

1969700 - SQL Statement Collection for SAP HANA

Version	367	Type	SAP Note
Language	English	Master Language	English
Priority	Recommendations / Additional Info	Category	Consulting
Release Status	Released for Customer	Released On	23.02.2019
Component	HAN-DB (SAP HANA Database)		

Please find the original document at <https://launchpad.support.sap.com/#/notes/1969700>

Symptom

During a SAP HANA database analysis it is required to collect special information.

Last update: 2019/02/23

Reason and Prerequisites

As part of SAP services like

- SAP Technical Performance Optimization - HANA (TPO - HANA; see SAP Note [2177604](#))
- SAP System Administration - HANA

and during the analysis of complex problems it can be required to determine special database information which is partially not available in standard functionalities like SAP HANA Studio (SAP Note [2073112](#)) or DBACOCKPIT (SAP Note [2222220](#)).

Solution

Attached to this SAP Note are two SQL statement collections:

- SQLStatements_1.00.99-.zip: Outdated SQL statements that are only useful for SAP HANA 1.0 SPS 09 and before
- SQLStatements.zip: SQL statements that are useful for SAP HANA 1.0 SPS 10 and above

Some statements are SAP HANA release and statistics server dependent, e.g.:

- *SQL: "HANA_Threads_CurrentThreads"*: Valid for all SAP HANA releases
- *SQL: "HANA_Configuration_MiniChecks_1.00.102.01+"*: Valid for SAP HANA Revisions starting with 1.00.102.01 (i.e. SAP HANA 1.0, Revision 102.01)
- *SQL: "HANA_IO_Savepoints_1.00.80+_ESS"*: Valid for SAP HANA Revisions starting with 1.00.80, restricted to the embedded statistics server (ESS)

In case multiple versions exist, you should always use the best fitting one. For example, if you are on SAP HANA 2.00.023 and there are versions 2.00.000+, 2.00.010+ and 2.00.030+ available, you should use the 2.00.010+ variant. The 2.00.000+ variant would also work, but the result may not be as good as with the 2.00.010+ variant. The 2.00.030+ variant won't work, e.g. because certain used monitoring views or columns don't exist with SAP HANA 2.00.023, yet. The same rule applies across SAP HANA Version levels. For example, variant 1.00.70+ can still be used on SAP HANA 2.00.023 (unless there is a newer fitting variant available).

The following name suffixes can be used to restrict context of specific SQL statements:

- ESS: Embedded statistics server
- CommandGenerator: SQL statement generates another SQL statement that needs to be executed in a second step
- Fallback: Fallback version if standard version doesn't work (e.g. due to permission or performance issues)
- Internal: Not available for public use
- MDC: System database of multitenant database container installation
- SSS: Standalone statistics server

Some of the SQL statements contain configuration settings marked with `"/ * Modification section */` that can be adapted to change the results.

The modification section typically contains settings of the following types:

Number parameters:

Used for various purposes, -1 disables the setting

```
30 NUM_RECORDS /* Limit result set to 30 records */
```

```
-1 NUM_RECORDS /* No limitation of the result set */
```

String parameters (fix set of possible values):

Used to choose from a fix set of different processing options, allowed values are provided in comment

```
'SUM' AGGREGATION_TYPE, /* MAX, AVG, SUM */
```

String parameters (variable values):

Used for various purposes, '%' is placeholder for an arbitrary string (comparable to '*' in other environments), '_' is placeholder for a single character (comparable to '?' in other environments), single '%' disables the setting

```
'indexserver' SERVICE_NAME, /* Perform analysis only for indexserver */
```

```
'%' SERVICE_NAME, /* Perform analysis for all services */
```

```
'%server' SERVICE_NAME, /* Perform analysis for all services ending with 'server' */
```

On / off switches:

Used to enable and disable a certain feature, 'X' for enabling the functionality, ' ' for disabling the functionality

```
'X' ONLY_ACTIVE_THREADS, /* Show only information for active threads */
```

```
' ' ONLY_ACTIVE_THREADS, /* Show information for all threads */
```

Be aware that the retention time of histories is typically related to the configuration of the statistics server. Normally histories of 42 days are configured, but on demand you can adjust it based on your needs. See SAP Note [2147247](#) for more information.

SAP doesn't take over responsibility for the correctness of the commands. The statements can put a significant load of the system, that can result in resource bottlenecks and terminations in the worst case. Therefore you should familiarize yourself with the commands in a non-production environment.

