Volume 4, Issue 1, January 2014 ISSN: 2277 128X

International Journal of Advanced Research in Computer Science and Software Engineering

Research Paper

Available online at: www.ijarcsse.com

Analysis of SAP HANA High Achievability Competence

Lalit Dole, Girish Talmale
Assistant Professor
Dept. of Computer Science
G.H.Raisoni College of Engg, Nagpur, India

Jayant Rajurkar
Dept.of Computer Science
G.H.Raisoni College of Engg
Nagpur, India

Abstract— One of the core principles of the SAP HANA database system is the comprehensive support of distributed query facility. Supporting scale-out scenarios was one of the major design principles of the system from the very beginning. Within this paper, we first give an overview of the overall functionality with respect to data allocation, metadata caching and query routing. We then dive into some level of detail for specific topics and explain features of SAP HANA high achievability. In summary, the paper provides a comprehensive overview of distributed query processing in SAP HANA database to High achievability competence.

Keywords— SAP HANA, HA Architecture**, **Scalability, **Disaster Recovery, Backup & Recovery,**

I. INTRODUCTION

In today's business world, High Availability (HA) is an extremely important consideration for any enterprise-level IT architecture. A technology platform missing on critical HA features are not enterprise-ready. HANA has been touted by SAP as the modern platform for real-time analytics and applications. The SAP HANA documentation [1] even states that "SAP HANA is fully designed for high availability." However – as the analysis in this whitepaper shows, this assertion is entirely incorrect!

This paper provides an in-depth analysis of the HA features available with SAP HANA Support Package Stack (SPS) 05 [2]. It also provides a comparison with HA features available with the Oracle Database, as part of Oracle Maximum Availability Architecture (MAA) [3]. As this paper demonstrates – as far as HA is concerned, there are serious drawbacks in HANA's technical architecture, and HANA lacks several key capabilities that are absolutely necessary to implement a robust HA architecture. SAP HANA is not enterprise-ready.The intended audience of this whitepaper are members of IT managerial and technical teams interested in finding out whether SAP HANA is viable for any of their critical applications, and/or who want to get a better understanding of how the HANA architecture compares with well-established high availability technologies.

II SAP-HANA-AN ARCHITECTURAL OVERVIEW

SAP uses the term "SAP HANA" to designate the SAP In-Memory Appliance which is aimed at data analytics. It is a combination of hardware and software, and it is delivered as an optimized appliance in cooperation with SAP's hardware partners for SAP HANA. SAP also uses the term to indicate the SAP in-memory database, which is a hybrid in-memory database that combines row-based, column-based, and object-based database technology. The following diagram provides a High-level overview of the SAP HANA architecture.

