Vol. 10 No. 1 (2019):753-787
DOI: https://doi.org/10.61841/turcomat.v10i1.14805

SAP Failover Setup for HANA Database 1.0 SP12 and SAP ASCS/ERS using Pacemaker Tool

Harikrishna Madathala

Senior SAP BASIS Architect

Abstract

This research paper explores the implementation of a high availability solution for SAP HANA Database 1.0 SP12 and SAP ASCS/ERS (ABAP Central Services/Enqueue Replication Server) using the Pacemaker clustering tool. The study focuses on the architecture, configuration, and best practices for setting up a robust failover mechanism to ensure continuous operation of critical SAP systems. By examining the integration of SAP HANA's system replication capabilities with Pacemaker's cluster management features, this paper provides insights into creating a resilient SAP environment capable of minimizing downtime and maintaining data consistency during failover events.

Keywords: SAP HANA, High Availability, Pacemaker, Failover, System Replication, ASCS, ERS, Cluster Management

1. Introduction

1.1 Overview of SAP HANA Database and ASCS/ERS

SAP HANA is an in memory, column orient database management system developed by SAP SE also referred to as High-Performance Analytic Appliance. SAP HANA 1. 0 This platform is SP12 which was introduced in 2016 and is richer in features as compared to previous release which provides greater performance along with scalability for enterprise level applications. ABAP

Vol. 10 No. 1 (2019):753-787
DOI: https://doi.org/10.61841/turcomat.v10i1.14805

Central Services (ASCS), and Enqueue Replication Server (ERS) are two important technical modules of SAP NetWeaver that take charge of distributing workloads and data synchronization across SAP environment.

1.2 Importance of High Availability in SAP Systems

SAP systems have today become very critical especially in today's environment where businesses are always on 24/7. There are several implications of unplanned downtime, they include; High financial implications, low productivity, and business reputation. HASP solutions for SAP environments are aimed at availability of service and fast restoration in case of hardware failures, software problems or planned maintenances.

1.3 Role of Pacemaker in SAP Failover Structures

Pacemaker is an open-source cluster resource manager which is used for the implementation of High Availability solutions in SAP environments. It offers the means to orchestrate and supervise the resources that are part of clusters such as SAP HANA databases and, ASCS/ERS instances. Such highly integrated failover schemes are easy to accomplish with the help of Pacemaker and allow for non-interruption of mission-critical SAP environments.

Vol. 10 No. 1 (2019):753-787
DOI: https://doi.org/10.61841/turcomat.v10i1.14805

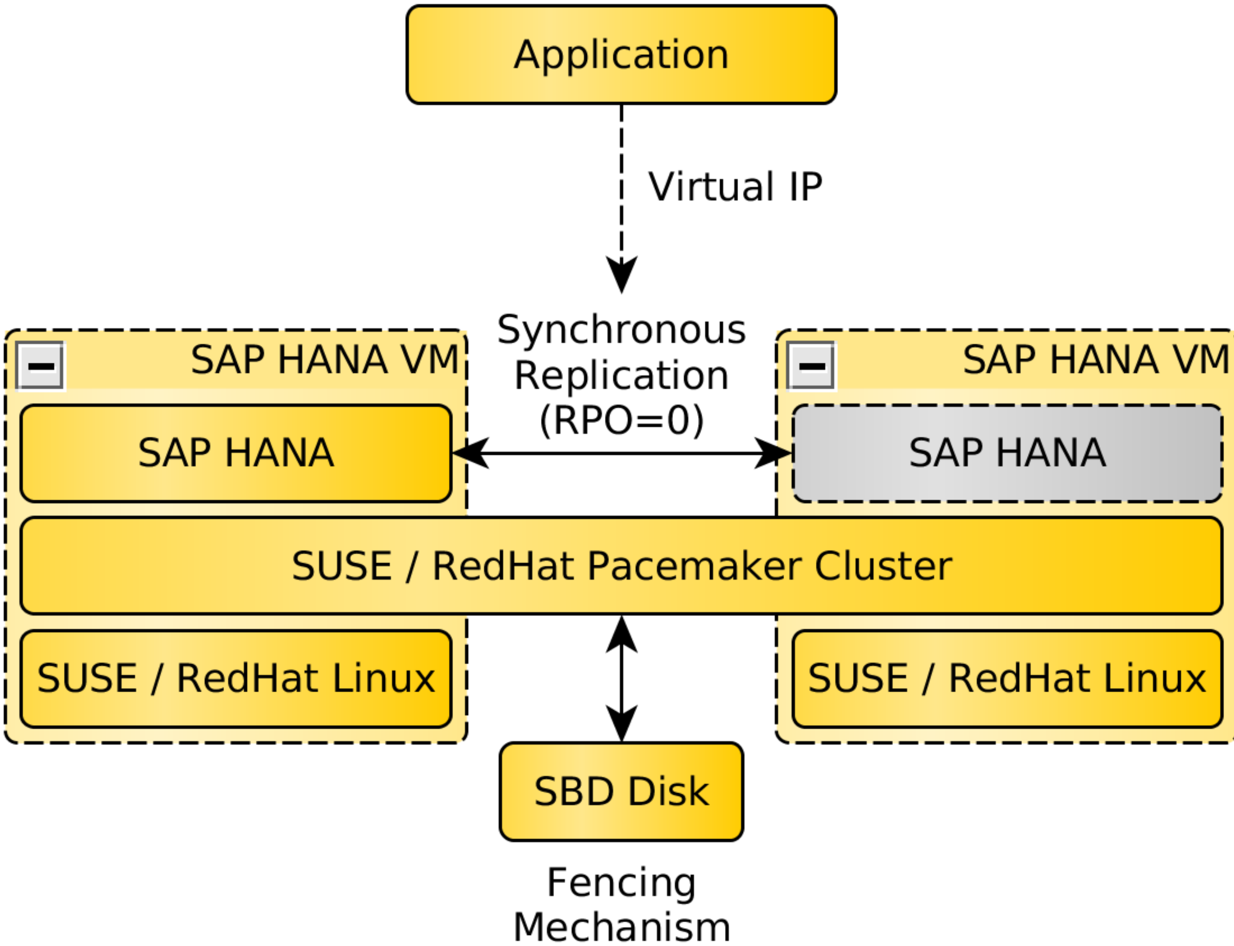

1.4 Research Aims and Objectives

This research aims to:

1. Sap HANA can be analyzed architecturally when we come to the first one as follows Therefore, SAP HANA can be analyzed architecturally when we come to 1 as follows: 0 SP12 and all of its high availability features.
2. Discuss how to configure SAP ASCS/ERS for proper high available solutions.
3. Consider the effectiveness of the utilizing of the Pacemaker tool in the context of SAP failover process.

Vol. 10 No. 1 (2019):753-787
DOI: https://doi.org/10.61841/turcomat.v10i1.14805

4. Offer step by step to conduct a failover setup with Pacemaker for SAP HANA and ASCS/ERS.
5. Exploring best practice, issues and trend of SAP high availability solutions.

2. SAP HANA 1.0 SP12 Architecture

2.1 Basic Modules of SAP HANA 1.0

SAP HANA 1. 0 As mentioned earlier, SP12 released in 2016 is a great advance of SAP in its in-memory database technology. Basically, SAP HANA 1 is a platform for database that has inbuilt analytic abilities which is used for the real-time analysis of large amount of data. 0 SP12 is made up of a number of related parts all of which combine to provide robust data throughput and analysis. The Index Server plays an important role of the nervous system in the HANA framework where it controls all the data tables, processes queries and handles transactions so impressively. This component makes use of column-oriented storage paradigms as well as advanced in-memory query optimizations in order to deliver unprecedented query throughput, especially when it comes to analytical queries.

The organization of the system topology is complemented by the second server, which is the Name Server. It can be described as the system that hold and/or distribute information about the SAP HANA environments in order to allow different components to exchange information while using resources optimally within distributed settings. Extending the physical HANA platform through an integrated application server, the XS Engine further broadens the scope of possibilities by providing native-development capabilities for XS applications that may leverage the full potential of in-memory computing technologies included natively in HANA.

In particular, enterprises that need to handle large amounts of unstructured data, it is possible to use more advanced text analysis and text mining capabilities of the Preprocessor Server. This component uses the advanced natural language processing to achieve the text mining that allows the organizations to get the maximum value of unstructured information. The Statistics Server

Vol. 10 No. 1 (2019):753-787
DOI: https://doi.org/10.61841/turcomat.v10i1.14805

completes the list of the main components of the RAC and is responsible for the collection of the system statistics that can be useful to the system administrators to diagnose the performance and potential problems of the system and its utilization of resources.