The following table lists typical questions and answers related to the SQL statement collection:

Problem	Details
Accessing or downloading of the attachment fails with an error like: Could not retrieve attachment Edit in CSS not allowed!	Please use the new Launchpad as described in SAP Note 2310449.
Executing a command fails with errors like: -709: Cannot connect to VolumeID=<volume_id>: Cannot connect to host <host>:<port> [Unknown host] -813: Cannot connect to VolumeID=<volume_id>: Cannot connect to host <host>:<port> [Connection refused] -813: Cannot connect to VolumeID=<volume_id>: Cannot connect to host <host>:<port> [Connection timed out]	These errors indicate that SAP HANA hosts and disks can't be reached. It is not a problem of the SQL statements, instead it is an infrastructure issue. See SAP Note 2222200 and take appropriate actions that the SAP HANA host and disk volumes are accessible.
2: general error: indexId=<id> is not found in indexMap 7: feature not supported: Cannot access RS table : <table> in secondary system [sic!]	These errors are issued when commands are run on the secondary system replication site with activated read access (logreplay_readaccess mode, SAP Note 1999880) and tables are queried that are located in row store (e.g. certain tables of the statistics server, SAP Note 2147247). This isn't supported on secondary system replication sites, only column store tables can be queried.
3: fatal error: Execution flow must not reach here 7: feature not supported: scalar subquery is not allowed	A bug with SAP HANA Rev. <= 112.02 can result in terminations with these error messages. As a workaround you can use the NO_JOIN_REMOVAL hint.
7: feature not supported: incorrect CESU-8 string	This error can be reported in the context of M_MERGED_TRACES accesses with SAP HANA <= SPS 09 if traces contain unusual characters like Chinese letters. As a workaround you can remove the M_MERGED_TRACES sections in the mini checks and directly look into the trace files rather than using queries based on M_MERGED_TRACES. If the problem still happens with SAP HANA Rev. >= 102.01, open a SAP incident for analysis purposes.
7: feature not supported: invalid CESU-8 encoding for Unicode string	The error can show up in the context of M_SQL_PLAN_CACHE accesses with SAP HANA Rev. 122.05 or lower and SAP HANA 2.00.000. See SAP Note 2418887 for more information.
7: feature not supported: invalid character encoding	Failing accesses to M_SERVICE_THREAD_SAMPLES with SAP HANA Rev. <= 1.00.122.09 and <= 2.00.002.01 can be caused by the bug described in SAP Note 2453766.
8: invalid argument: Error while parsing trace file	This error can be caused by a bug with SAP HANA <= 2.00.020 when parsing the unloads or loads trace files during M_CS_UNLOADS / M_CS_LOADS access (SAP Note 2127458). If this access is part of a more comprehensive analysis command (e.g. mini checks or timeframe report) you can manually remove the UNION ALL accessing M_CS_UNLOADS as a workaround. Alternatively you can delete the *.unloads.* / *.loads.* trace files from the SAP HANA trace directory. A fix is available with SAP HANA 2.00.021.

257: sql syntax error	This error can be a consequence of a restricted result set size (e.g. due to a "Rows:" restriction in transaction DBACOCKPIT) when running *CommandGenerator* commands. As a consequence the generated SQL statement is incomplete and so it can fail with different errors, in particular with syntax errors. Make sure the output is complete and increase result set restrictions sufficiently if required.
Cannot iterate over result set rows 257: sql syntax error: incorrect syntax near "WITH"	This error happens when a SQL statement starting with WITH is executed in the "System information" section of SAP HANA Studio (SAP Note 2073112). See SAP Note 2405756 and use the "SQL Editor" instead to execute these queries. Starting with SAP HANA Studio 2.3.33 this problem is fixed.
258: insufficient privilege: Not authorized	This error indicates insufficient privileges for the executing user. Consider the following details: Make sure that at least a role like MONITORING is assigned so that accesses to monitoring views are permitted. For queries accessing application tables it is additionally required that the executing user has the permission to access this data. Queries available via internal SAP Note 2082221 fail with this error until they are executed with user SYSTEM or with the SAP_INTERNAL_HANA_SUPPORT role.
259: invalid table name: Could not find table/view DD02L in schema	For technical reasons SQL: "HANA_ABAP_MiniChecks" fails with this error if the command is run in a schema that is different from the SAP ABAP schema.
259: invalid table name: Could not find table/view RSMDATASTATE_EXT in schema	For technical reasons SQL: "HANA_BW_DataTargets" fails with this error if the command is run in a schema that is different from the BW schema.
259: invalid table name: Could not find table/view RSODSO in schema	For technical reasons SQL: "HANA_BW_DSOTypes" and SQL: "HANA_InconsistentDSOTables" fail with this error if the command is run in a schema that is different from the BW schema. For the second command you may use the fallback option SQL: "HANA_InconsistentDSOTables_Fallback".
259: invalid table name: Could not find table/view TBTCO in schema	For technical reasons SQL: "HANA_ABAP_BatchJobs" fails with this error if the command is run in a schema that is different from the SAP ABAP schema.
259: invalid table name: Could not find table/view VBHDR in schema	For technical reasons SQL: "HANA_ABAP_UpdateErrors" fails with this error if the command is run in a schema that is different from the SAP ABAP schema.
259: invalid table name: Could not find table/view /SAPAPO/OMLOGHDR in schema	For technical reasons SQL: "HANA_liveCache_LCApps_Executions" fails with this error if the command is run in a schema that is different from the SAP ABAP schema.
259: invalid table name: Could not find table/view /SDF/SMON_WPINFO in schema	For technical reasons SQL: "HANA_ABAP_SDFSMON_FilterAndAggregation" fails with this error if the command is run in a schema that is different from the SAP ABAP schema.

