International Journal of Research in Engineering and Science (IJRES)
ISSN (Online): 2320-9364, ISSN (Print): 2320-9356
www.ijres.org Volume 11 Issue 10 ǁ October 2023 ǁ PP. 205-212

SAP HANA Database Restart At Blazingly Fast Speed: An Overview

Rajesh Azmeera, *(https://orcid.org/0009-0005-4643-1599)*
Rhea Khanna, *(https://orcid.org/0009-0001-4625-5100)*
Deepak Nanuru Yagamurthy, *(https://orcid.org/0009-0009-9546-6615)*

Abstract—*The year-on-year increase in the usage of data requirements for IT enterprises has paved the way for significant improvements in database technology for storing and retrieving data stored on them at lightning-fast speeds. For instance, 6 TB DB without fast restart on AWS takes 5 min for restart and 2 hours for loading the database into memory while 6 TB DB with fast restart takes 5 min to restart and 15 min for loading the database into memory. As a thumb rule, the DB size is directly proportional to the restart time; higher the size of the database, longer is the restart time. With SAP HANA DB, the entire database must be fully loaded into the memory for starting and running the SAP applications, otherwise enterprises won't be able to kick-start their operations. New-age enterprises are experiencing a lot of pain points associated with digital database restarts. The top pain points are high downtime and revenue losses. What is obvious is there is a rising challenge for enterprises in the way they deal with maintenance of the database and downtime during a database restart. SAP HANA DB has a time-tested and proven solution for this.*

Keywords—*SAP HANANA DB Restart, Fast Digital Database Restart, SAP In-memory DB Start Stop HANA SYSTEM*

Date of Submission: 05-10-2023 Date of acceptance: 19-10-2023

I. INTRODUCTION

SAP HANA (High Performance Analytic Appliance) is an in-memory RDBMS: multi-model database that is designed to store data in its memory rather than on the disk. This results in data processing speeds that are astronomically larger than the speeds attained through conventional disk-based database systems and the best part is it makes advanced analytics a reality for enterprises.

SAP 4/HANA system acts as a robust platform for businesses to run SAP applications and manage data in on-premise, cloud and hybrid environments. SAP HANA handles data from multiple areas of an organization.
Traditional business documents: SLAs, contracts & spreadsheets

UX/UI (User Experience/User Interface): web landing pages, Subscription forms, & emails

IoT: sensor data that resides on the mechanical and digital devices connected by the Internet of Things (IoT) in warehouses, workplaces, stores, & more

Mobile: mobile data of your workforce and customers that resides on smartphones

SAP HANA is not just designed to integrate data but it can also leverage the power of AI and ML to accelerate the process of decision making by offering key insights into a firm's day-to-day operations.

II. LITERATURE REVIEW

The global big data analytics is forecasted to touch a whopping $68 billion USD in revenue in 2025. This is an increase of $53 billion USD from the revenue generated in 2019—$15 billion USD, with a CAGR of 30% during the time frame.

The market trends and data suggest that an estimated 180 Zettabytes of data will be created, consumed, and stored by the world by the end of the year 2025. That number is 92 Zettabytes for the year 2022, setting a new high point for the rapid growth of data that the world has created, consumed and stored since the outbreak of the COVID-19 pandemic. That increase can be mainly attributed to three new global trends that were a result of the partial

shutdowns of businesses for staying safe from COVID-19: remote WFH jobs, cloud usage for remote work, and usage of home entertainment channels (OTT) like Netflix.

III. How to Start and Stop SAP HANA Systems

Step 1: Log into system with User: SIDADM
Step 2: Verify processes with HDB information