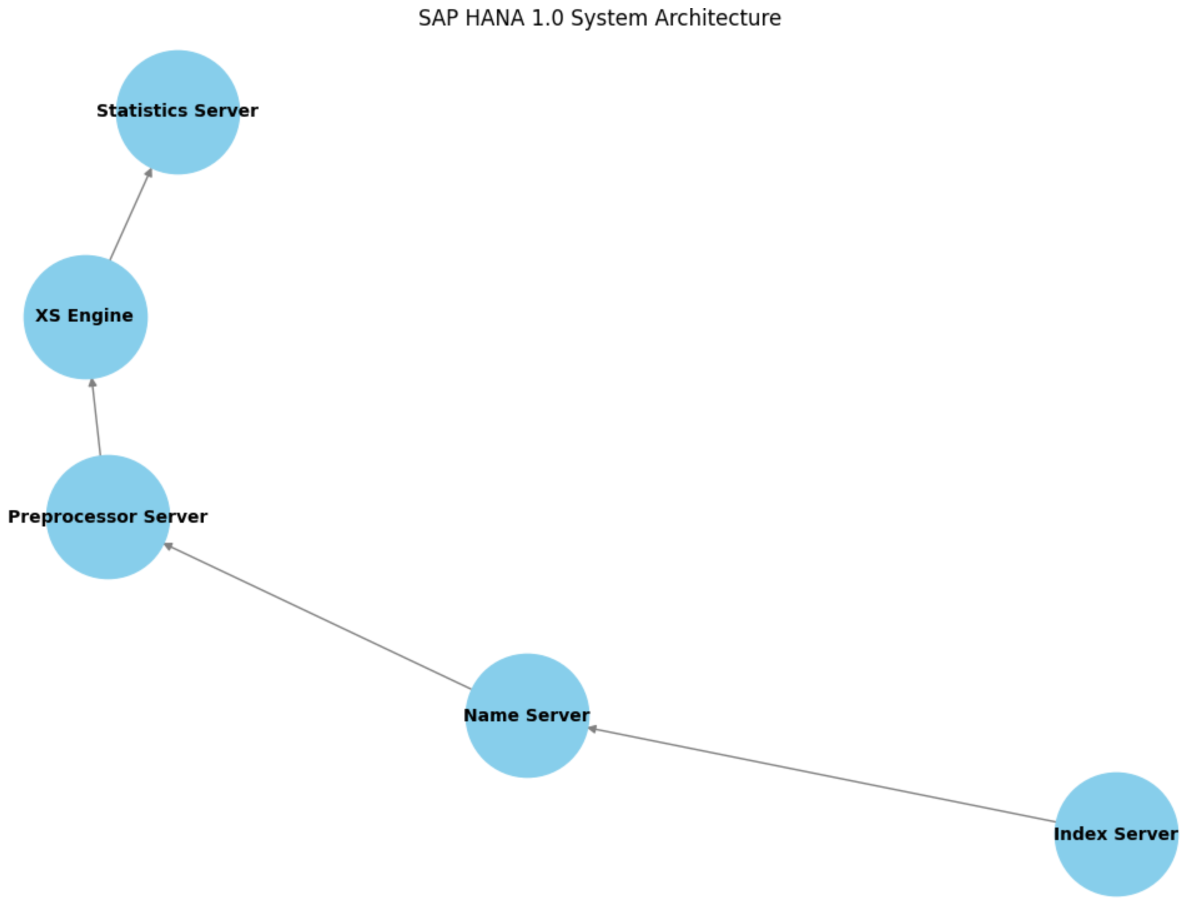

2.2 System Replication Mechanism in SAP HANA

SAP HANA 1. 0 SP12 brings in a major change to its system replication feature which is a critical element for business continuity and disaster recovery. The system replication capability in HANA enables production of a stand-by system which keeps an almost synchronous copy of the primary system. This mechanism is designed to support a range of replication modes that may be unique to a business' needs, or may depend on its infrastructure.

Vol. 10 No. 1 (2019):753-787
DOI: https://doi.org/10.61841/turcomat.v10i1.14805

In synchronous replication mode, changes are made simultaneously to both the primary as well as the secondary systems before the transaction is considered to be through. This approach results to a RPO of 0, but can entail some latency, especially where data replication is between two or more sites. Asynchronous replication, on the other hand, seeks to enhance performance since the primary system has to validate the transactions before they are mirrored in the secondary system. Long-distance replication is perhaps the most beneficial through this mode especially in situations where disaster recovery is required.

SAP HANA 1.0 SP12 also comes with the idea of multi-tier system replication, meaning one can create cascading replication settings. In this case, further replication can be run by a secondary system or distributed in parallel across several systems of a disaster-recovery tandem which can occupy other data centers or regions. The system replication feature is fully aligned with backup and recovery in HANA meaning that the state of different systems replicating the original system

Vol. 10 No. 1 (2019):753-787
DOI: https://doi.org/10.61841/turcomat.v10i1.14805

are readable and recoverable.

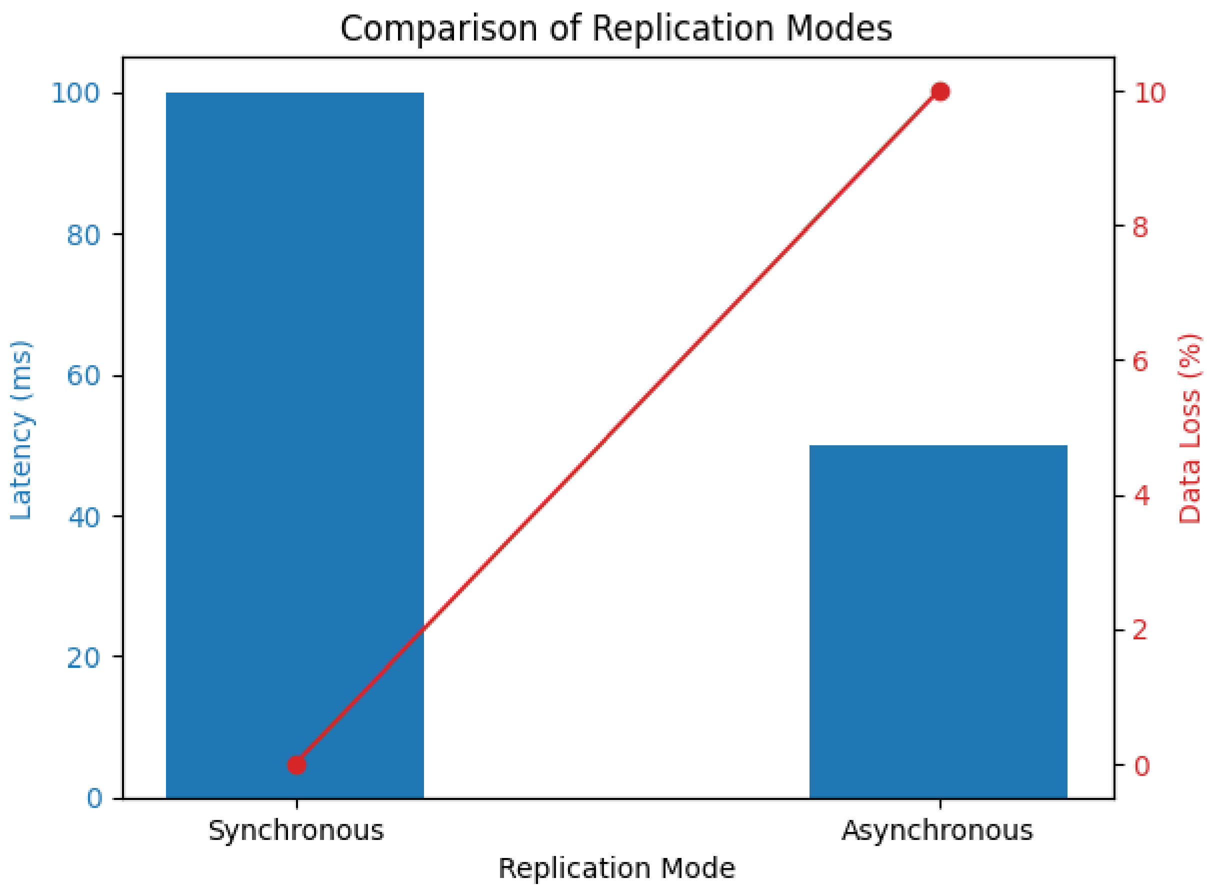

2.3 Availability Characteristics in SAP HANA

Since the modern systems of enterprises are to be highly available, SAP HANA 1. The table below shows that 0 SP12 provides an array of functions that can be used to meet this need fully. Auto-Host Failover feature allow a system to switch from the active host to a redundant host in cases of hardware crash or scheduled for maintenance works. This feature harnesses HANA structural characteristics by distributing workload to further nodes in case something happens to one of them, thus enabling uninterrupted performance of business transactions.

