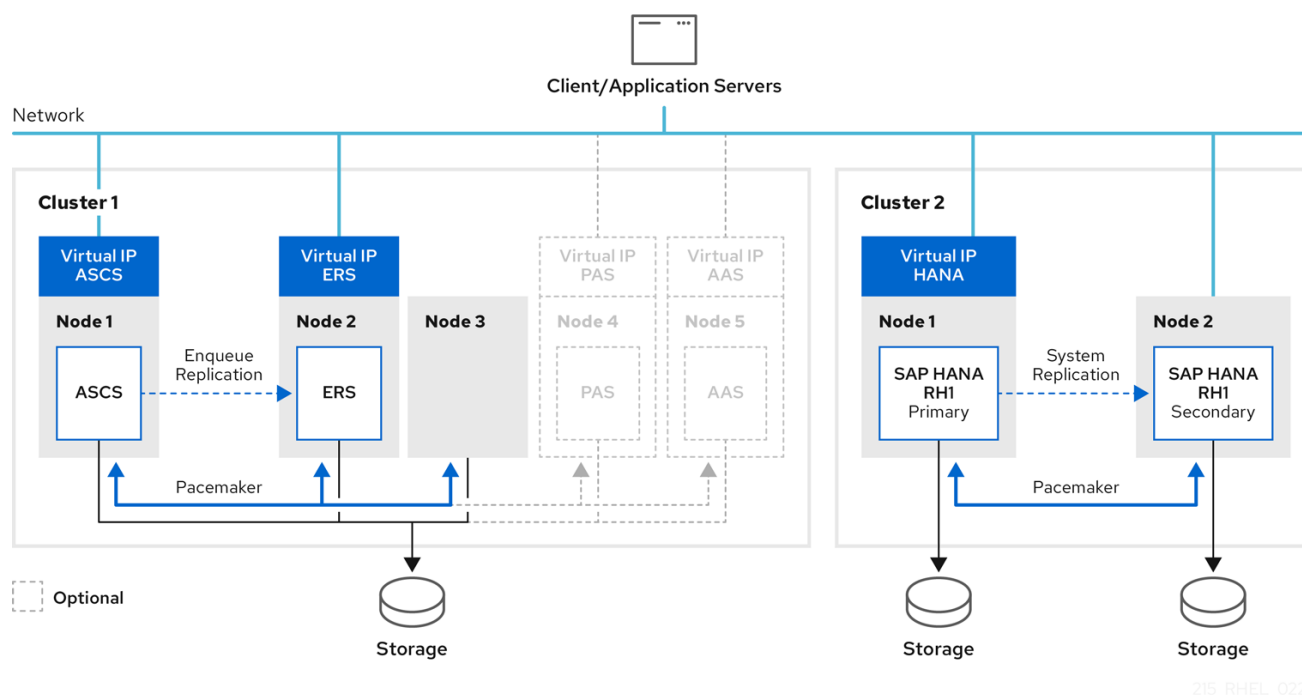


In ABAP platform 1809 or newer, the Standalone Enqueue Server 2 (ENSA2) has become the default installation.

3.2. SUPPORTED SCENARIOS

Supported Scenario	Description
Multi-node cluster	Because in ENSA2 ASCS doesn't have to "follow" ERS, that makes a multi-node cluster possible
Two-node cluster	Upgraded two-node cluster can easily adjust the setup to switch from ENSA1 to ENSA2

Brand new installations can take this chance to design the architecture, choosing between multi-node cluster or two-node cluster. Below is the architecture diagram of a typical 3-node cluster. Of course more nodes can be added upon customer's datacenter requirements or needs.



