International Journal of Science and Research (IJSR)
ISSN: 2319-7064
SJIF (2022): 7.942

Optimization of Memory: SAP HANA Database Perspective - Part 3

Rajesh Azmeera[1], Rhea Khanna[2], Deepak Nanuru Yagamurthy[3]

[1]Technology Professional, Department of Information and Technology, Stryker, US (Corresponding Author)

[2]Technology Professional, Department of Research, Microsoft, US

[3]Technology Professional, Department of Information and Technology, 7 Leven Inc, US

Abstract: ***Memory is a fundamental resource of the SAP HANA database. Understanding how the SAP HANA database requests, uses, and manages this resource is crucial to the understanding of SAP HANA. As Data is growing rapidly and these data is stored in memory to process. Increasing DB size is directly proportional to its memory consumed. SAP HANA provides a variety of memory usage indicators that allow for monitoring, tracking, and alerting. The most important indicators are used memory and peak used memory. Memory needs to be optimized, otherwise increasing Data size can't be accommodated by existing memory. Since SAP HANA contains its own memory manager and memory pool, external indicators such as the size of resident memory at host level and the size of virtual and resident memory at process level can be misleading when you are estimating the real memory requirements of an SAP HANA deployment. SAP HANA integrates data from multiple areas within an organization, for example: Traditional business documents – including contracts and spreadsheets. UX/UI (User Experience/User Interface)- including website forms, emails and other customer interactions Mobile - information from the mobile devices of customers and your workforce. IoT (Internet of Things) – data from the many sensors that run in every aspect of a business, from warehouses and trucks to stores and offices. The SAP HANA system not only integrates all of this data; it can also apply machine learning and AI to analyze it instantly and deeply, accelerating real-time decision-making by providing key insights into a company's operations.***

Keywords: SAP Memory Management, SAP HANA, Memory Optimization, Resource Utilization and SAP HANA Database

1. Introduction

SAP HANA (High-performance Analytic Appliance) is a multi-model database that stores data in its memory instead of keeping it on a disk. The column-oriented in-memory database design allows you to run advanced analytics alongside high-speed transactions – in a single system. Why is this so important? Because it lets companies process massive amounts of data with near-zero latency, query data in an instant, and become truly data-driven. By storing data in column-based tables in main memory and bringing online analytical processing (OLAP) and online transactional processing (OLTP) together, SAP HANA is unique – and significantly faster than other database management systems (DBMS) on the market today.

Pre-requisite: To understand this paper thoroughly pre requisite is Optimization of Memory: SAP HANA Database perspective – PART 1 and Optimization of Memory: SAP HANA Database perspective – PART 2

2. Discussion

Analytic Views performance

Analytic views are used to model data that includes measures. For example, transactional fact table representing sales order history would include measures for quantity, price, and so on. During the service the SAP experts can create these views via SAP HANA studio. After adding data sources, definition of central fact table, definition of output columns joins can be created in the star join node between data fields of the fact table to an interval field of the attribute view. After definition of the attributes and measures the analytic views activated.

After creating an analytic view, the SAP service experts can perform certain additional tasks to obtain the desired output. The table below lists the additional tasks that you can perform to enrich the analytic view.

Requirement	Task to Perform
If you want to filter the output of the data foundation node	Filter Output of Data Foundation Node

Attributes and Measures:

Requirement	Task to perform
If you want to count the count the number of distinct values for a set of attribute Columns	Create Counters
If you want to create new output columns and calculate its values at runtime using an expression.	Create Calculated Columns
If you want to restrict measure values based on attribute restrictions	Create Restricted Columns
If you want to assign semantic types to provide more meaning to attributes and measures in analytic views.	Assign Semantics
If you want to parameterize attribute views and execute them based on the values users provide at query runtime	Create Input Parameters
If you want to, for example, filter the results based on the values that users provide to attributes at runtime	Assign Variables
If you want associate measures with currency codes and perform currency conversions	Associate Measures with

	Currency
If you want associate measures with unit of measures and perform unit conversions.	Associate Measures with Unit of Measure
If you want to group related measures together in a folder	Group Related Measures