Vol. 10 No. 1 (2019):753-787
DOI: https://doi.org/10.61841/turcomat.v10i1.14805

Currently HANA 1.0 SP12 offers support towards Storage Replication is a plus for data protection since the added feature means compatibility with the storage-level replication technologies. This feature also allows utilization of organizations' current storage infrastructure investments while at the same time improving the high availability of SAP HANA implementations.

SAP HANA 1.0 SP12 also brings improvements to support for disaster recover configurations in multiple tiers. It enables organization to perform multi data center recovery which in turn help in protecting the organization from ruining failure modes and help the organization to continue operating in failure modes. Thus, the disaster recovery is closely linked with HANA, concerning the system replication and back up tools resulting in fast and coherent restoration of a business.

2.4 Brief Presentation of SAP ASCS (ABAP Central Services) and ERS (Enqueue Replication Server)

Though not intrinsically components of SAP HANA architecture, the ABAP Central Services (ASCS) and the Enqueue Replication Server (ERS) are very vital in providing availability in SAP environments that interface with HANA. The ASCS component is handling enqueue services and message servers to organize workload distribution and to guarantee data integrity for SAP landscapes divided across multiple instances. It performs important tasks like lock management and interposes communication and facilitates the operation of SAP applications which are based on the HANA platform.

The Enqueue Replication Server (ERS) supplements the ASCS by being a backup copy of the lock table. In case of failure of ASCS, ERS can easily get to manage the locks thereby reducing the impact of failure on the operation of other business processes. This replication mechanism guarantees the capability by the system to be able to uphold data integrity and consistency in case of failure of certain components or during the times that the components are being serviced.

In SAP HANA 1.0 SP12 environments, the interdependencies of ASCS and ERS with HANA architecture form the basis of implementable high availability. When integrated with HANA these

Vol. 10 No. 1 (2019):753-787
DOI: https://doi.org/10.61841/turcomat.v10i1.14805

components enhance the high availability native capabilities of HANA to deliver a solid platform for business-critical applications. HANA in-memory computing together with the lock administration and message dealing with services of ASCS and ERS consequently provides organizations with superior SAP efficiency and high degree of general availability.

3. Pacemaker Tool: An Overview

3.1 About Pacemaker in High Availability (HA) Clusters

As the key open-source cluster resource manager, Pacemaker was developed to manage high availability solutions in much larger IT contexts. It was primarily raised in response to the necessity for high-performance and compliance-driven clubbing solutions; Over time, it turned out to recognize that cluster administration of some of the most vital kinds of business apps, for instance, SAP HANA along with the ASCS/ERS components, cannot be possible without pacemaker interfacing. It has been designed to support integration with different enterprise systems and to meet the contemporary tendencies in IT environments flexibility.

On a basic level, Pacemaker works by keep a constant check on cluster nodes and the resources that reside in them. It also works in a vigilant manner, permits a quick identification of failures and starts pre-defined procedures to ensure service continuity. The design philosophy behind this tool is based on the policy-based resource management where the administrator is able to set different policies and constraints on how various resources in the cluster behave depending on the conditions that are met.

3.2 Major Components of the Pacemaker

Pacemaker also has a broad number of reservations which are best for increasing the high availability SAP surroundings. Its ability to work with resources does not limit to their startup and shutdown processes, but also includes powerful monitoring and recovery functions critical to each resource's requirements. The fencing mechanisms of the tool, which are very important in

Vol. 10 No. 1 (2019):753-787
DOI: https://doi.org/10.61841/turcomat.v10i1.14805

maintaining the data integrity of the shared storage, afford a high level of protection against the split-brain situation that is very detrimental to system synchronization.

A major value-add in Pacemaker is the constraint-based configuration model it has been developed with. This approach gives the administrators the ability to specify dependencies and constraints between the resources required by a service enabling it to start, stop or migrate to another host at the right time. Another pointed advantage is the ability of tool to scale with the overall size of the cluster, from the formation of a two-node cluster to a dense cluster that might have tens of node.

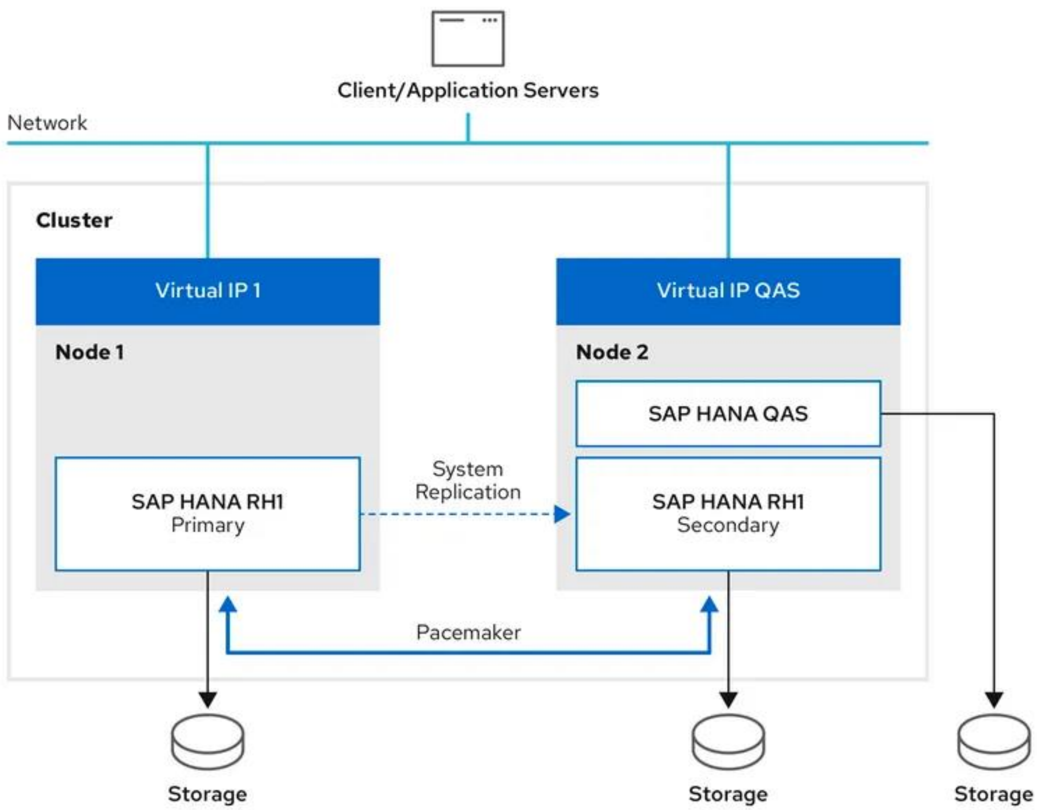

Policy- and resource-wise, Pacemaker is rather versatile, offering integration with a large number of resource agents. These agents work as the interface between the Pacemaker and the managed resources and offer common means of exerting control over the numerous applications and

Vol. 10 No. 1 (2019):753-787
DOI: https://doi.org/10.61841/turcomat.v10i1.14805

services. In SAP environments there are the enhanced developed resource agents able to handle with the HANA database and ASCS/ERS instances.

3.3 Pacemaker and Integration with SAP HANA and ASCS/ERS

There is a fundamental compatibility between PACEMAKER and SAP systems, primarily HANA and ASCS/ERS, and on this, much effort has been directed at improving the compatibility. SAP has officially endorsed Pacemaker for usage in High Availability environment that include these crucial parts stating that it is very reliable and effective when it comes to fail-over situations. This certification also covers different deployment of SAP HANA such as scale up and scale out.

Pacemaker's resource agents for SAP HANA are designed to be used in conjunction with the HANA system replication feature so that the primary and secondary systems may be easily transferred. These agents are capable of monitoring of the health of HANA instances, the promotion of secondary systems to primary and the general failover. In the same manner, for ASCS/ERS setups, there are specific agents that Pacemaker has that Has environment specifics knowledge of the different aspects of this component with regards to managing the virtual IP addresses, File systems, and SAP instance resources.

3.4 Advantages of Using Pacemaker for an SAP Failover

The integration of Pacemaker in handling SAP failover scenarios has many advantages to organizations that require stronger SAP systems' reliability. Among these are the impressive cuts in downtime that is occasioned by automated failover procedures. Design and implementation of a system like Pacemaker means that the health of a system is monitored at all times, and failures escalated quickly so that they do not affect business processes.

Pacemaker has the implementation of general and complex failover strategies adjusted to the particularity of the business case. This also makes it possible for the solution to support many SAP deployment models depending on the needs of a business; whether it will be installed on the local infrastructure or a hybrid or cloud infrastructure. This way, clustered resources are easily managed

Vol. 10 No. 1 (2019):753-787
DOI: https://doi.org/10.61841/turcomat.v10i1.14805

through a single point of contact all of which are governed by Pacemaker's centralized management approach. Moreover, it can also be said that with the integration of Pacemaker, the failover policies of the overall SAP landscape are more reliable. This help cut instances of human error when performing manual failover exercises and gives assurance that all the components of the SAP environment are being run in line with the laid down best practices.