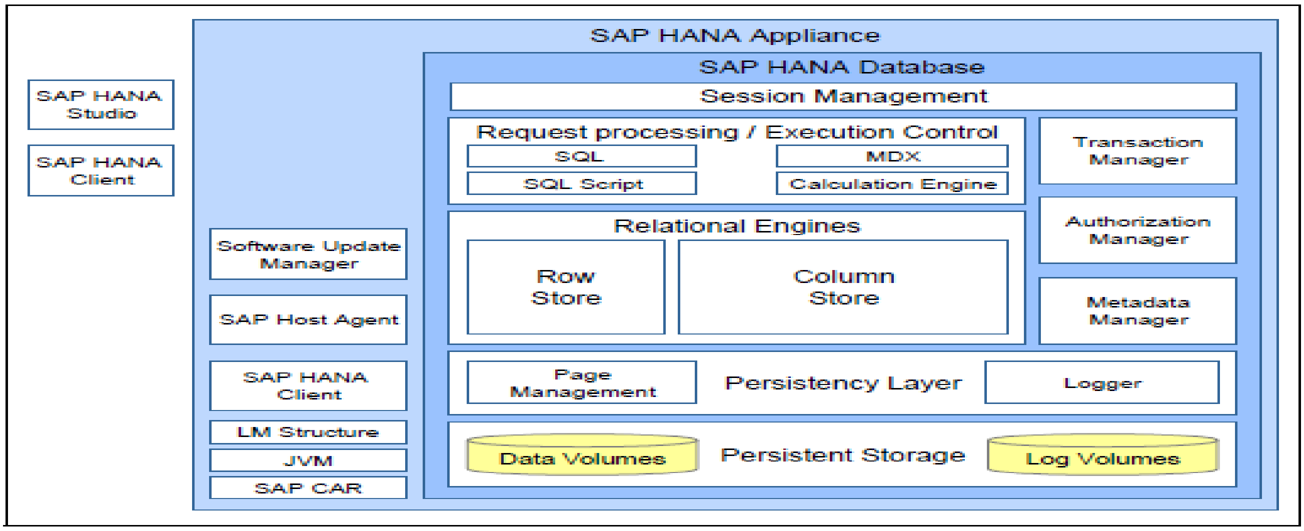

Fig: 1: SAP HANA Architecture

II. SAP HANA DATABASE

The SAP HANA database consists of two database engines:

- The column-based store, storing relational data in columns,Optimized for holding data mart tables with large amounts Of data, which are aggregated and used in analytical Operations.
- The row-based store, storing relational data in rows. This Row store is optimized for write operations and has a lower Compression rate, and its query performance is much lower Compared to the column-based store. The engine that is used to store data can be selected on a per-table basis at the time Of creation of a table (and changed subsequently). Tables in the row-store are loaded into memory at start-up time, Whereas tables in the column-store can be either loaded at Start up or on demand, during normal operation of the SAP HANA database. The engine that is used to store data can be selected on a per-table basis at the time of creation of a table (and changed subsequently). Tables in the row-store are loaded into memory at start-up time, whereas tables in the column-store can be either loaded at start-up or on demand, during normal operation of the SAP HANA database.

IV. SAP HANA High achievability (HA) Features

The HA analysis of SAP HANA can be done across the following three elements.

1. Scalability support.
2. Disaster Recovery (DR) support.
3. Backup & Recovery support.

The following sections will look into each of these in greater detail.

1.Scalability Features:

HANA implements scalability through the use of multiple servers in one SAP HANA cluster, using a shared nothing approach for the underlying data, which has to be partitioned across these servers (hosts). SAP often refers to this layout as the distributed SAP HANA database / system.

Distributed HANA System: Name, Index and Statistics Servers

The Index / Name / Statistics system components are integral to a distributed HANA system and deserve special mention. In a single-host system, all these components are installed on a single SAP HANA database instance on one host. In a distributed SAP HANA system, they are installed on multiple database instances on different hosts.