259: invalid table name: Could not find table/view <name> in schema <schema> 260: invalid column name: <column>	Make sure that you use the SQL command variant fitting to your system (proper Revision level and statistics server type).
303: invalid DATE, TIME or TIMESTAMP value	Make sure that you use valid time format for "Modification section" switches like BEGIN_TIME or END_TIME. Be aware that you can only reference the other timestamp via "B" or "E" when the other timestamp is explicitly specified. For example, if implicitly END_TIME is specified with "C" (for current timestamp), you can't specify BEGIN_TIME with "E-D7" (end time minus 7 days). Instead you have to use something like "C-D7" (current timestamp minus 7 days). Additionally this error can show up on SAP HANA 2.00.010 when the USERS view is accessed (e.g. as part of SQL: "HANA_Configuration_MiniChecks"). As a workaround you can remove the accesses to users. For a permanent solution you have to upgrade to SAP HANA 2.00.011 or higher.
339: invalid number: not a valid number string '<value>' 339: invalid number: '<value>' isn't a valid numeric value	These errors are typically a consequence of an invalid SAP HANA parameter setting. If a parameter that is supposed to contain a number value is configured with a character value, this termination is possible. See SAP Note 2186744 and use SQL: "HANA_Configuration_Parameters" (available as part of the SQL statement collection) in order to make sure that all SAP HANA parameters are set properly. Also manually check via M_INIFILE_CONTENTS if a parameter is set to the value <value>. Another reason is the reason of a non-supported Linux kernel version, for example 3.0.101-412.g1823839-bigmem on SLES 11.4. In this case you have to move to a standard Linux kernel version or - as a temporary workaround - manually remove the part of the SQL statement linked to the OS kernel check (e.g. OS_KERNEL_VERSION check in the mini checks).
339: invalid number: not a valid number string 'yes'	When running the mini checks with ONLY_POTENTIALLY_CRITICAL_RESULTS = 'X' terminates with this error, the problem is most likely caused by a SAP HANA bug and the NO_SUBPLAN_SHARING hint can be used as a workaround. Download the most recent statement collection from this SAP Note, because the hint is already added per default now.
362: invalid schema name: SAPSR3	For technical reasons SQL: "HANA_BankAnalyzer_IndexConsistency" has a hard-coded schema name 'SAPSR3' outside of the "Modification section" that has to be manually adjusted to the database schema name of the individual Bank Analyzer installation.
362: invalid schema name: SYS_DATABASES	This error appears if you execute a command related to multitenant database containers (MDC, SAP Note 2101244) in a non-MDC system or in a tenant of an MDC system. These commands can only run successfully when