Analytic view properties

Requirement	Task to perform
If you want to filter the view data either using a fixed client value or using a session client set for the user.	Filter Data for Specific Clients
If you want to execute time travel queries on analytic views	Enable Information Views for Time Travel Queries
If you want to invalidate or remove data from the cache after specific time intervals	Invalidate Cached Content
If you want to maintain object label texts in different languages	Maintain Modeler Objects in Multiple Languages
If you do not recommend using an analytic view	Deprecate Information Views

3. Application & Statement performance analysis and problems with Data models

View Doesn`T Return Data

In case that a view doesn`t return any data a check needs to be performed to see if the client is set correctly. Therefore open the view in the Modeler, check in the Properties if the Default Client is set. If the Default Client is set "dynamic" then the Session Client is used.

These are the steps to be proceeded in SAP HANA Studio: Catalog □ Authorization □ Users □ <user> check "Session Client".
In order to find out the correct session client it needs to be checked if the fact table of the view contains column mandt or client.

This can be done with the help of the following command select mandt, count (*) from <fact table> group by mandt order by count (*) desc;
ACTIVATION OF CALCULATION VIEW OR ANALYTIC VIEW FAILS

Activation says "Repository: Encountered an internal error in the repository code, this is most likely a bug in the implementation; Got error when trying to change object owner to _SYS_REPO, detailed error is: internal error"

The issue can be solved in the following way: drop view "_SYS_BIC"."<package>/<view name>"; before each activation.

SQL Trace Analysis

The first step in application performance analysis is to figure out if the database layer is causing performance problems for your application at all. During this analysis the experts will check how many and which database calls are made and what their contribution to the overall application performance is. This will be dome within the context of a given user interface step or transaction.

Example:

Start the tracing of database calls. Run the application from its user interface or with any other driver. Both, SAP HANA studio and SAP HANA Web-based Development Workbench provide two deep tracing (including complete execution plans) is provided by Plan Trace in SAP HANA studio.

Terminate the tracing and review aggregated and individual results. As a result of this investigation you might see some indicators for bad application logic creating excessive load on the database such as too many database calls (per transaction/UI step). Many identical executions, for example repeated identical selects. Too many records returned (per execution or in total). Too many columns or all columns of a row selected. Inefficient statement reuse, that is, statements that need to be optimized over and over again. One or more database calls with unexpected bad performance, so you should further investigate those calls.

Statement Measurement

Once it has been determined which SQL statements are problematic the SAP team will perform a sound measurement in order to get reliable performance numbers and to make sure that indeed your statement(s) are causing the issues and not the current state of your SAP HANA system

Execute your statement(s) and measure their performance (in particular response time). Both, SAP HANA studio and SAP HANA Web-based Development Workbench offer basic measurement of SQL statements. In addition, the SAP HANA Web-based. Development Workbench implicitly supports observing the system state and also executing repeated measurements. Checking the SAP HANA system status for disturbing conditions, such as high load, high resource usage and so on. The measurement will be repeated until stable results without major variations (for example, 3 stable executions in a row) can be encountered.

Once you have a stable result you may also acquire a detailed SAP HANA engine trace which will allow for a deep dive analysis. As a result of this activity reliable data for query performance, both for initial query execution performance (possibly cold execution) and stabilized execution performance (warm execution) are available.

Technical Analysis

The deepest level of performance analysis addresses the technical details of a database statement execution.

There are a number of tools that support the service team:

SAP HANA studio offers the Plan Visualizer perspective which allows for deep technical analysis. Dedicated views and filters support the analysis along numerous dimensions. Furthermore, there are multiple tools allowing for even deeper analysis such as Explain Plan and the numerous

tracing tools.

Possible Procedure:
Inspect aggregated execution KPIs (execution time, resource consumption, distribution characteristics) in order to figure out in which aspect or dimension you might look for possible issues. For any suspicious KPI, the team will track down the KPI to the deepest possible level manifesting the symptom.

From there, they will try to correlate the symptom with the cause, in higher level statement elements. Possible restriction of the analysis scope by focusing on the critical path, or on a specific time interval, system node, engine, or execution plan operator.