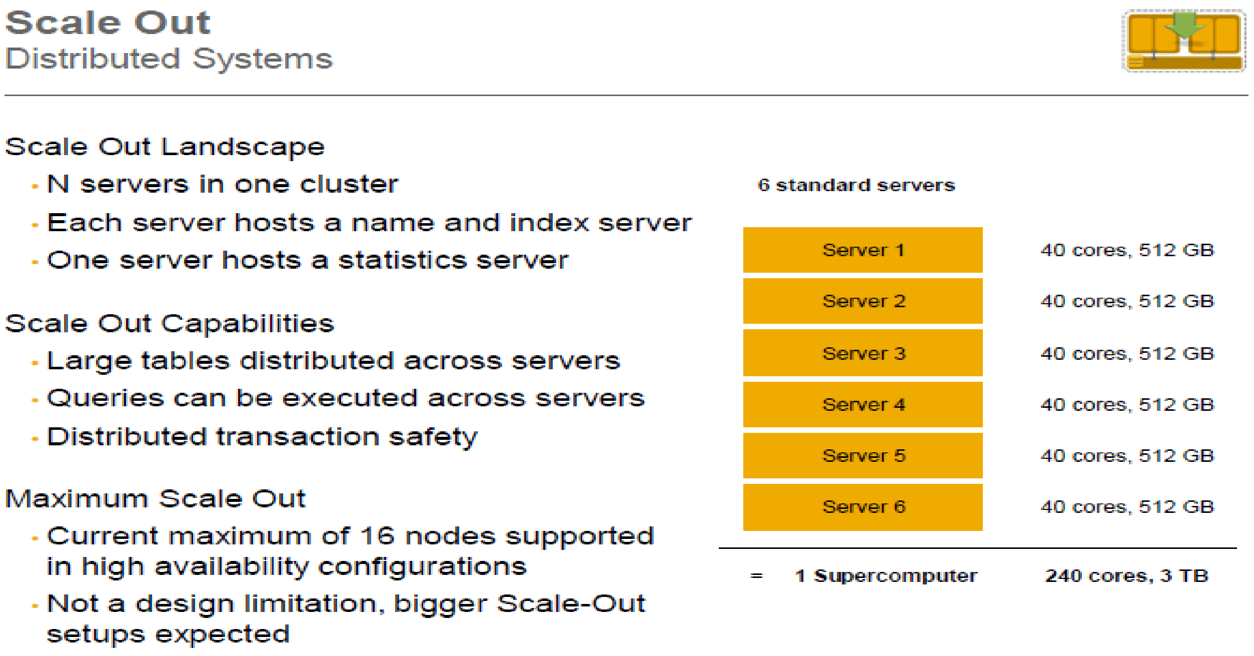

Scale Out
Distributed Systems

Scale Out Landscape

- N servers in one cluster
- Each server hosts a name and index server
- One server hosts a statistics server

Scale Out Capabilities

- Large tables distributed across servers
- Queries can be executed across servers
- Distributed transaction safety

Maximum Scale Out

- Current maximum of 16 nodes supported in high availability configurations
- Not a design limitation, bigger Scale-Out setups expected

6 standard servers

Server 1	40 cores, 512 GB
Server 2	40 cores, 512 GB
Server 3	40 cores, 512 GB
Server 4	40 cores, 512 GB
Server 5	40 cores, 512 GB
Server 6	40 cores, 512 GB
= 1 Supercomputer	240 cores, 3 TB

Fig: 2: SAP HANA Scale Out

Name server – The name server owns the information about the topology of the SAP HANA system. In a distributed system with instances of the SAP HANA database on multiple hosts, the name server knows where the components are running and which data is located on which server.

Index server – The index server contains the actual data stores and the engines for processing the data.

Statistics server – This statistics server collects information about status, performance and resource consumption from all components belonging to the SAP HANA system. Monitoring clients such as the SAP HANA studio access the statistics server to get the status of various alert monitors. The statistics server also provides a history of measurement data for further analysis.

The Index Server is the core database engine – following is the architecture diagram for the Index Server:

