Best-Practice Document

Building High Availability for SAP NetWeaver and SAP HANA on Linux

A Technical Design Guide

Version 2.1

Dietmar-Hopp-Allee 16
D-69190 Walldorf

DATE January 2018

SAP PRODUCT(S)	PRODUCT VERSION(S)
SAP NetWeaver Based Products	SAP NetWeaver with Kernel 720 and higher
OPERATING SYSTEM(S)	DATABASE(S)
Linux	SAP HANA
ALM PHASE(S)	
Plan, Design, Build, Test, Deploy	
SAP SOLUTION MANAGER 7.1 SP	SAP SOLUTION MANAGER WORK CENTER(S)
All	Implementation/Upgrade

Table of Content

List of Figures

Change History

Version 2.1 (January 2018) – Update relevant features up to SAP HANA 2.0 SPS 2, IaaS cloud scenarios, certified SAP NetWeaver HA-Interface information, supported features and implementation guides from cluster vendors, and some minor corrections

Version 2.0.1 (October 2017) – Corrected with some minor changes

Version 2.0 (March 2017) – Update relevant features for SAP HANA 2.0, and Linux cluster solutions for SAP HANA System Replication automation including some IaaS cloud scenarios

Version 1.2 (July 2016) – Update features supported by different cluster solutions for SAP HANA System Replication automation and some minor changes

Version 1.1 (May 2016) – Update certified HA-Interface information and some minor corrections

Version 1.0 (March 2016) – Initial version

1 Objectives and Summary

1.1 Objectives and Scope of the Document

The purpose of this document is to support the technical architecture design and implementation when deploying SAP NetWeaver (NW) based solutions with SAP HANA on Linux Operating Systems, including High Availability (HA) and virtualization[1]. The focus is on SAP applications using SAP HANA as underlying database – for example, SAP Business Suite or SAP BW systems. The content of this document is based on SAP product features as of January 2018.

The objective of HA is to improve the availability of a software solution, protecting single point of failures in the software itself, the server hardware failures and storage, within the same data center, Disaster Recovery (DR) is the technology for protection across data centers. HA and DR are often combined with each other. However, DR is not in the scope of this document, even if the concept in this document can also apply to build an active-active data center solution when using two nearby data centers with sufficient bandwidth and low latency in between.

This document is not a step-by-step installation guide but instead provides guidelines (and details for some important points) for setting up HA for SAP NetWeaver and SAP HANA systems. The referenced documents to specific vendor information about step-by-step installations are listed in the footnotes and appendix of this document[2].

1.2 How to Read this Document

This document helps you to select the technical architecture and deployment model best suited to your requirements. The sections in the document guide the reader through the design and implementation process. We start with a general SAP NetWeaver and SAP HANA architecture introduction and the required general technical infrastructure, then move to the reference architecture design and implementation tips to reach HA by integrating with third-party cluster solutions, and conclude with HA verification scenarios at different levels. In the appendix, we list all referenced SAP Notes, detailed implementation guides, and other resources on the Web.

1.3 Summary

HA for a complete SAP Landscape of SAP NetWeaver based systems and SAP HANA on Linux can be achieved by designing and implementing technical measures to utilize native features of SAP NetWeaver and SAP HANA, in combination with third-party cluster solutions.

- Technically, the HA needs cover the following component levels for each specific service: applications, database, server, storage, and network;

[1] Virtualization is limited to the scope when it's HA related
[2] As the vendor may also update the relevant documents frequently, while you can directly find the exact references used by this document, you are also encouraged to search the latest references using the keywords in document titles

- SAP NetWeaver architecture (SAP NW ABAP only, SAP NW Java only, and SAP NW Dual Stack) are ready to run an HA environment, with the support of SAP Central Services, SAP Enqueue Replication Services, and scalable SAP Application Servers;
- In general, SAP Central Services Instances, SAP HANA Databases, SAP Global File Systems, and SAP Web Dispatcher are the Single Point of Failures (SPOFs), which need to be protected using native or third party cluster solutions;
- SAP Web Dispatcher is a software based loading balancing mechanism for HTTP(S) based requests, it requires HA when deployed in critical production environments;
- SAP HANA is scalable via scale-up on large server hardware or via scale-out using multiple server nodes. HA can be achieved with standby server nodes (SAP HANA Host Auto Failover) or synchronized secondary databases (SAP HANA Synchronous System replication);
- To automate the failover of SAP HANA System Replication, several third-party solutions are available for SAP HANA and SAP NetWeaver. Certification programs for these cluster solutions ensure quality of service;
- Most of the discussed HA cluster solutions work for both physical and virtual environments, while some virtual environments have specific solutions;
- The designed and configured HA solutions need to be verified after implementation at application, cluster and hardware levels, to ensure operational readiness.

Feedback to this document is welcome. Please contact stephen.sun@sap.com.

2 High Availability Overview for SAP NetWeaver and SAP HANA

2.1 High Availability Introduction

System outages can cost companies millions of dollars. The outages derive from Single Point of Failures (SPOFs) like network, hardware, and software failures, or even human errors. High Availability (HA) is the requirement to maximize system availability. From an end-user perspective the technology behind HA is hidden. This document demonstrates how to mitigate SPOFs with redundancy or failover solutions to minimize the risk of system outages and to achieve HA from technical perspective.

Recovery Point Objective (RPO) and Recovery Time Objective (RTO) are common terms to specify requirements for HA/DR solutions. RPO defines how much business data can be lost in terms of business activity time - or phrased differently - how far you can go back in time to resume application processing. So it is the target tolerance threshold for lost business transactions. RTO focuses on how fast the application can be "restarted" and brought back online, so it is the target maximum time to resume application functionality. In an HA context, normally what we have is the demand of "zero" for RPO (no data loss) and several minutes for RTO. The HA solution is expected to be capable of a fully automated failover for an active-passive cluster, or for continuous availability for an active-active cluster while no data is lost and business can continue.

SAP addresses these challenges within the SAP NetWeaver (NW) platform. SAP NetWeaver is a combination of several technical components, forming a technology platform to run SAP Business Suite or SAP BW systems. This document focuses on HA solutions for the following SAP related components.

- SAP Central Services Instances
- SAP HANA Databases
- SAP Global File Systems
- SAP Web Dispatcher

For hardware-related HA-solutions, please refer to relevant hardware partners.

2.2 SAP NetWeaver Architecture

2.2.1 SAP NetWeaver Installation Variations

Depending on the SAP NetWeaver products, the SAP NetWeaver platform has three kinds of installation variants:

- SAP NetWeaver ABAP (ABAP execution environment only)
- SAP NetWeaver Java (Java / Java EE execution environment only)
- SAP NetWeaver Dual Stack (both execution environments in a single system)

The NetWeaver Dual Stack systems are only applicable to SAP NetWeaver Process Integration (SAP NW PI) and SAP Solution Manager[3].

[3] Please note that the Dual Stack is deprecated. New SAP products will use separate systems for ABAP and Java execution environments. SAP NetWeaver PI systems, as well as SAP Solution Manager systems have to be split into separate single stack systems in new product releases.

When multiple AS instances are used, usually SAP Web Dispatcher can be used for load balancing end-user requests (HTTP(S)/DIAG).

The following architectures are applicable to SAP NetWeaver version 7.1 and higher.

SAP NetWeaver ABAP

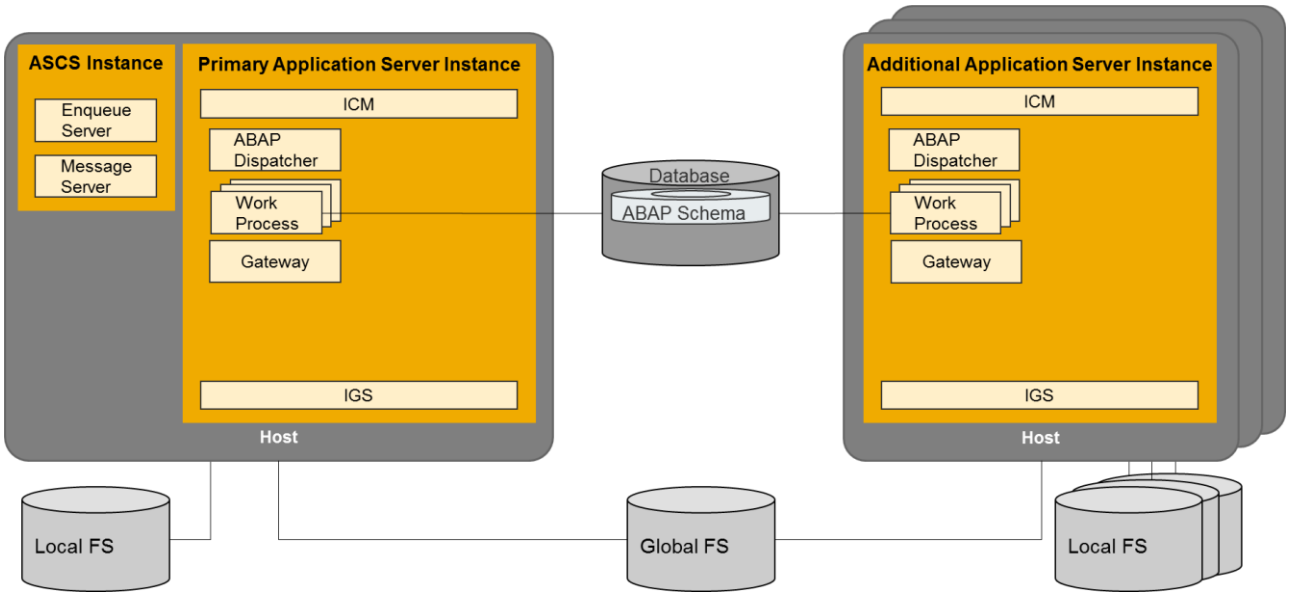

Figure 1 - General SAP NetWeaver ABAP Architecture

In a SAP NW ABAP system there exists

- A single Central Services (CS or ASCS) instance on local file systems. It consists of two processes: The Message Server (MSG) for the communication between application server instances, and the Enqueue Server (ENQ) holding a memory table with application locks.
- One or multiple Application Server (AS) instances on local file systems, serving user requests by running the related ABAP application. The first (or single) instance is called Primary Application Server (PAS). Each AS instance consists of multiple operating system processes:
  - ABAP dispatcher process, dispatching requests from users working with SAPGUI-based applications.
  - Multiple work processes for executing ABAP application code.
  - Internet Communication Manager (ICM) process, dispatching requests coming in via HTTP(S) protocol, e.g. from Web-based applications or for Web Service calls.
  - Gateway process (GW)[4], dispatching incoming requests and outgoing calls for Remote Function Calls (RFC).

[4] This is not the SAP NetWeaver Gateway, which is the technology providing a simple way to connect devices, environments and platforms to SAP software. SAP NetWeaver Gateway is needed for SAP Fiori.

  - Internet Graphic Server (IGS) processes for generating graphical output, e.g. for Web-based applications.

- One database management system.
- Global file systems for all instances, such as /sapmnt/<SID> and /usr/sap/trans.

SAP NetWeaver Java

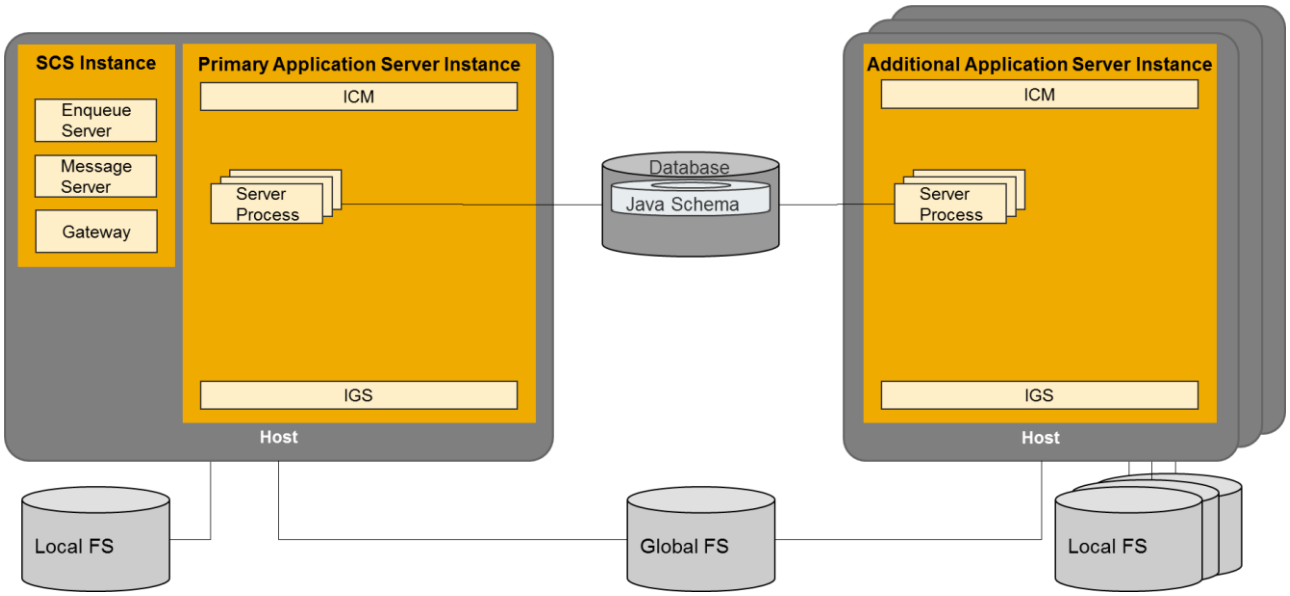

Figure 2 - General SAP NetWeaver Java Architecture

For each SAP NW Java system there exists

- A single Central Services (CS or SCS) instance on local file systems. It consists of three processes: The Message Server (MSG) for the communication between application server instances, the Enqueue Server (ENQ) holding a memory table with application locks, and Gateway (GW) process for and external communication.
- One or multiple Application Server (AS) instances on local file systems, serving user requests by running the related Java application. The first (or single) instance is called Primary Application Server (PAS). Each AS instance consists of multiple operating system processes:
  - Internet Communication Manager (ICM) handles requests from clients and dispatches them to the available server processes. ICM replaces the previously used Java Dispatcher process.
  - Multiple server processes for executing Java application code. The process is responsible for processing incoming requests, which are assigned by the ICM. Each server process is multi-threaded and can simultaneously process a large number of requests.
  - Internet Graphic Server (IGS) processes for generating graphical output, e.g. for Web-based applications.
- One database management system.

- Global file systems for all instances, such as /sapmnt/<SID> and /usr/sap/trans.

SAP NetWeaver Dual Stack

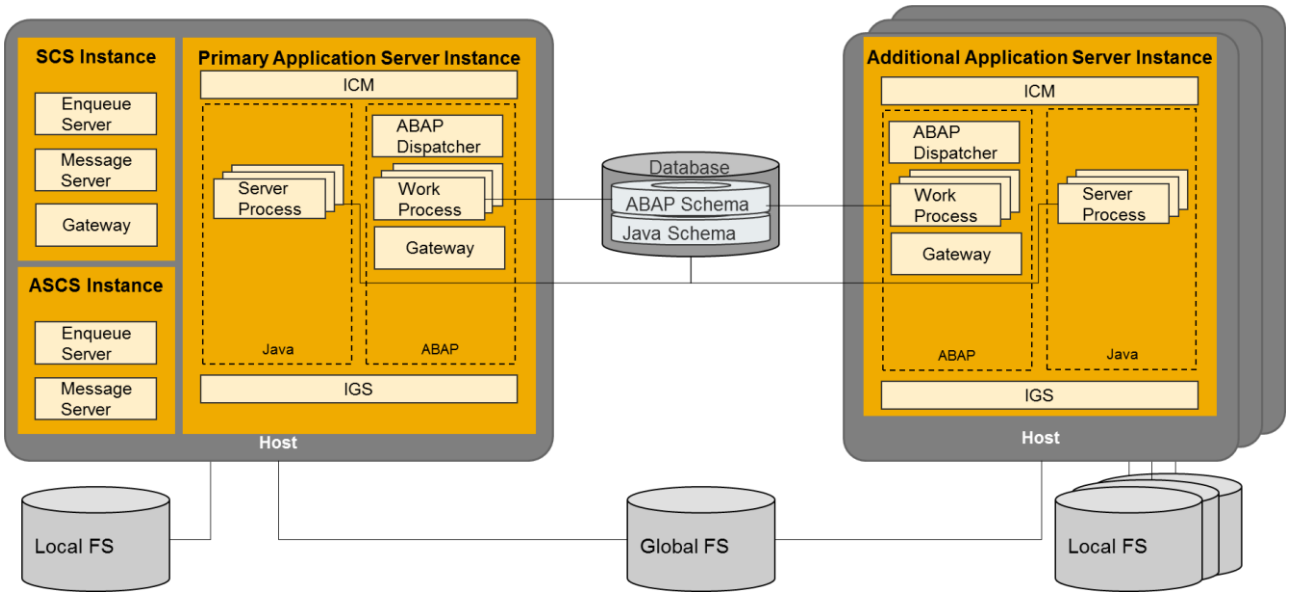

Figure 3 - General SAP NetWeaver Dual Stack Architecture

A SAP NW Dual Stack system is a combination of an SAP NW ABAP and SAP NW Java in one SAP system, using a single system identifier (SID):

- Central services are not shared or coupled. SAP NW AS ABAP and SAP NW AS Java both have an own instance for central services (ASCS for the AS ABAP, SCS for the AS Java).
- AS instances of the AS ABAP and AS Java part are coupled.
- Two separate database schemas are used in one shared database management system.
- Dual Stack installations are generally not recommended anymore, but as mentioned may be technically required for some SAP applications, e.g. SAP Solution Manager or SAP NW PI.

SAP Web Dispatcher

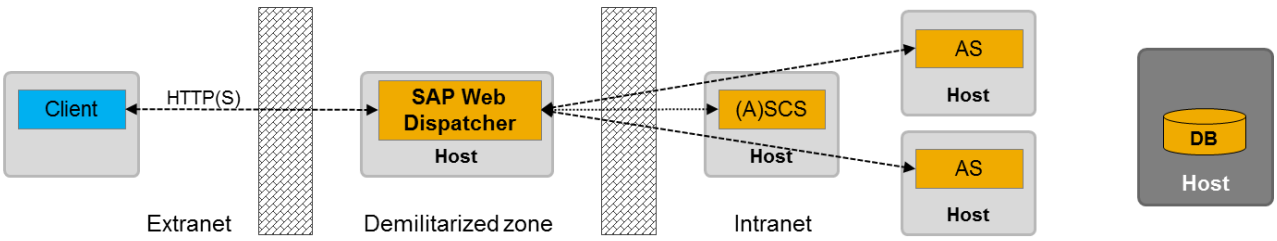

Figure 4 - General Architecture for HTTP(S) Load Balancing with SAP Web Dispatcher

There is also an option of using the SAP Web Dispatcher (WD) in the demilitarized zone (DMZ). This DMZ is located before the SAP system, where it distributes incoming HTTP(S) requests among the instances of the SAP system, according to specified rules.

2.2.2 Single Point of Failures of SAP NetWeaver

With SAP NW native capabilities, a system can scale by adding more AS instances to form an SAP NW AS cluster, so AS instances (PAS or AAS) are not Single Point of Failures (SPOF). However, to form a SAP NW AS cluster mitigating hardware failures, at least two AS instances need to run on different physical servers to protect from hardware failures.

All other components – namely SAP (A)SCS instance, database management system, global file systems – in the above architecture variants (Figure 1~3) are classified as SPOFs, as there exist one and only one copy of these components for each SAP NW system.

When an SAP Web Dispatcher (WD) is in use, it is also a SPOF and needs protection. Actually, there are three options[5] for SAP WD HA. In this document, we only describe how to use a third-party cluster solution for it, if you use other options, just follow SAP help document.

By default, with SAP NW ABAP stack, there is no longer an HA protected SAP gateway, if we only protect ASCS instance with an HA cluster. To solve this issue, it is possible to include SAP NW PAS in the cluster, but it is recommended to configure a standalone Gateway in an HA protected ASCS instance[6].

From simplification perspective, we only recommend to protect components classified as SPOFs using third-party cluster solution, either failover cluster or active-active cluster.

2.2.3 SAP Enqueue Server and Enqueue Replication Server

Unlike the Message Server in an (A)SCS instance, which can be restarted without concern as it's stateless, the Enqueue Lock Table information of Enqueue Server needs to be retained after unexpected Enqueue server failures. So when there is a failover of the cluster node, the (A)SCS services will switch to the surviving cluster node. In doing so, the Enqueue Lock Table will be lost. In order to preserve the Enqueue Locks, a duplicated Enqueue Table on the other cluster node is needed. The Enqueue Replication Server (ERS) does exactly this[7].

[5] See also High Availability of the SAP Web Dispatcher
[6] This is only needed if you use external programs that cannot be started on all application servers, e.g. registered server programs or programs installed on dedicated hosts only. See also SAP Note 1010990 - Configuring a Standalone Gateway in an HA ASCS instance
[7] See also High Availability with the Standalone Enqueue Server

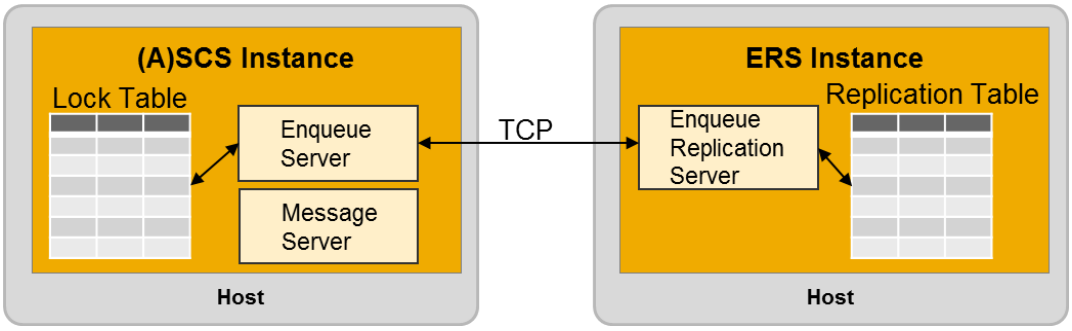

Figure 5 - SAP Enqueue Replication Server

For each change made to the Enqueue Lock table (for example, if the enqueue server receives an enqueue or dequeue request), the line changed in the Lock Table is passed to the ERS, which adds the line to the Replication Table. The response from the actual enqueue request is only sent to the enqueue (server) client once the replication is successfully completed[8] on the ERS.

2.2.4 SAP Global File Systems

A Global File System /sapmnt/<SID> needs to be available on all hosts with the same operating system[9].

Under /sapmnt/<SID>, there are three subfolders:

- global: contains globally shared data
- profile: contains the technical configuration profiles of all instances
- exe: contains executable programs (NetWeaver kernel and SAP JVM)

The directory /usr/sap/trans is the global transport directory shared within the transport landscape, such as development, quality assurance, and production systems. This also needs to be available on all hosts with an AS instance.

2.3 SAP HANA Overview

2.3.1 SAP HANA Introduction

SAP HANA is the Platform for Next-Generation Applications and Analytics. SAP HANA converges database and application platform capabilities in-memory to transform transactions, analytics, text analysis, predictive and spatial processing so businesses can operate in real-time. SAP HANA can be deployed on premise or cloud.

[8] See also Replication and Failover
[9] For a distributed SAP system (an SAP NetWeaver system with multiple instances on multiple hosts), including HA-enabled systems, a Global File System /sapmnt/<SID> needs to be available on all hosts with the same operating system. This document focuses on Linux based systems, so there is no need to discuss heterogeneous distributed system scenarios.

Looking at the SAP HANA architecture layers, SAP HANA is much more than just a database. However, this document focuses on High Availability, so SAP HANA is seen mainly as a database management system, which is classified as one of SPOFs in section 2.2.2.

The SAP HANA database supports both RowStore and ColumnStore storage. However, it is optimized for ColumnStore storage. Tables that are organized in columns are read optimized and have better compression rates than tables organized in rows. ColumnStore is typically suitable for big tables with bulk updates. However, update and insert performance is better on RowStore. RowStore is typically suitable for small tables with frequent single updates.

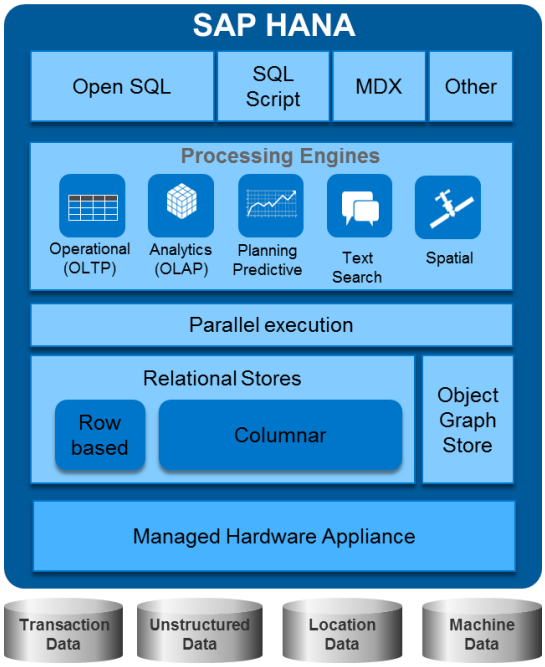

Figure 6 - SAP HANA Architecture Layers

The following is a brief overview of the most important server components of the SAP HANA system and the corresponding OS processes[10]:

Server Component	OS Process	Description
SAP start service	Sapstartsrv	Trigger the startup of HANA daemon process, and also be used to read log or trace files and determine the system status.
Daemon	Hdbdaemon	Responsible for starting and stopping the other services in the correct order, and keeping the core process running (restart the core process if in case they get killed)
Index server	hdbindexserver	Contains the actual data stores and the engines for processing the data.