Case study:

Troubleshoot System Replication
If system replication appears to slow down transaction processing you can check the network and disk I/O on the secondary site. Often due to network related issues or a disk I/O related issue on the secondary site, system replication can slow down transaction processing in the following cases:

Asynchronous replication is configured over long distances. Multitier system replication is configured and a tier 3 system is attached.

SYNC/SYNCMEM replication mode is configured over short distances. The following troubleshooting steps can help you determine and resolve the underlying cause:

Check If Log Can Be Shipped In Time
Check if log shipping is significantly slower than local log write (SYNC/SYNCMEM). Check Async Buffer Full Count (ASYNC). These checks can be used to verify that the problem is really related to system replication. For analyses of system replication KPIs use HANA_Replication_SystemReplication_Overview.txt which is attached to the OSS note 1969700 - SQL Statement Collection for SAP HANA.

Check If Data Load Can Be Handled By Network Link
Estimate the amount of data/log shipped from the primary site. Compare this with the available bandwidth (the recommended bandwidth is 10 Gbit/s). Do a network performance test.

For analyses of bandwidth for Data/Log shipping use *HANA_Replication_SystemReplication_Bandwidth.txt* which is attached to the OSS note 1969700 - SQL Statement Collection for SAP HANA. It is recommended to use this SQL statement when system replication is disabled. Collect network information on bandwidth and latency. If the bandwidth can handle load check if the network is shared and whether other applications may be interfering with performance. Check the network utilization profile for the network link to see if the maximum capacity of the network has been reached. Slow disk I/O on the secondary can postpone releasing log buffers on primary, which results in wait situations on the primary. Perform disk performance test. Check the Monitoring view via command line tool – run command on secondary site: hdbcons "statreg print -n M_VOLUME_IO_TOTAL_STATISTICS -h"

For overall status of replication check as adm-user on primary node via: python systemReplicationStatus.py (you have to be in destination of alias cdpy). For troubleshooting check log in destination /usr/sap/<DBSID>/HDB<instance #>/<node name>/trace. Check also related HOW document of SUD OPR.HDB.18 - High_Availability_of_HDB

Network Performance and Connectivity Problems
In cases where a subjectively slow performing system behavior is experienced, but a first analysis of the SAP HANA resource utilization does not reveal any obvious culprits, it is often necessary to analyze the network performance between the SAP HANA server host(s) and SAP Application Server(s) / Non-ABAP clients, SAP HANA nodes (inter-node communication in SAP HANA scale-out environments), or, in an SAP HANA system replication scenario, between primary and secondary site.

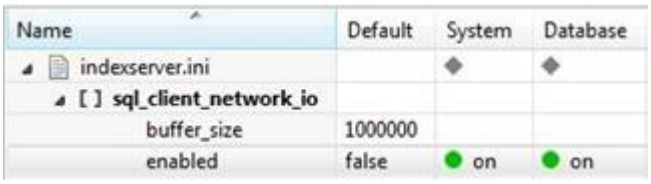

this parameter change implies a certain performance overhead and should only be active for the duration of the troubleshooting activity

To check the statistic call via SQL Console SQL command: SELECT * FROM M_SQL_CLIENT_NETWORK_IO

HOST	PORT	CLIENT_HOST	CONNECTION_ID	MESSAGE_ID	CLIENT_DURATION	SERVER_DURATION	SERVER_RECEIVED_TIME	RECEIVED_MESSAGE_SIZE	SEND_MESSAGE_SIZE
d	30.040	D 5	214.859	3.240	854	455	29.02.2016 15:48:53.520621	256	656
d	30.040	D 5	214.858	6.171	816	109	29.02.2016 15:48:53.520378	192	280
d	30.040	D 5	214.859	3.241	744	312	29.02.2016 15:48:53.521513	232	168
d	30.040	D 5	214.872	2.897	1.033	293	29.02.2016 15:48:53.521145	224	280
d	30.040	D 5	214.859	3.242	3.135	365	29.02.2016 15:48:53.524752	232	640
d	30.040	D. 5	214.872	2.898	3.058	258	29.02.2016 15:48:53.525187	224	280
d	30.040	D 5	214.872	2.899	585	298	29.02.2016 15:48:53.525872	504	400
d	30.040	D 5	214.872	2.900	726	467	29.02.2016 15:48:53.526534	320	176
d	30.040	D. 5	214.872	2.901	1.804	1.517	29.02.2016 15:48:53.527445	176	152
d	30.040	D 5	214.872	2.902	644	347	29.02.2016 15:48:53.529759	456	512
d	30.040	D. 5	214.872	2.903	782	545	29.02.2016 15:48:53.530457	480	272