4. SAP HANA and ASCS/ERS Failover Strategies

4.1 Understanding Failover & Fallback Mechanisms

When using SAP HANA or ASCS/ERS high availability configuration, failover and fallback are the basic foundation that is used to guarantee availability of the SAP services. Failover means the transfer of the system or its portion to redundant or standby subsystem, element or component whenever the latter stops functioning. This transition is supposed to be smooth in a way that it either creates or eliminates the do-not-disturb time or period for the end-users as well as business processes. On the other hand, fallback entails the intentional shifting back to the main primary system as soon as the system has been rectified fully.

The various aspects that need to be taken into consideration when considering failover procedures regulate chiefly by far data coherency, network delays and application specificity in terms of its behavior during the failover. In SAP HANA failover means making the secondary node a primary node, which is a very sensitive operation that has to be carried out promptly and with no data loss. While ASCS/ERS failover is more concerned with maintaining the availability of primary SAP services through handing over their control to the active stand-by instances/nodes.

4.2 SAP HANA Failover with System Replication

SAP HANA failover using system replication can be configured in several steps that when consolidated give a strong replication connection between the principal and the backup HANA nodes. This setup is often set up to synchronizes the secondary system to pulling and integrating

Vol. 10 No. 1 (2019):753-787
DOI: https://doi.org/10.61841/turcomat.v10i1.14805

data changes from the primary system every few seconds – thus creating a near real-time copy of the database. Pacemaker has a great importance in this configuration to always check the health of these two systems, and facilitate the failover in case of issues.

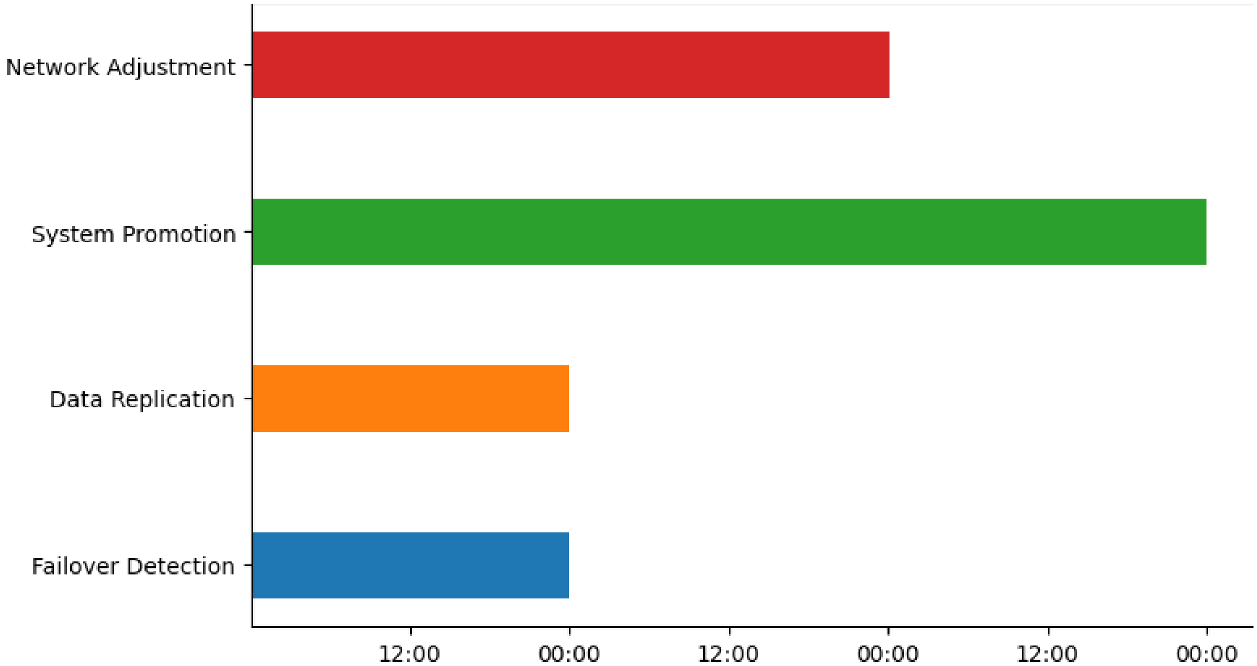

Decisions to make in this configuration include identification of the correct replication mode, for instance, synchronous or asynchronous replication depending on objects such as distance in geography and network lag. Although synchronous replication grants zero data loss, performance can be affected especially when the distance between nodes is large while asynchronous replication is the reverse; often introducing minimal data loss but with superior performance. The second example shows that in this setup Pacemaker is also responsible for the promotion of the secondary system to the primary role during the failover which usually includes many complex actions such as swap replication directions and changes connection details of the clients.

Example Pacemaker configuration for SAP HANA:

Vol. 10 No. 1 (2019):753-787
DOI: https://doi.org/10.61841/turcomat.v10i1.14805

```
<clone id="SAPHanaTopology_HANACluster_HDB00" multi_state="false">
  <primitive class="ocf" id="SAPHanaTopology_HANACluster_HDB00" provider
      ="heartbeat" type="SAPHanaTopology">
    <instance_attributes id="SAPHanaTopology_HANACluster_HDB00
        -instance_attributes">
      <nvpair id="SAPHanaTopology_HANACluster_HDB00-instance_attributes-SID"
          name="SID" value="HDB"/>
      <nvpair id="SAPHanaTopology_HANACluster_HDB00-instance_attributes
          -InstanceNumber" name="InstanceNumber" value="00"/>
    </instance_attributes>
    <operations id="SAPHanaTopology_HANACluster_HDB00-operations">
      <op id="SAPHanaTopology_HANACluster_HDB00-operations-monitor" interval
          ="10" name="monitor" timeout="600"/>
      <op id="SAPHanaTopology_HANACluster_HDB00-operations-start" interval
          ="0" name="start" timeout="600"/>
      <op id="SAPHanaTopology_HANACluster_HDB00-operations-stop" interval
          ="0" name="stop" timeout="300"/>
    </operations>
  </primitive>
</clone>
```

4.3 SAP ASCS/ERS Failover with Pacemaker Configuration

This guide examines how to configure ASCS/ERS failover with Pacemaker with the aim of making these SAP services always available. Such configuration normally entails running of ASCS and ERS instances on different nodes in a cluster that is managed by Pacemaker. As part of the configuration client-side virtual IP addresses are configured and the shared file systems required for SAP resources are established, in addition, Pacemaker is used to configure the above components as cluster resources.

Pacemaker's resource agents for ASCS and ERS are instrumental in this configuration, providing the necessary logic to start, stop, and monitor these SAP components. The configuration also involves setting up constraints within Pacemaker to ensure that ASCS and ERS instances are

Vol. 10 No. 1 (2019):753-787
DOI: https://doi.org/10.61841/turcomat.v10i1.14805

always running on separate nodes, a critical requirement for maintaining service availability in the event of node failures.

Example Pacemaker configuration for ASCS/ERS:

```
<group id="g-ASCS">
  <primitive class="ocf" id="rsc_ip_ASCS" provider="heartbeat" type="IPaddr2">
    <instance_attributes id="rsc_ip_ASCS-instance_attributes">
      <nvpair id="rsc_ip_ASCS-instance_attributes-ip" name="ip" value="192.168.1.10"
          />
    </instance_attributes>
  </primitive>
  <primitive class="ocf" id="rsc_fs_ASCS" provider="heartbeat" type="Filesystem">
    <instance_attributes id="rsc_fs_ASCS-instance_attributes">
      <nvpair id="rsc_fs_ASCS-instance_attributes-device" name="device" value="/dev
          /vg_sap/lv_ASCS"/>
      <nvpair id="rsc_fs_ASCS-instance_attributes-directory" name="directory" value
          ="/usr/sap/PRD/ASCS00"/>
      <nvpair id="rsc_fs_ASCS-instance_attributes-fstype" name="fstype" value="xfs"
          />
    </instance_attributes>
  </primitive>
  <primitive class="ocf" id="rsc_sap_ASCS" provider="heartbeat" type="SAPInstance">
    <instance_attributes id="rsc_sap_ASCS-instance_attributes">
      <nvpair id="rsc_sap_ASCS-instance_attributes-InstanceName" name="InstanceName"
          value="PRD_ASCS00_ascsvm"/>
    </instance_attributes>
  </primitive>
</group>
```

4.4 Data Consistency Challenges During Failover

Therefore, managing data synchronization during failover events represents the key challenge in SAP environments to keep data accurate and guard the integrity of business processes. For SAP HANA this entails the use of such features as synchronous replication so that no data is lost in case

Vol. 10 No. 1 (2019):753-787
DOI: https://doi.org/10.61841/turcomat.v10i1.14805

of a failover. In making for consistency, Pacemaker has appropriate fencing mechanisms which STONITH: Shoot the Other Node in The Head means that several nodes are likely to attempt to take the primary role at once since during split-brain, more than one node will think that it is possible to become the stand-in for the pacemaker.