[10] Also see SAP HANA Administration Guide

Server Component	OS Process	Description
Preprocessor server	hdbpreprocessor	Analyzes text data using an index server and extracts information based on the text-search capabilities.
Name server	hdbnameserver	Owns information about the topology of the SAP HANA system.
XS classic server	hdbxsengine	Provides application and application developers with an access to the SAP HANA Database using a consumption model, which is exposed through HTTP.
XS advanced runtime	hdbxscontroller hdbxsexeagent hdixsuaaserver	As of SAP HANA 1.0 SPS 11, SAP HANA includes an additional run-time environment for application development: SAP HANA extended application services (XS), advanced model. The SAP HANA XS Advanced runtime runs either on dedicated hosts or together with other SAP HANA components on the same host.[11]
Extended store server	hdbesserver	It is part of the SAP HANA dynamic tiering option of SAP HANA and provides a high-performance disk-based column store for very big data up to the petabyte range.[12]
Data provisioning server	hdbdpserver	It is part of the SAP HANA smart data integration option of SAP HANA and provides capabilities such as data provisioning in real time and batch mode, real-time data transformations, data quality functions, adapters for various types of remote sources, and an adapter SDK.[13]
Streaming cluster	hdbstreamingserver	Part of the SAP HANA smart data streaming option of SAP HANA. Smart data streaming extends SAP HANA with capabilities of SAP Event Stream Processor for consuming data streams and complex event processing.
SAP HANA Deployment Infrastructure (HDI) server	hdbdiserver	HDI handles the deployment of design-time artifacts into SAP HANA.
Compile server	hdbcompileserver	Compiles stored procedures and programs.
Script server	hdbscriptserver	The script server is used to execute application function libraries written in C++. The script server is optional and must be started manually[14].
Statistics service	hdbstatisticsserver	Collects information related to system performance, status, and resource usage[15].

[11] Also see SAP Note 2233866 - SAP HANA XS Advanced Model SPS11 Release Note
[12] Also see SAP Note 2140959 - SAP HANA Dynamic Tiering - Additional Information
[13] Also see SAP Note 2091095 - SAP HANA Enterprise Information Management
[14] Also see SAP Note 1650957 - SAP HANA Database: Starting the Script Server
[15] Also see SAP Note 2091313 - HANA Statistics Server - changed standard setting of the statistics server as of Revision 93

Server Component	OS Process	Description
SAP Web Dispatcher	hdbwebdispatcher	The native SAP HANA Web Dispatcher processes inbound HTTP and HTTPS connections to XS service. It has replaced the previous standalone executable sapwebdisp_hdb as of SAP HANA 1.0 SPS 9[16].

2.3.2 Native High Availability Options for SAP HANA[17]

For HA, SAP HANA provides two native options[18]:

1. Synchronous System Replication[19]: a synchronous standby database, implemented with the SAP HANA system replication feature, can be used to failover in case of hardware failures[20]. It provides short failover times, since database data is preloaded into main memory already. It additionally enables recovery from technical data corruptions on storage level. Technically, SAP HANA System Replication can also be used without preloading tables into memory, which allows to use the secondary system, for a non-production system. However, this will lead to longer failover/takeover time, so consequently it may not be considered as an HA implementation option when the RTO requirement is high.

   SAP HANA System Replication provides two operation modes for synchronization. **Delta data shipping** mode frequently transfers increments of the data (data modifications) from the primary to the secondary system (operationMode=delta_datashipping); **Continuous log replay** mode uses pure log-based transfer, and redo log is processed immediately on the secondary system (transaction boundary, operationMode=logreplay), which can reduce failover times and network traffic significantly. This continuous log replay is also called HotStandby and in this mode, the use of the secondary system for non-production system would need a bigger memory footprint than delta data shipping.

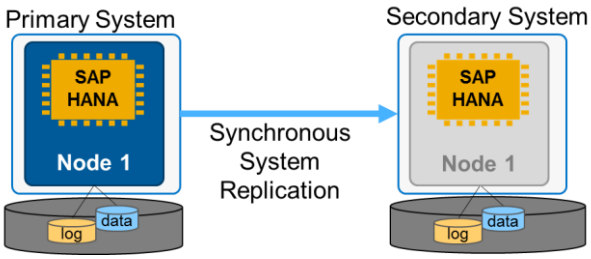

Figure 7 - SAP HANA Synchronous System Replication for Single Node (scale-up) System

[16] Also see SAP Note 1990354 - Changes in internal HANA Web Dispatcher as of HANA SP9
[17] Also see Technical Deployment Options for SAP Systems with SAP HANA
[18] Also see SAP Note 2057595 - FAQ: SAP HANA High Availability
[19] Also see SAP Note 1999880 - FAQ: SAP HANA System Replication
[20] Also see How to Perform System Replication for SAP HANA

Based on continuous log replay, SAP HANA 2.0 introduces Active/Active (Read Enabled) capability to enable read-only access on the secondary system for SAP HANA System Replication[21]. Active/Active (read enabled) is integrated into the System Replication solution and gets activated with the operation mode logreplay_readaccess. Active/Active (read enabled) inherits characteristics from continuous log replay:

- fast take-overs;
- reduced need for bandwidth in continuous operation;
- compatible with the existing replication modes: SYNC (with or without the full sync option), SYNCMEM, ASYNC.

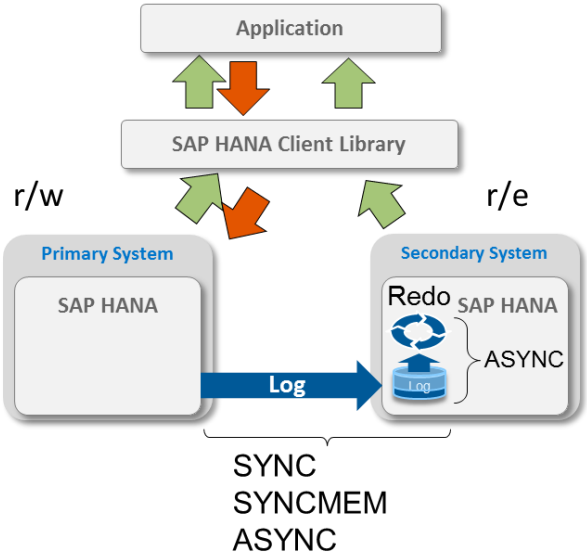

Figure 8 - SAP HANA Active/Active System Replication

There are several considerations related to the secondary read operations:

- the redo log replay runs as an asynchronous process on the secondary system. The secondary system provides statement level snapshot isolation with potentially delayed view on the data and no minimum delay guarantee;
- The secondary system allows read access if the primary system runs the same SAP HANA version;
- read-only access to the secondary system is not supported during rolling upgrade (zero downtime upgrade);
- the secondary system does not accept new connections if it cannot access the primary system;
- some internal processes for operations, not redo log relevant, like ColumnStore delta merges take place on the secondary system;

[21] Also see SAP HANA Administration Guide, SAP Note 2369981 - Required configuration steps for authentication with HANA System Replication

- both scale-up and scale-out setups are supported;
- there are two types of connections: explicit read-only connection to the secondary system, and implicit hint-based statement routing[22].

Especially, when working with third party cluster software, **the secondary system gets its own virtual IP addresses or host names** representing the secondary system.

Currently the following limitations apply:

- In a multi-tier SAP HANA system replication Active/Active (read enabled) is supported; however, the read access is limited to tier two and the connections cannot be opened to the tier three system;
- There is no Active/Active (read enabled) support for systems supporting Dynamic Tiering services, before SAP HANA 2.0 SPS 2[23];
- Before SAP HANA 2.0 SPS 2, the read access in the secondary system is not supported with row tables[24];
- The export of tables is possible with CSV as target. However, binary exports on the secondary system are not supported;
- Running further SAP HANA systems (SIDs) on the read enabled secondary is currently not supported[25].

2. Host Auto-Failover: This is a local fault recovery solution that can be used in addition or as an alternative measure to system replication. In a scale-out (multi-node) architecture, one or more server nodes can be configured as standby nodes, taking over data and processing in case of a hardware failure of another server node. Database data is not preloaded in the standby nodes, since a standby node should be able to replace different failed server nodes. This SAP HANA feature is called "host auto-failover". The standby nodes cannot be reused for different purposes except failover. A minimal highly available multi-node setup with two nodes, one active node plus one standby node, is theoretically possible. However, typically there are more active nodes than standby nodes. The minimal highly available scale-out setup with two server nodes is typically not attractive, since the similar SAP HANA system replication feature provides better failover times and more data protection. However, host auto-failover requires less hardware than system replication, since only one copy of data is kept. The host auto-failover is an SAP HANA feature that needs support from the underlying storage or file system to ensure avoidance of "split brain" situations. This capability is provided by certified SAP HANA storage hardware[26].

[22] Also see SAP Note 2392706 - HANA Active/Active (read enabled) hint based statement routing
[23] Also see SAP Note 2356851 – SAP HANA Dynamic Tiering Support for SAP HANA System Replication
[24] Also see SAP Note 2391079 - Access restrictions in Active/Active system setup, SAP Note 2402547 - Active/Active (read enabled) - monitoring of the secondary system with HANA Cockpit 2.0
[25] Also see SAP Note 2422925 - Restriction of secondary usage for other HANA systems (SIDs) in an Active/Active (read enabled) configuration
[26] Also see SAP Note 2062631 - high availability limitation for SAN storage

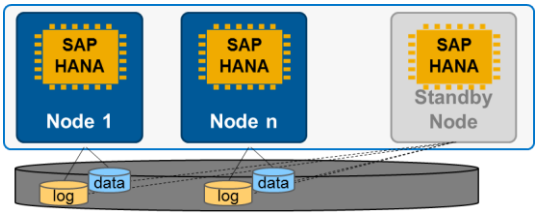

Figure 9 - Host Auto-Failover Normal Operation

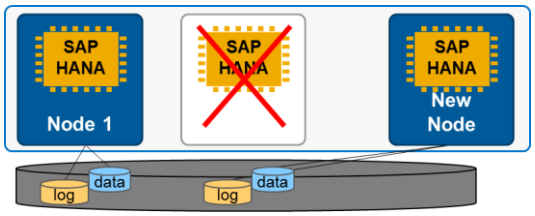

Figure 10 - Host Auto-Failover After Failover

Some hardware vendor also provide Disk or Storage based Replication for HA purpose, as mentioned previously, for HA scenarios, zero data loss (RPO=0) is expected, so only consider synchronous disk or storage replication[27]. This option replicates the data and the redo log of every committed and uncommitted transaction to a secondary hardware on storage level, so a crash recovery is needed for failover. Compared with SAP HANA system replication, this will again typically not attractive, since SAP HANA system replication feature provides much better failover times and more data protection.

SAP is continuously improving SAP HANA HA (and DR) capability:

- As of SAP HANA 1.0 SPS 5, SAP HANA introduces Synchronous System Replication (SYNC and SYNCMEM);
- in SAP HANA 1.0 SPS 6, Asynchronous System Replication is added;
- in SAP HANA 1.0 SPS7, SAP HANA is extended to support Multi-tier System Replication (cascading/chained systems)[28], Zero Downtime maintenance, and compressed log transfer;
- with SAP HANA 1.0 SPS 8, SAP HANA introduces a new option for the Synchronous Mode (SYNC) named Full Sync option and enhanced network management with separated network zones. Log write is successful when the log buffer has been written to the logfile of the primary and the secondary instance; in addition, when the secondary system is disconnected, such as because of network failure, the primary systems suspends transaction processing until the connection to the secondary system is re-established; no data loss occurs in this scenario;
- from SAP HANA 1.0 SPS 9, SAP HANA extends support for SAP HANA Multi-tenant Database Containers setups, improved take-over performance by preventing reload of RowStore during take-over, optimized delta data shipment in Multi-tier environments with both Log and Data Compression (according to compression algorithm LZ4) on transfer between sites, online Add Host & Remove Host, and HA/DR Provider Framework;
- in SAP HANA 1.0 SPS 10, SAP HANA improves support for SAP HANA Multitenant Database Containers setups, take-over times by parallelizing efforts with RowStore initialization;
- until SAP HANA 1.0 SPS 10, SAP HANA System Replication frequently transfers increments of the data with delta data shipping from the primary to the secondary system. From SAP HANA 1.0 SPS 11

[27] Also see SAP Note 1755396 - Released DT solutions for SAP HANA with disk replication
[28] Also see SAP Note 2303243 - SAP HANA Multitier System Replication – supported replication modes between sites

onwards and after initial data load, pure log-based transfer is possible, and redo log is processed immediately on secondary system, which can reduce failover times and network traffic significantly;

- in SAP HANA 1.0 SPS 12, snapshot data backups for initial data load for setting up SAP HANA System Replication is supported; and automated transfer of parameter changes to secondary instances at all sites can be enabled to minimize manual changes and risk of human errors;
- SAP HANA 2.0 SPS 0 introduces Active/Active (Read Enabled) capability to enable read-only access on the secondary system for SAP HANA System Replication, and supports Dynamic Tiering (DT) with two-tier synchronous replication mode [29], but Active/Active for DT is not available. Additional configuration steps are required to setup HANA System Replication, because replication connections now use certificate (SSFS) based authentication[30];
- in SAP HANA 2.0 SPS 1, SAP HANA System Replication for DT allows multiple tenants, but only one tenant with extended store; the performance is further improved: faster initial load from primary to secondary system with multi-streaming option per SAP HANA services (up to 32 streams); faster pre-load from disk to memory in secondary system with parallel (multi-threaded up to 32) table load;
- In SAP HANA 2.0 SPS 2, Row Store read access restriction is removed for Active/Active feature; two-tier asynchronous and three-tier replication for DT support are added with 'logreplay' and 'logreplay_readaccess' operation modes only; Active/Active for DT is supported on secondary site, but DT data will be excluded from query results; arbitrary number of tenants can be configured for System Replication with DT. Tenant copy/move into a primary of a active HANA System Replication as well as out of the primary to another target (other than the secondary) is possible; tenants that are copied/moved to the primary will be added to replication depending on the value of SECONDARY_FULLY_RECOVERABLE at the end of the operation. To avoid the changes in client connection string, after tenant move, a host-independent tenant access mechanism is introduced with HANA 2.0 SPS 2 accordingly, refer to Section "Multi-Tenant Database Containers" for more details.

2.3.3 SAP HANA HA Provider[31]

As of SAP HANA 1.0 SPS 9, the SAP HANA nameserver provides a Python-based API, which is called at important points of the host auto-failover and system replication takeover process.

These so called "hooks" or "HA/DR providers" can be used for arbitrary operations that need to be executed. One of the most important uses of the failover hooks is moving around a virtual IP address (in conjunction with fencing mechanism). Nevertheless, there are other purposes like starting tools and applications on certain hosts after failover or even stopping Development or Quality Assurance SAP HANA instances on secondary nodes before takeover.

[29] Also see SAP Note 2380229 - SAP HANA Platform 2.0 – Central Note, SAP Note 2380257 - SAP HANA Platform 2.0 SPS 00 Release Note, SAP Note 2420699 - Release of SAP HANA Database 2.0 for older SAP Versions
[30] Also see SAP Note 2369981 - Required configuration steps for authentication with HANA System Replication
[31] Also see Implementing a HA/DR Provider

2.3.4 Multi-Tenant Database Containers

SAP HANA introduces Multi-Tenant Database Containers (MDC) since SAP HANA 1.0 SPS 9. Starting with SAP HANA 2.0 SPS 1, MDC is the default setting and the only operational mode for SAP HANA systems[32]. MDC allows combining multiple databases on the same hardware. MDC uses strictly separated databases ("tenants") that share the same SAP HANA system: the SAP HANA software, the server hardware, and storage system. It allows an SAP HANA system to have control over the physical resources and balance workload between database tenants with SAP HANA-embedded heuristics. SAP HANA currently provides static CPU and memory resource assignment. Multi-tenancy is supported for both single-node (scale-up) and multi-node (scale-out) SAP HANA architectures. Multiple database tenants can share a single-node server hardware, but a mixture of single-node and multi-node databases is also possible in the same (scale-out) SAP HANA system.

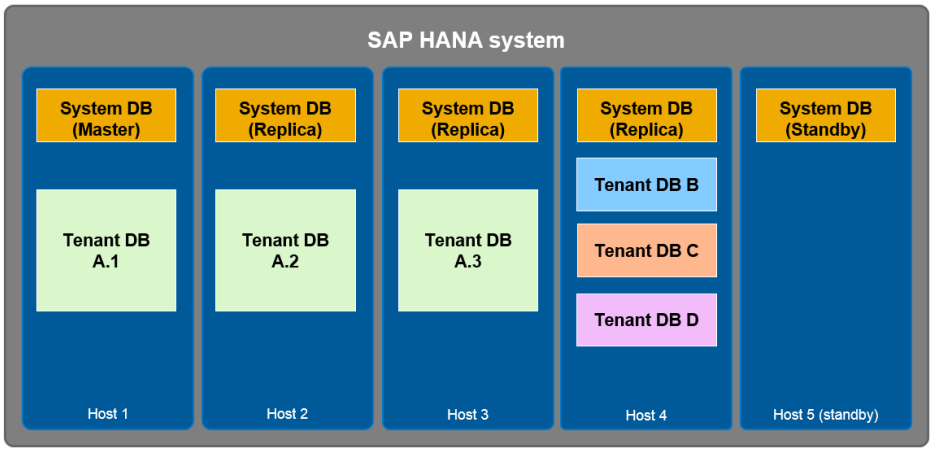

Figure 11 - SAP HANA Multi-Tenant Database Containers

Multi-tenancy is supported for production and non-production systems. Database lifecycle management (start/stop), backup, and restore are separate for each tenant, but can also be managed for the whole system with all tenants. MDC is the strategic deployment option for SAP. It combines flexibility and efficient hardware usage with the possibility to share data on the SAP HANA platform level.

The availability implication for MDC:

- One database software version for a SAP HANA system (all tenant databases), so if one tenant database requires certain software version, it has to be changed for all tenant databases
- One HA setting for a SAP HANA system: system replication setup for the whole SAP HANA system (all tenant databases), all tenants are included in one HA scenario[33]:
  - SAP HANA System Replication and storage replication can be used to replicate the whole system for HA

[32] Also see SAP Note 2423367 - Multitenant database containers will become the standard and only operation mode
[33] Also see SAP Notes 2092793 Reduced delivery with SPS09: Restrictions for the usage of SAP HANA system replication with multi-tenant database containers, and SAP Note 2096000 SAP HANA multitenant database containers - Additional Information

- Primary and secondary systems must be identical at the time of system replication setup, including tenant databases
- The replication process treats the complete collection of tenant containers as one
- All tenants fail over together
- Replication of a single tenant to an individual location is not possible as of today
- Adding, deleting, changing tenants requires setting up the replication newly

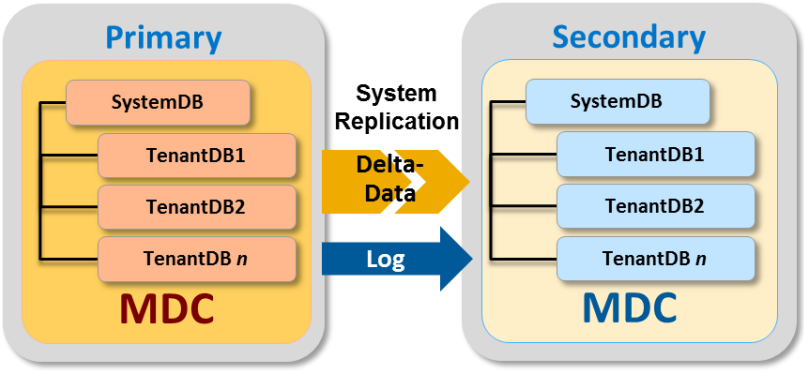

Figure 12 - SAP HANA System Replication with MDC delta_datashipping Operation Mode

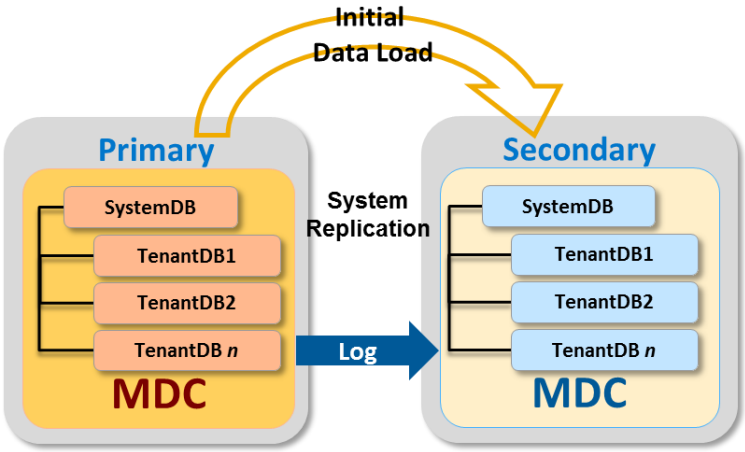

Figure 13 - SAP HANA System Replication with MDC logreplay Operation Mode

Host-independent tenant addresses for tenant move scenarios:

- Tenant database client connectivity methods
  - Client connect via tenant port (connect string <hostname>:<tenant port>): after tenant move, this does not work anymore since tenant port may change. It is hardly possible to manually keep all port numbers stable;

  - Client connect via System DB (connect string <hostname>:3<instance_number>13@<tenant name>): this does not work either because the target MDC system has a different hostname (physical or virtual) and a different System DB port.

- Host-independent SAP HANA Client connect supporting tenant move
  - Since a tenant's ports will change with a tenant move, the client should connect to a tenant via the System DB using the specific tenant database name;
  - To establish a connection to a tenant independent of its current host, the connection should use a virtual IP/hostname and an additional port number that is mapped to the tenant;
  - As of SAP HANA 2.0 SPS 2 revision 22, this request is addressed. The additional port number can be configured, and this is independent of the SAP HANA system number and can be chosen individually for each tenant or for a group of tenants in a group of MDC (MDC Group) enabled systems where tenants are supposed to move across[34]. This is actually to configure additional ports on which the System DB listens in addition to port 3<instance_number>13;
  - A additional unique virtual IP address/hostname is assigned for each individual tenant (or MDC group) and is used for the connection string instead of the already existing hostname (physical or virtual). This virtual IP address is following the tenant (or MDC group) when it is moved to another HANA system by relocating the virtual IP;
  - The connection string used by a HANA client to access a tenant thus looks as follows: <tenant specific virtual IP/hostname>:<landscape wide tenant port>@<tenant name>. The System DB knows the actual port of all its current tenants and returns the port number and actual hostname to the client to establish the final connection;
  - Since all additional port numbers can be added to the System DB of the candidate hosting SAP HANA systems, this part is fully transparent to tenant move. For the tenant specific virtual IP address, the relocation can not be fully automated since it involves different SAP HANA systems.

- Implications for SAP HANA System Replication failover automation with third party cluster solutions
  - Besides the relocatable virtual IP address for external accesses always stays within the primary SAP HANA System, and the additional relocatable virtual IP address for external read-only accesses to the secondary SAP HANA System for SAP HANA 2.0 when Active/Active feature enabled, all tenant specific virtual IPs/hostnames also need to be relocated/handled automatically; the third party cluster solutions should be able to handle these relocatable IP addresses automatically during failover since this is still working within the same SAP HANA system and the same cluster; when Active/Active feature is enabled, the required tenant specific virtual IPs/hostnames are actually needed to be doubled;
  - Consequences for a tenant move: all relevant tenant specific virtual IPs/hostnames must be moved/relocated from source SAP HANA system to the target SAP HANA system. For the

[34] Also see Configure Host-Independent Tenant Addresses

running Linux clusters, you can first stop/offline the relevant virtual IPs/hostnames on the source Linux cluster (you may need go to maintenance mode before doing that), and then manually add the virtual IPs/hostnames on the target hosts. And then in the next maintenance window, you can adjust the virtual IP resource and corresponding relationship configurations of the Linux clusters. The virtual IPs/hostnames resources should be able to be reconfigured online, but you may want to adjust and test them properly during the announced maintenance windows.

2.3.5 Multiple Components One System (MCOS or Multi-SID)

SAP HANA supports multiple SAP HANA systems (SIDs) share one server (also known as MCOS or "Multi-SID") deployment option for production system usage as of SAP HANA 1.0 SPS 9[35].

Do not confuse "Multi-SID" with SAP HANA MDC feature[36]. In the MDC architecture, there is one SAP HANA System (SID), with numerous tenant DBs running inside that one SAP HANA system, but each tenant DB does not have its own SID. Specifically, "Multi-SID" has the following characteristics:

- 1 x Hardware + Operating System
- n x SAP HANA systems
- n x DB schemas
- n x (SAP NetWeaver) Systems

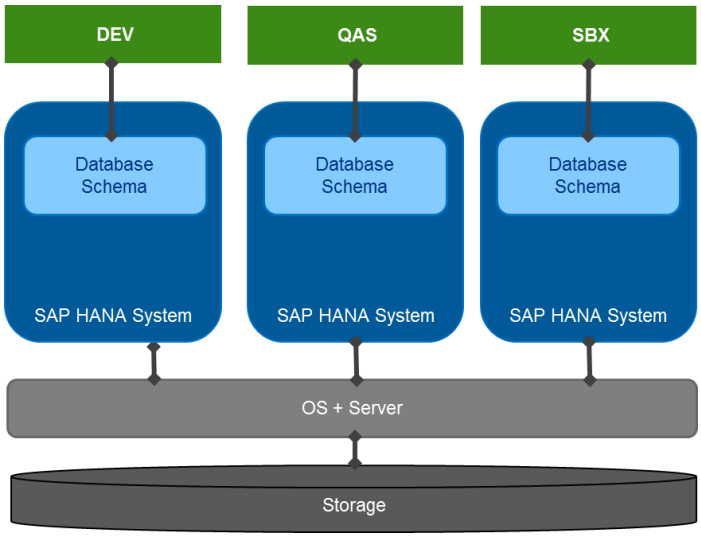

Figure 14 - SAP HANA Multiple Components One System (MCOS or Multi-SID) Deployment Option

[35] Also see SAP Note 1681092 - Multiple SAP HANA DBMSs (SIDs) on one SAP HANA system
[36] Also see SAP Note 1661202 - Support multiple applications one SAP HANA database / tenant DB, SAP Note 1826100 - Multiple applications SAP Business Suite powered by SAP HANA

For production systems, this is restricted to scale-up scenarios only; for non-production usages, scale-out setups are also supported. Check SAP Note 1681092 for more restrictions. You need also ensure that your SAP HANA system is sized appropriately for any additional SIDs that you wish to deploy[37].

