

# Optimization of Memory: SAP HANA Database Perspective - Part 1

Rajesh Azmeera<sup>1</sup>, Rhea Khanna<sup>2</sup>, Deepak Nanuru Yagamurthy<sup>3</sup>

<sup>1</sup>Technology Professional, Department of Information and Technology, Stryker, US (Corresponding Author)

<sup>2</sup>Technology Professional, Department of Research, Microsoft, US

<sup>3</sup>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.

## 2. Analysis

### Performance and High Resource Utilization

By observing the general symptoms shown by the system such as poor performance, high memory usage, paging or column store unloads we can start to narrow down the possible causes as a first step in analyzing the issue. Analysis of possible high memory consumption Analysis of possible out-of-memory situations

Analysis of possible Operating System Level Paging  
Analysis of Column Store Unload

Analysis of slow individual SQL-Statements or with increasingly long runtime

### Problem Analysis Principles

Each performance analysis has to be related to a valid record.

Validate the assignment when receiving a record. Review details captured in the incident record, and make corrections if needed. If description of the issue is not clear ask requestor for further information. Document all activities undertaken for resolution in the incident record to provide a complete and accurate account of events.

Perform a detailed investigation and diagnosis with the aim to resolve the incident and restore normal service operation as quickly as possible. This may include the employment of more technical scripts, diagnostic tools, knowledge bases, support documentation and change/release schedules available to deeper levels of support, or the engagement of other internal and or external service providers for assistance to identify the series of events, failing components and a resolution.

If required SAP platform support or Network infrastructure support has to be involved, if problem is not related to SAP product. If the Incident's duration is extended during investigation and diagnosis, it may be necessary to perform functional escalation to the next level support team.

Throughout this activity, notify the end user of progress.

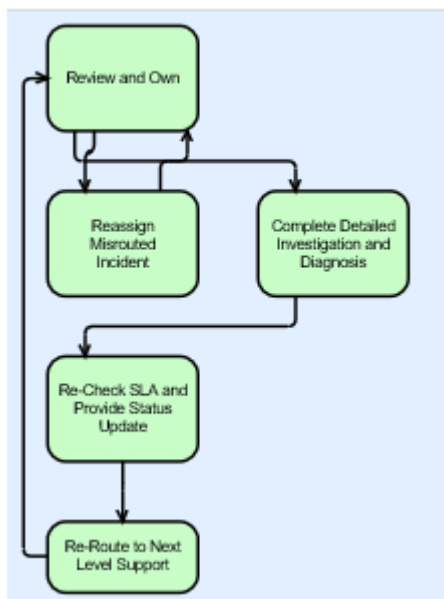
If problem cannot be solved and SAP support has to be involved, creation of OSS incident is necessary. This incident should contain all information which describe actual problem (with possible procedure how to reproduce the problem), open required connection to the affected instance and if possible also collect information for SAP Support described in chapter Collecting Diagnosis Information for SAP support

Volume 11 Issue 4, April 2022

[www.ijsr.net](http://www.ijsr.net)

Licensed Under Creative Commons Attribution CC BY

Support Specialist



### 3. Case Study

#### Database or Service Crash

The following main reasons exist for restarts and crashes:

Explicit restart by administrator or software component  
 Software bugs (SAP HANA, operating system) Hardware problems