As in any well-designed ASCS/ERS setups, there is always a strict enforcement of consistency in the replication process of the enqueue. Pacemaker makes sure that failover of ASCS instances is synchronized with ERS to make copy of lock table always up-to-date. This is very important to maintain the consistency of the on-going transactions and to avoid the problems that result from inconsistent lock information on each node.

5. Pacemaker Configuration with SAP HANA and ASCS/EPR Failover

5. 1 Required Pacemaker Packages and Dependencies

When installing the solution based on Pacemaker, for SAP HANA and ASCS/ERS failover it is necessary to have a set of packages and dependencies on them that are chosen by selecting specific parameters with high priority. Generally, Pacemaker consists of the Pacemaker cluster manager and Corosync cluster engine, which is used to provide low level node-to-node communication, and a set of resource agents intended for controlling SAP elements. Other packages can be system-specific enhancements: for example, sap-sues-cluster-connector allows SUSE Linux Enterprise servers to strengthen the connection with SAP applications and the clustering system. It’s not only sufficient to rely on dependencies of the software packages, but also the systems configuration and network settings. For example, the correct synchronization of time in all the cluster nodes is important in their proper functioning and correct logging activity. Likewise, the firewall settings have to be changed to enable cluster communications and traffic for resource management while maintaining required level of security to the cluster.

5. 2 The Cluster Environment

Vol. 10 No. 1 (2019):753-787
DOI: https://doi.org/10.61841/turcomat.v10i1.14805

Building the solid cluster environment constitutes the basis of the failover based on the Pacemaker software. Then, the initial setup of the Pacemaker and Corosync is performed on all nodes to be used in a cluster. Configuration of Corosync is all the more important; it outlines the paths and protocols by which nodes in a cluster perceive the state of the cluster.

The cluster initialization process involves defining the cluster name, setting up authentication mechanisms to secure inter-node communication, and establishing quorum settings to prevent split-brain scenarios in multi-node clusters. For SAP HANA and ASCS/ERS setups, special consideration must be given to configuring appropriate timeout values and failure detection mechanisms to align with SAP's specific requirements for system responsiveness and failover timing.

Vol. 10 No. 1 (2019):753-787
DOI: https://doi.org/10.61841/turcomat.v10i1.14805

```
totem {
    version: 2
    cluster_name: sap_cluster
    transport: udpu
    interface {
        ringnumber: 0
        bindnetaddr: 192.168.1.0
        mcastport: 5405
    }
}

nodelist {
    node {
        ring0_addr: 192.168.1.1
        nodeid: 1
    }
    node {
        ring0_addr: 192.168.1.2
        nodeid: 2
    }
}

quorum {
    provider: corosync_votequorum
    two_node: 1
}
```

Vol. 10 No. 1 (2019):753-787
DOI: https://doi.org/10.61841/turcomat.v10i1.14805

5. 3 Buffering Fencing Mechanisms for Node Recovery

Fencing also known as STONITH: Shoot the Other Node in The Head is very important for any cluster especially where there is Shared Storage involved. In SAP HANA and ASCS/ERS setups, correct fencing makes sure that failed or unresponsive node is out of the cluster to prevent the recovery or access of its resource on another node. This mechanism is very important because corrupting the data is the easiest way to compromise the SAP environment. Pacemaker currently supports a number of fencing types some of which are: IPMI fencing which is a hardware level fencing and SBD (STONITH Block Device) which provides software-based fencing. Decision about which type of fencing should be implemented is based on the underlying hardware support and particularities of SAP configuration. Fencing devices configuration in Pacemaker include declaration of the fencing resources, setting the ways of the access to the fencing devices and control of the fencing mechanisms, and the monitoring of the fencing mechanisms which ensures that the fencing mechanisms are functioning properly.

```
pcs stonith create fence_ipmi_node1 fence_ipmilan pcmk_host_list="node1" ipaddr="192
    .168.1.101" login="admin" passwd="password" op monitor interval="60s"
pcs stonith create fence_ipmi_node2 fence_ipmilan pcmk_host_list="node2" ipaddr="192
    .168.1.102" login="admin" passwd="password" op monitor interval="60s"
```

5.4 SAP HANA and ASCS/ERS certification of Pacemaker

SAP HANA as well as ASCS/ERS components are configured as part of Pacemaker cluster which involves defining them as manageable resources in the cluster environment. For SAP HANA, this often means setting up resources for the SAP HANA topology, an object that controls the replication status and roles of HANA instances, as well as the SAP HANA database instances. These resources use dedicated SAP HANA resource agents that are familiar with the HANA processes including replication and instances.

Vol. 10 No. 1 (2019):753-787
DOI: https://doi.org/10.61841/turcomat.v10i1.14805

Like with integration of ASCS and ERS, resources exist for virtual IP addresses, file systems and the instances of SAP. The configuration should make sure that each ASCS instance and ERS instance is located in two different nodes for the sake of high availability. Such setup usually requires making of resource groups as well as defining ordering when it comes to starting and stopping resources.

```
pcs resource create SAPHanaTopology SAPHanaTopology SID=HDB InstanceNumber=00 op
    start timeout=600 stop timeout=300 monitor interval=10 timeout=600

pcs resource create SAPHana SAPHana SID=HDB InstanceNumber=00 PREFER_SITE_TAKEOVER
    =true DUPLICATE_PRIMARY_TIMEOUT=7200 AUTOMATED_REGISTER=false op start timeout
    =3600 stop timeout=3600 monitor interval=60 role=Master timeout=700 op monitor
    interval=61 role=Slave timeout=700

pcs resource create vip_ASCS IPaddr2 ip=192.168.1.10
pcs resource create fs_ASCS Filesystem device="/dev/vg_sap/lv_ASCS" directory="/usr
    /sap/PRD/ASCS00" fstype="xfs"
pcs resource create sap_ASCS SAPInstance InstanceName=PRD_ASCS00_ascsvm
```

5.5 Setting Up Resource Constraints and Dependencies\

Specifically, resource constraints and dependencies fix the behavior of SAP components inside of Pacemaker cluster. These configurations make it possible to have the right placement of resources in the available nodes and also have a good way of handling resources that are dependent on each other during failure and recovery.

General constraints usually include co-location constraints, which specify that the primary and the secondary SAP HANA instances cannot reside on the same nodes and ordering constraints that define the order in which the actions have to be taken during failure situations. In ASCS/ERS setups, constraints are given to define the limitation of ASCS and ERS instance distribution between several nodes besides the control of the relations between virtual IP, file systems and SAP instance resources. Pacemaker has a constraint system for the setting of complex

Vol. 10 No. 1 (2019):753-787
DOI: https://doi.org/10.61841/turcomat.v10i1.14805

configurations, which can work with many failure chicks to keep SAP environment highly performant and available.

```
pcs constraint colocation add SAPHana with SAPHanaTopology 2000
pcs constraint order SAPHanaTopology then SAPHana

pcs resource group add g-ASCS vip_ASCS fs_ASCS sap_ASCS
pcs constraint colocation add g-ASCS with SAPHana
```

5.6 Testing and Validating the Failover Configuration

To maintain the reliability and effectiveness of the high availability setup, it should be executed after passing different test and validation procedures. This process entails emulating different forms of failures on nodes, in partitions as well as resource-specific failures to be certain that the corresponding cluster behaves as expected with respect to resultant service availability. SAP HANA testing should include the areas like primary instance failure, network failure in replication environment and also planned maintenance where controlled failover is desired. The primary goal of ASCS/ERS failover testing is to check if services are switched between nodes properly, if clients' connections are being continued or if they are recovered quickly and whether the enqueue replication works properly in failover scenario or not. Thorough testing also involves ensuring that the behavior of the cluster transforms back to the initial state when the failure conditions are corrected In Simple terms, stress testing entails indicate how systems can gracefully roll back to specific states in case the failure conditions resolve.

6. Pacemaker: Techniques for SAP Failover

6.1 Performance Optimization for SAP HANA Failover

Vol. 10 No. 1 (2019):753-787
DOI: https://doi.org/10.61841/turcomat.v10i1.14805

The main steps which can be taken in order to optimize the performance of SAP HANA failover in the environment managed by Pacemaker include the following: Some are the mirroring parameters that have to be customized to reduce latency but at the same time maintaining data synchronization, the network settings to enable efficient data transfer from the primary to the secondary servers, and balancing the Pacemaker resource agents' parameters to give a responsive configuration with stability for the system.