2.3.6 Dynamic Tiering

SAP HANA Dynamic Tiering (DT)[38] is an optional add-on to the SAP HANA database for managing historical data. Its purpose is to extend SAP HANA memory with a disk-centric columnar store (as opposed to SAP HANA's in-memory store) for managing less frequently accessed warm data. Warm data has relaxed performance requirements compared to highly active "hot" data. Data in the extended store is on line, and available for both queries and updates. Although it is possible to add dynamic tiering to small SAP HANA databases, dynamic tiering is targeted at SAP HANA database sizes of 512GB and larger, where large data volumes begin to necessitate a data lifecycle management solution.

Specially, DT provides the ability to create and process disk-based columnar database tables, called extended tables, in addition to the traditional column- or row-oriented in-memory tables of the SAP HANA database. Like in-memory tables (Hot Store), extended tables (Warm Store) are first-class objects of the database with full ACID compliance. All database requests regarding extended tables are integrated into SAP HANA's transactional context. With SAP HANA 2.0 SPS 0, SAP HANA introduces a new database object called a multistore table. A multistore table is a type of partitioned SAP HANA column table that has partitions in both SAP HANA default column storage, and dynamic tiering extended storage.

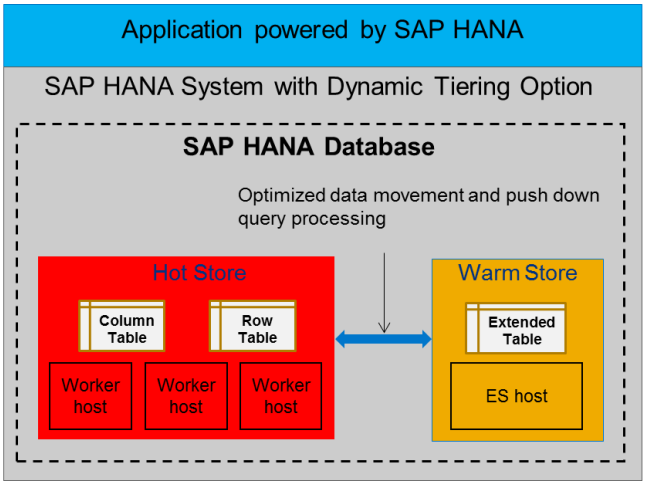

Figure 15 - SAP HANA System with Dynamic Tiering Option

[37] See also SAP Note 2422925 - Restriction of secondary usage for other HANA systems (SIDs) in an Active/Active (read enabled) configuration
[38] See also SAP Note 2140959 - SAP HANA Dynamic Tiering - Additional Information, SAP Note 2394124 - SAP HANA 2.0 Dynamic Tiering - Additional Information

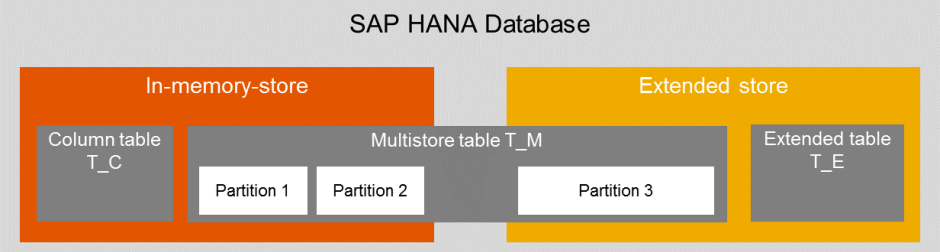

Figure 16 - SAP HANA Multistore Partitioned Column Table

With SAP HANA 2.0 SPS 0, there are some architecture improvements for DT:

- Co-deployment: the SAP HANA and DT servers may be installed on the same host for production scale-up systems. However, compute resource usage between the SAP HANA and DT servers is not completely isolated at this time, even you can limit the amount of memory used by DT[39]. Separate hosts deployments are still recommended for optimum performance; on storage layer, different physical units (disk, flash) should be used for SAP HANA and DT.
- SAP HANA System Replication support: DT can be added to an SAP HANA system running two-tier synchronous System Replication for HA.

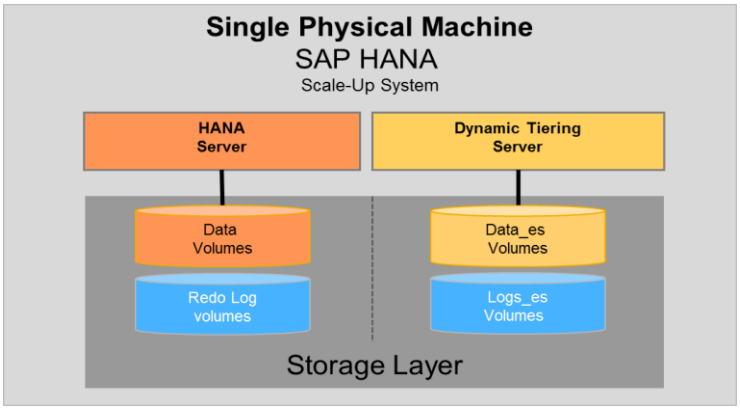

Figure 17 - Co-Deployment of SAP HANA and Dynamic Tiering on single Host

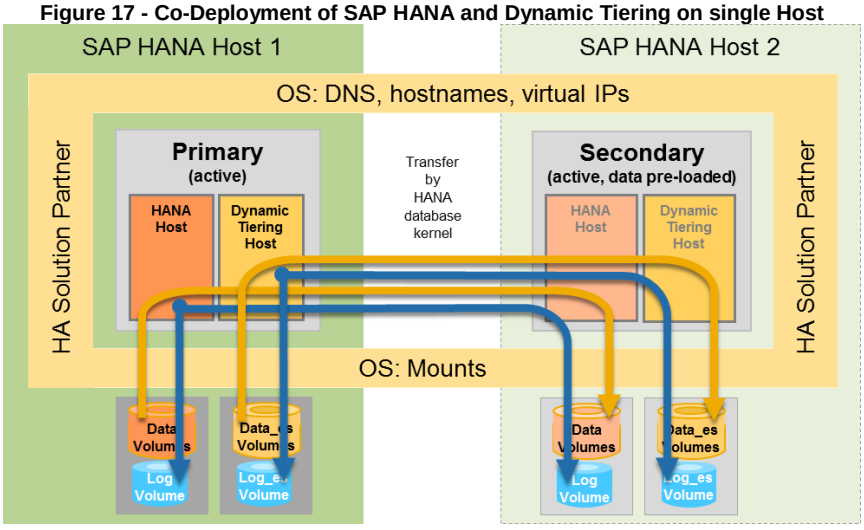

Figure 18 - SAP HANA System Replication with Dynamic Tiering

[39] Also see SAP Note 2086829 - SAP HANA dynamic tiering sizing ratios

With SAP HANA 2.0 SPS 1, SAP HANA System Replication for DT allows multiple tenants, but only one tenant with extended store; and adds BACKINT support with Dynamic Tiering in MDC set up.

In SAP HANA 2.0 SPS 2, two-tier asynchronous and three-tier replication for DT support are added, with 'logreplay' and 'logreplay_readaccess' operation modes only; Active/Active for DT is supported on secondary site, but DT data will be excluded from query results; arbitrary number of tenants can be configured for System Replication with DT.

3 Technical Infrastructure for High Availability

A high availability architecture employs fault resilient hardware, standby hardware and software solutions, automated failure detection, and recovery functions to ensure availability of business operations. From a technology perspective, the architecture is dependent on the business requirements and the products used to implement these business functions within the infrastructure.

HA starts by focusing on and preventing failures from occurring in the first place. Eliminating Single Points of Failures with redundancy is the most important aspect in the planning phase.

Building a solution for system wide HA involves designing each of the following components for availability. In this document, the requirements for involved technical infrastructure components are discussed.

Area	Protection requirements	Protection Measures	HA	DR
Data Center[40]	Loss of data center or data center facilities	Multiple data centers, and data replication between data centers		✓
Applications	Failure of application server components	Standby/failover setup or redundancy (active-active setup)	✓	
Database	Failure of database components	Standby/failover setup or redundancy (active-active setup)	✓	
Server	Hardware failure, Operating System service failure, network disconnect	Standby hardware, redundant server components	✓	
Storage	Hardware failure, data corruption	Replication (multiple data copies), redundant storage device components	✓	✓
Network	Hardware or service failure, connection interruption	Redundant network infrastructure, redundant network device components	✓	

3.1 Server

A hardware platform decision depends on many factors. Required hardware capacity, workload characteristics, existing corporate infrastructure, partner relationship, application compatibility, and availability requirements are just a few factors in a hardware platform selection process.

The SAP Installation Guides[41] describe the basic hardware requirements for each kind of product component. To size the hardware for specific applications and workloads, SAP provides the SAP Quicksizer[42] tool and standardized SAP performance benchmarks[43].

[40] Not discussed in this document
[41] Find the guides at http://service.sap.com/instguides
[42] Available at http://service.sap.com/quicksizer
[43] Also see SAP Standard Application Benchmarks

For an HA implementation, redundant and reliable servers are needed. In an HA cluster, the first level of redundancy is to have at least one additional server to build an active-active or active-standby cluster. This second server provides the capability to take over resources even if the primary server fails. In addition, redundancy should also be included for the components within the servers to make each server reliable. These components include communication networks, power, memory, and storage connectivity.

3.2 Operating Systems

The Linux operating system has gained popularity over the last several years in Data Centers. For SAP HANA, Linux is still the only supported platform. There are many Linux distributions, three of them are supported to run for SAP environments – SUSE Linux Enterprise Server, Red Hat Enterprise Linux, and Oracle Linux[44].

SUSE Linux Enterprise Server (SLES) is a Linux based operating system developed by SUSE. SLES for SAP Applications is based on SUSE Linux Enterprise Server and is optimized for SAP workloads. It contains the SUSE Linux Enterprise High Availability Extension (SLE HAE), which is an integrated suite of open source clustering technologies[45]. Based on different usages, SLES for SAP Applications has provided preconfigured patterns – "SAP Application Server Base", "SAP NetWeaver Server Base", "SAP HANA Server Base". SAP has provided Best Practices SAP Notes for different supported versions[46]. SLES for SAP Applications offers the following HA related important components: Pacemaker, DRBD (Distributed Replicated Block Device), OCFS2 (Oracle Cluster File System ver 2), ReaR (Relax and Recover), LVS (Linux Virtual Server) and Resource Agents for SAP application servers and SAP HANA System Replication.

Red Hat Enterprise Linux (RHEL) is a Linux distribution developed by Red Hat and targeted toward the commercial market. To run SAP Application Servers on Red Hat Enterprise Linux (RHEL), RHEL for SAP Business Applications[47] is recommended. It offers the following additional HA related important components: GFS/GFS2 (Global File System), LVS, Resource agents and scripts for Pacemaker based High Availability setups on RHEL 6.5 and later[48]. SAP has also provided Best Practices SAP Notes for different supported versions[49].

Oracle Linux, also known as Oracle Enterprise Linux, is a Linux distribution based on RHEL, repackaged and distributed by Oracle[50]. Oracle Clusterware is the native High Availability Extension from Oracle. Currently Oracle Linux can be used for SAP Application Servers, but not for SAP HANA[51].

---

[44] Also see SAP Note 171356 - SAP Software on Linux: General information, and SAP Note 936887 - End of maintenance for Linux distributions
[45] Also see SAP Note 1763512 - Support details for SUSE Linux Enterprise High Availability
[46] Also see SAP Note 1310037 - SUSE LINUX Enterprise Server 11: Installation notes, SAP Note 1984787 - SUSE LINUX Enterprise Server 12: Installation notes, and SAP Note 2001528 - Linux: SAP HANA Database SPS 08, SPS 09 and SPS 10 on RHEL 6 or SLES 11
[47] Also see SAP Note 1631106 - Red Hat Enterprise Linux for SAP Business Applications, SAP Note 2397039 - FAQ: SAP on RHEL, and Red Hat Solution https://access.redhat.com/solutions/34169
[48] Also see SAP Note 1908655 - Support details for Red Hat Enterprise Linux HA Add-On
[49] Also see SAP Note 1496410 - Red Hat Enterprise Linux 6.x: Installation and Upgrade, and SAP Note 2002167 - Red Hat Enterprise Linux 7.x: Installation and Upgrade
[50] Also see SAP Note 1565179 - SAP software and Oracle Linux
[51] Also see SAP Note 2235581 - SAP HANA: Supported Operating Systems

3.3 Network

Data centers typically run multiple separate networks, including an Ethernet network for client-to-server and server-to-server communications and a Storage Area Network (SAN, Fibre Channel is the predominant for now). Network-Attached Storage (NAS) could also be used as storage solution. Here storage devices are connected to a NAS server that makes the storage available at a file system level to the other computers across the Local Area Network (LAN, over Ethernet).

To support various types of networks, data centers use separate redundant interface modules for each network: Ethernet network interface cards (NICs) and Fibre Channel interfaces (Host Bus Adapters, HBAs) in their servers, and redundant pairs of switches at each layer in the network architecture.

A network interface card is the point of interconnection between a computer and a private or public network. To achieve HA at network level, redundant NICs are needed.

The term Link Aggregation applies to various methods of combining (aggregating) multiple network connections in parallel in order to provide redundancy in case one of the links fails and in addition to increase throughput beyond what a single connection could provide.

Linux uses a special kernel module called Bonding to allow users to join multiple NICs into a single channel. When bonded, two NICs appear to be the same logical device having the same MAC address[52].The terms bonding and teaming are quite often used interchangeably, but in Linux world, teaming is a new functionality, which is working a bit differently to bonding. Bonding works in Linux Kernel space, while teaming is a User space bonding variant with the help of the *teamd* daemon to scale better.

SAP HANA uses multiple network connections to transfer data, for example, from clients to the database during standard operations, between the nodes in a scale-out solution, between primary and secondary systems and to persist data on storage. The components belonging to SAP HANA communicate over several logical network zones[53]:

- HANA Client zone: Different clients, such as SQL clients on SAP application servers, browser applications using HTTP/S to the SAP HANA XS server and other data sources (such as BI) need a network communication channel to the SAP HANA database.
- Internal zone: The internal zone covers the communication between hosts in a distributed SAP HANA system as well as the communication used by SAP HANA system replication[54] between two SAP HANA systems.
- Storage zone: Although SAP HANA holds the bulk of its data in memory, the data is also stored in persistent storage locations – this layer is accessed via a network – to provide protection against power failures or other events that cause hosts to become unavailable.

[52] Also see SUSE Linux Administration Guide or Red Hat documentation
[53] Also see SAP HANA Network Requirements Whitepaper
[54] Also see Network Recommendations for SAP HANA System Replication

- Admin zone: used for administrator and/or external administration tools, it is possible to be configured on a separate network with low-bandwidth requirements.

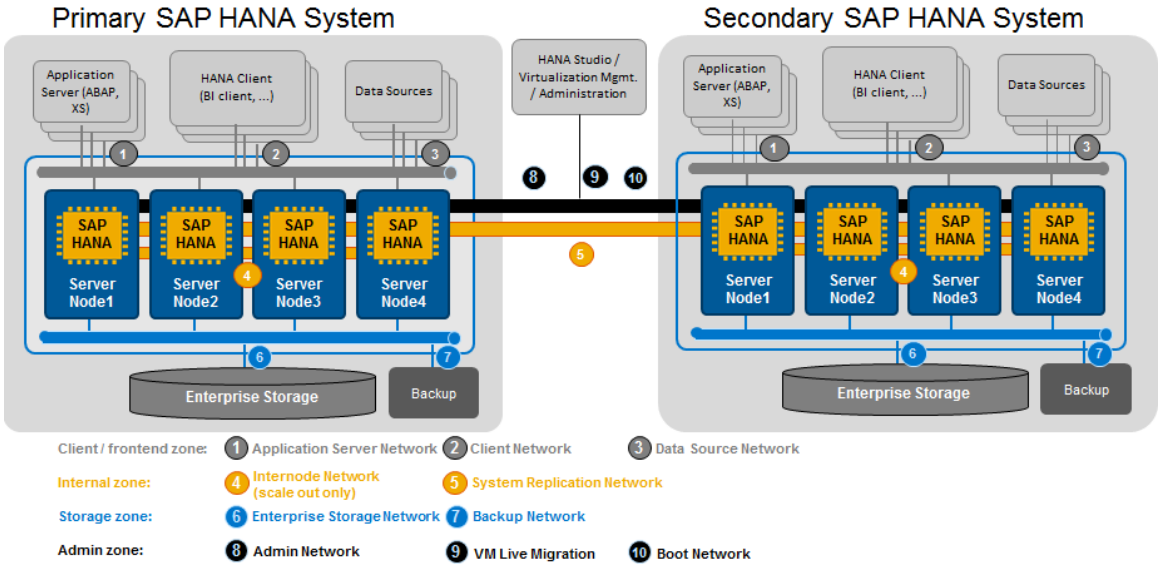

Figure 19 - SAP HANA Logical Network Zones

The below table lists the recommendation network settings for each kind of SAP HANA network:

Name	Function	Bandwidth	Quality of Service	Protection Measures
HANA Client / Frontend Zone				
Application Server Network	SAP AS and Database	10 Gigabit Ethernet		NIC Bonding
Client Network	DB Client/Client Apps and Database	1 / 10 Gigabit Ethernet		NIC Bonding
Data Source Network	Data import, replication and integration between Database	1 / 10 Gigabit Ethernet		NIC Bonding
Internal Zone				
Internode Network	Built-in node to node communication for scale-out SAP HANA system	10 Gigabit Ethernet	Non Blocking[55]	NIC Bonding
System Replication Network	Use of SAP HANA built-in system replication	10 Gigabit Ethernet for local data center	Non Blocking	NIC Bonding
Storage Zone				
Backup Network	Backup and Restore	10 Gigabit Ethernet		NIC Bonding

[55] Non-Blocking guarantees for every network port full performance even when network switch is fully utilized

Name	Function	Bandwidth	Quality of Service	Protection Measures
Storage Network	SAP HANA Node to Storage communication	10 Gigabit Ethernet or >= 8 Gigabit Fibre Channel	Non Blocking	NIC Bonding
Admin Zone				
Admin Network	SAP HANA administration	1 / 10 Gigabit Ethernet		NIC Bonding
Boot Network	Boot operating system from remote source	1 / 10 Gigabit Ethernet		NIC Bonding
VM Live Migration Network	Virtualization (Network between virtualized hosts)	10 Gigabit Ethernet	Non Blocking	NIC Bonding

3.4 Cluster Software

Cluster software helps ease the complexity of a cluster environment by automating both the administration and monitoring. It provides a centralized console to quickly review and analyze resources. Several cluster software packages can be used to make an SAP solution highly available.

SAP does not support the third-party cluster software itself, but rather makes the platform ready to run in a properly set up cluster environment. SAP helps cluster software vendors in the development of guidelines for setting up an SAP HA environment and in the process of solving problems observed in SAP customer systems.

3.4.1 Avoiding "Split Brain" Situations - Quorum, STONITH or Fencing Mechanisms[56]

In a cluster environment for HA, Quorum/STONITH/Fencing is a very important concept. Basically, it is a method to bring an HA cluster to a known state avoiding a so-called "split-brain" status: A computer cluster is a collection of cooperating computers or nodes, that talk to each other over communication channels to understand the node and resource status. There could be a situation when the nodes could not reach each other to understand with certainty the state of some node or resource. With a fencing mechanism, even when the cluster doesn't know what is happening on some node, it is possible to make sure that certain critical resources could not run on more than one node.

Fencing mechanisms can be categorized into resource level and node level. Resource level fencing makes sure that a node cannot access one or more resources, such as a SAN, where a fencing operation changes rules on the SAN switch to deny access from a node. The node level fencing assures that a node does not access any resources at all, which is usually done by a power reset using the power switch.

[56] Also see Fencing and Stonith

3.4.2 Linux Cluster Solutions for SAP NetWeaver

Most Linux cluster solutions can be used for all SAP Business Suite / SAP NetWeaver releases, such as SUSE Linux Enterprise Server for SAP applications, Red Hat Cluster Suite or RHEL HA add-on, HPE Serviceguard Extension for SAP on Linux, and so on[57].

However, there is one classical problem running SAP instances in a high available environment: if an SAP administrator changes the status of an SAP instance without using the cluster management software, then the cluster framework will interpret this as an error. It will automatically bring the SAP instance into the previous status by either starting or stopping the SAP instance. This can result in critical situations, if the cluster changes the status of a SAP instance during critical SAP maintenance tasks, such as during an SAP system upgrade[58]. Therefore, for SAP products based on NetWeaver Kernel 7.20 and higher, SAP offers a tighter integration of HA solutions with the SAP HA Script Connector Library and the HA-interface (Cluster API) certification[59] for the control of component sapstartsrv to achieve the following purposes[60]:

- Certify Start/Stop infrastructure within HA-Setups
- Establish a single SAP HA setup scenario
- Unify HA setups for SAP NetWeaver 7.x based on 7.20 DCK availability (ABAP, Java)
- Improve transparency about OS dependencies and 3rd party failover software
- Improve transparency about available DB-specific HA-Solutions of 3rd party vendors
- Clarify responsibilities of SAP and its partners, especially for support

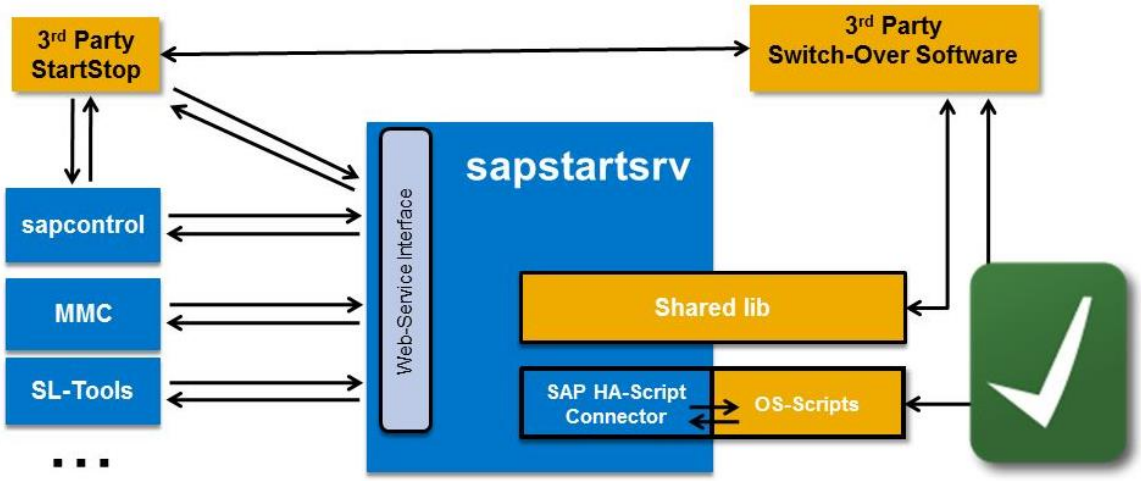

Figure 20 - Cluster-aware sapstartsrv Solution

Involved components of the cluster-aware sapstartsrv solution:

[57] See also SAP Note 1552925 - Linux: High Availability Cluster Solutions
[58] See also Using sap_vendor_cluster_connector for interaction between cluster framework and sapstartsrv
[59] See also SAP Note 1693245 - SAP HA Script Connector Library
[60] See also Manual Page of sap_vendor_cluster_connector

- Client (responsible: SAP): The clients like sapcontrol, startsap, SAP NetWeaver Adaptive Computing Controller (ACC), SAP Management Console and others are the same as in the standard scenario. They are part of the SAP software delivery. All clients must use the web service interface of sapstartsrv. No client should work around this procedure.
- Server (responsible: SAP): The server process sapstartsrv is providing web service interfaces for starting and stopping of SAP instances, and also for getting information about the status of an SAP instance. This program is provided by SAP.
- Library (responsible: SAP): The library saphascriptco.so for Unix is loaded by sapstartsrv, if the parameter service/halib is used in the SAP instance profile. Sapstartsrv uses the library as a generic interface to a set of cluster products. The library saphascriptco.so takes the calls by sapstartsrv and matches them to calls to an external helper program (the cluster connector).
- Cluster connector (responsible: cluster vendor): the cluster connector sap_vendor_cluster_connector (such as sap_suse_cluster_connector, or sap_redhat_cluster_connector) is called by the library saphascriptco.so and is an abstraction layer between the SAP sapstartsrv and the cluster product. While sapstartsrv and sapscriptco.so do not need anything to know about the cluster product, the cluster connector needs to answer all queries or to trigger cluster commands to change the status of cluster resources.
- Cluster product (responsible: cluster vendor): the cluster product must be able to answer the queries of the cluster connector and processes the triggered changes of cluster resources.