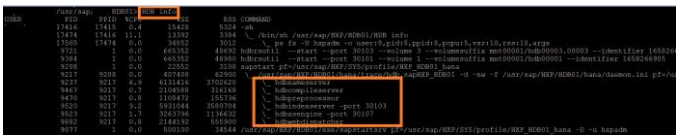

Step 3: Ensure all DB processes are running

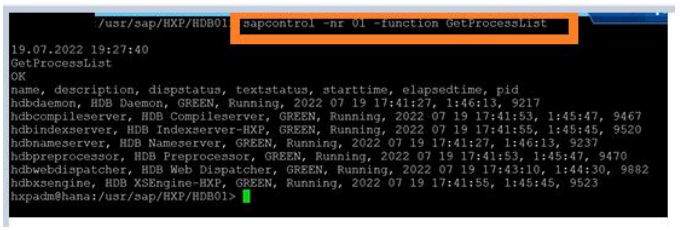

Fig 1. SAP Running Processes

Step 4: Stop HANA DB with SAP Control -nr - function stop system

01 is instance number

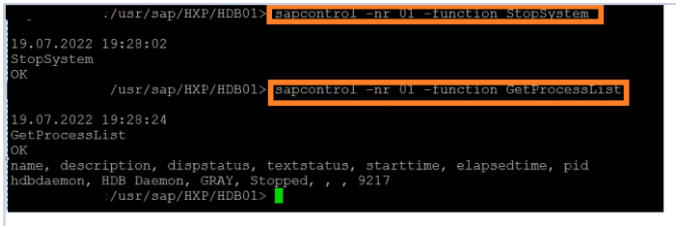

Fig 2. Stopped Processes

GetProcessList shows all processes that are stopped
Step 5: Start HANA DB with command SAP control -nr o1- function start system
01 is instance number

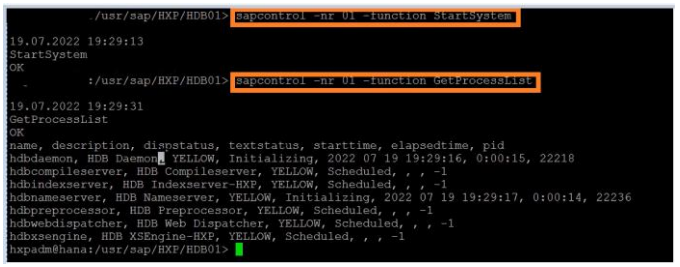

Fig 3. Started Processes

HANA processes that are fully started

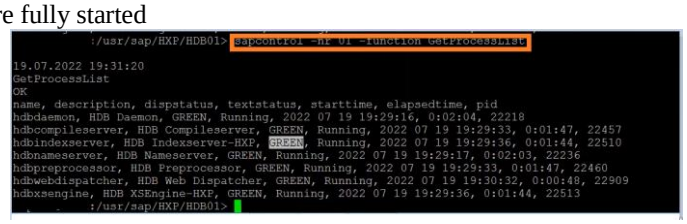

Fig 4. GetProcess List output

DB load takes time and unless finished SAP applications cannot run.

It is interesting to note that 6 TB of HANA DB on AWS took 5 mins to restart and 2 hours to load the entire database into memory while 6 TB of HANA DB using the SAP fast restart option took 5 min to restart all the DB processes and 15 min for loading the database into memory.

IV. Digital Database Fast Restart in SAP HANA

4.1 Solution Overview

Before we take a dive into how SAP fast restart happens, let us first understand the concept SAP Memory Management. The DB processes of SAP HANA run on Linux/Unix OS environment. In case of Linux, it is the OS that is responsible for allocating the required memory for the DB processes. After SAP HANA starts up, the Linux OS allocates memory for program code, program stack, and static data. It is also the responsibility of the OS to allocate memory dynamically when a request is placed by the SAP Memory Manager during the run time. Dynamically allocated memory can be of two types: heap memory and shared memory.
The infographic below depicts memory of code, stack and used table data:

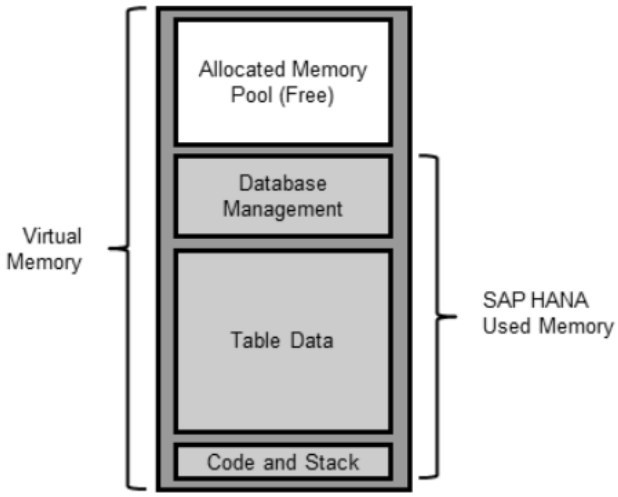

Fig 5. HANA Memory

It may be noted that the memory size allocated to code and stack is just about 6 GB, so most of the SAP HANA used memory space is taken up by table data, database management & computations.
Two (2) types of table storage systems are supported by the SAP HANA DB system— column store and row store, the default table storage is column store.
As far as the column store is concerned, it is made up of two data structures—Main and Delta. Keep in mind that the Main data structure is optimized and compressed for all Read operations while the Delta data structure is used for all Write operations. Note that a Delta Merge operation is used for moving Delta changes from Delta storage to Main storage,
A volatile temporary file system, abbreviated as TMPF, residing in the virtual memory (RAM) is used by the SAP HANA Fast Restart as a measure for prevention and reusability of Main data storage. This is a time-tested and proven method for minimizing the memory load time to a great extent, used in situations when the OS is not started, holds good for the following scenarios.