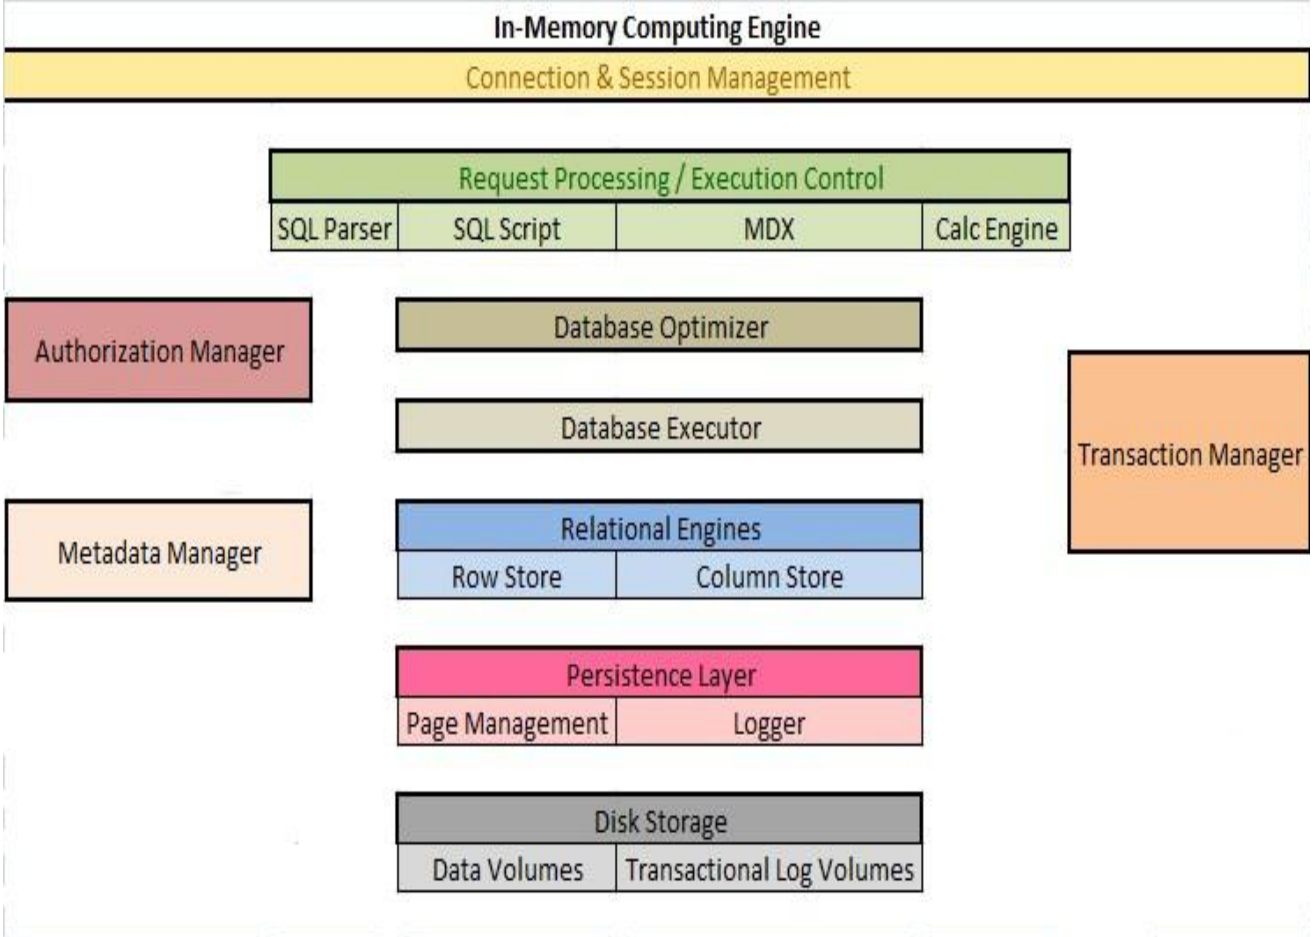

Fig 3: SAP HANA Index Server Architecture

Setting up a distributed HANA system involves a complex configuration of Name Server and Index Server roles for the hosts implementing the distributed system. It works as follows.

Name Server

Up to 3 hosts can be configured as MASTER Name Servers (using a “MASTER 1, MASTER 2, MASTER 3” naming designation), while others have to be configured as “SLAVE”. However, in steady state, only one of the “MASTER n” configured hosts has its actual Name Server role as “MASTER” while all other hosts have their actual Name Server roles as “SLAVE”. The MASTER Name Server is responsible for assigning storage/data volumes to active Index Servers at the time they start up. [Note: From SAP Collateral, the functionality provided by the SLAVE Name Server is not clear. If the MASTER Name Server host fails, one of the remaining hosts configured as a “MASTER n” Name Server becomes the active MASTER Name Server.

Index Server

During setup, hosts can be configured as “WORKER” or “STANDBY” Index Servers. In steady state, one WORKER host gets the actual Index Server role as “MASTER” – this is assigned on the same host as the MASTER Name Server. Other WORKER hosts get the actual Index Server role of “SLAVE”. The STANDBY Index Server gets the actual Index Server role of “STANDBY”. The MASTER Index Server provides metadata for the other SLAVE Index Servers. The SLAVE Index Server is an active database server, is assigned a data volume, and accepts database connections. A host configured as a STANDBY Index Server is not used for database processing – all database processes run on this host, but they are idle and do not allow SQL connections. If an active Index Server fails, the active MASTER Name Server assigns its volume to one of the hosts in the STANDBY Index Server role.

Storage Layout: Single-Instance HANA System:

To better understand how a distributed HANA system with multiple hosts persists data and accesses storage, it helps to see what the storage layout of a single-instance HANA system looks. The following diagram shows this layout from an IBM systems/storage standpoint.