The below table lists the status of certified HA-Interface partners and solutions for Linux platform[61]:

Partner Name	HA Solution Name	Certified Operating Systems	SAP Support Note	Certified Until[62]
Hewlett Packard Enterprise Co.	HPE Serviceguard A.12.00	Linux on x86_64	1779681	03 Jul. 2019
NEC Corporation	EXPRESSCLUSTER 3.3	Linux on x86_64	1768213	07 Mar. 2019
NEC Corporation	EXPRESSCLUSTER 3.3	Linux on Power	1768213	14 Mar. 2020
Oracle Deutschland B.V. & Co. KG	Oracle Clusterware 12c (Rel. 1 for SAP)	Linux on x86_64	1496927	22 Mar. 2019
Sios Technology Corp.	SIOS Protection Suite for Linux v9	Linux on x86_64	1662610	23 Feb. 2019

[61] See also Certified HA-Interface Partners and High Availability Partner Information
[62] Some certifications has already expired or are close to expiration date, regularly check Certified HA-Interface Partners for status update. All certifications last for 3 years; if a certification expires it does not mean the solution does not work anymore, but is still fully supported by the cluster vendors

Partner Name	HA Solution Name	Certified Operating Systems	SAP Support Note	Certified Until[62]
SUSE LINUX Products GmbH	SUSE Linux Enterprise High Availability 12	Linux on x86_64	1763512	7 Aug. 2020
SUSE LINUX Products GmbH	SUSE Linux Enterprise High Availability 12	Linux on Power	1763512	22 Jan. 2021
Veritas Technologies LLC	InfoScale Availability / InfoScale Enterprise 7.0	Linux on x86_64	1906423	18 Jan. 2020
Red Hat, Inc.	Red Hat Enterprise Linux 6	Linux on x86_64	1908655	09 Dec. 2016

3.4.3 SAP HANA System Replication Automation with Cluster Software

The SAP HANA System Replication feature has no dependencies on the chosen SAP HANA hardware server model, but is a native feature of the SAP HANA software. It can be combined with a cluster management solution from partners to automate failover management, and then reduce the technical downtime in case of hardware failure. In this document, we describe five SAP HANA System Replication automation solutions based on SLES for SAP Applications, RHEL for SAP HANA with High Availability, Veritas InfoScale Availability, HPE Serviceguard Extension for SAP, and IBM Tivoli System Automation for Multiplatforms for SAP HANA System Replication.

3.4.3.1 SLES for SAP Applications with "SAPHana"/"SAPHanaController" and "SAPHanaTopology" Resource Agents[63]

SUSE created a new solution under "SLES for SAP Applications" to automate the takeover in SAP HANA System Replication setups. For scale-up SAP HANA systems, SLES for SAP Applications uses the new SAPHana Resource Agent (RA) to check the SAP HANA replication. The master node (master/slave here is the status for the special multi-state resource under Pacemaker based cluster environments) assumes responsibility for the SAP HANA databases running in primary mode, and the slave is responsible for instances operated in synchronous (secondary) status. The master/slave resource also bundles the client access virtual IP address to the primary SAP HANA system; for scale-up scenarios, it's supported to provide one additional relocatable virtual IP address for external read-only accesses for SAP HANA 2.0 when Active/Active feature enabled.

[63] Also see Fail-Safe Operation of SAP HANA: SUSE Extends Its High-Availability Solution, Automate SAP HANA System Replication with SLES for SAP Applications, Automate your SAP HANA System Replication Failover Setup-guide, SAP HANA SR Cost Optimized Scenario: SUSE Linux Enterprise Server for SAP Applications 11 SP4, SAP HANA SR Performance Optimized Scenario: SUSE Linux Enterprise Server for SAP Applications 11 SP4, SAP HANA SR Cost Optimized Scenario: SUSE Linux Enterprise Server for SAP Applications 12 SP1, SAP HANA SR Performance Optimized Scenario: SUSE Linux Enterprise Server for SAP Applications 12 SP1, Best Practices for Mission-Critical SAP Applications, SAP HANA System Replication Automation (HanaSR) for HANA Scale-Out now officially available with SLES for SAP 12 SP2, SAPHanaSR-ScaleOut: Automating SAP HANA System Replication for Scale-Out Installations with SLES for SAP Applications

SUSE offers within the solution a second RA, the SAPHanaTopology RA running on all cluster nodes, gathers information about the status and configuration of SAP HANA System Replications. This RA analyzes the SAP HANA topology and sends all findings via the node status attributes to all nodes in the cluster. These attributes are taken by the SAPHana RA to control the SAP HANA Systems. In addition it starts and monitors the local saphostagent. These RAs are provided in SAPHanaSR-*.noarch.rpm RPM package.

SUSE also provides RAs and tools for setting up and managing automation of SAP HANA system replication for scale-out setups. With this mode of operation, internal SAP HANA HA mechanisms and the SUSE cluster have to work together. SAP HANA is able to replace failing nodes with standby nodes or to restart certain subsystems on other nodes. As long as the SAP HANA landscape status is not "ERROR", the SUSE cluster will not act. From SUSE cluster, there are two RAs (SAPHanaController and SAPHanaTopology) running on each SAP HANA node. SAPHanaController is designed as Pacemaker master/slave resource. The Pacemaker master status will be assigned to the SAP HANA Active Master Nameserver of the primary site. All other SAP HANA nodes are running with the Pacemaker slave resource status. SAPHanaController is able to handle the takeover to the other site, and restart former failed worker nodes as standby. Additionally, a dedicated majority-maker node helps to make the right decision, which of the two SAP HANA sites will survive. In addition to the SAPHanaTopology RA, who analyzed the current status and shares it's findings with SAPHanaController, the SUSE cluster for scale-out setups also uses a "HA/DR providers" API provided by SAP HANA to get informed about the current state of the system replication. These RAs are provided in SAPHanaSR-ScaleOut-*.noarch.rpm RPM package.

If the parameter "AUTOMATED_REGISTER" of the cluster configuration is activated[64], the former (failed) primary SAP HANA system can be automatically registered to be a new secondary system. This enables the System Replication resumes automatically after a successful takeover.

Until SLES for SAP Application 12 SP1, only cluster with replicated single-node (scale-up) SAP HANA instances is officially supported. It allows the operation of a replicated production SAP HANA instance together with non-production SAP HANA instances (cost optimized scenarios), which will be halted as part of a takeover operation, with certain restrictions[65]. If the secondary SAP HANA instance is not co-deployed with non-production instances (performance optimized scenarios), the secondary instance should be configured to preload the column tables into memory to shorten failover time. Multi-tier (three-tier for now) SAP HANA System Replication is also supported, but the cluster is only be defined for a leading (between the first and second SAP HANA systems) synchronous SAP HANA System Replication. Failover of primary to secondary SAP HANA instance works automatically, but later on when the failed system rejoin the replication, the asynchronous replication relationship to the third system must be manually removed. The feature of automation of SAP HANA System Replication of SAP HANA scale-out systems is introduced as of SLES for SAP Application 12 SP2. Currently, only single-tier System Replication for the scale-out setups, and there is no other SAP HANA systems on the replicating node which needs to be stopped during takeover.

[64] By default, it is not activated, because if a cluster node goes down, an administrator should have a detailed look what happens and then take actions accordingly to bring it back

[65] Additional storage hardware is required to create independent I/O channels and subsystems that do not interfere with the replication infrastructure; do not pre-load column tables of the replication system into SAP HANA system memory; all additional instances on the secondary node must be stopped before triggering a takeover

3.4.3.2 RHEL for SAP HANA with SAPHana and SAPHanaTopology Resource Agents[66]

Red Hat Enterprise Linux (RHEL) for SAP HANA[67] with HA meets stability, reliability and availability requirements by enhancing the native SAP HANA System Replication and failover technology to automate the takeover process.

RHEL for SAP HANA with HA provides pacemaker resource agents (SAPHana and SAPHanaTopology, the same as SLES for SAP Applications) to automate the takeover from primary system to secondary system. These agents are provided in resource-agents-sap-hana*.x86_64.rpm RPM package.

Today, only clusters with replicated single-node (scale-up) SAP HANA instances are officially supported.

The client access virtual IP address is bundled with the primary SAP HANA system; for scale-up scenarios, it's also supported to provide one additional relocatable virtual IP address for external read-only accesses for SAP HANA 2.0 when Active/Active feature enabled.

3.4.3.3 SAPDatabase Resource Agent for SAP HANA under Pacemaker based Clusters

The purpose of the SAPDatabase resource agent[68] is to start, stop, and monitor the database instance of an SAP system. The operation of the resource agent can test the availability of the database by using SAP tools (R3trans or jdbcconnect). With that, it ensures that the database is truly accessible by the SAP system. Since February 2012, this resource agent also supports SAP HANA[69].

This is a cold failover solution similar as SAP Host Auto-Failover, but also manages the takeover of the virtual service IP address. Compared to SAP HANA Synchronous System Replication, this is again not attractive from failover times and data protection perspectives.

3.4.3.4 Veritas InfoScale Availability with Cluster Server Agent for SAP HANA[70]

Veritas InfoScale Availability (formly Symantec Veritas Cluster Server or VCS) can connect multiple independent systems into a management framework for increased availability. Each system, or node, runs its own operating system (Linux, Linux or Microsoft Windows) and cooperates at the software level to form a High Availability cluster.

Together with cluster framework, Veritas InfoScale Availability uses HA agents for application monitoring, failover and control. An "Agent Pack" is the release vehicle for Veritas HA agents. The HA agents can be classified as Application agents, Database agents, and Replication agents.

The High Availability agent for SAP HANA (application agent sap_agent[71], SAPHDB resource type) provides HA for SAP HANA instances, replicated with SAP HANA System Replication. The agent brings the SAP HANA

[66] Also see Automated SAP HANA System Replication with Pacemaker on RHEL Setup Guide, Supported HANA HA Scenarios on Red Hat Enterprise Linux
[67] Also see Red Hat Enterprise Linux (RHEL) 6.5 Configuration Guide for SAP HANA
[68] Also see Using SAPInstance to control non-productive SAP HANA database in the cost-optimized scenario
[69] Also see https://github.com/ClusterLabs/resource-agents/commits/master/heartbeat/SAPDatabase
[70] Also see Cluster Server Agent for SAP HANA Database Installation and Configuration Guide, SAPNW.7.0.8.0.SUSE Linux Enterprise Server, and SAPNW.7.0.8.0.Red Hat Enterprise Linux
[71] Also see Cluster Agent Pack Getting Started Guide

instance online, monitors the instance, and takes the instance offline. The agent monitors system processes and server state, and shuts down the server in case of a failover.

The agent currently supports the following configurations:

- Single-tier SAP HANA System Replication
  - Two single-node SAP HANA instances
  - Two scale-out SAP HANA instances
- Multi-tier SAP HANA System Replication
  - Three single-node SAP HANA instances

  One instance is in VCS with Global Cluster Option (GCO) and the other two instances are in the local cluster. Nevertheless, since VCS with GCO is for Disaster Recovery purpose, this is not discussed in this document.

The application agent (sap_agent) is also used to protect SAP NetWeaver (A)SCS and AS instances (SAP NetWeaver with HA-API support, SAPNW resource type) [72].

3.4.3.5 HPE Serviceguard Extension for SAP - SGeSAP[73]

HPE Serviceguard, formerly known as MC/Serviceguard, is a HA cluster software produced by HP that runs on HP-UX and Linux. HPE Serviceguard Extension for SAP on Linux (SGeSAP/LX) extends HPE Serviceguard's failover cluster capabilities to SAP application environments. It is intended to be used in conjunction with the HPE Serviceguard Linux product and the HPE Serviceguard Toolkit for NFS on Linux.

SGeSAP supports single-node and multiple-nodes synchronous SAP HANA System Replication for HA. The cluster automates SAP HANA instance startup, shutdown, monitoring, restart, and re-initialization of corrupted replications, takeover, and role-reversal operations. SAP HANA System Replication and dual-purpose configurations allow the operation of a production SAP HANA instance together with a non-production SAP HANA instance, which will be halted as part of a takeover operation, with the same restrictions as SLES for SAP Applications for dual-purpose scenarios. A non-production HANA instance can be installed on the secondary (replication) SAP HANA node. To support SAP HANA clusters, SGeSAP provides the cluster modules sgesap/hdbinstance, sgesap/hdbprimary, and sgesap/hdbdualpurpose (for reusing secondary SAP HANA system hardware for non-production purposes), the SAP HANA service monitor saphdbsys.mon, SAP cluster easy deployment tool deploysappkgs, and SAP cluster API library saphpsghalib.so.1[74]. A relocatable virtual IP address for external accesses always stays within the primary SAP HANA System, and one additional relocatable virtual IP address for external read-only accesses could be created to stay within the secondary SAP HANA System for SAP HANA 2.0 when Active/Active feature enabled.

[72] Also see SAP Note 1906423 - Support details for Veritas
[73] Also see HPE Serviceguard Extension for SAP Version B.12.10.00 Release Notes for Linux
[74] Also see Managing HPE Serviceguard Extension for SAP for Linux Version B.12.10.00

SAP HANA scale-out systems can be clustered (with HPE Serviceguard for Linux Enterprise Version) to realize many-to-many failover. The solution can be combined with cold standby instances and SAP HANA Host Auto-Failover to add a layer of local HA with the specific Quorum mechanism - Smart Quorum[75].

SGeSAP HA clusters also support a Multi-tier (three-tier for now) SAP HANA System Replication setup, but only be defined for a leading synchronous SAP HANA System Replication. Failover of primary to secondary SAP HANA instance works automatically. Subsequently, the failed over primary can become the new secondary (demotion operation) during the secondary package startup routine, but first, the asynchronous replication relationship to the third system must be manually removed. The SGeSAP HA Multi-tier feature is not available for scale-out SAP HANA systems.

3.4.3.6 IBM Tivoli System Automation for Multiplatforms for SAP HANA System Replication[76]

IBM Tivoli System Automation for Multiplatforms (TSAMP) manages the availability of applications running on Linux and AIX clusters. Reliable Scalable Cluster Technology (RSCT) is a product that is fully integrated into TSAMP. RSCT is a set of software products that together provide a comprehensive clustering environment. TSAMP uses the RSCT infrastructure to provide clusters with improved system availability, scalability, and ease of use.

The TSAMP SAP HANA System Replication policies use TSAMP to automate SAP HANA components and detect failed components and restarts them or initiates a failover/takeover. It utilizes SAP HANA primary instance, SAP HANA secondary instance, ServiceIP, and various status resources to maintain the HA of SAP HANA instances.

Only clusters with two replicated single-node (scale-up) SAP HANA instances are supported at this time.

3.4.3.7 Summary of Cluster Solution Features

The table below lists the features officially supported by different cluster solutions for SAP HANA System Replication automation[77]:

[75] In Serviceguard for Linux 12.00.30, Smart Quorum is introduced to handle the split brain DR scenario. One additional Quorum Serviceguard Quorum Server is added to provide arbitration services for Serviceguard clusters when a split brain is discovered (two equal-sized groups of nodes become separated from each other), the Quorum Server allows one group to achieve quorum and form the cluster, while the other group is denied quorum and cannot start a cluster

[76] Also see IBM Tivoli System Automation for Multiplatforms, Version 4.1, SAP HANA System Replication high availability policy

[77] This is the latest support status when the document is written, but may be changed any time with further development, so check the relevant vendors from time to time. **Please note:** this is not official support statement from SAP, but rather a summary of current development status snapshot, as the partner is responsible for the support of the relevant HA solution and the supported features

Features	SLES for SAP Applications	RHEL for SAP HANA with High Availability	Veritas InfoScale Availability	HPE Serviceguard Extension for SAP	IBM Tivoli System Automation for Multiplatforms [78]
Single-node (scale-up)	✓	✓	✓	✓	✓
Multiple-nodes (scale-out)	✓①	-	✓	✓	-
Automatic Register original primary as new secondary	✓	✓	✓②	✓②	✓
MDC (Multi-Tenant Database Containers)	✓	✓	-③	✓	-③
Reuse Secondary SAP HANA Hardware for additional SAP HANA instances④	✓	✓	✓	✓	-③
Single-tier System Replication (like A->B)	✓	✓	✓	✓	✓
Multi-tier System Replication (like A->B->C)	✓⑤	✓⑤	✓⑤	✓⑤	-
MCOS (Multiple Components One System / Multi-SID) for scale-up SAP HANA system	✓	✓	-	✓	-
Separate virtual IP address for secondary SAP HANA system with SAP HANA 2.0 Active/Active Read Enabled	✓⑤	✓⑤	-⑥	✓	-
Separate virtual IP addresses for SAP HANA tenants to support host-independent SAP HANA Client accesses for tenant move scenarios as of SAP HANA 2.0 SPS 2	✓⑦	✓⑦	-⑥	✓⑦	-
Dynamic Tiering	-	-	-	-	-

[78] SAP HANA on Power only at this moment

① staring from SLES for SAP Application 12 SP2
② only available for single-tier System Replication
③ not specified clearly, but this is crucial from SAP HANA 2.0 SPS 1, since then MDC is the default setting and the only operational mode
④ only useful for an HA solution when the RTO requirement is not high and cost is important factor
⑤ not available for scale-out SAP HANA systems
⑥ the relevant SAP application agent from Veritas does not support SAP HANA 2.0 yet
⑦ there is no technical limitation to support this

3.4.3.8 Other Cluster Solutions for Specific Scenarios or Environments

Fijitsu FlexFrame Orchestrator Solution

Some vendor, such as Fujitsu, provides highly integrated solution, which can make you perform hardware configuration, SAP system installation, SAP HANA System Replication failover automation with specifically designed procedure[79]. By leveraging SLES for SAP Application, Fujitsu FlexFrame Orchestrator (FFO, formly FlexFrame for SAP) solution is able to automate SAP HANA System Replication failover for both scale-up and scale-out setups. Please note Fujitsu FFO solution support two flavors: cold standby failover via shared storage for SAP HANA without System Replication, and failover automation for SAP HANA with System Replication.

SIOS Protection Suite for Linux SAP HANA Recovery Kit

SIOS has provided two brief guides (SIOS Protection Suite for Linux SAP HANA 1.0 Recovery Kit and SIOS Protection Suite for Linux SAP HANA 2.0 Recovery Kit) for automating SAP HANA System Replication, but the detailed supported scenarios and configurations are not documented yet.

Cluster Solutions for IaaS Cloud Environments

Cloud environments are now becoming more and more popular, and many companies are deploying their critical applications on Infrastructure as a Service environments. Companies are using SAP HANA not only for fast analysis of big data but also for their mission critical systems. SAP also supports to run SAP HANA on IaaS environments with certain IaaS providers, specifically Amazon Web Services (AWS), Huawei Technologies Co. Ltd[80], IBM Cloud, Alibaba Cloud Computing Limited, Google Cloud Platform, and Microsoft Azure[81] at this moment[82].

[79] Also see White paper Safeguarding Business Continuity on Critical SAP HANA Systems, SAP Note 2127006 – SAP HANA DB used in Fujitsu FlexFrame Orchestrator, Fujitsu FlexFrame Orchestrator Manuals
[80] for native HANA Data Mart scenarios or standalone only at this moment
[81] Also see SAP HANA (large instances) high availability and disaster recovery on Azure
[82] Also see Certified Iaas Platforms for SAP HANA, SAP Note 1380654 - SAP support in public cloud environments, SAP Note 1964437 - SAP HANA on AWS: Supported AWS EC2 products, SAP Note 2316233 - SAP HANA on Microsoft Azure

Within these IaaS providers for SAP HANA, currently AWS has two available cluster solutions to automate SAP HANA System Replication failover for scale-up setups: NEC EXPRESSCLUSTER[83] and SLE HAE[84]. IBM Cloud mentioned SAP HANA System Replication can be configured with an automated failover from one server to a replica, with implementation of high availability solutions for SAP HANA through Red Hat or SUSE Linux[85]. Microsoft Azure also provides extensive documentation about available HA configurations for SAP NFS, SAP Central Services, and SAP HANA System Replication[86]. Google Cloud Platform and Alibaba Cloud provide automatic HA functions based on Virtual Machine (Compute Engine or Elastic Compute Service / ECS) and instance restart capabilities[87].

3.4.3.9 Implementation notice: Hostname Case-Sensitivity for SAP HANA Servers Protected with Cluster Software

Some components within SAP HANA automatically convert hostnames into lower case, such as in SAP HANA topology. Since Linux operating systems and cluster software does not do this, an exact matching of the hostname strings are not possible any more if uppercase hostnames are used. Then subsequently, the missing variables in the Resource Agents result in unexpected behaviors.

So even if both uppercase and lowercase hostnames are allowed[88], it is best practice to use lowercase hostnames when working with Linux clusters for SAP HANA, even in general UNIX/Linux based SAP environments[89].

3.5 Virtualization and Host Failover Clustering

Virtualization can increase IT agility, flexibility, and scalability while creating significant cost savings. It is supported by SAP to run SAP software on Linux and certain combinations of suitable servers with the virtualization technologies including Xen, KVM, VMware, LPAR technology, and IBM Power and System z

[83] Also see SAP Business Suite Powered by SAP HANA High Availability with NEC EXPRESSCLUSTER, the corresponding White Paper, An Integrated High Availability and Disaster Recovery Solution - Version 3.3 New Enhancement, and SAP Note 2302728 - Supported scenarios with NEC Expresscluster on Amazon Web Services; NEC also describes one failover automation scenario for SAP HANA System Replication on virtualized IBM Power Systems with NEC EXPRESSCLUSTER at White Paper A High Availability Solution for SAP S/4 HANA Running on IBM Power Systems Implemented by EXPRESSCLUSTER
[84] Also see SAP Note 2309342 - SUSE Linux Enterprise High Availability Extension on AWS for SAP HANA
[85] Also see SAP on IBM Bluemix infrastructure and SAP HANA on IBM Bluemix Infrastructure Implementation Guide
[86] Also see SAP Note 2513384 - SUSE Linux Enterprise Server for SAP Applications on Azure, High Availability of SAP HANA on Azure Virtual Machines (VMs), High availability set up in SUSE using the STONITH, and High availability for SAP NetWeaver on Azure VMs on SUSE Linux Enterprise Server for SAP applications
[87] Both Google Cloud Platform and Alibaba Cloud support SAP HANA System Replication, but how to automate the takeover is not mentioned. Also see Setting Instance Availability Policies on Google Cloud Platform, SAP HANA High Availability and Disaster Recovery on Alibaba Cloud
[88] Also see SAP Note 611361 - Hostnames of SAP servers
[89] Also see SAP Note 1386233 - TREX or BWA installation on hosts with uppercase hostnames

based products[90]. Today within them, VMware vSphere 5.x and 6.x, Hitachi LPAR 2.0, Huawei FusionSphere[91] 3.1, 5.1 and 6.0, IBM PowerVM LPAR, KVM and Xen are certified to virtualize SAP HANA System[92].

Virtualization products can also provide HA features for virtual machines running SAP applications, that is, when the virtual machine failures are detected, they could be restarted on surviving hosts by certain policies, such as VMware HA[93], SLE HAE[94], RHEV HA[95], Citrix XenServer HA[96], Xen HA[97], and so on.

In this case, the failover software is installed on the hypervisor inside a physical machine, the failover takes place between physical machines, while the virtual machines are treated as a kind of resource. We can take this as Host Failover Clustering provided by virtualization.

With VMware HA, hardware and virtual machine level failures can be protected. Several products are available to complement VMware HA and make it application aware. To enable SAP application level monitoring and HA, products like Veritas/Symantec ApplicationHA[98], Neverfail vAppHA[99], SLE HAE or SIOS LifeKeeper Single Server Protection (SSP)[100] can help by introducing application monitoring and restart capabilities[101]. Currently Veritas/Symantec ApplicationHA can also protect single-node SAP HANA deployment on VMware: ApplicationHA Heartbeat agent sends the application heartbeat to the VMware HA. Veritas/Symantec ApplicationHA uses the application heartbeat to communicate the status of the application to the VMware HA. If a SAP HANA host fails, the ApplicationHA may perform the following actions: (1) restart the SAP HANA system for a configurable number of times; (2) if is unable to start the SAP HANA system, VMware HA takes the necessary corrective action[102], after the virtual machine is restarted, Veritas/Symantec agent for SAP HANA attempts to start the HANA and its dependent components in a predefined order[103].