#### Configuration issues

Error	Reason	Troubleshooting Steps
Error during asynchronous file transfer rc=5: Input/output error	Error accessing a file on disk	The input/output error can also happen if a file in persistence is no longer accessible for a service because a SAP HANA failover mechanism has moved it away from the host. See OSS note 2062631- that describes ping-pong failover situations that can be responsible.  Check in collaboration with your hardware and OS partners why the mentioned file can't be accessed sporadically or permanently.
Cannot lock file "" for write access rc=11: Resource temporarily unavailable	Error locking a file on disk	Another process may have locked the file (backup, cluster, virus scan, ...) See OSS note 1880382 - HANA: persistence startup error: Cannot lock file and use lsof to check for processes having set a lock on the file  If the error happens during a SAP HANA failover, it might be a timing issue or a consequence of another problem. Example: The nameserver fails over and the new nameserver tries to set a lock on its DATA file although the previous nameserver is somehow still running and locking the file. In this case please also check which error resulted in the failover and analyze that problem to avoid it in the future.  The LINUX aio-max-nr is not set large enough.  Check in collaboration with your hardware and OS partners why the mentioned file can't be locked sporadically or permanently
Error during asynchronous file transfer rc=28: No space left on device	No disk quota or disk space available	Make sure that sufficient space is available on filesystem level.  If you have configured disk quota limitations, make sure that they aren't hit See OSS note 1921354 - RC=28: No space left on device error, Startup of HANA database with IBM GPFS filesystem fails and OSS note 2051052 -  "No space left on device" when df shows free space on Lenovo eX5 with GPFS for GPFS and OSS note 2251334 - "No space left on device" when df shows free space on Lenovo x86_64 for GPFS and XFS.  Use GPFS version 3.4.0.23 or higher where a related GPFS bug is fixed (OSS note 1846872 - "No space left on device" error reported from HANA)
ljit_server_Compilervice.cpp: Emergency shutdown because of out-of-memory	Lack of available memory	Usually an out-of-memory (OOM) situation doesn't result in a crash, but if the compile server process runs out of memory, it may perform an emergency shutdown (OSS note 2261388 - SAP HANA Database: Compile server restarts during out of memory situations). In this case only a OOM trace file is generated, but no crashdump.  See OSS note 1999997 - FAQ: SAP HANA Memory in order to analyze and eliminate memory bottlenecks.  In the worst case the compileserver will not restart after the emergency shutdown with SAP HANA SPS <= 09. See OSS note 2229604 - Compileserver crashes and remains offline after OOM event for more information.

**Disk Full Error of Log Volume**

In a log volume full situation, no free log-segments are available for re-use & no more new log segments can be allocated due to limited disk quota/size, thus the database cannot be started or stops accepting requests. The database tracefile of a particular service contain "rc=24 no space left on device errors" for the basepath\_logvolumes or basepath\_logbackup.

Do not remove either data files or log files using operating system tools as this will corrupt the database! Follow OSS Note 1679938 to temporarily free up space in the log volume, this way you should be able to start up the database for root cause analysis and problem resolution (after the situation is resolved, undo the workaround). For analyses of root cause of full log situation continue according to the OSS note 2083715 - Analyzing log volume full situations

**Corrupt Log Segments**

The trace file will say something like cannot find or cannot read a log segment at a hexadecimal address, the only resolution to a corrupt log segment is to do a recovery that does not involve that log segment.

**Authorization Issues**

In the trace we will see the message 'not authorized' in the trace, in this scenario check as the <SID>adm user and see if that user can make a file in the location specified in the trace to verify this. If you cannot create the file run the chmod command on the folder to allow reading and writing (ie chmod 764). For future analyses check chapter Security-Related Issues with related HOW documents which are in this chapter.

**Hardware Issue**

There is no generic line in the trace that would point to hardware, but if the issue is OS related or a disk cannot mount please follow the hardware portion of the survival guide.

**Case Study 2:****System Hanging Situations:**

This section explains what checks to carry out and how to perform those checks if your SAP HANA instance is hanging. The database is said to be hanging when it no longer responds to queries that are executed against it.

The source of the system standstill might be related to any of the components involved, for example, the storage, OS and hardware, network, SAP HANA database or the application layer. For troubleshooting it is essential to collect information about the context of the active threads in the SAP HANA database. As SQL statements cannot usually be executed for analysis, the SAP service team might perform the following steps if it is still possible to log on to the OS of the master host.

Collect a runtime dump

Collect CPU, memory and other information about threads currently running in the system by executing the command top -H and taking a screenshot of the output

OSS note 1813020 - How to generate a runtime dump on SAP HANA

**4. Collecting Diagnosis Information for SAP Support**

**Note:** If there is no connection to the database, the Administration editor opens in diagnosis mode and you will be prompted to enter the credentials of the <sid>adm user. If you are a tenant database administrator and there is no connection to your tenant database, you cannot proceed. Only the system administrator can collect diagnosis information from the system database.