In the diagram, IBM’s General Parallel File System (GPFS), which is a shared-disk clustered file system providing concurrent file access to applications executing on multiple nodes of clusters, is implemented on two types of storage:

- The data storage (on SAS disks), here referred to as HDD, which holds the save points.
- The log storage (on SSD drives or PCIe Flash devices), here referred to as Flash, which holds the database logs .

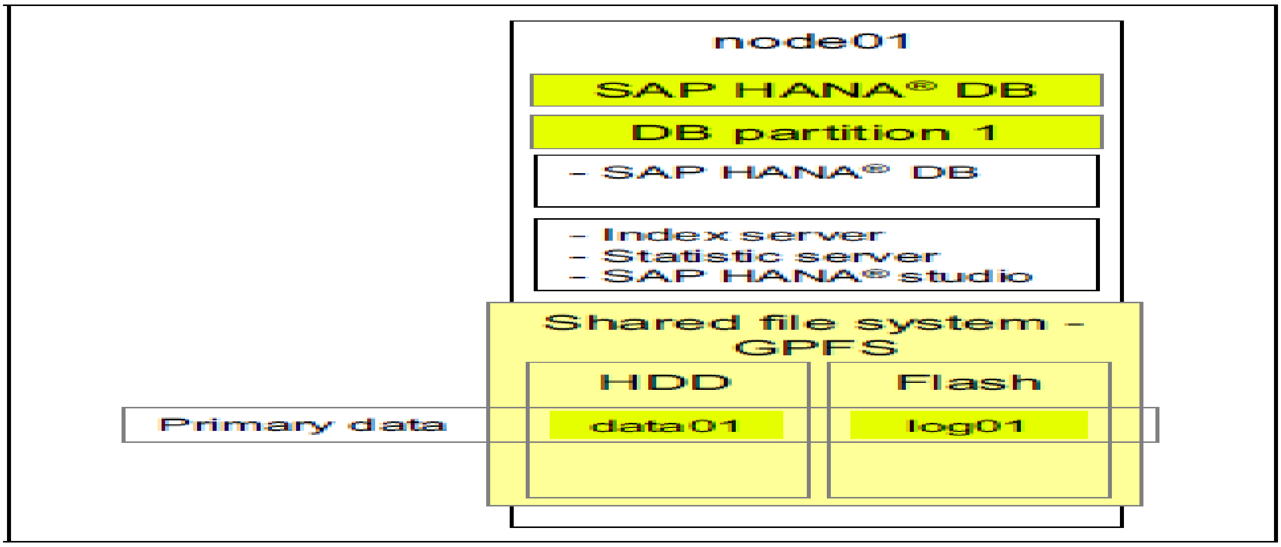

Fig: 4: Storage Layout: Single Instance SAP HANA System

This single node represents one single SAP HANA database consisting of one single database partition. Both the savepoints (data01) and the logs (log01) are stored once (that is, they are not replicated), denoted as being primary data in the above diagram.

Storage Layout: Distributed HANA System

The scale-out distributed HANA system differs from a single server solution in a number of ways:

- The solution consists of a cluster of building blocks,Interconnected with two separate 10 Gb Ethernet Networks, one for the SAP HANA application and one for the file system communication.
- The SAP HANA database is split into partitions, forming a single instance of the SAP HANA database.
- Each node of the cluster holds its own save points and database logs on the local storage devices of the server.
- The standby nodes run the SAP HANA application, but do not hold any data or take an active part in the processing. If one of the active nodes fails, a standby node will take over the role of the failed node, including the data (that is, the database partition) of the failed node. This takeover process is automatic, but orchestrated through the MASTER Name Server.
- The GPFS file system spans all nodes of the cluster, making the data of each node available to all other nodes of the cluster, including the standby node.

The following diagram illustrates this solution, showing a 4-node configuration as an example, in which the node04 serves as the standby.