- SAP HANA Restart
- SAP HANA Service Restart, including index server crash
- SAP HANA Upgrade/Service Pack

The infographic below illustrates a memory use case and how a TMPF can shrink or grow dynamically: The three parameters outlined below are of great relevance for this step.

- Basepath_persistent_memory_volumes: the location of the TMPF file system
- Persistent_memory_global_allocation_limit: by default this parameter is not set to any value. The maximum value for this parameter is the max size of the persistent memory on the host.
- Table_default: this value is set to ON by default but you can also set it to OFF for manually taking control of the persistent memory usage at these 3 levels: table partition or column using the persistent memory switch.

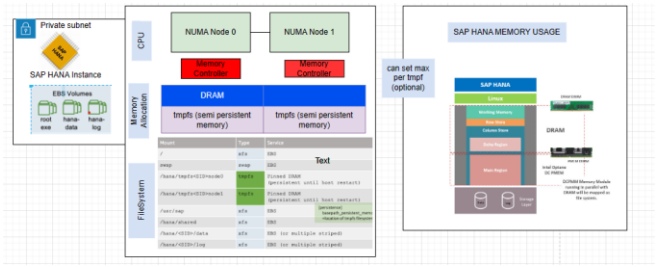

Fig 6: SAP HANA Fast Restart

SAP HANA Instances with required volumes root, hana data and hana log volumes with DRAM memory and necessary file systems / swap /hana/tmpfs /usr/sap /hana/shared /hana/<SID>/data /hana/<SID>/log.
Data loads from delta merge to main memory as Delta is read + write optimized where as main memory is fully read optimized this way DDL and DML statements are won't get conflict each other. These are requirements for HANA restart.
SAP highly recommends that businesses run SAP HANA 2.0 SP4 or higher as it doesn't impact SAP HANA online performance and sizing KPIs. Additionally, to maintain consistency, you must ensure that the SAP HANA Fast Restart option is enabled across your production and non-production system environments.

V. Test Results with and without SAP HANA Fast Restart Enabled

For testing purposes, data sample is taken with size of 1TB. I created multiple tables with test data in the DBACOCKPIT schema. Let's look at the restart times for a database size of 1,062 GB on an x2idn.32xlarge EC2 instance with 128 vCPUs, and 2048 (GiB) of memory with and without SAP HANA Fast Restart.
Operating System: SUSE Linux Enterprise Server 12 SP4
SAP HANA Version: 2.00.050.00.1592305219 (fa/hana2sp05)
HANA DB Size: 1,062 GB
Instance Type: x2idn.32xlarge
Storage Type: gp2/gp3

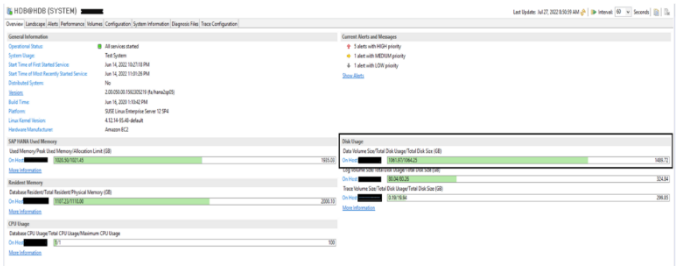

Fig 7. Disk size

5.1 Start up Load Time without Fast Restart

Index Server Trace:

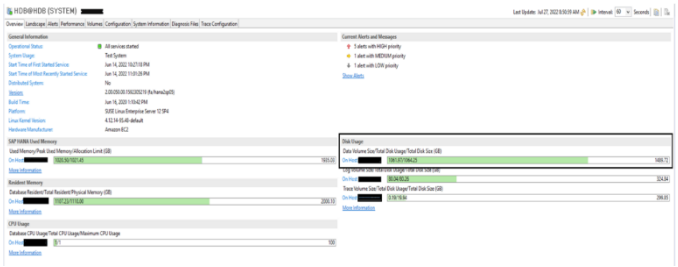

Fig 8. Data volume size

```
[57833]{-1}[12/-1] 2022-08-03 08:53:44.790361 i
TableReload       TRexApiSystem.cpp(00376) : Starting preload of table
HDB::DBACOCKPIT:LOADGENen
[57831]{-1}[-1/-1] 2022-08-03 09:16:28.450727 i
Service_Startup  TREXIndexServer.cpp(02059) : Pre-/Re-Loading of column store tables
finished
```

Fig 9. Load Tables

To fully load the column store, it almost took 23 minutes.