Optimizations in this area also include the appropriate calculation of the needed amount of system resources and the proper distribution of basic and additional HANA instances which should be capable enough to successfully meet the maximum demands to the CPU, memory utilization and I/O operations. Some fine-tuning of the configurations with the help of SAP HANA as well as general system monitoring activities using Pacemaker commands which helps to review the current state and characteristics of active and standy replication members, should be performed periodically to ensure that long-term failover performance remains optimal

6.2 Maximum Availability for SAP ASCS/ERS Failover

Keeping the downtime during ASCS/ERS failover to the minimum is paramount, with respect to SAP applications and business process in question. The efficiency enhancement in this area is directed to the optimization of the failover process itself, such as fine tuning of the start and stop timeouts of resources, usage of an effective health check mechanism and an optimization of the order of execution of the used resources during a failover. By insisting on utilizing such elements of Pacemaker resources stickiness and migration thresholds gesture which aimed at minimizing unnecessary fails and thus causing the reduction of the service transition frequency on the whole. Another important aspect regarding the reduction of the time spent in being offline is the proper techniques for clients' reconnection. As for the AS ABAP configuration, it is needed to set SAP systems and application servers to recognize ASCS/ERS failovers and change client connection routes towards the new active instances. Essentials in this process include Virtual IP addresses'

Vol. 10 No. 1 (2019):753-787
DOI: https://doi.org/10.61841/turcomat.v10i1.14805

deployment as well as the accurate mechanism for updating the Domain Naming System.

6.3 Monitoring and Troubleshooting of Pacemaker Clusters

Thus, monitoring and troubleshooting help to keep Pacemaker-managed SAP environments in optimal condition. In this area practices that are considered good are the usage of cluster logging and monitoring systems as well as localization of the informational is on the cluster activity and the state of the specific SAP components. Common tool which can be found in Pacemaker and other SAP's tools must be combined in order to provide an adequate picture of the state of the environment.

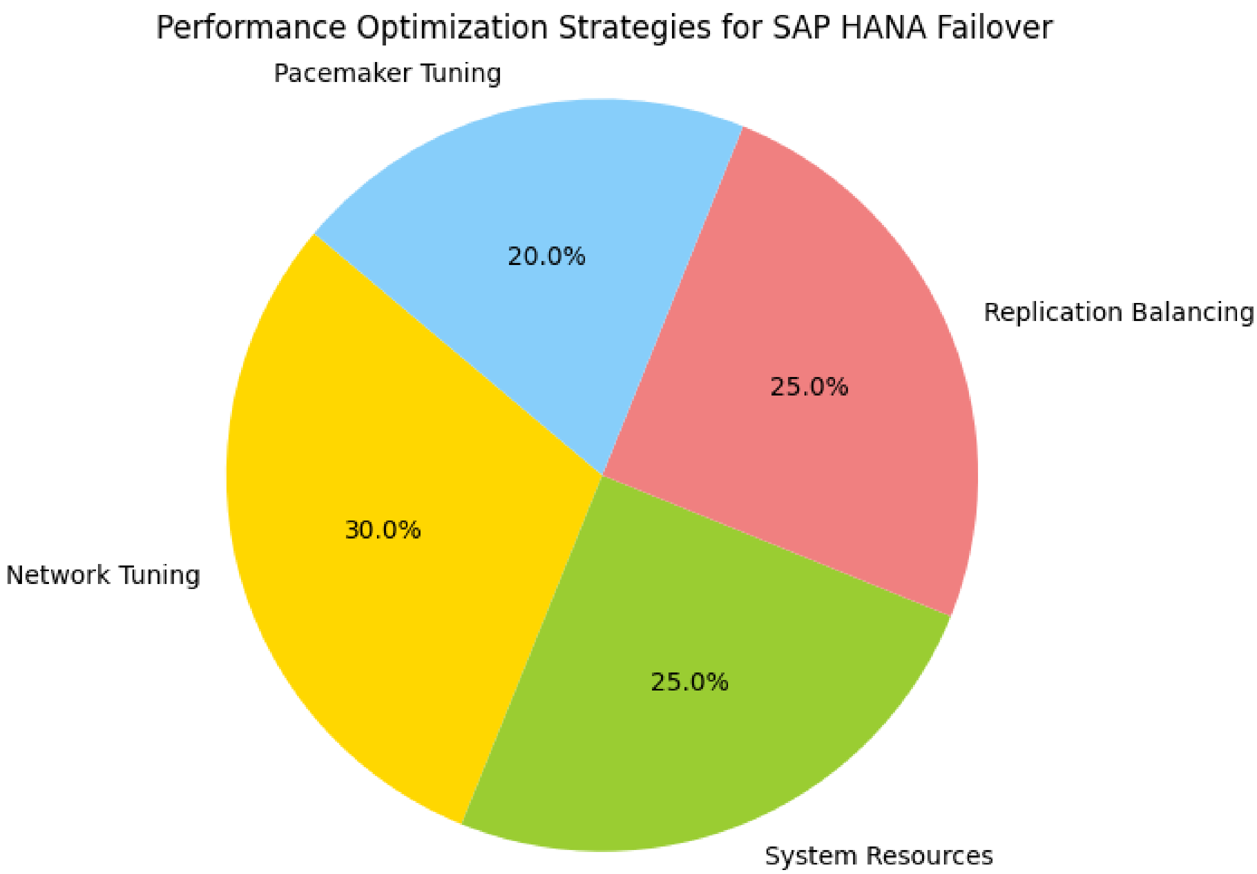

The use of monitoring techniques that allows one to predict problems such as failed resources and nodes, communication issues as well as reaching the levels of capacity make it easy to solve

Vol. 10 No. 1 (2019):753-787
DOI: https://doi.org/10.61841/turcomat.v10i1.14805

problems before they affect service availability. Detailing and documenting standard and referred troubleshooting procedures which are basically checklists on how to handle failure cases is very important as it enables quick action to be made the moment a failure occurs.

6.4 Security Concerns to Take When Setting Up the Failover

SAP failover involves security considerations especially because of the role played by the data and process involved. Each of these areas of best practice covers a broad area of considerations, some of which are using encryption and authentications on communication channels of the cluster as well as have good access control policies on the interfaces used to manage the cluster.

Particular attention should be used to protect the fencing mechanisms, as they have a significant input in the integrity of data during the failover scenarios. This is through establishing measures for securing the fencing devices intercommunications as well as conducting routine audits and reconfigurations of fencing. Further, the coordination with other safety mechanisms of an organization including the SIEM which is generally known as Security Information and Event Management systems allows in having overall safety picture of the cluster SAP.

7. Implementation issues and possible solutions in SAP failover configuration

7.1 Some of the Major Complexities in SAP HANA Failover executions

SAP HANA failover solution has a number of complexities that organizations face to have reliable high availability system. Another problem is data synchronization between the primary and secondary systems which can be problematic in the cases when the systems are geographically distributed and network synchronization of the Secondary System with the Primary System can be slow due to network latency. This challenge is made worse by the fact that data integrity must be achieved in the face of performance considerations, which arise out of stringent consistency demands on the system.

The second important issue that may be considered relevant at the moment is related to the complexity of SAP HANA landscapes and the ability to manage them, including in cases where it

Vol. 10 No. 1 (2019):753-787
DOI: https://doi.org/10.61841/turcomat.v10i1.14805

is required to address several tenants or complex distributed environments. By its very nature, regular and predictable failover behavior within such environments proves challenging to assure across all the various types of components. Also, managing SAP HANA failover with existing disaster recovery and business continuity management plans may involve a number of issues, which need to be solved in terms of risk management and system reliability.

7. 2 Managing Pacemaker and SAP Resource Failure

This is true for failures concerning the Pacemaker components as well as failures of SAP-specific resources which complicates the management of high availability. Principal concerns are false positive fail detection, where there will be instigation of a failover even when it is not required, and resource start / stop problems in failovers. Mitigating these challenges requires working on the following areas: enhancement of the failure identification of failure in Pacemaker, improvement of the resource monitoring model, and designing of the failure recovery strategies in different failure contexts.

These challenges could be solved by using some of the features of Pacemaker such as the resource operation history, failure tracking and other features that would enable one to make better decisions about the utilization of the resources in question. The management decrees staged recovery where the resources are grouped in stages and recovered controlled with some checks made at each stage

Vol. 10 No. 1 (2019):753-787
DOI: https://doi.org/10.61841/turcomat.v10i1.14805

as one of a solution to reducing effects of complicated failure scenarios.