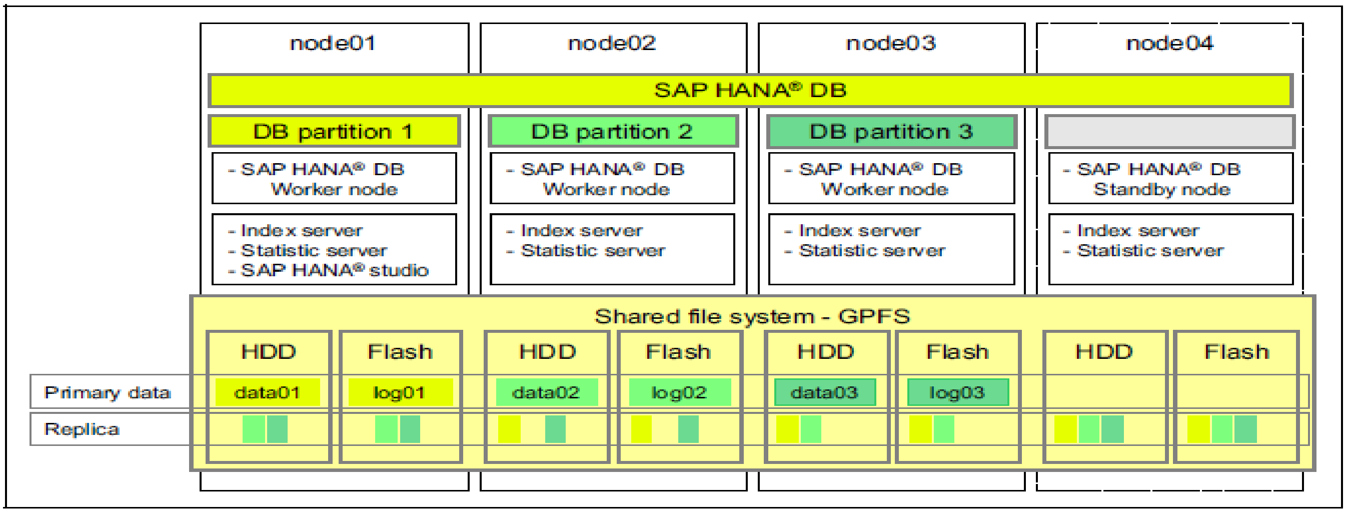

Fig: 5: Storage Layout: Distributed SAP HANA System

The SAP HANA software distributes application requests internally across the cluster to the individual worker nodes, which process the data and exchange intermediate results, which are then combined and sent back to the requestor. Each node maintains its own set of data, persisting it with save points and logging data changes to its database log.

With respect to the system components described previously, only the Name Server service on the active MASTER host persists data. SLAVE Name Sever hosts communicate with the MASTER, but do not persist data. The Index Server service on all hosts except STANDBY hosts persists data. The Statistics Server service can run only on one host and persists data on this host.GPFS combines the storage devices of the individual nodes into one file system, making sure that the SAP HANA software has access to all data regardless of its location in the cluster, while making sure that savepoints and database logs of an individual database partition are stored on the appropriate locally attached storage device (disk and flash) of the node on which the partition is located.

Distributed HANA System: Node Failover

To be able to take over the database partition from a failed node, the standby node has to load the savepoints and database logs of the failed node to recover the database partition and resume operation in place of the failed node. Using the previous diagram as an example, assume node03 experiences a problem and fails. The master node (node01) recognizes this and directs the standby node, node04, to take over from the failed node. To recreate database partition 3 in memory to be able to take over the role of node03 within the cluster, node04 reads the savepoints and database logs of node03 from the GPFS file system, reconstructs the savepoint data in memory, and re-applies the logs so that the partition data in memory is exactly like it was before node03 failed. After this is complete, Node04 is in operation as an active Index Server, and the database cluster has recovered.

After fixing the cause for the failure of node03, it can be reintegrated into the cluster as the new standby system. Data access in non-failed nodes can continue, although it is not clear from HANA collateral whether there is any brown-out impact.