5.2 Start Up Load Time With Fast Restart

Index Server Trace:

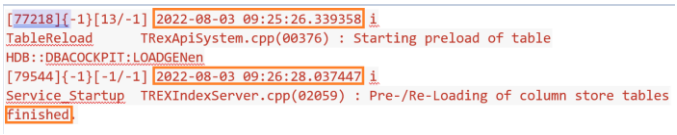

Fig 10. Completely Loaded

To fully load the column store, it only took 1 minute!

5.3 Result:

When the SAP HANA Fast Restart option is used, the startup load time of the DB got reduced drastically. In this illustrative example, from 23 minutes without Fast Restart to 1 minute with Fast Restart.

6. Step-by-Step Procedure to Implement SAP HANA Fast Restart

Here are the steps involved in implementing SAP HANA FAST Restart in an SAP certified EC2 instance.

Step 1 – Determine the memory for each CPU socket

cat /sys/devices/system/node/node*/meminfo | grep MemTotal | awsk 'BEGIN {printf %10s | %20s\n", "NUMA NODE", "MEMORY GB"; while i++ < 33) printf "-"; printf"\n"} {printf "%10d | %20.3f\n", $2, $4/1048576}'

```
NUMA NODE |          MEMORY .L
-------------------------------------------
         0 |           1000.034
         1 |           1000.067
```

Fig 13. Output

Step 2 – Create the mount points. Create 1 TMPFS per NUMA node. I am creating 2 mount points as x2idn.32xlarge has 2 NUMA nodes with 1000GB memory each.

```
mkdir -p /hana/tmpfs0/<SID>
mkdir -p /hana/tmpfs1/<SID>
chown -R <sid>adm:sapsys /hana/tmpfs*/<SID>
chmod 777 -R /hana/tmpfs*/<SID>
```

Fig 11. Permissions

Step 3 – Add the following lines to /etc/fstab

```
tmpfs<SID>0 /hana/tmpfs0/<SID> tmpfs rw,relatime,mpol=prefer:0
tmpfs<SID>1 /hana/tmpfs1/<SID> tmpfs rw,relatime,mpol=prefer:1
```

Fig 12. Temp FS

Step 4 (Optional) – To limit the memory allocated to the TMPFS filesystems, it is possible by passing the parameter "size". In the below example I am limiting the memory to 250G.

tmpfs<SID>0 /hana/tmpfs0/<SID> tmpfs rw,relatime,mpol=prefer:0,size=250G
tmpfs<SID>1 /hana/tmpfs1/<SID> tmpfs rw,relatime,mpol=prefer:1,size=250G

Step 5 – Mount the filesystems you have just added to /etc/fstab

Mount -a

Step 6 – Alter the following parameters using HANA Studio, or hdbsql (run as the user <dbsid>adm).

hdbsql -u system -p <password> "ALTER SYSTEM ALTER CONFIGURATION ('global.ini','SYSTEM') set ('persistence','basepath_persistent_memory_volumes') = '/hana/tmpfs0/<SID>;/hana/tmpfs1/<SID>' with reconfigure;"

hdbsql -u system -p <password> "ALTER SYSTEM ALTER CONFIGURATION ('indexserver.ini', 'SYSTEM') SET ('persistent_memory', 'table_default') = 'on' WITH RECONFIGURE;"

Step 7 (Optional) – Alter (or use default) parameters related to persistent_memory_global_allocation_limit and table_default. I am using default values for both the parameters.

persistent_memory_global_allocation_limit = Max Size (no limit is specified – default)
table_default = ON (default)
Step 8 – Restart SAP HANA in order to activate the changes
HDB start
Step 9 – Check tmpfs consumption.
Note: The first restart after implementing Steps 1 to 6 will take the same amount of time as without tmpfs. The subsequent restarts will be faster.