As mentioned earlier, the Enqueue Lock Table information of the Enqueue Server needs to be retained after an unexpected loss of the Enqueue server. Simply restarting SAP Central Services (CS) using VMware HA is not a solution for HA, since it will lose the content of the Enqueue Lock Table. Together with VMware HA, VMware also provides the "VMware Fault Tolerance (FT)" solution[104] to protect SAP Central Services instances. VMware FT protects a virtual machine by maintaining a second virtual machine that runs in lockstep

---

[90] Also see SAP Note 1122387 - Linux: SAP Support in virtualized environments, and SAP Note 1492000 - General Support Statement for Virtual Environments
[91] Also see SAP Note 2186187 - SAP HANA on Huawei FusionSphere in production (controlled availability), SAP Note 2279020 - SAP HANA on Huawei FusionSphere 5.1 and 6.0 in production
[92] Also see SAP Note 1788665 - SAP HANA Support for virtualized / partitioned (multi-tenant) environments
[93] Also see https://www.vmware.com/in/products/vsphere/features/high-availability
[94] SLE HAE is virtualization agnostic
[95] Also see Setting up a RHEV-M on a Highly Available Cluster
[96] Also see How to Configure High Availability Feature in XenServer
[97] Also see Xen FAQ High Availability
[98] Also see Veritas/Symantec ApplicationHA Solution for SAP NetWeaver
[99] Also see Neverfail Solutions for Virtualization
[100] Also see LifeKeeper Single Server Protection for Linux and Documentation for Linux Recovery Kits for v9.2.1 and SIOS Protection Suite for Linux SAP Recovery Kit Administration Guide
[101] Also see High Availability Solutions for SAP on VMware
[102] If VM.GracefulRebootPolicy policy is disabled, and SAP HANA system or its component fails, then ApplicationHA stops sending the heartbeat to VMware HA. As a result VMware HA can then restart the virtual machine. If this policy is enabled, ApplicationHA itself invokes a native operating system command to restart the virtual machine. The default value of this policy is Disabled
[103] Also see Veritas/Symantec ApplicationHA Solution for SAP HANA
[104] Also see VMware KB: VMware Fault Tolerance FAQ

mode with the primary virtual machine, i.e. the secondary virtual machine is always an identical copy of the primary virtual machine. If the primary virtual machine becomes unavailable due to a hardware issue, the secondary machine takes over with minimal downtime[105]. Working with SAP CS, VMware FT has the following limitations that need to be considered:

- Limited sizing to 1 vCPU (up to vSphere 5.5) or 4 vCPU (from vSphere 6.0), 64GB memory[106]
- Maximum 4 FT virtual machines per host up to vSphere 6.5
- Up to 98 FT virtual machines or 256 FT virtual machines vCPU per vSphere cluster in vSphere 6.5

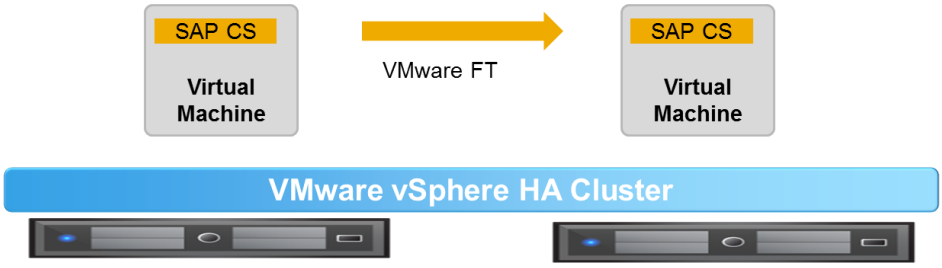

Figure 21 - VMware HA and VMware FT for SAP CS

3.5.1 Linux Clustering combined with Host Failover Clustering in Virtualized Environment

Besides the Host Failover Clustering from the virtualization platform, the native Linux cluster solutions discussed in section "Linux Cluster Solutions for SAP NetWeaver" can also be used similar as in a physical environment, in case an appropriate Quorum, STONITH or Fencing mechanism can be provided.

Normally power based fencing mechanisms cannot be used like in a physical environment, as in most cases there will be multiple virtual machines on the same physical box, and it is not desired to power them on and off at the same time due to an issue with one virtual machine. Both SUSE Linux Enterprise Server (SLES) and Red Hat Enterprise Linux (RHEL) support specific virtualization fencing mechanisms. Some examples are external/vcenter for VMware, external/libvirt for KVM, external/xen0 and external/xen0-ha for Xen on SLES cluster[107], and fence_vmware_soap for VMware, fence_virt or fence_xvm for KVM, fence_xvm for Xen on RHEL cluster[108]. The shared storage fencing mechanism can be used for a virtualization environment the same as in a physical environment, such as external/sbd[109] for SLES cluster, and fence_scsi for RHEL cluster[110].

---

[105] Also see SAP Solutions on VMware Best Practices Guide
[106] Also see The Case for SAP Central Services and VMware Fault Tolerance
[107] Also see SUSE Linux Enterprise High Availability Extension Documentation
[108] Also see Virtualization Support for RHEL High Availability and Resilient Storage Clusters, How to configure VMware fencing using fence_vmware_soap in Red Hat Enterprise Linux High Availability Add On, Fence Device and Agent Information for Red Hat Enterprise Linux
[109] Also see SUSE Linux Enterprise High Availability Extension Documentation - Storage-based Fencing and SBD Fencing
[110] Also see Using SCSI Persistent Reservation Fencing (fence_scsi) in RHEL 6, Using SCSI Persistent Reservation Fencing (fence_scsi) with pacemaker in a RHEL 6 or 7 High Availability cluster, When using fence_scsi with multipath devices and specifying a "devices" attribute in RHEL, reservation conflicts are logged

Clustered SAP ASCS/ERS instances that run in a Host Failover Cluster should not run with both cluster nodes on the same physical host if possible. The reason is if physical host running both cluster nodes fails, the Enqueue table gets completely lost. VMware HA can respect VM anti-affinity rules, which force specified virtual machines to remain apart during failover[111].

3.6 Storage

Fencing mechanisms protect data from being destroyed by unintended accesses. To keep data consistent and highly available in combination with clustering, two categories of techniques are developed – data redundancy[112] and access layer redundancy.

Data redundancy techniques:

1. RAID[113]: now commonly Redundant Array of Independent Disks is a storage virtualization technology that combines multiple physical disk drive components into a single logical unit for the purposes of data redundancy, performance improvement, or both. This works within one storage array.
2. SAN based replication/mirroring: enterprise storage arrays typically offer a mechanism to mirror all data from one storage array to one or more other arrays. However, the cluster uses only one array at one time. E.g., EMC SRDF, Hitachi HDS UR.
3. Host based mirroring: enterprise Linux offers the possibility to assemble independent storage devices (LUNs) on separate storage arrays to a soft-raid logical volume in order to prevent data loss and ensure availability in case of a failure on one of the physical arrays. The Logical Volume Manager (LVM) simultaneously reads and writes to two or more LUNs on separate storage arrays and keeps them synchronized. E.g., Multiple Device (MD) and LVM
4. File system based replication: Some file system solutions allows synchronizing data in in more than one location. This allows applications to run in Active-Active mode cross nearby data centers as if database is on one copy of the file systems. E.g., IBM GPFS-FPO shared nothing cluster.
5. Storage federation: some advanced storage virtualization can also keep the data in synchronization in more than one location. This can allow application to run in Active-Active mode cross nearby data centers. Comparing with file system based replication; this can reduce the OS level resource consumption. E.g., EMC VPLEX.
6. Application level replication: some applications, such as databases, can continuously transmit delta changes to the target in more than one location. The ongoing application level verification can detect corruption in early stage. E.g., SAP HANA synchronous system replication.

Access layer redundancy techniques:

1. Redundant network: as described in section 3.3, there are basically two types of networks Ethernet and Fibre Channel based on SAN. Redundancy and fault resilience should be designed: redundant

[111] Also see vSphere HA and DRS Affinity Rules
[112] Also see https://en.wikipedia.org/wiki/Data_redundancy
[113] Also see https://en.wikipedia.org/wiki/RAID

network interface cards (NICs) and Fibre Channel interfaces in their servers, and redundant pairs of switches at each layer in the network architecture. For SAN, Fibre Channel over Ethernet (FCoE) and iSCSI protocols are finally over Ethernet network.

2. Multipath I/O (MPIO)[114]: MPIO is a fault-tolerance and performance-enhancement technique by using more than one physical path between a computing server and its devices through the buses, controllers, switches, and bridge devices connecting them. If one controller, port or switch fails, the operating system can route the I/O through the remaining controller, port or switch transparently and with no changes visible to the applications, although perhaps resulting in increased latency. MPIO can also leverage the redundant paths to provide performance-enhancing features, including dynamic load balancing, traffic shaping, automatic path management, and dynamic reconfiguration. E.g., EMC PowerPath, Veritas Dynamic Multipathing (VxDMP), Linux kernel based Device-Mapper Multipath (DM-MPIO).
3. Cluster file systems: file system shared by being simultaneously mounted on multiple servers accessing the same storage. It can provide features like location-independent addressing and redundancy, which improve reliability and availability. The shared-disk cluster file systems use concurrency control mechanisms to provide a consistent and serializable view of the file systems, avoiding corruption and unintended data loss even when multiple clients try to access the same files at the same time. Some sort of a fencing mechanism are employed to prevent data corruption in case of node failures, because an unfenced device can cause data corruption if it loses communication with other cluster nodes, and tries to access the same information other nodes are accessing. E.g., shared-disk IBM GPFS file system, Oracle Cluster File System (OCFS2) under SLES, Global File System (GFS2) under RHEL.
4. File systems protected by failover cluster solutions: unlike cluster file systems which can provide concurrent file system accesses to multiple cluster nodes, these file systems are locally mounted in a cluster environment. It is only accessible to the cluster node where it is mounted. In order to access this file system after a failover or switchover, this file system has to be switched to the backup node. Fencing mechanisms prevent data corruption in case of node failures. This is normally file system protected by cluster software, such as EXTx, XFS, NFS file systems protected by Linux cluster software.

As discussed, SAP Global File Systems are classified as SPOFs. There are several options to make them high available. The below table lists options, and Pros and Cons.

HA Protection Options Comparison for SAP Global File Systems

[114] Also see https://en.wikipedia.org/wiki/Multipath_I/O

	Option	Preconditions	Pro		Con	
1	Cluster file systems - OCFS2, GFS2	Shared external storage – SAN	+	No additional nodes needed	–	Additional file system cluster introduce complexity
					–	All AS instance hosts need to be included in the cluster
2	Internal NFS server[115] with HA cluster (such as SLES HAE, RHEL HA)	Shared external storage – SAN	+	No additional NFS server nodes needed	–	Mounting NFS shares over loopback is a configuration, which is known to cause deadlocks under heavy load[116]
3	External NFS with HA cluster (such as SLES HAE, RHEL HA)	Shared external storage – SAN, and additional NFS nodes	+	Avoid loopback NFS issues under Linux	–	Need additional NFS nodes
4	External NFS and DRBD[117] with HA cluster(such as SLES HAE, RHEL HA)	Additional NFS nodes	+	Avoid loopback NFS issues under Linux	–	Need additional NFS nodes
			+	the storage is duplicated for corruption protection		
5	External NAS server	External NAS server	+	External NAS provide HA functionality and easy configuration	–	Need additional external NAS server

[115] The internal NFS server means the NFS server is deployed in the same cluster as the resources accessing the NFS server, otherwise, it's called external NFS server. Internal NFS server needs mount NFS shares over loopback.
[116] See also Loopback NFS: theory and practice and Hangs occur if a Red Hat Enterprise Linux system is used as both NFS server and NFS client for the same mount. The Operation System vendors may solve the issue in the subsequent versions or patches, such as SUSE now supports the internal NFS scenario in SLES 12
[117] different vendors have different packaging license terms

4 High Availability Design and Implementation

4.1 Design and Implementation for SAP NetWeaver with High Availably

4.1.1 Technical Architecture for SAP NetWeaver with Pacemaker Cluster Resource Manager

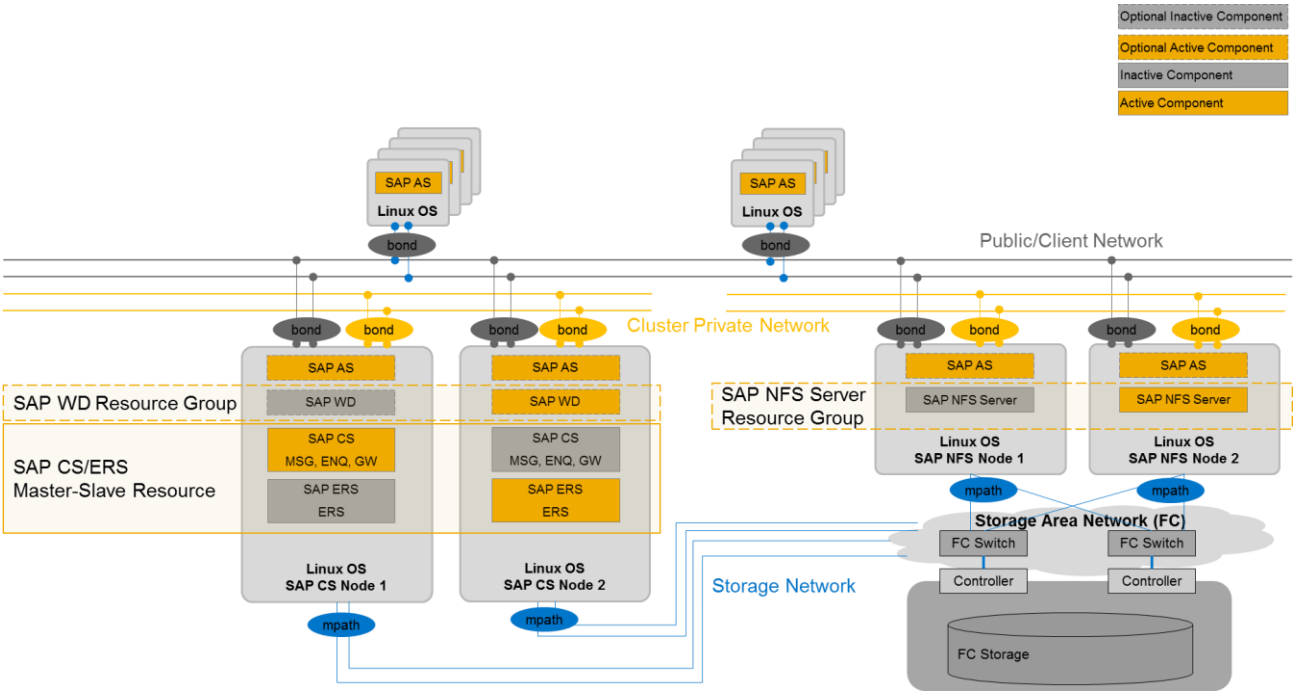

Figure 22 - Reference Architecture for SAP NetWeaver with Pacemaker Cluster Resource Manager

For simplification purposes, in the above reference architecture, we protect all discussed SPOFs using cluster software only at minimum level, while SAP AS (PAS and AAS) instance availability is ensured by multiple SAP AS installations.

As mentioned, in a Linux environment, generally, we do not recommend using the integrated or internal NFS server[118] for the SAP Global File Systems, so two clusters are designed: one for the SAP CS, and the other for the SAP NFS service. In the SAP CS cluster, the main components are SAP CS and SAP ERS instances. SAP CS and SAP ERS should never run on the same node, and for this SAP certified Pacemaker based cluster solutions (SLES 11, RHEL 6) provide a Master-Slave resource. If the original master (where the SAP CS instance running) fails, the cluster software will promote the slave (where the SAP ERS instance is running) to the new master, and make the original master to the new slave - when possible. There are the following general recommendations:

(1) For SAP dual-stack systems that have both SAP ABAP and SAP Java CS instances, set up separate Master-Slave resources for each SAP CS instance;

[118] Some operating system vendor may have different support statements for different operating system versions, so check with your vendor for supportability on integrated NFS server, such as in SLES 12, this is supported

(2) For SAP ABAP CS (ASCS) instances, customize the instance profile to configure a standalone Gateway in an HA protected ASCS instance, the example profile can be found in appendix;

(3) When SAP Web Dispatcher (WD) is used for HTTP(S) request load balancing, protect it as a separate resource group (one special SAPInstance resource). You can deploy multiple NFS resource groups on the NFS cluster to make it a central NFS service provider, in case there is enough capacity. If you want to use an external HA NAS for the SAP Global File Systems, then the NFS server cluster is not needed.

All SAP NetWeaver components can be deployed on physical servers and in virtualized environments. When deployed in virtualized environment, make sure the cluster nodes of individual cluster run on different physical servers. When physical servers are used, an SAP CS instances may not need the full available server capacity, so additional SAP AS instances can be installed within the cluster nodes (without third-party cluster protection) to make full use of the resources.

Three kinds of network may be involved:

- Storage network (FC is most popular) to connect to external storage subsystem: Cluster protected NFS and SAP WD will need shared storage available from external storage subsystem. If storage based fencing is used, then this also needs a storage network. To protect the storage network, multi-path I/O and redundant SAN switches will be used.
- Public network for client-server and server-server communications: On network infrastructure level, redundant switches are needed, while on server level network bonding is recommended for HA.
- Cluster private network: cluster heart-beating communication. This should be separate network segments from public network, and network bonding or teaming is recommended for HA.

4.1.2 Implementation Tips for SAP NetWeaver with Pacemaker Cluster Resource Manager

Detailed Implementation Guides from SUSE and Red Hat

Several detailed implementation guides for configuration HA for SAP on SUSE Linux Enterprise Server can be found at https://www.suse.com/products/sles-for-sap/resource-library/sap-best-practices.html [119], especially the guide "Best practices for running SAP NetWeaver on SUSE Linux Enterprise Server with High Availability—Enqueue Replication Server and sap_suse_cluster_connector Integration" and "Best Practice for SAP Enqueue Replication - SAP NetWeaver High Availability on SUSE Linux Enterprise 12" can be directly applied to the recommended approach.

Red Hat also provides one detailed guide "Deploying Highly Available SAP NetWeaver-based Servers Using Red Hat Enterprise Linux HA add-on with Pacemaker"[120] and the updated version.

[119] Also see Addendum for SUSE Whitepapers and Best Practices - SAP NetWeaver HA
[120] Red Hat also provided another method "Deploying Highly Available SAP Servers using Red Hat Clustering" with rgmanager as cluster resource manager, but that is not certified scenarios discussed in this document.

File System Layout and Virtual Hostnames

Single Point of Failure	SAP Global File System	SAP Central Services	SAP Web Dispatcher
Resource / Resource Group	rg_sap_<SID>_nfs	msl_sap_<SID>_cs	rg_sap_<SID>_wd
Volume Group	vg_<SID>_nfs	vg_<SID>_cs, vg_<SID>_ers	vg_<SID>_wd
Logical Volume	lv_<SID>_sapmnt, lv_<WD_SID>_sapmnt, lv_<SID>_saptrans	lv_<SID>_cs, lv_<SID>_ers	lv_<SID>_wd
File System Mount Points[121]	/export/sapmnt/<SID>[122], /export/sapmnt/WD[123] /export/saptrans/<SID>[124]	/usr/sap/<SID>/ASCS<Inst_Nr>, /usr/sap/<SID>/ERS<Inst_Nr>[125]	/usr/sap/<WD_SID>
File System Local or Shared[126]	Shared	Local	Shared
Virtual Hostname	<sid>vnfs	<sid>vascs, <sid>vaers[127]	<sid>vwd

The above table lists the planning information for resource groups and file systems, and a recommended naming convention if you have not yet have one.

SAP NetWeaver System Installation

After the SAP HANA instance is installed, the SAP NetWeaver system installation can be started following the SAP NetWeaver installation Guide using SAP Software Provisioning Manager (SWPM).

The overall complete general steps are as follows:

- Installing and configuring operating systems on all nodes according to the SAP guidelines for specific operating system version[128];

[121] For the file system sizing requirement, check with the SAP NetWeaver installation guides available at http://service.sap.com/instguides
[122] This will be mounted at /sapmnt/<SID> at all hosts with SAP instances
[123] This will be mounted at /sapmnt/<WD_SID> at all SAP Web Dispatcher cluster nodes
[124] This will be mounted at /usr/sap/trans at all hosts with SAP AS instances
[125] This is SAP NetWeaver ABAP stack format, for SAP NetWeaver Java stack, the correspond file systems will be /usr/sap/<SID>/SCS<Inst_Nr>, /usr/sap/<SID>/ERS<Inst_Nr>
[126] Local file system can be allocated from local disks or from external storage, but do not share with other cluster nodes; shared file systems must be allocated from external storage, and could be failed and switched over between cluster nodes
[127] For SAP NetWeaver Java stack, <sid>vscs and <sid>vjers can be used instead
[128] See SAP Note 171356 - SAP Software on Linux: General information, SAP Note 1631106 - Red Hat Enterprise Linux for SAP Business Applications, SAP Note 1565179 - SAP software and Oracle Linux, SAP Note 1310037 - SUSE LINUX Enterprise Server 11: Installation notes, SAP Note 1984787 - SUSE LINUX Enterprise Server 12: Installation notes, SAP Note 1944799 - SAP HANA Guidelines for SLES Operating

- Install SAP HANA instance[129]. Then you immediately configure the SAP HANA System Replication Failover Automation using cluster software. You can use the virtual hostname of the SAP HANA instance in the following installation; otherwise you can use first the physical hostname to finish the installation, and then use the tool ***hdbuserstore*** to reconfigure the SAP AS instance connectivity to SAP HANA instance later on;
- Configure NFS resource group (or other options) for SAP Global File Systems, and mount them on all hosts of SAP NetWeaver instances;
- Manually start virtual IPs on the hosts where to install the SAP CS, SAP ERS and SAP WD, such as using ***ifconfig*** or ***ip*** command; or bind virtual hostnames using SAP Host Agent (saphostctrl) function;
- Install SAP CS instances on the first node of the cluster using SWPM with parameter SAPINST_USE_HOSTNAME=<SAP CS virtual hostname>, such as /<path to sapinst>/sapinst SAPINST_USE_HOSTNAME=<sid>vascs;
- Install SAP ERS instances on the first node of the cluster using SWPM with parameter SAPINST_USE_HOSTNAME=<SAP ERS virtual hostname>, such as /<path to sapinst>/sapinst SAPINST_USE_HOSTNAME=<sid>vaers;
- Manually move virtual IPs to the second node of the cluster. Use ***tar*** command to copy the directories /home and /usr/sap, and files /etc/hosts, /etc/group, /etc/password, /etc/services, /etc/rc.d/sapinit. In the home directory, adapt the filenames containing hostnames[130]. Manually install SAP Host Agent on the second node also. And then try to start and stop the SAP CS and SAP ERS instances on the second node to see whether there is any error;
- Install SAP PAS and AAS instances using SWPM without parameter SAPINST_USE_HOSTNAME;
- When needed, install SAP WD using SWPM with parameter SAPINST_USE_HOSTNAME=<SAP Web Dispatcher virtual hostname>;
- Stop the SAP NetWeaver system and the manually started virtual IPs;
- Configure HA for SAP CS and SAP WD (if there is any) by following the relevant implementation guides;
- Unlike SAP HANA where a license is only required to be installed once on the primary SAP HANA instance, you need install official SAP licenses twice, that is, for both primary and standby SAP CS instances.

SAPInstance Configuration for Gateway and SAP Web Dispatcher

System Installation, SAP Note 2009879 - SAP HANA Guidelines for RedHat Enterprise Linux (RHEL) Operating System, SAP Note 1496410 - Red Hat Enterprise Linux 6.x: Installation and Upgrade SAP Note 2002167 - Red Hat Enterprise Linux 7.x: Installation and Upgrade, and SAP Note 2001528 - Linux: SAP HANA Database SPS 08, SPS 09 and SPS 10 on RHEL 6 or SLES 11

[129] Installation and configuration guides for SAP HANA are available from SAP Help

[130] Red Hat guide recommends to remove all user login profiles (that include a hostname in the file name) for RHEL cluster

Resource agent SAPInstance has parameter MONITOR_SERVICES to customize the services monitored by the cluster[131]. The default is: default="disp+work|msg_server|enserver|enrepserver|jcontrol|jstart"

Use "sapcontrol -nr XX -function GetProcessList" to check, which SAP service names are valid for your environment. For example, in below environment, we can use gwrd for standalone Gateway and sapwebdisp for SAP Web Dispatcher monitoring

dsthost1:dstadm 55> **sapcontrol -nr 00 -function GetProcessList**

08.09.2015 15:37:39

GetProcessList

OK

name, description, dispstatus, textstatus, starttime, elapsedtime, pid

msg_server, MessageServer, GREEN, Running, 2015 09 08 15:35:44, 0:01:55, 8148

enserver, EnqueueServer, GREEN, Running, 2015 09 08 15:35:44, 0:01:55, 8149

gwrd, Gateway, GREEN, Running, 2015 09 08 15:35:44, 0:01:55, 8150

gw1app:gw1adm 49> **sapcontrol -nr 55 -function GetProcessList**

08.09.2015 16:41:00

GetProcessList

OK

name, description, dispstatus, textstatus, starttime, elapsedtime, pid

sapwebdisp, Web Dispatcher, GREEN, Running, 2015 08 13 12:29:53, 628:11:07, 16451

SAP CS Instance Profile Adjustment

To leverage the replication table in the SAP ERS, in the Master-Slave scenario, the Enqueue Server must not be restarted locally. We need to change the start directive for the enqueue server in the SAP CS instance profiles, that is, use of Start_Program_nn for the enqueue server instead of Restart_Program_nn to avoid the local restart of the Enqueue Server after the process has failed, for example:

```
#--------------------------------------------------------------------
# Start SAP enqueue server
#--------------------------------------------------------------------
_EN = en.sap$(SAPSYSTEMNAME)_$(INSTANCE_NAME)
Execute_04 = local rm -f $(_EN)
```

[131] Also see http://linux-ha.org/doc/man-pages/re-ra-SAPInstance.html

```
Execute_05 = local ln -s -f $(DIR_EXECUTABLE)/enserver$(FT_EXE) $(_EN)
#Restart_Program_01 = local $(_EN) pf=$(_PF) # <---- original line
Start_Program_01 = local $(_EN) pf=$(_PF) # <---- modified line
```

SUSE also recommends adding one line

```
Max_Program_Restart = 03
```

to define that sapstart only restarts processes three times for takeover of the ASCS instance if the message server is failing multiple times.

Integrate SAP HA Library and Vendor Cluster Connector

To integrate SAP HA Library (saphascriptco.so loaded by sapstartsrv) and vendor cluster connector, the following steps are needed[132]:

- allow the <sid>adm user to communicate with the cluster
  groupmod -A <sid>adm haclient
- the SAP command sapstartsrv must be instructed to load and use the saphascriptco library. Add the following to the instance profile of each SAP instance that should be managed by cluster
  `service/halib = $(DIR_CT_RUN)/saphascriptco.so`
  `service/halib_cluster_connector = /usr/bin/sap_suse_cluster_connector` (for SLES)
  `service/halib_cluster_connector = /usr/sbin/sap_redhat_cluster_connector` (for RHEL)

For Red Hat, it is also required to switch on the recording of pending commands in the cluster, that is set record-pending=true for defaults resource operations. And also set below instance parameter

```
service/halib_tmp_prefix = <prefix (e.g. /tmp)>
```

Other Pacemaker Cluster Considerations

- Cluster resource constraints, relationship or dependency: There are rich methods – scores, location rules, collocation rules, order constraints, and implicit dependencies – to control where the resource should be placed, and when and how the resource status should be changed. In the proposed simplified reference architecture, Master-Slave, location, and collocation relationships are enough, except that you want to define other resource groups and relationship in between.
- STONITH: fencing mechanism must be enabled (stonith-enabled="true")[133]. You need adjust the STONITH timeout values as per your hardware and SAN environment during volume/stress tests.