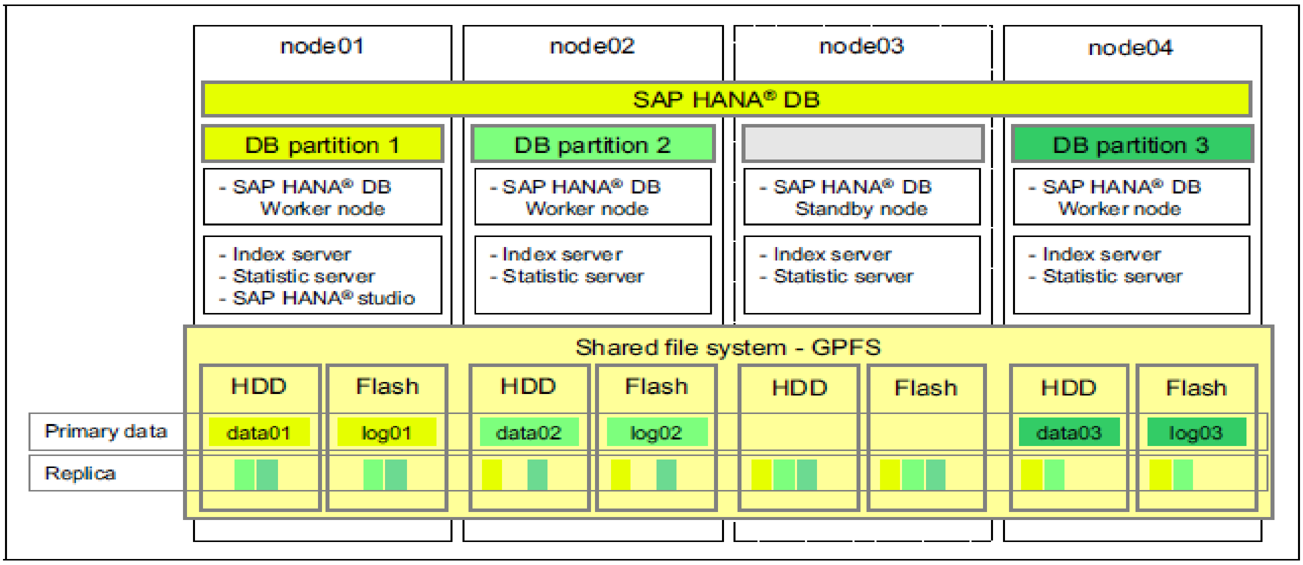

Fig: 6: Failover: Distributed SAP HANA System

When a node has an unrecoverable hardware error, the storage devices holding the node's data might become unavailable or even destroyed. To mitigate this, the GPFS file system can be configured to replicate the data of each node to the other nodes, to prevent data loss if one of the nodes (with its local storage) goes down. The replication is done by GPFS synchronously, i.e. each write operation only finishes when the data has been both written locally and replicated.

2.SAP HANA: Analysis of Disaster Recovery Features :

HANA's Disaster Recovery (DR) support hinges on two capabilities:

• Storage replication: Continuous replication (mirroring) between primary storage and backup storage over a network (may be synchronous). This solution is provided by storage partners (e.g. HP, HDS), and was theonly DR support available in earlier HANA releases.

• System replication: Continuous (currently: synchronous) update of secondary system by primary system, including in-memory table loading. This capability is new in SPS05, and in earlier roadmap discussions, was referred to as warm standby.

Storage Replication

SAP relies on specific certified hardware partners to offer a storage-level replication solution, which delivers a copy of the persisted volumes or file-system to a remote, networked storage system. For synchronous replication, the SAP HANA transaction only completes when the locally persisted transaction log has been replicated remotely. SAP suggests synchronous storage replication can be used only where the primary-DR site distance is 100 kilometers or less. In any case, this solution requires a reliable, high bandwidth and low latency connection between the primary site and the secondary site.

Following are two examples of how this configuration look As evident from the previous diagrams, while the storage-bits are propagated to the remote storage arrays, no servers can be mounted on those remote arrays while mirroring is in progress.

Among the hardware vendor partners, HP has made a big deal [10] about validating a hardware solution to provide DR for HANA.