**Procedure**

Choose Diagnosis Information -> Collect Specify the scope of information to be collected:

**Collect all diagnosis information**

Select this option if you want to collect all diagnosis information for a specific time period, by default the last 7 days. If you also want information from system views, then select Include system views.

**Note:** If you are connected to the system database of a multiple-container system, only information from the system views of the system database will be collected. Information from the system views of tenant databases will not be collected regardless of this option. Information from system views is collected through the execution of SQL statements, which may impact performance. In addition, the database must be online, so this option is not available in diagnosis mode.

**Create and collect one or multiple sets of runtime environment (RTE) dump files**

Select this option if you want to restrict the information collected to one or more RTE dump files.

You can configure the creation and collection of dump files by specifying the following additional information. The index server(s) from which RTE dump files are to be collected

The number of RTE dump file sets to be collected (possible values are 1, 2, 3, 4, and 5)

The interval (in minutes) at which RTE dump files are to be collected (possible values are 1, 5, 10, 15, and 30). The default value is 1.

Older systems do not support all of the above options. It may not be possible to exclude system views from collection or you may require operating system (<sid>adm) user access to do so.

The system collects the relevant information and saves it to a zip file. This may take some time and can be allowed to run in the background.

If you are connected to the system database of a multiple-container system, information from all tenant databases is collected and saved to separate zip files.

To download the zip file containing the collected diagnosis information, proceed as follows: Choose Diagnosis Files List


The Diagnosis Information dialog box opens. The zip file containing the collected diagnosis information is listed together with any other zip files of previously collected information.

Select the relevant zip file and choose Download Collection. Specify the download location.

**Root Causes and Solutions**

For standard scale-up configuration (with or without HA\DR) you can find information about memory consumption of server and HANA DB in SAP HANA Studio.

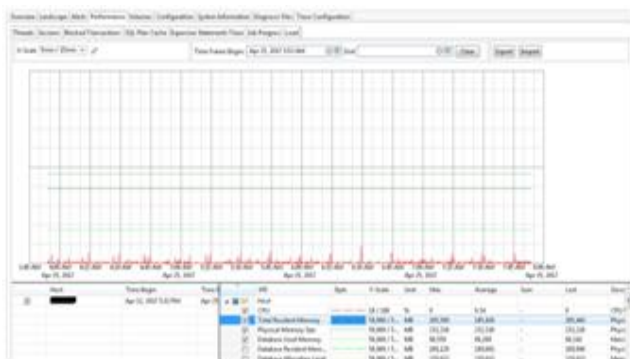
Below are only 3 examples of all available metrics. High Level information about physical memory:



Allocation limit and used memory for each service:

Service	Host	Start Time	Process ID	Memory	Used Memory (MB)	Peak Used Memory (MB)	Effective Allocation Limit (MB)	Usage (%)
hdi	ddhd078	Apr 11, 2022 08:15:01.000	2239	1,712	1,876	1,876	10,000	18.76%
hdi	ddhd078	Apr 11, 2022 08:15:01.000	2239	7	7	7	10,000	0.07%
hdi	ddhd078	Apr 11, 2022 08:15:01.000	2239	1,027	82,824	82,824	100,000	82.82%
hdi	ddhd078	Apr 11, 2022 08:15:01.000	2239	1,027	1,027	1,027	100,000	1.03%
hdi	ddhd078	Apr 11, 2022 08:15:01.000	2239	1,027	1,027	1,027	100,000	1.03%
hdi	ddhd078	Apr 11, 2022 08:15:01.000	2239	1,027	1,027	1,027	100,000	1.03%
hdi	ddhd078	Apr 11, 2022 08:15:01.000	2239	1,027	1,027	1,027	100,000	1.03%
hdi	ddhd078	Apr 11, 2022 08:15:01.000	2239	1,027	1,027	1,027	100,000	1.03%
hdi	ddhd078	Apr 11, 2022 08:15:01.000	2239	1,027	1,027	1,027	100,000	1.03%