[132] For details, check the relevant implementation guide.
[133] For two-node clusters, you may need special setting for quorum management, check with relevant vendor for latest supported configuration when you design your cluster as it may subject to changes. Currently SUSE recommends no-quorum-policy="ignore", while Red Hat proposes default pacemaker setting. This also applies to pacemaker based clusters for SAP HANA.

The setting "-INFINITY" should be used instead of INFINITY, if location constraints should be used to determine where the STONITH device should run[134].

4.1.3 Technical Architecture for SAP NetWeaver with other Cluster Resource Managers

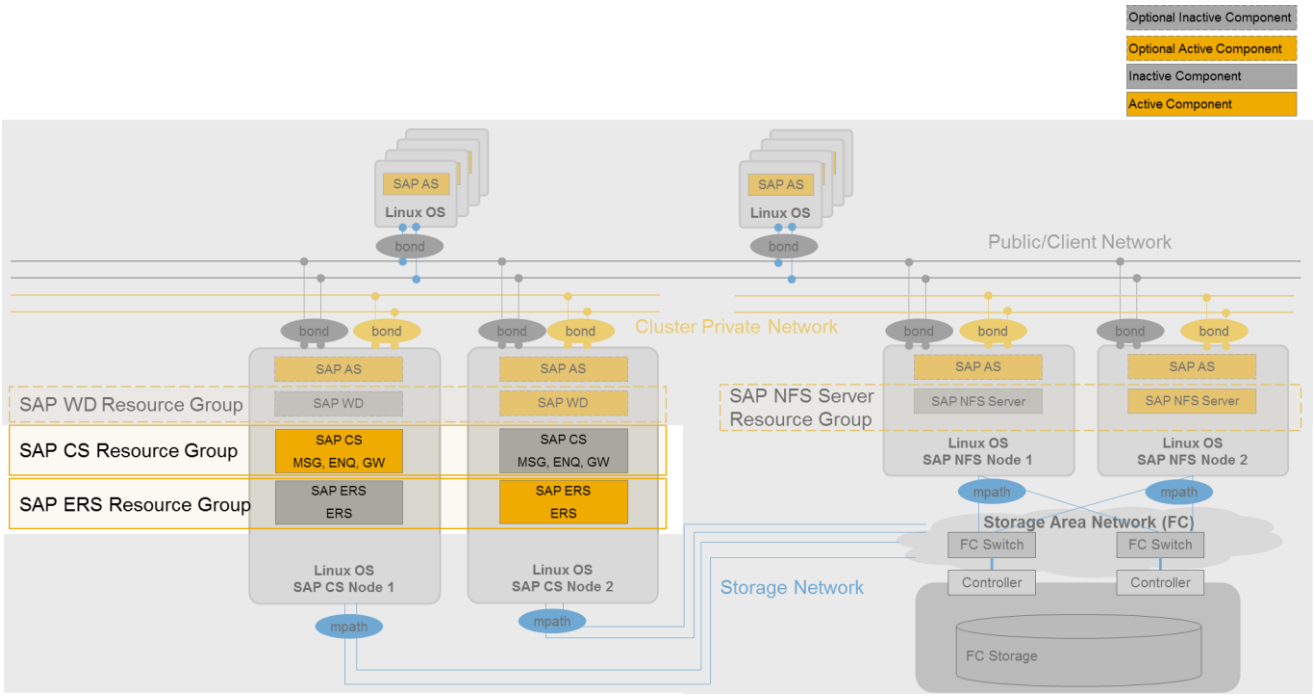

Figure 23 - Reference Architecture for SAP NetWeaver with non-Pacemaker Cluster Resource Manager[135]

In the above reference architecture, we again only protect the all discussed SPOFs using cluster software at minimum level, while SAP AS (PAS and AAS) instance availability is ensured by multiple SAP AS installations.

The main difference here is how to handle HA for SAP Central Services. Instead of one Master-Slave resource for each SAP CS and SAP ERS combination, SAP ERS instance is separated into its own resource group (service group/package). Then these cluster solutions use parent-child or follow-and-push dependency, or other similar mechanisms to control the relationship between SAP CS and SAP ERS instances. The principle is the same: the SAP ERS instance runs on the other node than the SAP CS instance, and holds a complete copy of the Enqueue lock table; when the SAP Enqueue server fails, the cluster moves the SAP CS instance (including SAP Enqueue server) to the node on which the corresponding SAP ERS process is currently running; the lock table (replication table) stored on the SAP ERS is transferred to the enqueue server process being recovered and the new lock table is created from it; after that, the active replication server is then deactivated by the cluster; the cluster will then also try to restart the SAP ERS on the node which ran former corresponding SAP CS instance if possible.

[134] See also http://clusterlabs.org/doc/crm_fencing.html. There is probability that the STONITH operation is going to fail. Hence, a STONITH operation on the node, which is the executioner, too is not reliable, and then the cluster is in trouble. By convention, STONITH refuses to kill its host.
[135] The grayed out part of the picture does not change compared to Figure 22

4.1.4 Implementation Tips for SAP NetWeaver with other Cluster Resource Managers

Detailed Implementation Guides from Different Cluster Vendors

HP has a detailed guide "Managing HPE Serviceguard Extension for SAP for Linux" describing how to plan, configure, and administer highly available SAP systems on Red Hat Enterprise Linux and SUSE Linux Enterprise Server systems using HPE Serviceguard.

NEC provides the guides "EXPRESSCLUSTER X 3.3 for Linux SAP NetWeaver System Configuration Guide" and "EXPRESSCLUSTER X 3.3 for Linux SAP NetWeaver Configuration Example"[136].

Oracle has the guide "Providing High Availability for SAP Resources with Oracle Clusterware 11g Release 2 and Clusterware 12c Release 1", which uses the Oracle Clusterware tool - SAP Control (SAPCTL) to manage SAP HA resources.

For SIOS Protection Suite for SAP, two guides "SIOS Protection Suite for Linux SAP Recovery Kit" and "SIOS Protection Suite for Linux - SAP Solution" are available; SIOS LifeKeeper Single Server Protection (SSP) for Linux technical documentation is for integrating with Virtual Environments (VMware, Citrix, KVM, Oracle VM, Amazon EC2, and so on)[137] to provide application monitoring.

Veritas provides the guide "Cluster Server Agent for SAP NetWeaver Installation and Configuration Guide 7.0" for failover cluster and the guide "Veritas/Symantec ApplicationHA Solution for SAP NetWeaver" specific for virtualized environments.

File System Layout and Virtual Hostnames

Single Point of Failure	SAP Global File System	SAP Central Services	SAP Enqueue Replication Server	SAP Web Dispatcher
Resource Group[138]	rg_sap_<SID>_nfs	rg_sap_<SID>_cs	rg_sap_<SID>_ers	rg_sap_<SID>_wd
Volume Group	vg_<SID>_nfs	vg_<SID>_cs	vg_<SID>_ers	vg_<SID>_wd
Logical Volume	lv_<SID>_sapmnt, lv_<WD_SID>_sapmnt, lv_<SID>_saptrans	lv_<SID>_cs	lv_<SID>_ers	lv_<SID>_wd

[136] You may find the updated guides at http://www.nec.com/en/global/prod/expresscluster/en/support/Setup.html
[137] Also see SIOS LifeKeeper Single Server Protection Release Notes at http://docs.us.sios.com/Linux/9.2.1/LifeKeeperSSP/TechDoc/Content/LifeKeeper_Single_Server_Protection_Release_Notes.htm
[138] This may also be called Service Group or Package by cluster software vendors

Single Point of Failure	SAP Global File System	SAP Central Services	SAP Enqueue Replication Server	SAP Web Dispatcher
File System Mount Points[139]	/export/sapmnt/<SID>, /export/sapmnt/WD /export/saptrans/<SID>	/usr/sap/<SID>/ASCS<Nr>[140]	/usr/sap/<SID>/ERS<Nr>	/usr/sap/<WD_SID>
File System Local or Shared[141]	Shared	Shared	Shared/Local	Shared
Virtual Hostname	<sid>vnfs	<sid>vascs	<sid>vaers[142]	<sid>vwd

The above table lists the planning information for resource groups and file systems, and a recommended naming convention if you have not yet have one.

SAP NetWeaver System Installation

After the SAP HANA instance is installed, the SAP NetWeaver system installation can be started by following the SAP NetWeaver installation Guide using SAP Software Provisioning Manager (SWPM).

The overall general steps are quite similar as with the Pacemaker cluster. The slight difference is with synchronization of the second cluster node: If the SAP CS and SAP ERS instance directories are shared, then during the file synchronization, you manually need to deal with the volume group, lvm and file systems to make them accessible to the second node - instead of copy those instance directories.

SAP CS Instance Profile Adjustment

To leverage the replication table in the SAP ERS, the Enqueue Server must not be restarted locally. We need to change the start directive for the enqueue server in the SAP CS instance profiles, that is, use of Start_Program_nn for the enqueue server instead of Restart_Program_nn to avoid the local restart of the Enqueue Server after the process has failed, for example:

```
#-----------------------------------------------------------------------
# Start SAP enqueue server
#-----------------------------------------------------------------------
_EN = en.sap$(SAPSYSTEMNAME)_$(INSTANCE_NAME)
```

[139] For the file system sizing requirement, check with the SAP NetWeaver installation guides available at http://service.sap.com/instguides
[140] This is SAP NetWeaver ABAP stack format, for SAP NetWeaver Java stack, the correspond file systems will be /usr/sap/<SID>/SCS<Nr>
[141] Local file system can be allocated from local disks or from external storage, but do not share with other cluster nodes; shared file systems must be allocated from external storage (or via external NFS service), and could be failed and switched over between cluster nodes. Refer to relevant implementation guides, such as SIOS LifeKeeper requires ERS instance directory mounted locally, while others use shared file system. Veritas VCS allow all instance directories to be either shared or local.
[142] For SAP NetWeaver Java stack, <sid>vscs and <sid>vjers can be used instead

```
Execute_04 = local rm -f $(_EN)
Execute_05 = local ln -s -f $(DIR_EXECUTABLE)/enserver$(FT_EXE) $(_EN)
#Restart_Program_01 = local $(_EN) pf=$(_PF) # <---- original line
Start_Program_01 = local $(_EN) pf=$(_PF) # <---- modified line
```

Some cluster solutions, such as Veritas InfoScale Availability, may also recommend to disable SAP Enqueue Replication Server (ERS) restart in its instance profile by changing Restart_Program_nn to Start_Program_nn.

Integrate SAP HA Library and Vendor Cluster Connector

Different cluster solutions vary slightly to integrate the SAP HA Library and vendor cluster connector, so refer to relevant implementation guide. For example, Veritas InfoScale Availability uses the below parameter settings in each instance profile that should be managed by cluster:

```
service/halib = $(DIR_CT_RUN)/saphascriptco.so
service/halib_cluster_connector = /opt/VRTSagents/ha/bin/SAPNW/sap_symc_cluster_connector
```

HPE Serviceguard Extension for SAP (SGeSAP) uses the setting below:

`service/halib = /opt/cmcluster/lib/saphpsghalib.so` (for SLES)

`service/halib = /usr/local/cmcluster/lib/saphpsghalib.so` (for RHEL)

Cluster Resource Constraints, Relationship or Dependency

Again different cluster solutions provide different methods to control where the resource should be placed, and when and how the resource status should be changed, such as HPE Serviceguard Extension for SAP (SGeSAP) uses a follow-and-push mechanism and same-node/up dependency, while SIOS LifeKeeper uses a parent-child dependency and resource hierarchy. So refer to relevant implement guides for detailed configuration options.

4.2 Design and Implementation for SAP HANA with System Replication Automation

4.2.1 Technical Architecture for SAP HANA System Replication Automation with Pacemaker Cluster Resource Manager

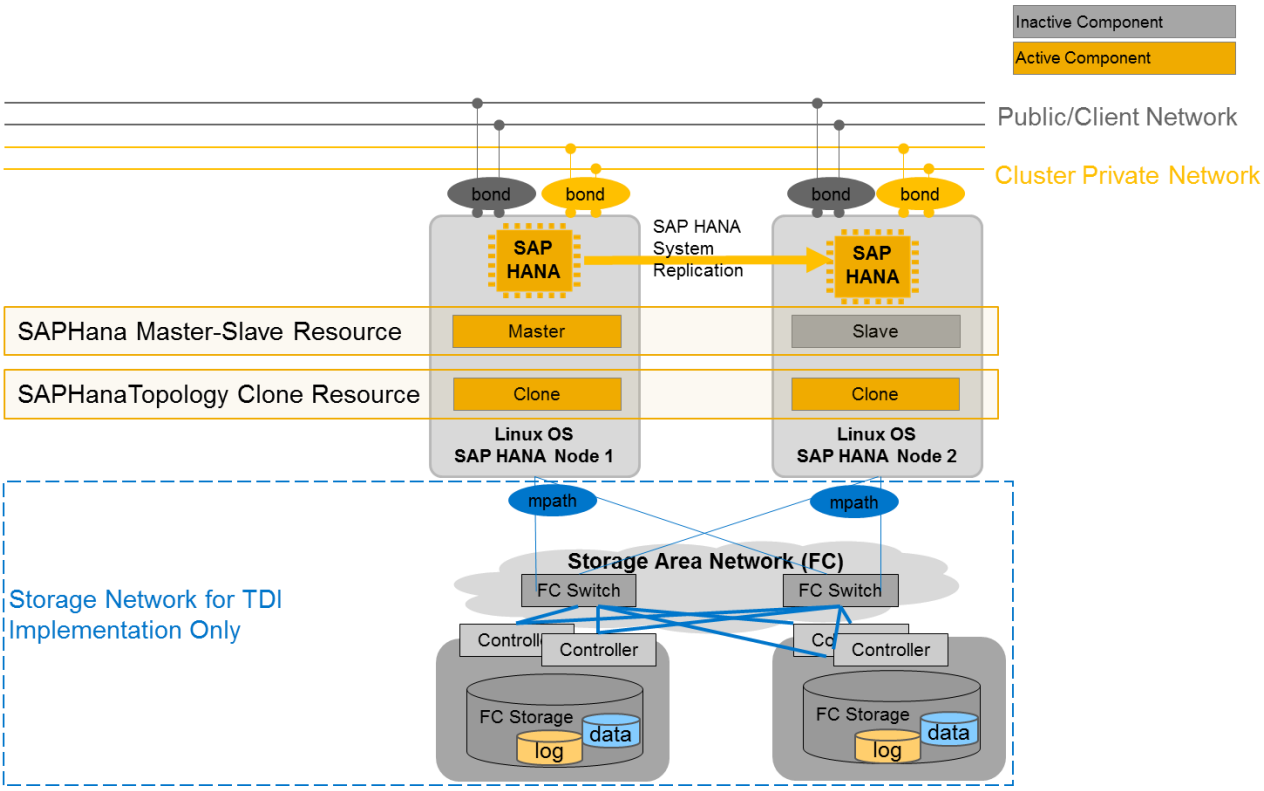

Figure 24 - Reference Architecture for SAP HANA System Replication Automation with Pacemaker Cluster Resource Manager

In the above reference architecture, there are two special resources:

- **SAPHana Master-Slave resource**: Main tasks
  - Manages the two SAP HANA instances (primary and secondary) controlling the start/stop of the SAP HANA instances
  - Monitors the SAP HANA databases to check their availability on landscape host configuration level
  - Checks the synchronization status of the two SAP HANA instances, so if the synchronization is not OK, then the cluster does not failover to the secondary side if the primary fails, which is to improve the data consistency.

  In general, the primary SAP HANA instance is promoted to be Master, while the secondary SAP HANA instance is Slave. One virtual IP address/hostname of SAP HANA is used to provide external accesses, and this virtual IP address always runs on the same node as the primary SAP HANA instance. For scale-up scenarios, one additional relocatable virtual IP address for external read-only accesses could be created to stay with secondary SAP HANA instance for SAP HANA 2.0

when Active/Active feature enabled. The status of virtual IP addresses will be controlled by Master-Slave resource.

- **SAPHanaTopology Clone resource**: The main tasks are to analyze the SAP HANA topology and send all findings via the node status attributes to all nodes in the cluster. These attributes are taken by the SAPHana resource to control the SAP HANA instances. Additionally, it starts and monitors the local SAP Host Agent. The SAPHanaTopology resource always starts before the SAPHana resource.

For a SAP HANA appliance-based solution, the appliance internal storage will be used, so two kinds of network will be involved:

- Public network for client-server communications. On network infrastructure level, redundant switches are needed, while one server level network bonding is recommended for HA. This is corresponding to Client / Frontend Zone of SAP HANA logical network zones.
- Cluster private network for cluster heart-beat communication. This should be separate network segments from public network, and network bonding or teaming is recommended for HA. This should be in the Internal Zone.

If Tailored Datacenter Integration (TDI)[143] for storage concept is employed, then separate Enterprise Storage will be used, requiring

- Storage network (FC is dominant) to connect to external storage subsystem. To protect storage network, redundant SAN switches and HBA card teaming is recommended. If the two SAP HANA systems span over two nearby data centers, then each data center need redundant SAN switches.

4.2.2 Implementation Tips for SAP HANA System Replication Automation with Pacemaker Cluster Resource Manager

Detailed Implementation Guides from SUSE and Red Hat

SUSE has detailed implementation guides available at https://www.suse.com/products/sles-for-sap/resource-library/sap-best-practices/ for different versions and scenarios; and Red Hat also provides one detailed guide "Automated SAP HANA System Replication with Pacemaker on RHEL Setup Guide" for RHEL 6.5 and later, and the updated version.

SAP HANA System Installation

SAP HANA installation should follow the SAP HANA installation guide. The whole process includes the following steps:

[143] Also see SAP HANA Tailored Datacenter Integration (TDI)

- Install and configure relevant Operating Systems for SAP HANA on all nodes according to the SAP guidelines[144]
- Install SAP HANA on all nodes using the same SID and instance number;
- Manually configure SAP HANA System Replication and verify that it works according to SAP guide "How to Perform System Replication for SAP HANA";
- Install and configure relevant clusters as per the detailed implementation guides;
- Perform cluster tests to verify that failover works under all circumstances.

Other SAP HANA Pacemaker Cluster Considerations

- Different vendors have slightly different terms regarding supported scenarios and prerequisites, so check the relevant implementation guides carefully;
- A working Fencing/STONITH mechanism must be available for all cluster nodes[145];
- The virtual IP address for external access and the primary SAP HANA instance always run on the same node. As of SAP HANA 2.0, the additional virtual IP address for read-only access are bounded to the secondary SAP HANA instance[146]. The SAPHanaTopology resource always starts before the SAPHANA resource, so a colocation and an order constraint are required;

---

[144] See SAP Note 1310037 - SUSE LINUX Enterprise Server 11: Installation notes, SAP Note 1944799 - SAP HANA Guidelines for SLES Operating System Installation, SAP Note 1954788 - SAP HANA DB: Recommended OS settings for SLES 11 / SLES for SAP Applications 11 SP3, SAP Note 2240716 - SAP HANA DB: Recommended OS settings for SLES 11 / SLES for SAP Applications 11 SP4, SAP Note 2205917 - SAP HANA DB: Recommended OS settings for SLES 12 / SLES for SAP Applications 12, SAP Note 1855805 - Recommended SLES 11 packages for HANA support on OS level, SAP Note 2009879 - SAP HANA Guidelines for RedHat Enterprise Linux (RHEL) Operating System, SAP Note 2013638 - SAP HANA DB: Recommended OS settings for RHEL 6.5, SAP Note 2136965 - SAP HANA DB: Recommended OS settings for RHEL 6.6, SAP Note 2247020 - SAP HANA DB: Recommended OS settings for RHEL 6.7, SAP Note 2292690 - SAP HANA DB: Recommended OS settings for RHEL 7, SAP Note 2001528 - Linux: SAP HANA Database SPS 08, SPS 09 and SPS 10 on RHEL 6 or SLES 11, and SAP Note 2228351 - Linux: SAP HANA Database SPS 11 revision 110 (or higher) on RHEL 6 or SLES 11, and SAP Note 2195019 - Linux: Using SAP kernel 748 and higher or SAPinst 749 and higher on RHEL 6, OL6 and SLES 11, and SAP Note 2338763 - Linux: Running SAP applications compiled with GCC 5.x

[145] Pay special attention to stonith-enabled and no-quorum-policy settings, refer to section 4.1.2.

[146] The second virtual IP address resource is managed with a colocation constraint with Slave status. For example, for SLES:

```
# IP address for r/o access
primitive rsc_ip_PRD_ro ocf:heartbeat:IPaddr2 params ip="xxx.xxx.xxx.xxx" \
 op monitor interval="60" timeout="20" on_fail="restart"
# colocation IP address for r/o with HANA secondary
colocation col_ip_ro_on_sr-secondary 2000: rsc_ip_PRD_ro:Started msl_SAPHana_PRD_HDB00:Slave
```

for RHEL:

```
# Create IP address for r/o access
pcs resource create rsc_ip_PRD_ro IPaddr2 ip="xxx.xxx.xxx.xxx"
# colocation IP address for r/o with HANA secondary
pcs constraint colocation add rsc_ip_PRD_ro with slave msl_rsc_SAPHana_PRD_HDB00 2000
```

- The timeouts parameters of operations of SAPHana and SAPHanaTopology resources should be adjusted based on volume tests, as large SAP HANA databases can take longer to start up, therefore the start timeout might have to be increased;
- The time on all nodes must be in sync, which can be achieved by using a time synchronization mechanism like NTP;
- SAP HANA System Replication will only work after an initial backup has been performed, and the SAP HANA "log_mode" variable must be set to "normal";
- Ensure that SAP HANA is not configured to automatically start after system boot, since startup and shutdown of SAP HANA will be controlled by the cluster;
- A technical user with "CATALOG READ" and "MONITOR ADMIN" privileges must exist in SAP HANA for the resource agents to be able to run queries on the System Replication status.

4.2.3 Technical Architecture for SAP HANA System Replication Automation with Veritas InfoScale Availability

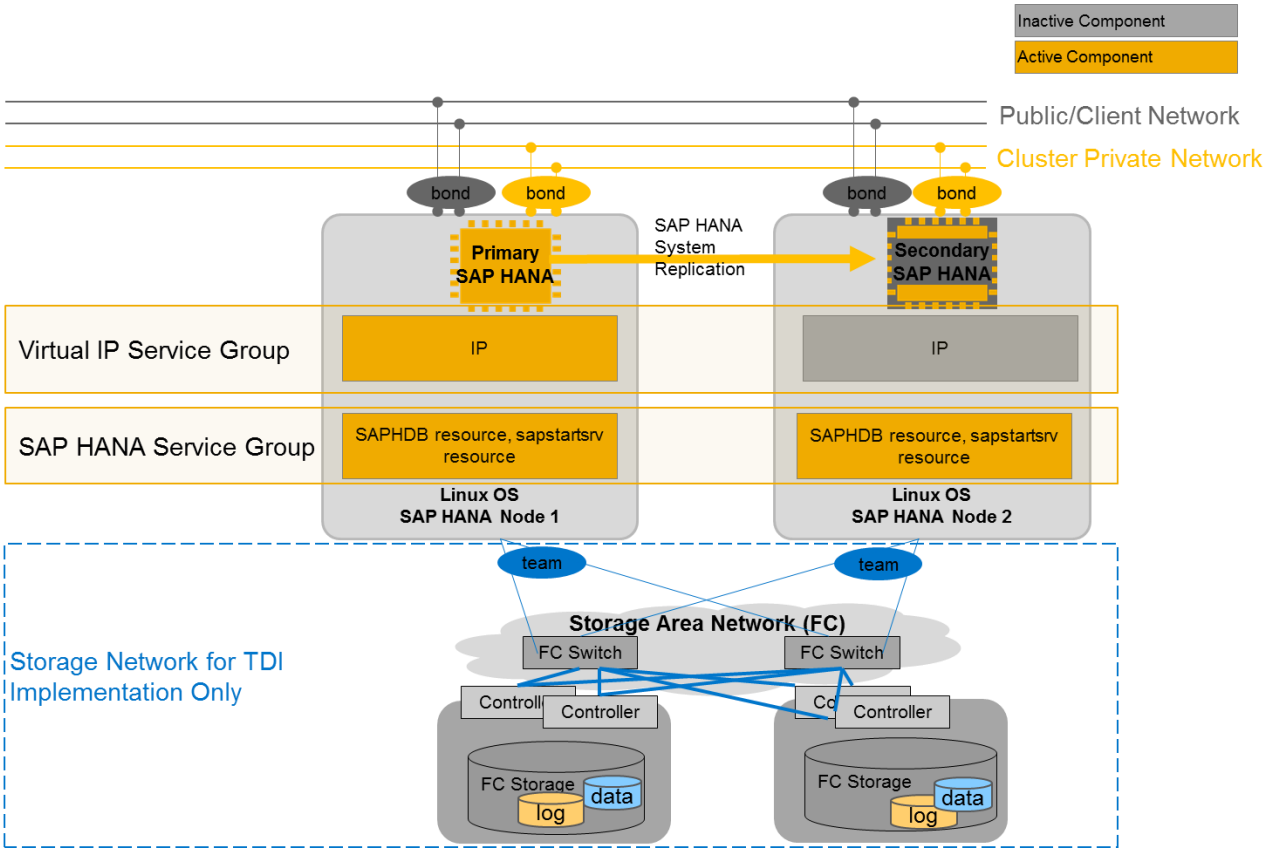

Figure 25 - Reference Architecture for SAP HANA System Replication Automation with Veritas InfoScale Availability

In the reference architecture, Veritas InfoScale Availability cluster has two Service Groups:

- **SAP HANA Service Group**: a parallel Service Group contains two resources SAPHDB resource and sapstartsrv resource. SAPHDB is parent, while sapstartsrv is child. Child resources must be online

before parent resources become online. Parent resources must be taken offline before child resources are taken offline[147]. This will maintain two SAP HANA systems in synchronous System Replication mode. Veritas InfoScale Availability cluster utilizes the SAP Start Service (sapstartsrv) for SAP HANA instance start, stop and monitoring operations.