df -h | head -1; df -h| grep tmpfs<SID>

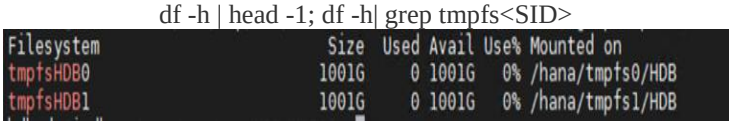

Fig 14. Tmpfs

Step 10 – Check table consistency (optional)

CALL CHECK_TABLE_CONSISTENCY('CHECK_PERSISTENT_MEMORY_CHECKSUM', NULL, NULL);

VI. Postscript

Once the setup of SAP HANA FSO is done successfully, a link is established between the TMPFS, CPU and memory specifications of the EC2 instance. If you desire to alter the flexibility of the instance sizes, then you must consider setting up a systemd service to resize TMPFS.
Step 1 – Navigate to /etc/systemd/system
Step 2 – Create your service. E.g. sap-hana-tmpfs.service. You can use this script as an example to create your service:

```
1   [Unit]
2   Description=SAP HANA TMPFS
3   ConditionFileIsExecutable=/etc/rc.d/sap-hana-tmpfs.sh
4   After=local-fs.target
5   OnFailure=
6
7   [Service]
8   Type=forking
9   ExecStart=/etc/rc.d/sap-hana-tmpfs.sh start
10  ExecStop=/etc/rc.d/sap-hana-tmpfs.sh stop
11  RemainAfterExit=yes
12
13  [Install]
14  WantedBy=multi-user.target
```

Fig 15. Tmps.sh

Step 3 – Create the script which is called by the service created in Step 2. E.g. sap-hana-tmpfs.sh script at /etc/rc.d/sap-hana-tmpfs.sh. You can use this script as an example:

```
#!/bin/bash
SID=HDB
COUNT=0
SIDADM=${SID,,}adm

start() {
        for each in $(ls -d /sys/devices/system/node/node*); do
                #
                mkdir -p /hana/tmpfs${COUNT}/${SID}
                mount tmpfs${SID}${COUNT} -t tmpfs -o mpol=prefer:${COUNT} /hana/tmpfs${COUNT}/${SID}
                COUNT=$((COUNT} + 1))
                chown -R $SIDADM:sapsys /hana/tmpfs*/${SID}
                chmod 777 -R /hana/tmpfs*/${SID}
        done
}

stop() {
        #
        for each in $(mount | grep '/hana/tmpfs' | awk '{ print $3 }'); do
                #
                umount $each
        done
}

# See how we were called.
case "$1" in
    start)
        start
        ;;
    stop)
        stop
        ;;
    restart|force-reload)
        stop
        start
        ;;
    reload)
        stop
        start
        ;;
    *)
        echo $"Usage: $0 {start|stop|restart|force-reload|reload}"
        exit 2
esac
```

Fig 16. Sap-hana-tmpfs.sh

Step 4 – Reload the service files to include the new service:

sudo systemctl daemon-reload

Step 5 – Start the service:

sudo systemctl start <your_service_name>

Step 6 – Check status of your service

sudo systemctl status <your_service_name>

Note: After changing the instance type, check the HANA parameters manually to ensure the configuration is according to your requirement.

VII. Conclusion

Regardless of the industry, technology is something that gets reinvented over time. What is latest today soon becomes obsolete! Usually, with greater costs and complexities in technology comes higher perks but this isn't the case with SAP HANA DB—the cost of setting up a HANA database for fast restart is very minimal and the perks of using it are very attractive. Businesses are leveraging the power of HANA DB for reduced downtime during a database restart without impacting the underlying infrastructure at low costs.

Database Size	Time Taken
1 TB DB without Fast restart config	1 minute to restart + 23 minutes to load DB into Memory
1 TB DB with Fast restart config	1 minute to restart + 1 minute to load DB into Memory

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
- Authors' contributions:

R.A. Topics Covered: Digital Database Fast restart in SAP HANA, Test results with and without SAP HANA Fast restart enabled, Step-by-Step Procedure to implement SAP HANA Fast restart
R.K. Topics Covered: Abstract, Introduction and How to Start and Stop HANA Systems
D.N.M. Topics Covered: Post script, conclusion and Declarations
All Authors have reviewed the manuscript.

Acknowledgements: Thank you co-authors Rhea Khanna and DeepakNanuru Yagamurthy for their expertise and assistance throughout all aspects of our study and for their help in covering few topics and reviewing the manuscript.

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
[4]. Memory Usage in the SAP HANA Database. SAP Help. [Online]. Available at: 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
[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