CLIENT_DURATION and SERVER_DURATION contain the values in microseconds. The difference between CLIENT_DURATION and SERVER_DURATION makes the total transfer time of the result (SEND_MESSAGE_SIZE in bytes). This allows you to see whether the transfer time from the SAP HANA server to the client host is exceptionally high. Another important KPI is the Round Trip Time (RTT) from server to client. For analyses of RTT KPI use HANA_Network_Clients.txt which is attached to the OSS note 1969700 - SQL Statement Collection for SAP HANA.

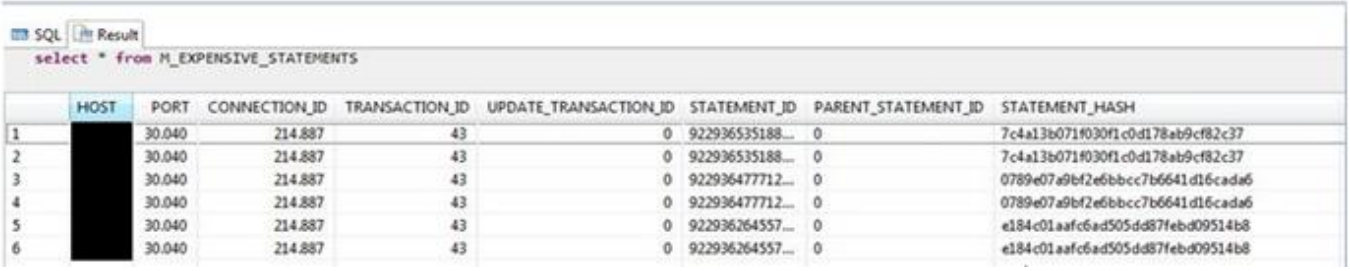

SQL Result

```
select * from M_EXPENSIVE_STATEMENTS
```

	HOST	PORT	CONNECTION_ID	TRANSACTION_ID	UPDATE_TRANSACTION_ID	STATEMENT_ID	PARENT_STATEMENT_ID	STATEMENT_HASH
1	[redacted]	30.040	214.887	43	0	922936535188...	0	7c4a13b071f030f1c0d178ab9cf82c37
2		30.040	214.887	43	0	922936535188...	0	7c4a13b071f030f1c0d178ab9cf82c37
3		30.040	214.887	43	0	922936477712...	0	0789e07a9bf2e6bbcc7b6641d16cada6
4		30.040	214.887	43	0	922936477712...	0	0789e07a9bf2e6bbcc7b6641d16cada6
5		30.040	214.887	43	0	922936264557...	0	e184c01aafc6ad505dd87febd09514b8
6		30.040	214.887	43	0	922936264557...	0	e184c01aafc6ad505dd87febd09514b8

Because of SQL statement is using the view M_EXPENSIVE_STATEMENTS, the expensive statements trace needs to be active in SAP HANA studio Administration editor Trace Configuration tab. Once expensive statement is running, it is required to identify STATEMENT_HASH from M_EXPENSIVE_STATEMENTS view via SQL Console SQL command: SELECT * FROM M_EXPENSIVE_STATEMENTS

4. Application and Database Connectivity Analysis

On an ABAP application server, Run transaction OS01 - Database - Ping x10. If a connection to the database cannot be established over a longer period of time by an SAP ABAP application work process, the work process is terminated. First, the work process enters the reconnect state in which it constantly tries to connect to the database, after a predefined amount of retries fail, the work process terminates. In this case the connectivity from the SAP application server to the SAP HANA server must be verified.