215_RHEL_0222

3.3. SUPPORT POLICIES

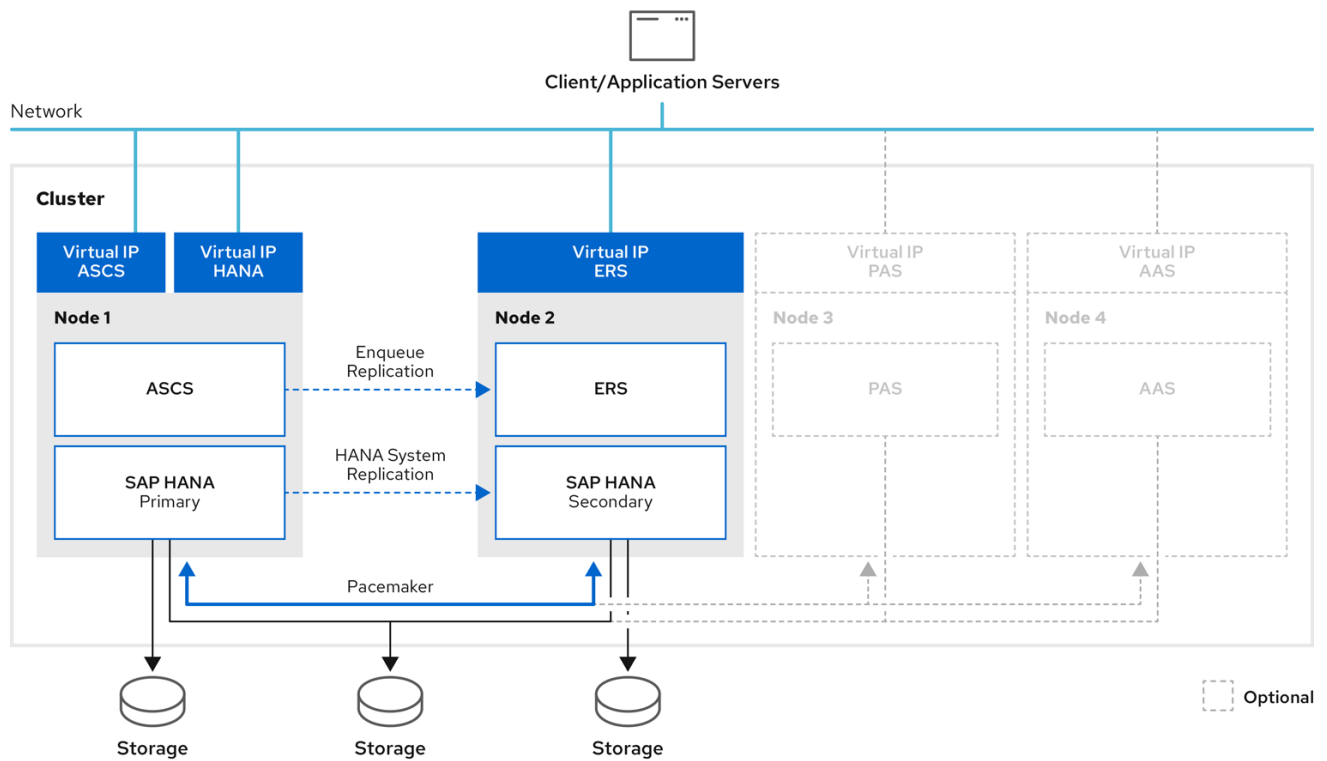
Please refer to [Support Policies for RHEL High Availability Clusters - Management of SAP S/4HANA](#).

3.4. CONFIGURATION GUIDES

- On-Premise: [Configuring SAP S/4HANA ASCS/ERS with Standalone Enqueue Server 2 \(ENSA2\) in Pacemaker](#)
- AWS: [Configure SAP S/4HANA ASCS/ERS ENSA2 on Amazon Web Services \(AWS\)](#)
- Azure with GlusterFS: [Azure Virtual Machines high availability for SAP NetWeaver on Red Hat Enterprise Linux with GlusterFS](#)
- Azure with NetApp: [Azure Virtual Machines high availability for SAP NetWeaver on Red Hat Enterprise Linux with Azure NetApp Files for SAP applications](#)
- Azure with NFS: [High availability for SAP NetWeaver on Azure VMs on Red Hat Enterprise Linux with NFS on Azure Files](#)
- GCP: [HA cluster configuration guide for SAP NetWeaver on RHEL](#)
- IBM Power System Virtual Server: [Configuring high availability for SAP S/4HANA \(ASCS and ERS\) in a RHEL HA Add-On cluster](#)