- **Virtual IP Service Group**: only one IP resource is in this failover Service Group to provide virtual IP address/hostname for external SAP HANA accesses. The virtual IP Service Group is the parent group of SAP HANA Service Group, there is "online local hard" dependency[148] between them.
- For the network, a public network for client-server communications and a cluster private network for a cluster heart-beat communication are needed. If Tailored Datacenter Integration (TDI) for storage concept is employed, a storage network is also involved. To achieve HA at network level, redundancy and interface card bonding or teaming is recommended.

4.2.4 Implementation Tips for SAP HANA System Replication with Veritas InfoScale Availability

Detailed Implementation Guide with Veritas InfoScale Availability

Symantec provides detailed guides "Cluster Server Agent for SAP HANA Database Installation and Configuration Guide" and "Veritas/Symantec ApplicationHA Solution for SAP HANA" for physical and virtual environment correspondingly[149].

SAP HANA System Installation

This are the same as steps listed for "SAP HANA System Replication Automation with Pacemaker Cluster Resource Manager".

[147] See also Resource and resource dependency in Veritas VCS
[148] If a child group faults, the parent is taken offline before the child group is taken offline; if the child group fails over, the parent fails over to the same node; if the child group cannot fail over or brought online, the parent group remains offline; if the parent faults, the child is taken offline. See also Service group dependency configurations in Veritas VCS
[149] Also see https://sort.veritas.com/agents

4.2.5 Technical Architecture for SAP HANA System Replication Automation with HPE Serviceguard

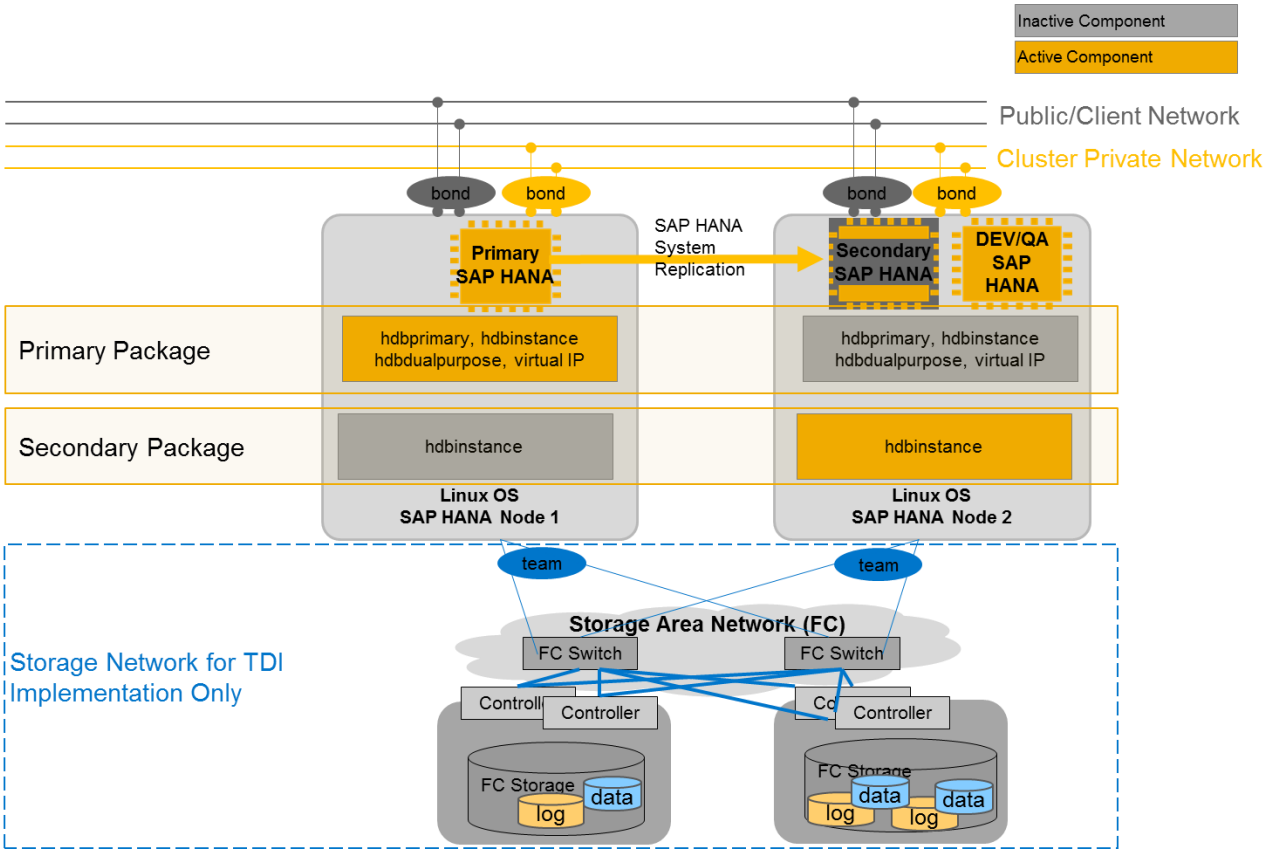

Figure 26 - Reference Architecture for SAP HANA System Replication Automation with HPE Serviceguard

In the above reference architecture, HPE Serviceguard cluster creates two packages:

- **Primary Package**: makes sure that one of the SAP HANA system works as SAP HANA primary system, with the help of HPE Serviceguard modules sg/package_ip, sg/dependency, sg/priority, sgesap/hdbprimary, and sgesap/hdbinstance for scale-up setups (or sgesap/hdbdistinstance for scale-out systems). A relocatable virtual IP address for external accesses also stays within the Primary Package.
- **Secondary Package**: registers the SAP HANA instance to the primary SAP HANA system as the secondary system and keeps the secondary SAP HANA replicating from primary SAP HANA system with modules sg/dependency, sg/priority, and sgesap/hdbinstance for scale-up setups (or sgesap/hdbdistinstance for scale-out systems). SAP HANA System Replication and dualpurpose configurations allow the operation of a production SAP HANA instance together with a non-production SAP HANA instance with certain restrictions. A non-production HANA instance can be installed on the replication node. However, installing and running additional SAP HANA instances on the primary node is not supported. Hence, it is not possible to do a complete role-reversal in which node one takes over the role of node two for both production and non-production instances. In case of a failover of the Primary Package to the node running the secondary SAP HANA instance on a non-production SAP HANA system, HPE Serviceguard ensures the non-production systems on the node are halted before

it triggers the takeover procedure, which promotes the secondary SAP HANA instance into the primary SAP HANA production instance. Although the non-production SAP HANA instance is running on the server of secondary SAP HANA instance, the HPE Serviceguard module sgesap/hdbdualpupose is configured in the Primary Package. To reuse the secondary SAP HANA server for non-production SAP HANA instance, make sure it does not use the preconfigured storage, but add additional DATA/LOG volumes. One additional relocatable virtual IP address for external read-only accesses could be created to stay within the Secondary Package for SAP HANA 2.0 when Active/Active feature enabled. The SGeSAP modules hdbinstance and hdbdistinstance can be used to define automatically handled virtual IP addresses for HANA read-only accesses. As part of a primary failover, these addresses become unconfigured, while during the role reversal automation, read-only access becomes re-configured on the previously failed primary instance after that instance is re-registered as new secondary.

- HPE Serviceguard interoperates with the SAP Start Service (sapstartsrv) agents during SAP HANA instance start, stop and monitoring operations.
- For the network, again public network for client-server communications and cluster private network for cluster heart-beat communication are needed. If Tailored Datacenter Integration (TDI) for storage concept is employed, storage network is also involved. To achieve HA at network level, redundancy and interface card bonding or teaming is recommended.

4.2.6 Implementation Tips for SAP HANA System Replication with HPE Serviceguard

Detailed Implementation Guide with HPE Serviceguard

The detailed HP guide "Managing HPE Serviceguard Extension for SAP for Linux" has specific sections for SAP HANA System Replication configuration.

SAP HANA System Installation

This is the same as steps listed for "SAP HANA System Replication Automation with Pacemaker Cluster Resource Manager".

Other HPE Serviceguard Cluster Considerations

- As HPE Serviceguard uses SAP Start Service, define the cluster awareness library in the SAP HANA instance profiles. Add the following line the SAP HANA instance profiles /usr/sap/<SID>/SYS/profile/<SID>_<INST_Nr>_<nodename>

  `service/halib = /opt/cmcluster/lib64/saphpsghalib.so` (for SLES)

  `service/halib = /usr/local/cmcluster/lib/saphpsghalib.so` (for RHEL)

- To prevent Serviceguard occupying the HANA SQL port temporarily, reserve the HANA SQL port explicitly in the SAP host controller on all the cluster nodes with the following line in /usr/sap/hostctrl/exe/host_profile file:

  ```
  ADDITIONAL_RESERVED_PORTS = 3<INSTNR>15
  ```

- If the specific fencing mechanism - Smart Quorum is used, Quorum networks must not be used for heartbeat exchange.

4.2.7 Technical Architecture for SAP HANA System Replication Automation with IBM TSAMP

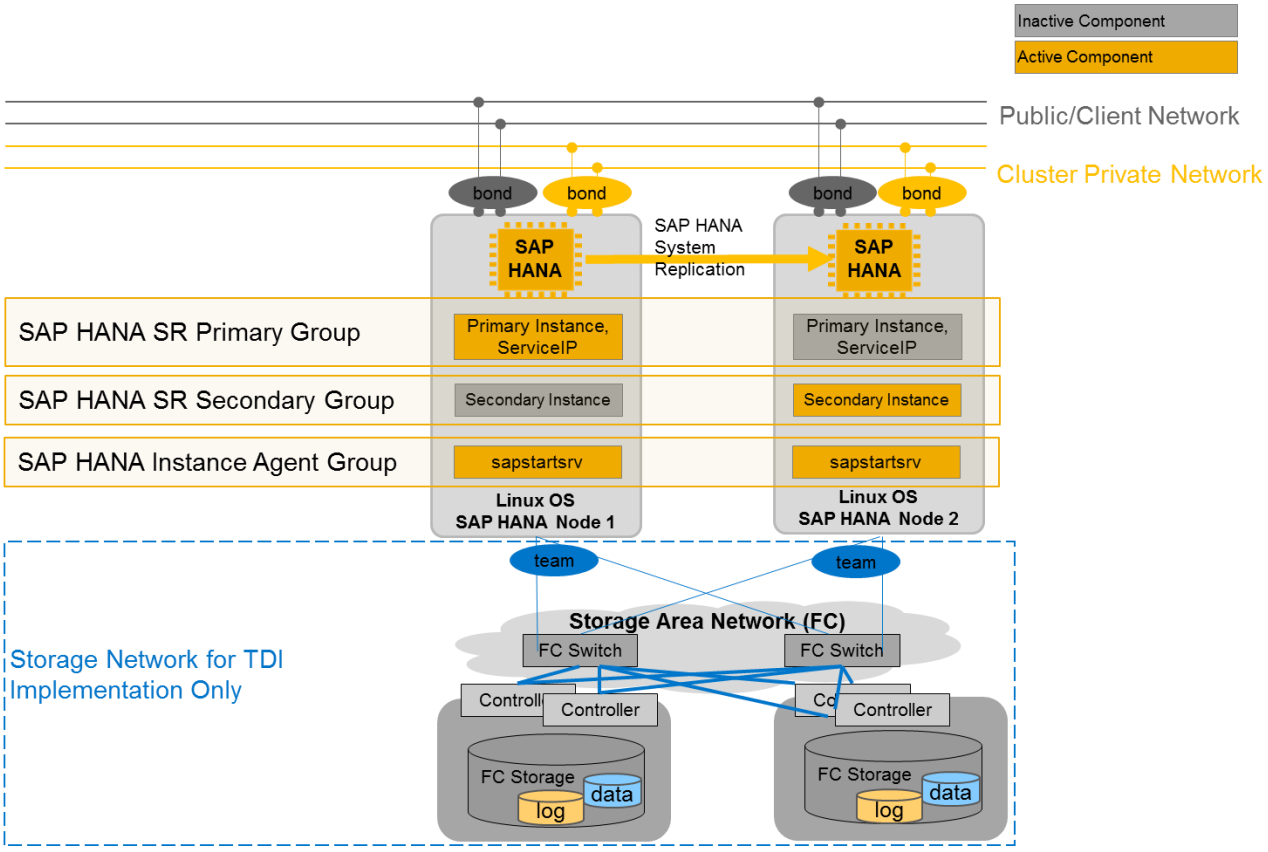

Figure 27 - Reference Architecture for SAP HANA System Replication Automation with IBM TSAMP

In the above reference architecture, the IBM TSAMP cluster has three main resource groups:

- **SAP HANA SR Primary Group**: includes components "HDB Primary instance" and "HDB ServiceIP". "HDB Primary instance" is the SAP HANA Primary System, which provides external services. Floating "HDB ServiceIP" for external accesses also stays within this resource group.
- **SAP HANA SR Secondary Group**: registers the SAP HANA instance to the primary SAP HANA system as the secondary system and keeps the secondary SAP HANA replicating from primary SAP HANA system with component "HDB Secondary instance"

- **SAP HANA Instance Agent Group:** interoperates with the SAP Start Service (sapstartsrv) agents during SAP HANA instance start, stop and monitoring operations.

4.2.8 Implementation Tips for SAP HANA System Replication with IBM TSAMP

Detailed Implementation Guide with IBM TSAMP

The detailed IBM TSAMP guide "Tivoli System Automation for Multiplatforms 4.1 Product manuals" has specific sections for SAP HANA System Replication configuration.

SAP HANA System Installation

This is the same as steps listed for "SAP HANA System Replication Automation with Pacemaker Cluster Resource Manager".

Other IBM TSAMP Cluster Considerations

- The TSAMP high availability policy manages the SAP HANA System Replication Single-Node setup. SAP HANA on Power is required.
- define the cluster awareness library in the SAP HANA default profile. Add the following line the SAP HANA default profile /usr/sap/<SID>/SYS/profile/DEFAULT.PFL

```
service/halib = /usr/sap/<SID>/SYS/exe/uc/<your platform>/saphascriptco.so

service/halib_cluster_connector = /usr/sbin/rsct/sapolicies/sap/bin/sap_tsamp_cluster_connector
```

- disable autostart of all SAP HANA instances in all their profiles by commenting the line

```
Autostart = 1
```

4.3 Design and Implementation for SAP NetWeaver and SAP HANA in one Cluster

4.3.1 Technical Architecture for SAP NetWeaver and SAP HANA in one Cluster

The initial offered appliance delivery of SAP HANA required a separated deployment of the SAP HANA and the SAP NetWeaver system using separate hardware. Now in order to simplify technical setup, reuse of free hardware resources, and reduce cost and operational effort, SAP supports to deploy SAP HANA and SAP NetWeaver AS on one server or cluster in all productive and non-productive systems with the following prerequisites:

- SAP HANA single-node installation (no scale-out)
- SAP NW AS ABAP Stack: SAP NetWeaver AS ABAP 7.4 with SAP HANA 1.0 SPS 7 and later[150]

[150] See also SAP Note 1953429 - SAP HANA and SAP NetWeaver AS ABAP on one Server

- SAP NW AS JAVA Stack: SAP NetWeaver AS Java 7.4 SPS 5 with SAP HANA 1.0 SPS 8 and later[151]
- Separate SIDs for SAP HANA and SAP NetWeaver
- Additive sizing for memory resources, correct configuration of memory allocation for SAP HANA (GLOBAL_ALLOCATION_LIMIT) and SAP NetWeaver AS

However, there are potential disadvantages for this kind of setups:

- Potential mutual impact due to shared hardware resources (CPU and memory)
- Potential conflicts in OS kernel requirements and OS configuration settings

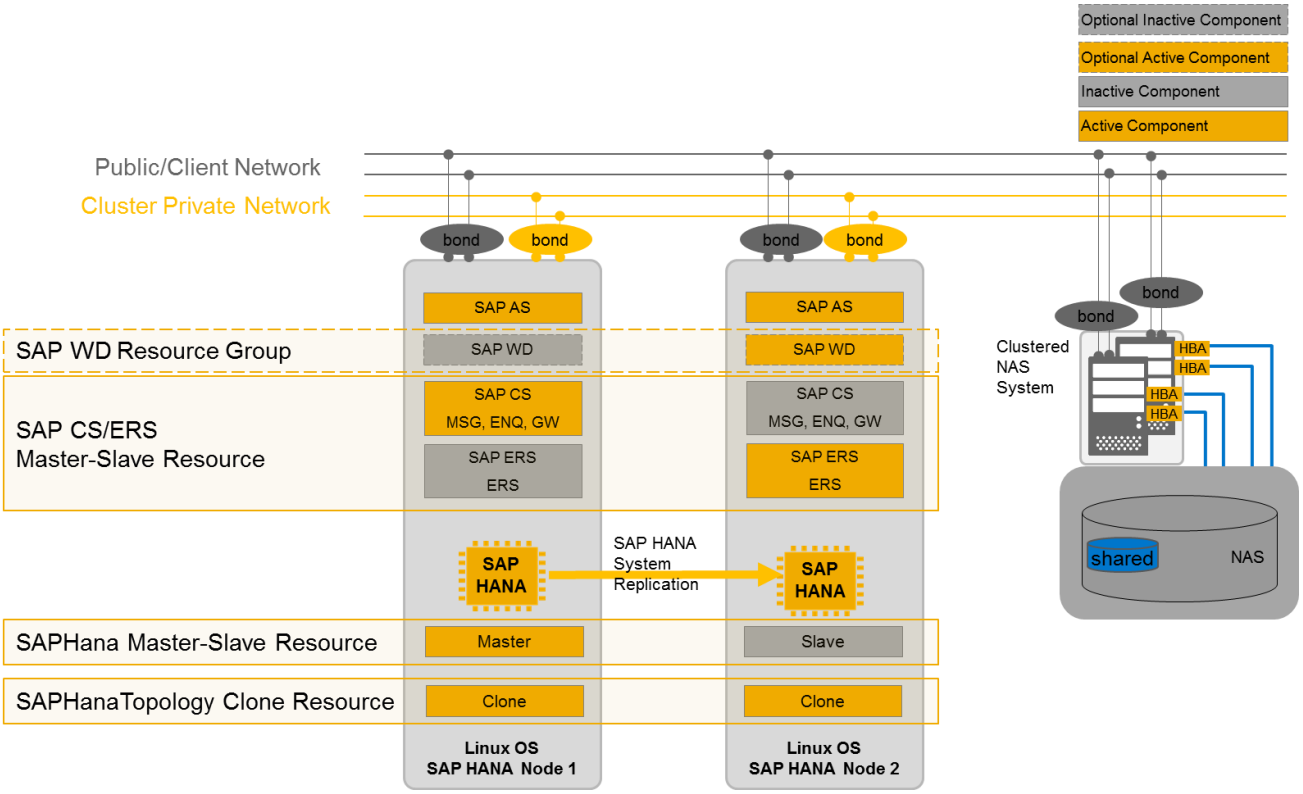

Figure 28 - Reference Architecture for SAP NetWeaver and SAP HANA System Replication in one Cluster

The above reference architecture takes Pacemaker based cluster solutions as example. It will combine the Resources, Resource Groups/Service Groups/Packages in one cluster. Since one cluster can only be managed by one cluster solution, it means the applicable cluster solution is limited to those, which both are SAP HA-interface certified and could automate SAP HANA System Replication.

4.3.2 Implementation Tips for SAP NetWeaver and SAP HANA in one Cluster

Detailed Implementation Guides

Also refer to the previous sections for the guides applicable to chosen cluster solution.

[151] See also SAP Note 2043509 - SAP HANA and SAP NetWeaver Java on a Single Host

Other Cluster Design Considerations

- As discussed, in a Linux environment it's not recommended to use the integrated NFS server for SAP Global File Systems, while only two nodes are available. Technically, cluster file systems such as GFS2 or OCFS2, can be used for a two-node cluster instead. If this additional cluster stack at file system level is considered too complex, external HA NAS system can be an alternative. In the above reference architecture, an external HA NAS system is adopted. If the chosen cluster solution needs the SAP Instances or SAP Web Dispatcher binary directories to be in shared file systems, this external NAS will also provide them.
- At this moment, the potential cluster solutions are SLES for SAP applications based cluster, RHEL HA cluster, HPE Serviceguard Extension for SAP on Linux, and Veritas InfoScale Availability. However, HPE Serviceguard has made clustered SAP HANA instances and clustered NetWeaver instances, both cannot be stacked on the same cluster nodes. Therefore, only SLES cluster, RHEL cluster and Veritas InfoScale Availability are possible options for now.
- To run SAP HANA and SAP NetWeaver in one cluster may not be a technical limit, but this scenario has not been described in cluster vendor guides, so it needs to be described and tested in detail. It is strongly recommended to define a Proof of Concept together with cluster vendor and SAP.
- As an alternative, SAP Note 1953429 has provided a method the install and configure SAP HANA and SAP NetWeaver in HA fashion, but the failover or switchover needs to be done manually.

5 Verification of High Availability Clusters

In the complete HA landscape, we could have a SAP NetWeaver cluster, a SAP HANA cluster, and an NFS cluster. After configuring these clusters, we need verify whether they are working as expected. Specifically for each cluster, all verifications need to be done at different levels: application level, operating systems and cluster level, and hardware level.

Below is the list of recommended minimum test scenarios to make sure the clusters work as designed.