Run Report ADBC_TEST_CONNECTION via SE38**.** If a specific database connection is failing, the report ADBC_TEST_CONNECTION offers a connectivity check for each defined database connection. OSS note 2213725. If an application is facing communication issues with the SAP HANA server, on client side the connectivity issue may be indicated by several 10709-errors, mostly short dumps. For an overview of the most common errors in this context and detailed explanations of how to resolve they check OSS note 2213725 - How-To: Troubleshooting of -10709 errors.

ON NON-ABAP APPLICATION

In case issues occur on non-ABAP client connections to a remote SAP HANA instance, it is of importance to make sure that a supported client is used. For that reason check related OSS note 1577128 - Supported clients for SAP HANA.

If no SAP Notes can be found which outline the root cause of a specific error message you can record an ODBC trace to gain more insight. To do that follow the procedure described in OSS note 1993254 - Collecting ODBC Trace.

Troubleshooting SAP HANA Network. In case the error occurs sporadically, it is useful to perform a long-term stress test between the client and SAP HANA server to confirm the network's stability. For more information check related chapters Stress Test with NIPING and OSS note 2081065.

5. Generic Smart Data Access Troubleshooting Steps

Verify that the issue might not be a flaw in the SAP HANA studio: try to connect via the 'isql' tool on the SAP HANA host directly as the SAP HANA<sid>adm. Verify that all libraries can be accessed on OS level by the SAP HANA <sid>adm user (PATH environment variable). In the SAP HANA <sid>adm home directory (cd $home) check that the correct host and port, and username and password combinations are used in the .odbc.ini file. Verify that the LD_LIBRARY_PATH environment variable of the SAP HANA <sid>adm user contains the unixODBC and remote DB driver libraries. Connections from the SAP HANA server to remote sources are established using the ODBC interface (unixODBC). In case issues occur on non-ABAP client connections to a remote SAP HANA instance, it is of importance to make sure that a supported client is used. For that reason check related OSS note 1577128 - Supported clients for SAP HANA. If no SAP Notes can be found which outline the root cause of a specific error message you can record an ODBC trace to gain more insight. To do that follow the procedure described in OSS note 1993254 - Collecting ODBC Trace. In case the error occurs sporadically, it is useful to perform a long-term stress test between the client and SAP HANA server to confirm the network's stability. For more information check related chapters Stress Test with NIPING and OSS note 2081065.

In the context of blocked transaction troubleshooting, the columns “Blocked by Connection Id” and “Blocks No. of Transactions” are of special interest. The first tells you whether the session is blocked by another session and identifies the ID of the blocking one. The latter gives you the corresponding information if a session blocks other sessions, and how many transactions are affected.

Cancel a session by right-clicking the session and choosing cancel Session. Jump to the related objects by right-clicking the session and choosing navigate To Activate the performance trace, SQL trace, or expensive statements trace by choosing configure Trace.

Trace Level	Description
NORMAL	All statements that have finished successfully are traced with detailed information such as executed timestamp, thread ID, connection ID, and statement ID.
ERROR	All statements that returned errors are traced with detailed information such as executed timestamp, thread ID, connection ID, and statement ID.

ERROR_ ROLL BACK	All statements that are rolled back are traced with detailed information such as executed timestamp, thread ID, connection ID and statement ID.
ALL	All statements including status of normal, error, and rollback are traced with detailed information such as executed timestamp, thread ID, connection ID and statement ID.
ALL_WITH_ R ESULTS	In addition to the trace generated with trace level ALL, the result returned by select statements is also included in the trace file

The SQL trace allows you to analyze the response time of SQL statements within an object. To perform an analyses, tracing has to be activated via: Trace Configuration dialog box => specify a name for the trace file => set the trace status to Active. Once SQL statement is finished deactivate tracing (Inactive).