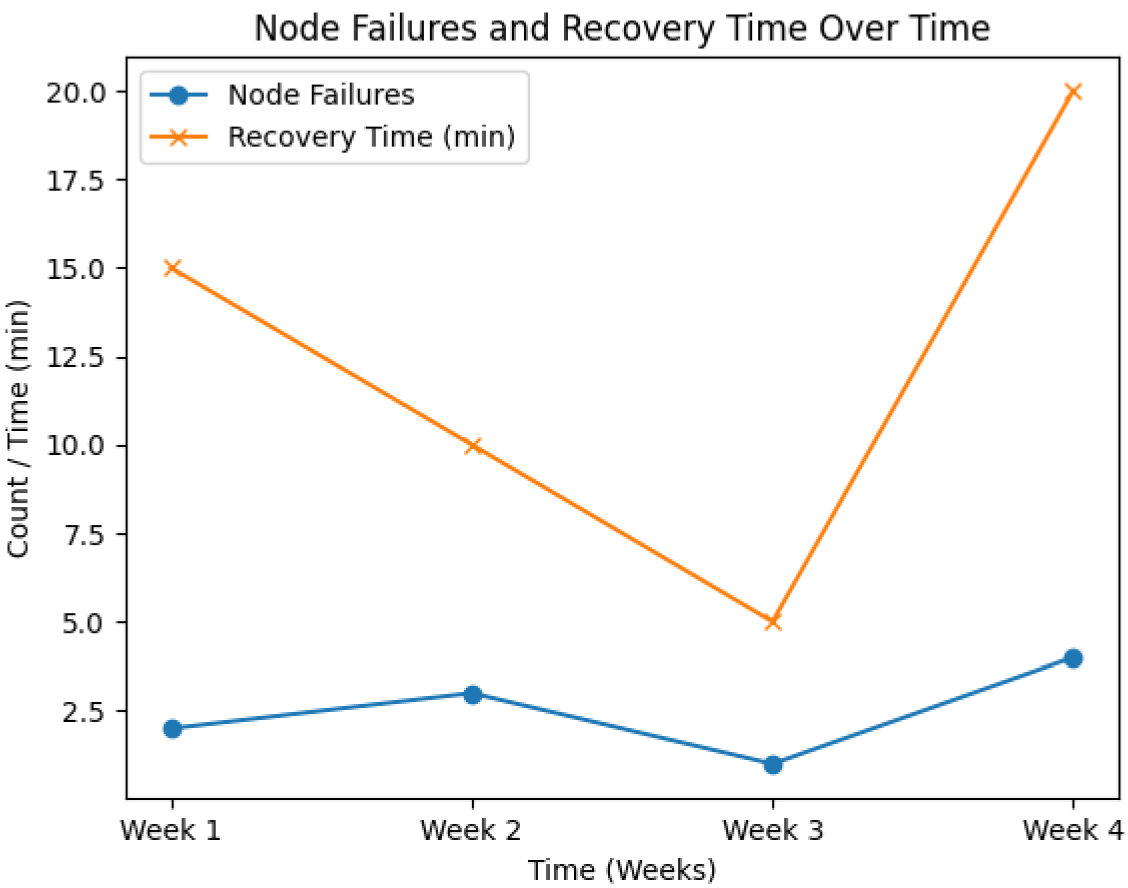

7.3 Tips for Solving High Availability Problems with SAP ASCS/ERS

SAP ASCS/ERS Setups usually involve High Availability issues, which demand a blend of understanding of SAP application and of Pacemaker cluster. In this area, the following problems are usually found: Enqueue replication does not work correctly; virtual IP addresses; shared file system access issue. These problems may occur in different levels ranging from marginal performances that slightly degrade the performance of the equipment's to total system failure in providing the services.

Some Reasoned approaches include the following ones; * Aproactive approach, this is, for example, monitoring the system where the problems are likely to occur * Log analysis, this is for

Vol. 10 No. 1 (2019):753-787
DOI: https://doi.org/10.61841/turcomat.v10i1.14805

instance using log files to identify where the problem could be arising from * Systematical Method, this is for example isolating a problem that is being experienced on the system to the root case. High-quality logging functionality should be applied to both the work of Pacemaker and the SAP processes themselves; this will help to identify the problem and respond to it as quickly as possible. It's thus necessary and helpful to manage common issues and solutions as knowledge items so that the same knowledge can be applied to similar problems in an efficient manner in production settings and with reduced mean time to recovery (MTTR). Moreover, using specific SAP support structures such as forums and webinars could go a long way in identifying how to approach complicated high availability problems.

7.4 Performance Consequences of Failover on SAP HANA Operations

One of the major challenges noticeable in high availability configurations is the effect of failover instances on the SAP HANA processes involved. Some of the problems in this area include how to move workloads from the primary system to a secondary system without interruption to normal transactions of analysis. Time which is needed for the promotion of a secondary system to the primary one, including operations like switching of replication directions and updating of the connection information for the clients, may result in short-term outages or decreased throughput. Managing these performance effects thus entails a number of technical improvements alongside a number of operational measures. On the technical side this involves fine tuning the system replication configuration to allow for fast data transfer, the configuration of the log shipping mechanisms to be efficient and last but not least tuning of the Pacemaker resource agent to ensure that failover periods are minimized. Functionally, some of the measures include, instituting application level retries, optimizing client connection management, ensuring that client-side failover is clearly understood by all concerned parties to reduce its disruption to business functions and usability to the end user.

8. Trends for SAP HANA and High Availability in the Future

Vol. 10 No. 1 (2019):753-787
DOI: https://doi.org/10.61841/turcomat.v10i1.14805

8.1 New Features in the SAP HANA 2.0 and Failover Technologies

The change in SAP HANA technology especially with the coming of SAP HANA 2. 0, opens new opportunities as well as new problems in the area of high availability and failover. SAP HANA 2. 0 brings new options for system replication with added functionality for replication of systems belonging to different tiers and more flexible failover procedures. They allow the means to build more intricate and fault tolerant high availability structures that are more likely to lower down time and increase overall system effectiveness.

The future trend in this area can be expected to be the further continuation of automating different failover processes to use machine learning and predictive analytics for failing over without waiting for a failure to occur. Moreover, SAP HANA's connectivity with cloud and container orchestrations such as Kubernetes for container orchestration are likely to bring new ideas about high availability management, and provide more detailed configuration of resource distribution and failover schemes.

8.2 The Role of Cloud-Based High Availability Solutions

As time proceeds, cloud computing becomes more and more important in SAP HANA, and it has the same prospect put in high availability solutions. Solutions from cloud platforms are distinguished by their flexibility for performing highly effective failover solutions, especially the speed of acquiring computing resources in different availability zones or regions. The above capability helps to build SAP HANA environments that are extremely robust for even large-scale infrastructural disaster.

The future of cloud high availability and disaster recovery of SAP HANA will involve enhancements in integration of on premise and cloud solutions to form the best of both worlds. Some of the scenarios for using multi-cloud failover may involve failing over SAP HANA instances to another cloud platform and not just to another cloud provider's data center.

8.3 Automation of Failover Processes using AI and Machine Learning

Vol. 10 No. 1 (2019):753-787
DOI: https://doi.org/10.61841/turcomat.v10i1.14805

It is noteworthy that, the implementation of AI and ML for SAP HANA failover extensification bears high potential for advanced high availability administration. These technologies have the following promise to increase the intelligence and self-governance of failover systems and go to the next level from simple rule-based management of high availability:

One can predict the following trends in this field: in the future, AI-based systems will be able to learn from failover experience and behaviors of systems in order to implement actions to failover more effectively. Additionally, AI-powered predictive maintenance capabilities could help identify potential hardware or software issues before they lead to failures, enabling proactive interventions to maintain system stability.

8.4 Predictive Failover Techniques in Future SAP Systems

Predictive failover can be classified as the direction of the development of HA solutions for SAP incarnation as promising. This approach is different from ordinary failover mechanisms whose alert is triggered by existing system failure and seeks to address it then and there, this strategy has a proactive approach, and its health check mechanism is designed to check and possibly avoid failure. There are numerous forms of failover techniques, with predictive failover techniques relying on complex analysis, monitoring, and machine learning to determine problems before they occur.

The future work in this context may also consist in extending predictive failover right into the SAP HANA and similar systems where proactive decisions could be made to trigger controlled failovers or other resource redistribution initiatives based on expected state observations. This could include such factors as low-level hardware log information, tens of system level metrics or higher-level application-level metrics to develop rich models of system behavior and use it as a baseline to recognize symptoms that lead to failures.

9. Conclusion

Vol. 10 No. 1 (2019):753-787
DOI: https://doi.org/10.61841/turcomat.v10i1.14805

9. 1 Summary of Key Findings

This research has looked into complicated issue of SAP failover setups in relation to HANA Database 1. In version 0 SP12 and SAP ASCS/ERS with tools such as the Pacemaker tool. Some of the findings key concerns revolve around failover mechanisms as a fundamental necessity for any organization with SAP environments as they are a necessity for the overall continuity of the business and a direct help for the avoidance of downtime. It has demonstrated that the Pacemaker with SAP HANA and ASCS/ERS components can be contributed to offering the flexible and powerful infrastructure to control the high availability which can cover the needs of the contemporary enterprise systems.