5.1 Application Level Verification

#	Category	Description	Action	Prerequisites	Result
A1	SAP NW	Kill Message Server (Recoverable, no failover)	Identify Message Servers PID, and then terminate the process	Cluster in origin state (both nodes up and active, all instances up)	Message server is restarted by cluster; Killing multiple times may cause failover SAP CS instance to other node
A2	SAP NW	Kill Enqueue Server(irrecoverable, failover required)	Identify Enqueue Server PID, and then terminate the process	Cluster in origin state (both nodes up and active, all instances up)	Cluster failover SAP CS instance to the node where active SAP ERS instance running, and no Enqueue information lost
A3	SAP NW	Kill Enqueue Replication Server (Recoverable, no failover)	Identify Enqueue Replication Server PID, and then terminate the process	Cluster in origin state (both nodes up and active, all instances up)	Enqueue replication server is restarted by cluster, and get synchronized with Enqueue server very fast
A4	SAP NW	Kill SAP CS sapstartsrv service(Recoverable, no failover)	Identify sapstartsrv PID: pgrep -fl "SCS.*sapstartsrv" then terminate the process	Cluster in origin state (both nodes up and active, all instances up)	Next monitor restarts sapstartsrv.
A5	SAP NW	Kill SAP ERS sapstartsrv service(Recoverable, no failover)	Identify sapstartsrv PID: pgrep -fl "ERS.*sapstartsrv" then terminate the process	Cluster in origin state (both nodes up and active, all instances up)	Next monitor restarts sapstartsrv.
A6	SAP NW	Kill SAP Gateway Server (Recoverable, no failover)	Identify SAP Gateway Server PID, and then terminate the process	Cluster in origin state (both nodes up and active, all instances up)	SAP Gateway server is restarted by cluster

#	Category	Description	Action	Prerequisites	Result
A7	SAP NW	Move the SAP CS Resource Group/Resource/Service Group/Package to candidate node where SAP ERS running	Use cluster administration procedures to move resources	Cluster in origin state (both nodes up and active)	Cluster moves SAP CS instance to the node where active SAP ERS instance running, and no Enqueue information lost
A8	SAP NW	Kill SAP Web Dispatcher (Recoverable, no failover)	Identify SAP Web Dispatcher PID, and then terminate the process	Cluster in origin state (both nodes up and active, all instances up)	SAP Web Dispatcher server is restarted by cluster
A9	SAP NW	Move the SAP Web Dispatcher Resource Group/Resource/Service Group/Package to candidate node	Use cluster administration procedures to move resources	Cluster in origin state (both nodes up and active)	Cluster moves SAP Web Dispatcher to the other node
A10	SAP NW	Start and stop SAP instances	sapcontrol -<nr> -function "Start|Stop"	Cluster in origin state (both nodes up and active, all instances up)	SAP instance addressed by the sapcontrol call is started/stopped
A11	SAP NW	Apply kernel patch in $DIR_CT_RUN and restart system (no failover required)	Update SAP Kernel	Cluster in origin state (both nodes up and active, all instances up)	SAP Instances started after the SAP kernel update are running the new kernel code
D1	SAP HANA	Kill Index Server on primary SAP HANA instance (Recoverable, no failover)	Identify hdbindexserver PID, then terminate process	Cluster in origin state (both nodes up and active, all instances up)	The killed Index Server is restarted by daemon
D2	SAP HANA	Kill Index Server on secondary SAP HANA instance (Recoverable, no failover)	Identify hdbindexserver PID, then terminate process	Cluster in origin state (both nodes up and active, all instances up)	The killed Index Server is restarted by daemon, and the Replication status go back to normal (Active) soon
D3	SAP HANA	Kill Name Server on primary SAP HANA instance (Recoverable, no failover)	Identify hdbnameserver PID, then terminate process	Cluster in origin state (both nodes up and active, all instances up)	The killed Name Server is restarted by daemon

#	Category	Description	Action	Prerequisites	Result
D4	SAP HANA	Kill Name Server on secondary SAP HANA instance (Recoverable, no failover)	Identify hdbnameserver PID, then terminate process	Cluster in origin state (both nodes up and active, all instances up)	The killed Name Server is restarted by daemon, and the Replication status go back to normal (Active) soon
D5	SAP HANA	Kill XS server on primary SAP HANA instance (Recoverable, no failover)	Identify hdbxsengine PID, then terminate process	Cluster in origin state (both nodes up and active, all instances up)	The killed XS server is restarted by daemon
D6	SAP HANA	Kill XS server on secondary SAP HANA instance (Recoverable, no failover)	Identify hdbxsengine PID, then terminate process	Cluster in origin state (both nodes up and active, all instances up)	The killed XS server is restarted by daemon, and the Replication status go back to normal (Active) soon
D7	SAP HANA	Kill Web Dispatcher on primary SAP HANA instance (Recoverable, no failover)	Identify hdbwebdispatcher PID, then terminate process	Cluster in origin state (both nodes up and active, all instances up)	The killed Web Dispatcher is restarted by daemon
D8	SAP HANA	Kill Web Dispatcher on secondary SAP HANA instance (Recoverable, no failover)	Identify hdbwebdispatcher PID, then terminate process	Cluster in origin state (both nodes up and active, all instances up)	The killed Web Dispatcher is restarted by daemon
D9	SAP HANA	Kill daemon on primary SAP HANA instance (Irrecoverable)	Identify daemon PID (pgrep -fl hdb.sap<SID>), then terminate process	Cluster in origin state (both nodes up and active, all instances up)	If takeover is preferred in SAP HANA resource configuration, takeover happens; if automatic role rehearsal is configured, the current primary becomes secondary
D10	SAP HANA	Kill daemon on secondary SAP HANA instance (Recoverable, no failover)	Identify daemon PID, then terminate process	Cluster in origin state (both nodes up and active, all instances up)	Secondary SAP HANA instance is restarted by cluster

#	Category	Description	Action	Prerequisites	Result
D11	SAP HANA	Kill sapstartsrv service on primary SAP HANA instance (Recoverable, no failover)	Identify sapstartsrv PID: pgrep -fl "<SID>.*sapstartsrv", then terminate the process	Cluster in origin state (both nodes up and active, all instances up)	Next monitor restarts sapstartsrv.
D12	SAP HANA	Kill sapstartsrv service on secondary SAP HANA instance (Recoverable, no failover)	Identify sapstartsrv PID: pgrep -fl "<SID>.*sapstartsrv", then terminate the process	Cluster in origin state (both nodes up and active, all instances up)	Next monitor restarts sapstartsrv.
D13	SAP HANA	Move the primary SAP HANA Resource Group/Resource/Service Group/Package to secondary SAP HANA node	Use cluster administration procedures to move resources	Cluster in origin state (both nodes up and active)	Primary SAP HANA instance moves to the node where secondary SAP HANA instance was running, and no transaction data lost; if automatic role rehearsal is configured, the current primary becomes secondary
N1	NFS	Kill NFS server process on active NFS cluster node	Identify NFS server process PID: pgrep -fl nfsd, then terminate the process	Cluster in origin state (both nodes up and active)	NFS server process is restarted by cluster, and the clients can still accesses the shared file systems
N2	NFS	Move the NFS Resource Group/Resource/Service Group/Package to candidate node	Use cluster administration procedures to move resources	Cluster in origin state (both nodes up and active)	Cluster moves NFS service to the other node, and the clients can still accesses the shared file systems

5.2 Cluster Level Verification

#	Description	Action	Prerequisites	Result
C1	Simulate cluster failure by terminating cluster software processes	Identify cluster software process(es), such as corosync then terminate the process(es)	Cluster in origin state (both nodes up and active)	The chosen node is fenced; all resources on that node are migrated to the remaining nodes

#	Description	Action	Prerequisites	Result
C2	Test cluster maintenance mode	Use cluster administration commands to enter cluster into maintenance mode	Cluster in origin state (both nodes up and active)	The cluster enters maintenance mode, while the running services are not impacted (still running)
C3	Fence one node	Use cluster administration commands to fence one node	Cluster in origin state (both nodes up and active)	The fenced node is rebooted, and the cluster move all resources to the other node
C4	Rejoin the failed node	Use cluster administration commands to start cluster service on the newly started the node	Cluster is running in single-node mode	Fenced node successfully rejoins cluster. Running resources are unaffected, and the resources expected to be slave start on the newly joined node
C5	Start the whole cluster	Use cluster administration commands to start cluster service on both nodes	Cluster is stopped	Both nodes join cluster, and resources start on their preferred nodes as per configuration
C6	Stop the whole cluster	Use cluster administration commands to stop cluster service on both nodes	Cluster in origin state (both nodes up and active)	Both nodes exit cluster, and resources stop on their respective nodes

5.3 Hardware Level Verification

#	Description	Action	Prerequisites	Result
H1	Power off one node	Power off one node	Cluster in origin state (both nodes up and active)	The cluster move resources to the other node
H2	Power on the stopped node	Power on the stopped node, and use cluster administration commands to start cluster service on the newly started the node	Cluster is running in single-node mode	Newly started node successfully rejoins cluster. Running resources are unaffected, and the resources expected to be slave start on the newly joined node
H3	Disable fencing/quorum/STONITH devices	Use Operating Systems, or hardware method to disable fencing/quorum/STONITH devices	Cluster in origin state (both nodes up and active)	Both nodes reset themselves, and cluster is down.

#	Description	Action	Prerequisites	Result
H4	Disable one of cluster interconnect link	Disable one of cluster interconnect link	Cluster in origin state (both nodes up and active)	Cluster and resources are not impacted.
H5	Disable all cluster interconnect links	Disable all cluster interconnect links	Cluster in origin state (both nodes up and active)	Cluster resets one node and moves resources to the surviving node

6 Appendix. Further Information

- Updates to this document <https://support.sap.com/content/dam/SAAP/SAP_Activate/AGS_70.pdf>
- SAP Notes <http://support.sap.com/notes>
    - Linux
        - SAP Note 171356 - SAP Software on Linux: General information
        - SAP Note 611361 - Hostnames of SAP servers
        - SAP Note 1386233 - TREX or BWA installation on hosts with uppercase hostnames
        - SAP Note 936887 - End of maintenance for Linux distributions
        - SAP Note 1056161 - SUSE Priority Support for SAP applications
        - SAP Note 1310037 - SUSE LINUX Enterprise Server 11: Installation notes
        - SAP Note 1984787 - SUSE LINUX Enterprise Server 12: Installation notes
        - SAP Note 2235581 - SAP HANA: Supported Operating Systems
        - SAP Note 1944799 - SAP HANA Guidelines for SLES Operating System Installation
        - SAP Note 1954788 - SAP HANA DB: Recommended OS settings for SLES 11 / SLES for SAP Applications 11 SP3
        - SAP Note 2240716 - SAP HANA DB: Recommended OS settings for SLES 11 / SLES for SAP Applications 11 SP4
        - SAP Note 2205917 - SAP HANA DB: Recommended OS settings for SLES 12 / SLES for SAP Applications 12
        - SAP Note 1855805 - Recommended SLES 11 packages for HANA support on OS level
        - SAP Note 2397039 - FAQ: SAP on RHEL
        - SAP Note 1631106 - Red Hat Enterprise Linux for SAP Business Applications
        - SAP Note 1496410 - Red Hat Enterprise Linux 6.x: Installation and Upgrade
        - SAP Note 2002167 - Red Hat Enterprise Linux 7.x: Installation and Upgrade
        - SAP Note 2009879 - SAP HANA Guidelines for Red Hat Enterprise Linux (RHEL) Operating System
        - SAP Note 2013638 - SAP HANA DB: Recommended OS settings for RHEL 6.5
        - SAP Note 2136965 - SAP HANA DB: Recommended OS settings for RHEL 6.6
        - SAP Note 2247020 - SAP HANA DB: Recommended OS settings for RHEL 6.7
        - SAP Note 2292690 - SAP HANA DB: Recommended OS settings for RHEL 7

- SAP Note 2001528 - Linux: SAP HANA Database SPS 08, SPS 09 and SPS 10 on RHEL 6 or SLES 11
- SAP Note 2228351 - Linux: SAP HANA Database SPS 11 revision 110 (or higher) on RHEL 6 or SLES 11
- SAP Note 2195019 - Linux: Using SAP kernel 748 and higher or SAPinst 749 and higher on RHEL 6, OL6 and SLES 11
- SAP Note 2338763 - Linux: Running SAP applications compiled with GCC 5.x
- SAP Note 1565179 - SAP software and Oracle Linux

- SAP HANA
  - SAP Note 1650957 - SAP HANA Database: Starting the Script Server
  - SAP Note 2091313 - HANA Statistics Server - changed standard setting of the statistics server as of Revision 93
  - SAP Note 1990354 - Changes in internal HANA Web Dispatcher as of HANA SP9
  - SAP Notes 2092793 Reduced delivery with SPS09: Restrictions for the usage of SAP HANA system replication with multi-tenant database containers
  - SAP Note 2096000 SAP HANA multitenant database containers - Additional Information
  - SAP Note 2423367 - Multitenant database containers will become the standard and only operation mode
  - SAP Note 1681092 - Multiple SAP HANA DBMSs (SIDs) on one SAP HANA system
  - SAP Note 1661202 - Support multiple applications one SAP HANA database / tenant DB
  - SAP Note 1826100 - Multiple applications SAP Business Suite powered by SAP HANA
  - SAP Note 2233866 - SAP HANA XS Advanced Model SPS11 Release Note
  - SAP Note 2140959 - SAP HANA Dynamic Tiering - Additional Information
  - SAP Note 2394124 - SAP HANA 2.0 Dynamic Tiering - Additional Information
  - SAP Note 2380229 - SAP HANA Platform 2.0 – Central Note
  - SAP Note 2380257 - SAP HANA Platform 2.0 SPS 00 Release Note
  - SAP Note 2420699 - Release of SAP HANA Database 2.0 for older SAP Versions
  - SAP Note 2086829 - SAP HANA dynamic tiering sizing ratios
  - SAP Note 2091095 - SAP HANA Enterprise Information Management
  - SAP Note 1872170 - Suite on HANA sizing report

- SAP Note 1736976 - Sizing Report for BW on HANA
- SAP Note 1953429 - SAP HANA and SAP NetWeaver AS ABAP on one Server
- SAP Note 2043509 - SAP HANA and SAP NetWeaver Java on a Single Host

- High Availability (HA)
  - Note 1552925 - Linux: High Availability Cluster Solutions
  - Note 1374671 - High Availability in Virtual Environment on Windows
  - Note 1678705 - Installation scenarios for a standalone ASCS instance
  - Note 1064290 - Standalone enqueue server and standalone gateway
  - Note 1010990 - Configuring a Standalone Gateway in an HA ASCS instance
  - SAP Note 1693245 - SAP HA Script Connector Library
  - SAP Note 1763512 - Support details for SUSE Linux Enterprise High Availability
  - SAP Note 2309342 - SUSE Linux Enterprise High Availability Extension on AWS for SAP HANA
  - SAP Note 2513384 - SUSE Linux Enterprise Server for SAP Applications on Azure
  - SAP Note 1908655 - Support details for Red Hat Enterprise Linux HA Add-On
  - SAP Note 1906423 - Support details for Veritas
  - SAP Note 2057595 - FAQ: SAP HANA High Availability
  - SAP Note 1999880 - FAQ: SAP HANA System Replication
  - SAP Note 2303243 - SAP HANA Multitier System Replication – supported replication modes between sites
  - SAP Note 2369981 - Required configuration steps for authentication with HANA System Replication
  - SAP Note 2356851 – SAP HANA Dynamic Tiering Support for SAP HANA System Replication
  - SAP Note 2391079 - Access restrictions in Active/Active system setup
  - SAP Note 2392706 - HANA Active/Active (read enabled) hint based statement routing
  - SAP Note 2402547 - Active/Active (read enabled) - monitoring of the secondary system with HANA Cockpit 2.0
  - SAP Note 2422925 - Restriction of secondary usage for other HANA systems (SIDs) in an Active/Active (read enabled) configuration
  - SAP Note 2062631 - high availability limitation for SAN storage
  - SAP Note 1755396 - Released DT solutions for SAP HANA with disk replication

        - SAP Note 2127006 – SAP HANA DB used in Fujitsu FlexFrame Orchestrator
        - SAP Note 1768213 - Support details for NEC EXPRESSCLUSTER
        - SAP Note 2302728 - Supported scenarios with NEC Expresscluster on Amazon Web Services
    - Partitioning and virtualization
        - SAP Note 1122387 - Linux: SAP Support in virtualized environments
        - SAP Note 1492000 - General Support Statement for Virtual Environments
        - Note 1788665 - SAP HANA Support for Virtualized Environments
        - Note 2077823 - SAP HANA DB: Supported Hardware partitioned environments
        - Note 1995460 - Single SAP HANA VM on VMware vSphere in production
        - Note 2024433 - Multiple SAP HANA VMs on VMware vSphere in production (controlled availability)
        - Note 2063057-SAP HANA VMs on Hitachi LPAR in production (controlled availability)
        - Note 2103848 - SAP HANA on HP Hardware Partitions (nPAR) in production
        - SAP Note 1380654 - SAP support in public cloud environments
        - SAP Note 1964437 - SAP HANA on AWS: Supported AWS EC2 products
        - SAP Note 2316233 - SAP HANA on Microsoft Azure
- Best-practice documents/white papers/implementation guides:
    - SAP HA Best Practices on SUSE Linux Enterprise Server <https://www.suse.com/products/sles-for-sap/>
    - Best Practices for Mission-Critical SAP Applications <https://www.suse.com/products/sles-for-sap/resource-library/sap-best-practices/>
    - Addendum for SUSE Whitepapers and Best Practices - SAP NetWeaver HA <https://wiki.scn.sap.com/wiki/display/ATopics/Addendum+for+SUSE+Whitepapers+and+Best+Practices+-+SAP+NetWeaver+HA>
    - Deploying Highly Available SAP NetWeaver-based Servers Using Red Hat Enterprise Linux HA add-on with Pacemaker <https://access.redhat.com/solutions/963123>
    - SAP NetWeaver in pacemaker cluster (Updated document from Red Hat) <https://access.redhat.com/articles/3150081>
    - EXPRESSCLUSTER X 3.3 for Linux SAP NetWeaver System Configuration Guide <http://www.nec.com/en/global/prod/expresscluster/en/support/setup/HOWTO_SAP_CG_Linux_EN_08.pdf>
    - EXPRESSCLUSTER X 3.3 for Linux SAP NetWeaver Configuration Example <http://www.nec.com/en/global/prod/expresscluster/en/support/setup/HOWTO_SAP_CE_Linux_EN_05.pdf>

- White Paper SAP Business Suite Powered by SAP HANA High Availability with NEC EXPRESSCLUSTER
  <http://www.nec.com/en/global/prod/expresscluster/en/collaterals/pdf/WP_SAPHANA_HA_EXPRESSCLUSTER_EN_03.pdf>
- White Paper A High Availability Solution for SAP S/4 HANA Running on IBM Power Systems Implemented by EXPRESSCLUSTER
  <http://www.nec.com/en/global/prod/expresscluster/en/collaterals/pdf/WP_SAPS4HANA_IBMPOWER_EXPRESSCLUSTER_EN_01.pdf>
- Providing High Availability for SAP Resources with Oracle Clusterware 11g Release 2 and Clusterware 12c Release 1 <http://www.oracle.com/us/solutions/sap/database/high-availability-for-sap-resources-3812782.pdf>
- SIOS Protection Suite for Linux SAP Recovery Kit
  <http://docs.us.sios.com/Linux/9.2.1/LK4L/SAPPDF/SAP.pdf>
- SIOS Protection Suite for Linux - SAP Solution
  <http://docs.us.sios.com/Linux/9.2.1/LK4L/SAPSolutionPDF/SAPSolution.pdf>
- SIOS LifeKeeper Single Server Protection (SSP) for Linux Technical Documentation
  <http://docs.us.sios.com/Linux/9.2.1/LifeKeeperSSP/TechDoc/index.htm>
- SIOS Protection Suite for Linux SAP HANA 1.0 Recovery Kit
  <http://docs.us.sios.com/Linux/9.2.1/LK4L/QuickStartPDFs/Content/Resources/pdfs/quick_start_guides/SAPHANA1.pdf>
- SIOS Protection Suite for Linux SAP HANA 2.0 Recovery Kit
  <http://docs.us.sios.com/Linux/9.2.1/LK4L/QuickStartPDFs/Content/Resources/pdfs/quick_start_guides/SAPHANA2.pdf>
- Cluster Server Agent for SAP NetWeaver Installation and Configuration Guide 7.0
  <https://sort.veritas.com/agents/download_docs/7201/vcs_sapnw_install>
- Veritas/Symantec ApplicationHA Solution for SAP NetWeaver
  <https://support.symantec.com/en_US/article.TECH213348.html>
- Automate your SAP HANA System Replication Failover Setup-guide
  <https://www.suse.com/promo/sap/hana/replication.html>
- SAP HANA SR Cost Optimized Scenario: SUSE Linux Enterprise Server for SAP Applications 11 SP4 <https://www.suse.com/docrepcontent/container.jsp?containerId=9109>
- SAP HANA SR Performance Optimized Scenario: SUSE Linux Enterprise Server for SAP Applications 11 SP4 <https://www.suse.com/docrepcontent/container.jsp?containerId=9111>
- SAP HANA SR Cost Optimized Scenario: SUSE Linux Enterprise Server for SAP Applications 12 SP1 <https://www.suse.com/docrepcontent/container.jsp?containerId=9110>

- SAP HANA SR Performance Optimized Scenario: SUSE Linux Enterprise Server for SAP Applications 12 SP1 <https://www.suse.com/docrep/documents/ir8w88iwu7/suse_linux_enterprise_server_for_sap_applications_12_sp1.pdf>
- Automated SAP HANA System Replication with Pacemaker on RHEL Setup Guide <https://access.redhat.com/articles/1466063>
- SAP HANA system replication in pacemaker cluster <https://access.redhat.com/articles/3004101>
- Cluster Server Agent for SAP HANA Database Installation and Configuration Guide <https://sort.veritas.com/agents/download_docs/7032/vcs_saphdb_install>
- Veritas/Symantec ApplicationHA Solution for SAP HANA <https://www.veritas.com/support/en_US/article.000023431>
- Managing HPE Serviceguard Extension for SAP for Linux Version B.12.10.00 <https://support.hpe.com/hpsc/doc/public/display?sp4ts.oid=null&docLocale=en_US&docId=emr_na-a00018047en_us>
- Tivoli System Automation for Multiplatforms 4.1 Product manuals <http://www.ibm.com/support/knowledgecenter/en/SSRM2X_4.1.0/com.ibm.samp.doc_4.1/samp_IC_productmanuals.html?lang=en>
- White paper Safeguarding Business Continuity on Critical SAP HANA Systems <https://docs.ts.fujitsu.com/dl.aspx?id=dd1d66ac-a1ab-472e-a297-bec5d23b3c7a>
- SAP HANA on IBM Bluemix Infrastructure Implementation Guide <http://wpc.c320.edgecastcdn.net/00C320/SAP%20HANA%20Implementation%20Guide.pdf>

- Web resources:
  - High Availability of the SAP Web Dispatcher <http://help.sap.com/saphelp_nw74/helpdata/en/48/9a9a6b48c673e8e10000000a42189b/content.htm>
  - High Availability with the Standalone Enqueue Server <http://help.sap.com/saphelp_nw74/helpdata/en/47/e023f3bf423c83e10000000a42189c/frameset.htm>
  - Replication and Failover <http://help.sap.com/saphelp_nw74/helpdata/en/47/e01bb686983c85e10000000a42189c/content.htm>
  - SAP Installation Guides <http://service.sap.com/instguides>
  - SAP Quicksizer <http://service.sap.com/quicksizer>
  - SAP Standard Application Benchmarks <http://www.sap.com/benchmark>
  - Using sap_vendor_cluster_connector for interaction between cluster framework and sapstartsrv <http://scn.sap.com/community/business-continuity/blog/2014/05/08/using-sapvendorclusterconnector-for-interaction-between-cluster-framework-and-sapstartsrv>

- Manual Page of sap_vendor_cluster_connector <http://scn.sap.com/docs/DOC-54845>
- Certified HA-Interface Partners <http://scn.sap.com/docs/DOC-31701>
- High Availability Partner Information <http://scn.sap.com/docs/DOC-8541>
- Fencing and Stonith <http://clusterlabs.org/doc/crm_fencing.html>
- Loopback NFS: theory and practice <http://lwn.net/Articles/595652/>
- Hangs occur if a Red Hat Enterprise Linux system is used as both NFS server and NFS client for the same mount <https://access.redhat.com/solutions/22231>
- High Availability Solutions for SAP on VMware <http://scn.sap.com/docs/DOC-32023>
- VMware KB: VMware Fault Tolerance FAQ <http://kb.vmware.com/kb/1013428>
- SAP HANA Platform <http://help.sap.com/hana_appliance>
- SAP HANA Administration Guide <http://help.sap.com/hana/sap_hana_administration_guide_en.pdf>
- Certified SAP HANA Hardware Directory (Appliances, Server and Storage Systems, IaaS Platforms)
  <http://global.sap.com/community/ebook/2014-09-02-hana-hardware/enEN/appliances.html>
- SAP HANA Tailored Data Center Integration (FAQ) <http://scn.sap.com/docs/DOC-60327>
- SAP HANA virtualized - Overview <http://scn.sap.com/docs/DOC-60329>
- Technical Deployment Options for SAP Systems with SAP HANA
  <https://service.sap.com/sap/bc/bsp/spn/esa_redirect/index.htm?gotocourse=X&courseid=70229650>
- How to Perform System Replication for SAP HANA <http://scn.sap.com/docs/DOC-47702>
- Implementing a HA/DR Provider
  <http://help.sap.com/saphelp_hanaplatform/helpdata/en/13/67c8fdefaa4808a7485b09815ae0f3/content.htm>
- SAP Certified Appliance Hardware for SAP HANA <http://scn.sap.com/docs/DOC-52522>
- SAP HANA Hardware Directory
  <http://global.sap.com/community/ebook/2014-09-02-hana-hardware/enEN/index.html>
- Approaches to SAP HANA sizing <http://service.sap.com/~sapidb/011000358700000050632013E>
- SAP HANA Tailored Datacenter Integration (TDI) <http://scn.sap.com/docs/DOC-59755>
- Sizing for Tailored Datacenter Integration (TDI) Storage <http://scn.sap.com/docs/DOC-60313>
- SAP HANA Network Requirements Whitepaper <http://scn.sap.com/docs/DOC-62437 >
- Network Recommendations for SAP HANA System Replication
  <https://scn.sap.com/docs/DOC-56044>

- Fail-Safe Operation of SAP HANA: SUSE Extends Its High-Availability Solution <http://scn.sap.com/community/hana-in-memory/blog/2014/04/04/fail-safe-operation-of-sap-hana-suse-extends-its-high-availability-solution>
- Automate SAP HANA System Replication with SLES for SAP Applications <http://scn.sap.com/docs/DOC-56278>
- SAP HANA System Replication Automation (HanaSR) for HANA Scale-Out now officially available with SLES for SAP 12 SP2 <https://www.suse.com/communities/blog/sap-hana-system-replication-automation-hanasr-hana-scale-now-officially-available-sles-sap-12-sp2/>
- SAPHanaSR-ScaleOut: Automating SAP HANA System Replication for Scale-Out Installations with SLES for SAP Applications <https://www.suse.com/communities/blog/saphanasr-scaleout-automating-sap-hana-system-replication-scale-installations-sles-sap-applications/>
- Supported HANA HA Scenarios on Red Hat Enterprise Linux <https://blogs.sap.com/2016/09/28/supported-hana-ha-scenarios-on-red-hat-enterprise-linux/>
- Fujitsu FlexFrame Orchestrator Manuals <http://manuals.ts.fujitsu.com/index.php?id=9215-9217>
- SAP Business Suite Powered by SAP HANA High Availability with NEC EXPRESSCLUSTER <http://scn.sap.com/docs/DOC-65242>
- SAP on IBM Bluemix infrastructure <https://knowledgelayer.softlayer.com/procedure/sap-ibm-bluemix-infrastructure>
- SAP HANA (large instances) high availability and disaster recovery on Azure <https://docs.microsoft.com/en-us/azure/virtual-machines/linux/sap-hana-overview-high-availability-disaster-recovery>
- High Availability of SAP HANA on Azure Virtual Machines (VMs) <https://docs.microsoft.com/en-us/azure/virtual-machines/workloads/sap/sap-hana-high-availability>
- High availability set up in SUSE using the STONITH <https://docs.microsoft.com/en-us/azure/virtual-machines/workloads/sap/ha-setup-with-stonith>
- High availability for SAP NetWeaver on Azure VMs on SUSE Linux Enterprise Server for SAP applications <https://docs.microsoft.com/en-us/azure/virtual-machines/workloads/sap/high-availability-guide-suse>
- Setting Instance Availability Policies on Google Cloud Platform <https://cloud.google.com/compute/docs/instances/setting-instance-scheduling-options#autorestart>
- SAP HANA Deployment Guide on Google Cloud Platform <https://cloud.google.com/solutions/partners/sap/sap-hana-deployment-guide>

- SAP HANA High Availability and Disaster Recovery on Alibaba Cloud
  <https://www.alibabacloud.com/help/doc-detail/57235.htm>