Information collected by the SQL trace includes overall execution time of each statement, the number of records affected, potential errors (for example, unique constraint violations) that were reported, the database connection being used, and so on. So the SQL trace is a good starting point for understanding executed statements and their potential effect on the overall application and system performance, as well as for identifying potential performance bottlenecks at statement level. SQL trace information is saved as an executable python program that you can access on the Diagnosis Files tab of the Administration editor.

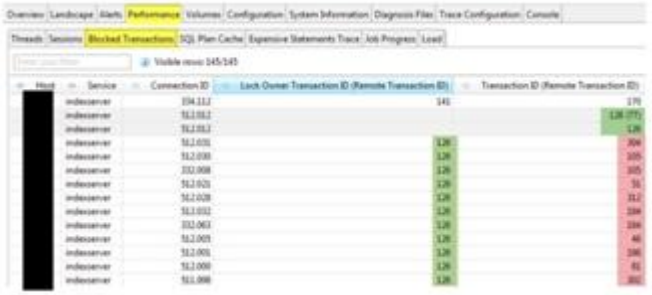

You can monitor all sessions in your landscape in the Administration editor on the Performance > Sessions sub-tab

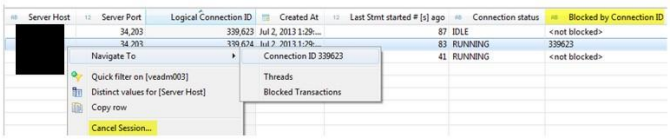

6. Load Monitoring

A graphical display of a range of system performance indicators is available in the Administration editor on the Performance > Load sub-tab. You can use the load graph for performance monitoring and analysis. For example, you can use it to get a general idea about how many blocked transactions exist now and in the past, or troubleshoot the root cause of slow statement performance

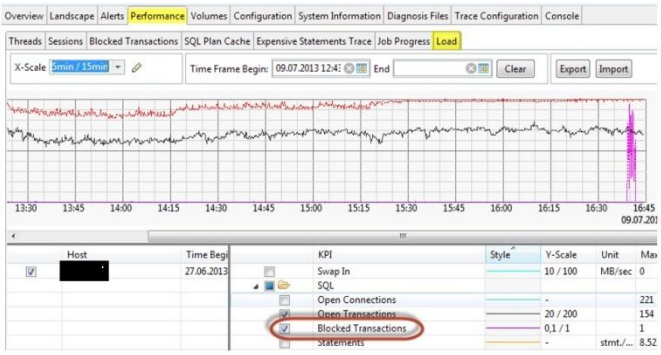

6.1 Plan Visualizer

To help you understand and analyze the execution plan of an SQL statement, you can generate a graphical view of the plan. This graphic is a tool for studying performance of queries on SAP HANA databases. You can explore the graphic further, for example, you can expand, collapse, or rearrange nodes on the screen. You can also save the graphic as an image or XML file, for example, so you can submit it as part of a support query.
A graphical representation of the query, with estimated performance, is displayed:

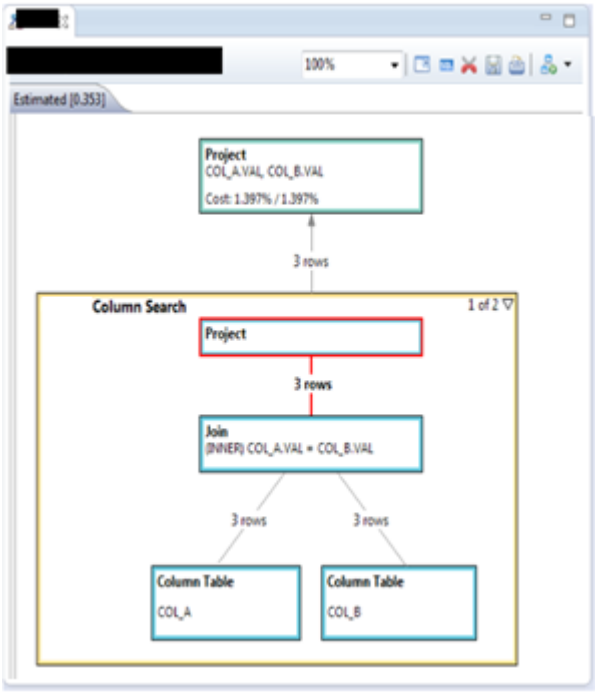

7. Conclusion

SAP HANA is a in memory complex database which requires to analyze thoroughly for large enterprise database systems. Being columnar Database for OLTP and OLAP, processing speed is much higher than traditional RDMS databases. A BI report used to take 10 minutes to process on

a Oracle database, consuming less resources and less than 1 minute (50 milli seconds)