3.5. COST-OPTIMIZED SAP S/4HANA HA SETUP (HANA SYSTEM REPLICATION AND ENSA2 COMBINED)

With current versions of S/4HANA, it is also possible to run HANA and ABAP Application Server instances on the same system. This makes it possible to have a "cost-optimized" S/4HANA setup where both HANA System Replication and ENSA2 are managed by a single cluster running on the same cluster nodes.

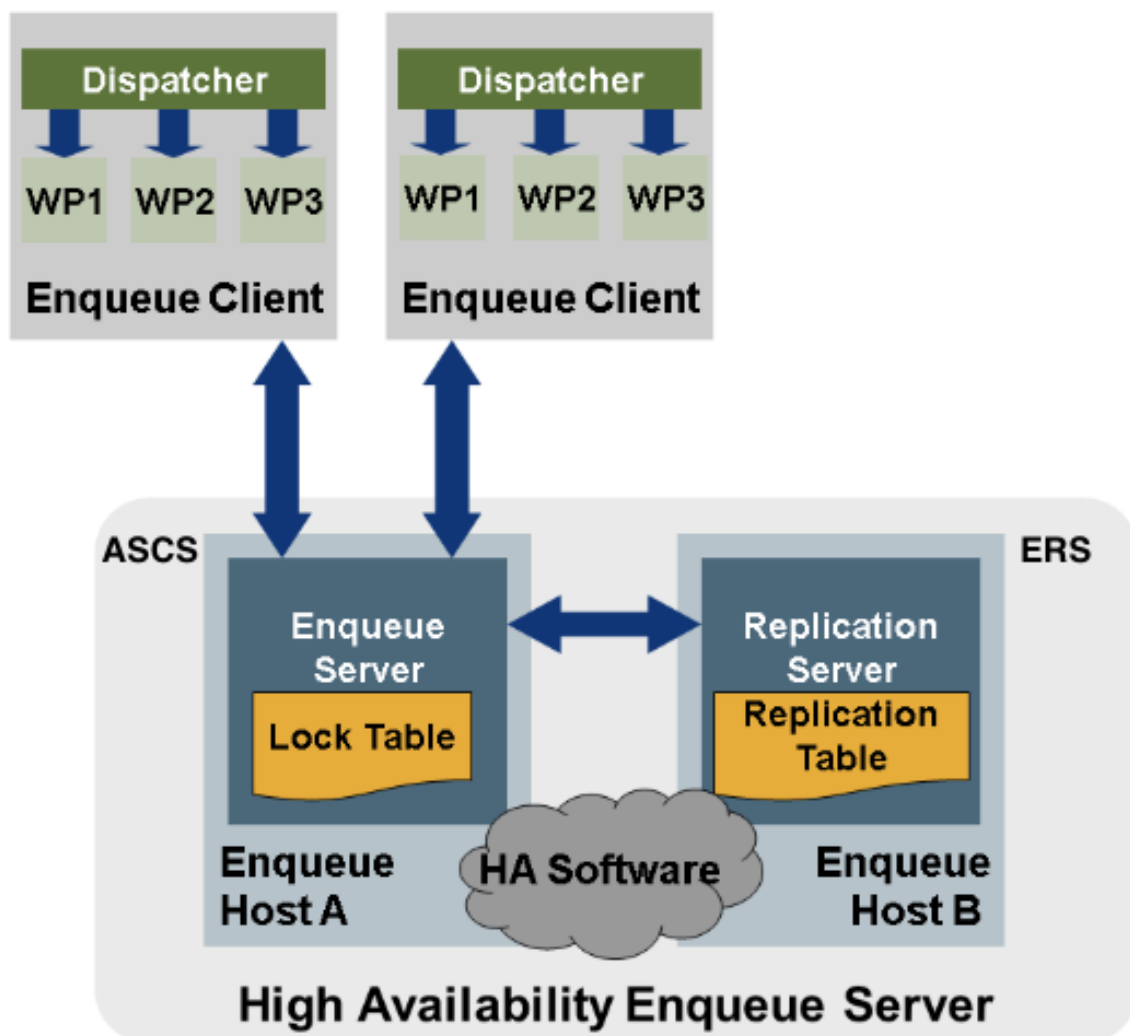


For more information, please refer to [Configuring a Cost-Optimized SAP S/4HANA HA cluster \(HANA System Replication + ENSA2\) using the RHEL HA Add-On](#).

CHAPTER 4. HA SOLUTION FOR NETWEAVER OR S/4 BASED ON ABAP PLATFORM 1709 OR OLDER

4.1. STANDALONE ENQUEUE SERVER 1 (ENSA1)

When using the old Standalone Enqueue Server (ENSA1), the ASCS instance has to fail over the cluster node where the active ERS instance is running and it must be ensured that the ERS instance is shut down and moved to the node where the ASCS instance was running when it was online. This is because the ASCS instance has to access the shared memory where the ERS instance maintains a copy of the enqueue lock table to continue to keep track of the enqueue locks of active transactions. ENSA1 is supported in pacemaker as a two-node cluster configuration, mainly because of the restriction that the ASCS instance must “follow” the ERS instance.



4.2. SUPPORTED SCENARIOS

Supported Scenario	Description
Two node cluster	For ENSA1 the ASCS instance must always move to the node where the ERS instance is running

Supported Scenario	Description
ABAP/Java Dual-Stack	Supported by Master/Slave Resources in all RHEL 7.x releases

4.2.1. ABAP/Java Dual-Stack