Load on historical data with different filters:



**Memory Information from Logs and Traces**

In case of critical memory issues more detailed information can be found in logs and trace files. In this case the service team can identify possible memory-related errors for example in the SAP HANA system alert trace file. In addition it will be checked if an out of memory trace file has been created. Via error message analysis on application side during high memory usage further detailed analysis can be conducted. Search for the strings “memory”, “allocate”, or “OOM” (case-insensitive).

Investigate error messages seen on the application side that occurred at times of high memory usage. If the application is an SAP NetWeaver system, good starting points for analysis are System Log (SM21), ABAP Runtime Error (ST22), and Job Selection (SM37).

In-depth analysis of a memory-intensive SQL statement can

be addressed within this service. In this case diagnosis information (full system info dump) and performance trace details can be included accordingly. For more details about it check chapter Tools and Tracing.

**Memory Information from SQL Commands**



If the database resident memory of all SAP HANA databases on the same host is significantly smaller than the total resident memory the service team will check which processes outside of the SAP HANA database are responsible for the additional memory requirements.

**Typical memory consumers are:**

- Operating system (for example: caches, mapping structures)
- Third party tools (for example: backup, virus scanner)

**Expensive SQL Statements**

SQL statements processing a high amount of data or using inefficient processing strategies can be responsible for increased memory requirements. In this case the service team will analyze together with the customer and partner teams to optimize the SQL statements.

For next analyses use chapters SQL Statement Analysis or Setting a Memory Limit for SQL Statements.

**Transactional Problems**

High memory consumption can be caused by problems with transactions. In some cases, high memory consumption is caused by wait situations, which can have different reasons such as long-running or unclosed cursors, blocked transactions, or hanging threads.

As one of the negative impacts, used memory is not released any more. In particular, the number of table versions can grow up to more than 8,000,000 which is considered the amount where an action is required. Within this service the SAP support team can analyze these problems and provide the required recommendations.

For next analyses use chapter Transactional Problems.


**4.8 Used Space Much Smaller than Allocated Space**

In order to optimize performance by minimizing the memory management overhead or due to fragmentation, SAP HANA may allocate additional memory rather than reusing free space within the already allocated memory. This can lead to undesired effects that the SAP HANA memory footprint increases without apparent need. The SAP HANA license checks against allocated space, so from a licensing

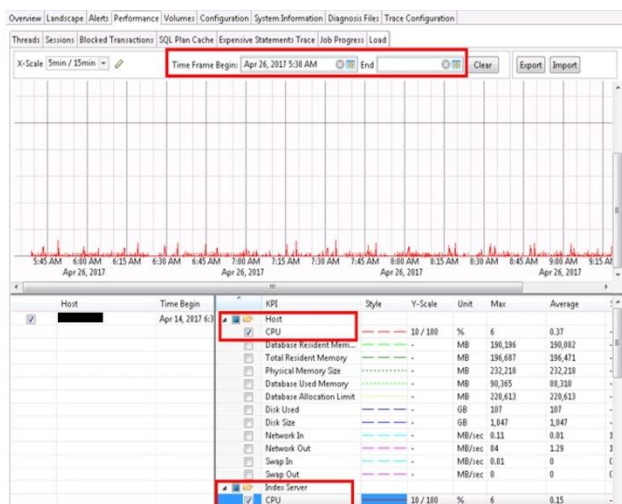
perspective it is important to keep the allocated space below the license limit. In order to limit the amount of allocated space you can set the parameter `global_allocation_limit` to a value not larger than the maximum memory that should be allocated

## 5. CPU Related Issues

CPU related issues are indicated by alerts issued or in views in the SAP HANA Studio. The following alerts may indicate CPU resource problems and will be analyzed during the service accordingly: Host CPU Usage (Alert 5), Most recent save point operation (Alert 28) and Save point duration (Alert 54).