Declarations

Ethics approval and consent to participate: Not Applicable

Consent for publication: All authors have consent to submit this paper to Journal of Cloud Computing. Also, we confirm that this paper or any part of this paper did not submit any where

Availability of data and materials: Not Applicable

Competing interests: Not Applicable

Funding: Not Applicable

References

[1] What is SAP HANA. [Online]. Available at: https://www.ibm.com/topics/sap-hana
[2] SAP HANA Installing and administering. SAP TRAINING. [Online]. Available at: https://training.sap.com/course/ha200-sap-hana-installing-and-administering-classroom-019-g-en/?
[3] SAP HANA Administration Guide. SAP Help. [Online]. Available at: https://help.sap.com/docs/SAP_HANA_PLATFORM/6b94445c94ae495c83a19646e7c3fd56/bde79b28bb5710149d6eee5e 75fe7f17.html
[4] Memory Usage in the SAP HANA Database. SAP Help. [Online]. Available at: https://help.sap.com/docs/SAP_HANA_PLATFORM/6b94445c94ae495c83a19646e7c3fd56/bde79b28bb5710149d6eee5e 75fe7f17.html
[5] How to Start and Stop an Instance of SAP HANA, SAP PRESS. [Online]. Available at: https://blog.sap-press.com/how-to-start-and-stop-an-instance-of-sap-hana#:~:text=With%20SAP%20HANA%20Studio%201%20Starting%20Choose%20Start,to%20look%20at%20the%20instance%20trace%20files.%20
[6] Sap-hana-tmpfs.service, Github. [Online]. Available at: https://github.com/aws-samples/aws-sap-hana-fast-restart-scripts/blob/main/sap-hana-tmpfs.service
[7] Sap-hana-tmpfs.sh, Github. [Online]. Available at: https://github.com/aws-samples/aws-sap-hana-fast-restart-scripts/blob/main/sap-hana-tmpfs.sh
[8] SAP HANA on AWS. Amazon Doc. [Online]. Available at: https://aws.amazon.com/sap/solutions/saphana/
[9] Help Portal Documentation, SAP HANA Platform. SAP Help. [Online]. Available at: https://help.sap.com/docs/SAP_HANA_PLATFORM
[10] Gilbert Wong, Data Management for Technical Tables for SAP S/4HANA, SAP Blogs. [Online]. Available at: https://blogs.sap.com/2021/02/09/data-management-for-technical-tables-for-sap-s-4hana/
[11] Big Data, Google Trends. [Online]. Available at: https://trends.google.com/trends/explore?date=all&q=Big%20Data
[12] Richard Bremer (Author), Lars Breddemann (Author), SAP HANA Administration, SAP PRESS [Online]. Available at: https://www.sap-press.com/sap-hana-administration_3506/
[13] TADM10_1 & TADM10_2 SAP Books, Technical Implementation and Operation I of SAP S/4HANA and SAP Business Suite
[14] TADM_51, SAP NetWeaver AS - DB Operation (Oracle), [Online}. Available at: https://cdn20.training.sap.com/cdn/course-pdf/TADM51_EN_Col15_ILT_FV_CO_A4.pdf/G/EN/TADM51/015
[15] Memory Usage in the SAP HANA Database. SAP Help. [online]. Available at: https://help.sap.com/docs/SAP_HANA_PLATFORM/6b94445c94ae495c83a19646e7c3fd56/bde79b28bb5710149d6eee5e 75fe7f17.html?version=2.0.04
[16] What is SAP HANA. Sap.com. [online]. Available at: https://www.sap.com/products/technology-platform/hana/what-is-sap-hana.html