3.SAP HANA: Analysis of Backup & Recovery Features :

SAP HANA uses in-memory technology, but it fully persists two types of data to storage: transaction logs, and data changes in the form of savepoints. A transaction redo log is used to record a change at commit time. Upon an outage, the most recent consistent state of the database can be restored by restoring the savepoint and replaying the changes recorded in the log, redoing completed transactions and rolling back incomplete ones. At a savepoint, all the changed data are written to storage, in the form of pages (i.e. blocks). When starting up the system, logs can be processed from the last savepoint position. Save points are coordinated across all processes and instances of the database to ensure transaction consistency. By default, save points are performed every five minutes, but this can be configured. HANA supports two kinds of backups [9] and both can be done only when the database is online:

- Backing up the save points in the form of full data backups, which can be used to restore a database to a specific point in time
- Smaller periodic log backups to recover from fatal storage faults to minimize the loss of data.

V. SAP HANA: Summary Analysis of all HA Features:

- Besides support for Scalability, Disaster Recovery, and Backup & Recovery, a well-rounded HA solution must include support for a variety of other capabilities that address various aspects of unplanned and planned downtime. For example – human errors is one of the biggest causes of downtime – when an administrator accidentally deletes important database tables or removes important database files, a well-rounded HA solution must include the capability to efficiently undo such errors. HANA has none. Similarly – reducing various aspects of planned downtime – e.g. software or hardware upgrades, technology refreshes, etc. should also be supported. HANA has none.
- In contrast, Oracle Maximum Availability Architecture (MAA), integrated with the Oracle Database, provides customers with a comprehensive and integrated set of industry-leading HA technologies that can be deployed at minimal cost, help avoid downtime, enable rapid recovery from unplanned failures, and minimize the impact of planned outages.
- Following is a list of critical HA features that comprise Oracle MAA, along with the benefits they provide, and also how the equivalent solutions from SAP HANA compare.

VI Conclusion

Successful enterprises understand the vital importance of maintaining highly available technology infrastructures to protect critical data and information systems. At the core of many mission critical information systems is the database system, responsible for the availability, security, and reliability of the information technology infrastructure. If this database system lacks any of these fundamental capabilities, it is simply not enterprise-ready – regardless of other features it may have. Building on decades of innovation, Oracle Database offers the industry's gold standard for integrated high availability and data protection that ensure customers can maximize their application uptime in the event of both planned maintenance activities and unexpected failures. In contrast – as this document demonstrates, SAP HANA, even though touted as the next-generation database architecture, lacks many of the fundamental high availability capabilities needed to run today's 24x7 enterprises.

References

[1] SAP HANA Administration Guide: http://help.sap.com/hana/hana_admin_en.pdf

[2] SAP HANA Documentation: http://help.sap.com/hana_appliance/

[3] Oracle Database High Availability: http://www.oracle.com/technetwork/database/features/availability/index.html

[5] SAP Collateral: SAP HANA In Data Centers,

[6] SAP Collateral: SAP HANA Scale-Out, High Availability & Disaster Tolerance, http://www.experiencehana.com/servlet/JiveServlet/download/1221-5-7948/SAP%20HANA-HighAvailability%26DisasterTolerance.pdf

[7] Blog: SAPA HANA Architecture, http://www.dwbiconcepts.com/database/28-hana/105-sap-hana- architecture.html

[8] IBM Redbook: In-memory Computing with SAP HANA on IBM eX5 Systems, http://www.redbooks.ibm.com/abstracts/sg248086.html

[9] SAP Collateral: SAP HANA – High Availability, http://www.experiencehana.com/servlet/JiveServlet/downloadBody/2775-102-1- 5238/HANA_HA1.0b.pdf

[10] HP AppSystems for SAP HANA, http://h20219.www2.hp.com/enterprise/us/en/partners/sap-high-performance-analytic-appliance.html

[11] HP Collateral: Disaster-Tolerant solution for HP AppSystems for SAP HANA, http://h20195.www2.hp.com/V2/GetPDF.aspx/4AA14846ENW.pdf

[12] Oracle HA Customer Case Studies: http://www.oracle.com/technetwork/database/features/ha-casestudies 098033.html

[13] IDC Insight: SAP and Sybase: A Marriage Made In Database Heaven, Junehttp://www.sybase.com/files/White_Papers/IDC_SAP_Sybase_DbHeaven_report.pdf.