The strategy for choosing and deploying the failover solutions has next shown that it is crucial to take several aspects into account and to distinguish between various behaviors of SAP systems and the principles of managing a Pacemaker cluster. Some of the challenges established especially in aspects of data consistency as well as other intricate failure scenarios contribute to the fact that high availability setups require better management solutions and techniques in the future.

9.2 Recommendations for Improving SAP Failover Configuration

Based on the findings of this research, several recommendations can be made for optimizing SAP failover setups:

1. Provide a thorough test plan that includes a number of failures that should turn the failover setup on to maximize its reliability.
2. IT staff should be provided with the opportunity to undergo formal training to update their knowledge in both, SAP systems and Pacemaker cluster management. It is necessary to review and modify failover configurations from time to time in order to suit the changing business needs or development in technology.
3. Use techniques such as monitoring and PdM to pick on problematic conditions that may affect the system and its availability.

Vol. 10 No. 1 (2019):753-787
DOI: https://doi.org/10.61841/turcomat.v10i1.14805

4. Cloud-based or hybrid architectures should be considered to take advantage of abilities to configure sophisticated failover systems and enhance the general system stability.
5. Be aware of the new trends in Artificial Intelligence and Machine learning that can be utilized for the enhancement of failover mechanisms and should be researched specifically for their implementation on SAP systems.

9.3 Future Research Directions

There are many directions for further research in the field of SAP technologies and high availability solutions due to the fact that the development pace is very high. Potential areas of investigation include:

1. Consequently, the effect of SAP HANA 2.0 and these and other subsequent releases on failover strategies and high availability architectures.
2. Non-supervised and supervised machine learning, and AI specifically in identifying and improving failover methodologies and cyclical system failures in environments that utilize SAP.
3. Research on premise option and cloud-based option and such comparative studies including hybrid failover facility for SAP systems in terms of efficiency, cost, reliability.
4. Container technologies and discussion with microservices architectures and SAP failover and hearing and availability management new paradigms.
5. More long-term research work to identify the consequences of the failover strategies on performance and reliability of giant SAP environments.
6. The creation of framework-supported parameters for the actual assessment of the quality of failover solutions in SAP scenarios.

Thus, it can be stated that while the topic of failover for SAP systems has been actively researched and developed throughout recent years, it is still a rather vibrant area. Further investigations in this

Vol. 10 No. 1 (2019):753-787
DOI: https://doi.org/10.61841/turcomat.v10i1.14805

sphere will be critical for responding to new problems associated with sustaining high availability in complicated and important enterprise infrastructures.

Vol. 10 No. 1 (2019):753-787
DOI: https://doi.org/10.61841/turcomat.v10i1.14805

References

1. Bartolini, C., Giurgiu, I., Tauber, M., & Wiesmann, M. (2017). Automated Application Recovery in the Cloud: Empirical Evaluation of SAP HANA on AWS. In 2017 IEEE 10th International Conference on Cloud Computing (CLOUD) (pp. 252-259). IEEE.
2. Beine, M., Herschel, R., & Castellanos, M. (2016). Enabling Real-Time Business Intelligence on SAP HANA. In Data Management Technologies and Applications, 201-221. Springer.
3. Bog, A., & Plattner, H. (2015). Interactive performance management for SAP HANA. In Proceedings of the 2015 ACM SIGMOD International Conference on Management of Data, 1119-1124.
4. Cheng, B., & Zhao, S. (2017). Pacemaker-based High Availability Cluster Implementation in OpenStack. In Proceedings of the International Conference on Computer Network, Electronic and Automation (ICCNEA), 164-168.
5. Daidone, A., Di Giandomenico, F., Chiaradonna, S., & Bondavalli, A. (2018). Hidden Markov Models as a Support for Diagnosis: Formalization of the Problem and Synthesis of the Solution. In 2018 48th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN) (pp. 1-12). IEEE.
6. Daidone, A., Di Giandomenico, F., Chiaradonna, S., & Bondavalli, A. (2018). Hidden Markov Models as a Support for Diagnosis: Formalization of the Problem and Synthesis of the Solution. In Proceedings of the 48th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN), 1-12.
7. Färber, F., Cha, S. K., Primsch, J., Bornhövd, C., Sigg, S., & Lehner, W. (2015). SAP HANA database: data management for modern business applications. ACM Sigmod Record, 40(4), 45-51.
8. Hasso Plattner Institute. (2016). In-Memory Data Management Technology. Springer-Verlag Berlin Heidelberg.

Vol. 10 No. 1 (2019):753-787
DOI: https://doi.org/10.61841/turcomat.v10i1.14805

9. Kemper, A., & Neumann, T. (2014). HyPer: A hybrid OLTP&OLAP main memory database system based on virtual memory snapshots. In Proceedings of the 27th IEEE International Conference on Data Engineering, 195-206.
10. Krüger, J., Grund, M., Zeier, A., & Plattner, H. (2015). Enterprise application-specific data management. In Proceedings of the 14th IEEE International Enterprise Distributed Object Computing Conference (EDOC 2010) (pp. 131-140). IEEE.
11. Lee, J., Kim, Y. S., & Kim, H. (2017). Failure prediction for autonomic live migration for high availability. Cluster Computing, 20(1), 587-599.
12. Liang, F., Feng, C., Lu, X., & Xu, Z. (2017). Performance benefits of DataNucleus over Hibernate for object-relational mapping. Future Generation Computer Systems, 76, 1-13.
13. Makris, A., Tserpes, K., Spiliopoulos, G., & Anagnostopoulos, D. (2018). Performance Evaluation of MongoDB and PostgreSQL for spatio-temporal data. In EDBT/ICDT Workshops.
14. Márton, G., & Varró, D. (2016). Formalizing and automating domain-specific refactorings. Software & Systems Modeling, 15(4), 1201-1222..
15. Mühe, H., Kemper, A., & Neumann, T. (2018). Executing long-running transactions in synchronization-free main memory database systems. In Proceedings of the 2018 International Conference on Management of Data, 135-148.
16. Plattner, H. (2014). The impact of columnar in-memory databases on enterprise systems: implications of eliminating transaction-maintained aggregates. Proceedings of the VLDB Endowment, 7(13), 1722-1729.
17. Plattner, H., & Zeier, A. (2016). In-Memory Data Management Technology. Springer-Verlag Berlin Heidelberg.
18. Qu, L., Wang, Y., & Orgun, M. A. (2018). Cloud service selection based on the aggregation of user feedback and quantitative performance assessment. In 2013 IEEE International Conference on Services Computing (pp. 152-159). IEEE.

Vol. 10 No. 1 (2019):753-787
DOI: https://doi.org/10.61841/turcomat.v10i1.14805

19. SAP SE. (2017). SAP HANA Administration Guide for SAP HANA Platform 2.0 SPS 02. SAP Help Portal.
20. Schaffner, J., Bog, A., Krüger, J., & Zeier, A. (2015). A hybrid row-column OLTP database architecture for operational reporting. In Business Intelligence for the Real-Time Enterprise (pp. 61-74). Springer, Berlin, Heidelberg.
21. Sharma, Y., Javadi, B., Si, W., & Sun, D. (2016). Reliability and energy efficiency in cloud computing systems: Survey and taxonomy. Journal of Network and Computer Applications, 74, 66-85.
22. Silberschatz, A., Korth, H. F., & Sudarshan, S. (2016). Database system concepts. McGraw-Hill Education.
23. Singh, A., & Singla, A. R. (2015). Performance optimization of fault-tolerant pacemaker cluster. International Journal of Computer Applications, 115(16), 34-38.
24. Stonebraker, M., & Cetintemel, U. (2015). "One size fits all": an idea whose time has come and gone. In Making Databases Work: the Pragmatic Wisdom of Michael Stonebraker (pp. 441-462).
25. Tan, Z., & Babu, S. (2016). Tempo: robust and self-tuning resource management in multi-tenant parallel databases. Proceedings of the VLDB Endowment, 9(10), 720-731.
26. Willhalm, T., Popovici, N., Boshmaf, Y., Plattner, H., Zeier, A., & Schaffner, J. (2015). SIMD-scan: ultra fast in-memory table scan using on-chip vector processing units. Proceedings of the VLDB Endowment, 2(1), 385-394.
27. Zhang, Q., Cheng, L., & Boutaba, R. (2015). Cloud computing: state-of-the-art and research challenges. Journal of Internet Services and Applications, 1(1), 7-18.