ABAP/JAVA Dual-Stack is supported using the Master/Slave methodology, that's supported in all RHEL 7.x minor releases. Please follow the configuration guide: [Configure \(A\)SCS/ERS SAPInstance cluster resource](#).

However, since SAP no longer recommends dual-stack setups where both ABAP and JAVA instances share the same SID, it would be better to consider doing a [Dual-Stack Split](#), so that the ABAP and JAVA stack will use separate SIDs. With such a setup you then could use the Standalone approach for managing the ASCS/ERS instances for the ABAP stack and the SCS/ERS instances for the JAVA stack. You need to configure separate resource groups for each instance and set up appropriate constraints based on the separate SIDs for each part of the environment. Please follow configuration guide: [Configuring SAP NetWeaver ASCS/ERS ENSA1 with Standalone Resources in RHEL 7.5+ and RHEL 8](#) .

4.3. SUPPORT POLICIES

Please refer to [Support Policies for RHEL High Availability Clusters - Management of SAP NetWeaver in a Cluster](#).

4.4. CONFIGURATION GUIDES

- On-Premise: [Configuring SAP NetWeaver ASCS/ERS ENSA1 with Standalone Resources in RHEL 7.5+ and RHEL 8](#)
- AWS: [Configure SAP NetWeaver ASCS/ERS ENSA1 on Amazon Web Services \(AWS\)](#)
- Azure: [Azure Virtual Machines high availability for SAP NetWeaver on Red Hat Enterprise Linux with GlusterFS](#)
- Azure: [Azure Virtual Machines high availability for SAP NetWeaver on Red Hat Enterprise Linux with Azure NetApp Files for SAP applications](#)
- Azure: [High availability for SAP NetWeaver on Azure VMs on Red Hat Enterprise Linux with NFS on Azure Files](#)
- GCP: [HA cluster configuration guide for SAP NetWeaver on RHEL](#)