The Load graph is currently showing actual CPU consumption/ or shows high consumption in the past:



## Bind Processes to CPUs

If the physical hardware on a host needs to be shared with other processes, it may be useful to assign a set of cores to an SAP HANA process. This is achieved by assigning an affinity to logical cores. Note that the logical cores used in the configuration do not have a fixed association to the physical cores on the hardware. The commands listed below can be used to discover this setting. The information on the cores provided by `sysfs` is used to restrict the cores available to SAP HANA in the `daemon.ini` file. `Sysfs` is a virtual file system and is used to access information about hardware. General information on `sysfs` is available in the Linux kernel documentation.

For Xen and VmWare, the users in the VM guest system see what is configured in the VM host. So the quality of the reported information depends on the configuration of the VM guest. Therefore SAP cannot give any performance

guarantees in this case. For Linux containers (like Docker) `/proc/cpuinfo` reports more cores than are actually available. As Linux containers use the same mechanism to limit CPU resources per container as SAP does for the affinity, the affinity for the container also applies to the complete SAP HANA system. This means that if the Linux container is configured to use logical cores 0-4, then only those will be available to SAP HANA processes.

There are a number of steps to analyze the topology of sockets and cores. This information can be used to define the desired setting for the affinity of SAP HANA processes.

Analyze physical and logical system details of CPU architecture using the `lscpu` command. This command returns a listing of details of the system architecture.

## 6. Limitations and Future Study

Memory is a vast topic and issues related to them can't be predicted. It depends on the behavior of Hardware, Database, OS and applications. Few cases are covered in this article. In another article few more cases will be analyzed. Hardware vendors do not provide SAS for SAP applications but limited to its sizes. Also, disk, CPU, I/O case studies, delta, main merge, MVCC will be covered.

## 7. Conclusion

SAP HANA is a memory complex database which requires to be analyzed thoroughly for large enterprise database systems. Being a columnar Database for OLTP and OLAP, processing speed is much higher than traditional RDBMS databases. A BI report used to take 10 minutes to process on an Oracle database, consuming less resources and less than 1 minute (50 milliseconds)

## Declarations

**Ethics approval and consent to participate:** Not Applicable

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

**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/bde79b28bb5710149d6eee5e75fe7f17.html](https://help.sap.com/docs/SAP_HANA_PLATFORM/6b94445c94ae495c83a19646e7c3fd56/bde79b28bb5710149d6eee5e75fe7f17.html)
- [4] Memory Usage in the SAP HANA Database. SAP Help. [Online]. Available at:  
[https://help.sap.com/docs/SAP\\_HANA\\_PLATFORM/6b94445c94ae495c83a19646e7c3fd56/bde79b28bb5710149d6eee5e75fe7f17.html](https://help.sap.com/docs/SAP_HANA_PLATFORM/6b94445c94ae495c83a19646e7c3fd56/bde79b28bb5710149d6eee5e75fe7f17.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](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/](https://www.sap-press.com/sap-hana-administration_3506/)
- [12] TADM10\_1 & TADM10\_2 SAP Books, Technical Implementation and Operation I of SAP S/4HANA and SAP Business Suite
- [13] 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](https://cdn20.training.sap.com/cdn/course-pdf/TADM51_EN_Col15_ILT_FV_CO_A4.pdf/G/EN/TADM51/015)
- [14] Memory Usage in the SAP HANA Database. SAP Help. [online]. Available at:  
[https://help.sap.com/docs/SAP\\_HANA\\_PLATFORM/6b94445c94ae495c83a19646e7c3fd56/bde79b28bb5710149d6eee5e75fe7f17.html?version=2.0.04](https://help.sap.com/docs/SAP_HANA_PLATFORM/6b94445c94ae495c83a19646e7c3fd56/bde79b28bb5710149d6eee5e75fe7f17.html?version=2.0.04)
- [15] What is SAP HANA. Sap.com. [online]. Available at:  
<https://www.sap.com/products/technology-platform/hana/what-is-sap-hana.html>