	being executed in a system database of a MDC system.
463: number of tables exceeds its maximum: 4095	This error happens when the amount of tables involved in a database request exceeds 4095 (SAP Note 2154870). It can happen in context of particularly large SQL statements like the mini checks (SAP Note 1999993). Also internal tables accessed in views are considered, so the termination can depend on factors like SAP HANA release level or configuration. Mini checks for SAP HANA 2.0 are already adjusted to minimize the risk of running into this termination. If you still face this termination you can set the following SAP HANA parameter: indexserver.ini -> [sql] -> table_size_check = false Attention: This setting is potentially dangerous because highly complex database requests involving a particularly high amount of tables are no longer blocked. Therefore this setting should only be implemented on a temporary bases in order to complete the failing analysis command successfully.
2048: column store error: search table error: [6850] Syntax error in format string;string [here]format(longdate "SERVER_TIMESTAMP", string '<string>') ...	This error can be a consequence of a wrong value <string> maintained in the TIME_AGGREGATE_BY input parameter in the "Modification section". Make sure that you only maintain supported values or NONE.
3584: distributed SQL error: [12027] ApplicationException : the application threw an exception info: ltt exception in thread 10791: exception 1: no.1000087 (lts/ios.hpp:550) basic_ios::clear Call stack module: TrexService::CoreHandler::handle_osNetStats	This error can happen on SAP HANA 2.00.020 due to a bug when accessing view M_HOST_NETWORK_STATISTICS. Upgrade to 2.00.021 or higher or avoid querying M_HOST_NETWORK_STATISTICS (e.g. by removing the related UNION ALL sections from the Mini Checks commands).
Parameter validation failed: data.body.statements[0].statement should NOT be longer than <length> characters	The SQL console in SAP HANA Cockpit (SAP Note 2185556) has a limited statement size based on the version level of the used SAP HANA Runtime Tools (HRTT): HRTT <= 2.5.51: 128,000 HRTT 2.5.52 - 2.5.59: 250,000 HRT >= 2.5.60: 300,000 Queries with a longer statement text fail with this error. You have to use a more recent HRTT version or another SAP HANA client to execute these commands successfully. See SAP Note 2621976 for more information.
Could not split statements; SQL console content is too large	This warning is issued by the SAP HANA Studio (SAP Note 2073112) SQL editor when SQL statements with a length of more than 200,000 characters are executed. It means that the statement is executed as a whole and SAP HANA Studio doesn't search for semi-colon (";") characters as statement separators. This is an acceptable behavior for the statements being part of this collection, so you can ignore the warning. It is possible to adjust the threshold of 200,000 characters in SAP HANA Studio via "Window" -> "Preferences" -> "SAP HANA" -> "Runtime" -> "SQL" -> "Maximum Numbers of Characters for Multiple Statement Execution".
SQL statements don't return historic data.	Historic data in HOST* and GLOBAL* tables is only

	<p>available if the statistics server collects this data and the retention time is set sufficiently high. Usually a retention of 42 days is used. If you are missing data, please check and adjust the data collection and retention of the statistics server accordingly. See SAP Note 2147247 for more information related to the statistics server. SQL: "HANA_StatisticsServer_Histories_RetentionTime" provides an overview about currently configured retention times and data availability.</p>
<p>The granularity of historic data returned by SQL statements isn't as precise as required.</p>	<p>Historic data in HOST* and GLOBAL* tables is only collected in certain intervals. Per default this interval varies between one minute and one day for different histories. You can adjust the intervals based on your requirements, but you have to make sure that an increased collection frequency doesn't put unnecessary load on the system. See SAP Note 2147247 for more information related to the statistics server. SQL: "HANA_StatisticsServer_Histories_Schedule" provides an overview about currently configured collection intervals.</p>
<p>The values in different historic intervals are inconsistent with each other.</p>	<p>Some SQL statements like SQL: "HANA_Memory_TopConsumers_Overview" gather historic data from multiple sources that are historicized in different intervals (e.g. hourly vs. every minute). If you now use an aggregation that is more precise than the least precise history collection, some lines of the output will contain the least precise information, others not. As a consequence consecutive lines in the output can appear to be inconsistent. If you want to avoid this scenario you shouldn't use a time granularity that is more precise than the least precise history collection. In case of SQL: "HANA_Memory_TopConsumers_Overview" you have to use at least an hourly aggregation.</p>
<p>The output is incomplete or empty</p>	<p>"Incomplete" or empty output can be caused by different reasons: RESULT_ROWS in "Modification section" (available for some SQL statements) set to a value different from -1, so that the output is limited the value defined in RESULT_ROWS DBACOCKPIT: "Rows:" is limited to a number of output rows (default: 250; increased to 1000 with SAP BASIS 7.02 SP21 / 7.30 SP19 / 7.31 SP23 / 7.40 SP20 / 7.50 SP11 / 7.51 SP06 / 7.52 SP02) and sometimes needs to be increased to see the complete output (e.g. when executing the mini checks, SAP Note 1999993) SAP HANA Studio: "Windows" -> "Preferences" -> "SAP HANA" -> "Runtime" -> "Result" -> "Maximum Number of Rows Displayed in Result" (default: 1000) can be increased when longer output is required (e.g. when executing the mini checks, SAP Note 1999993) DBACOCKPIT: Statements starting with WITH don't work with older DBACOCKPIT transactions before SAP BASIS 7.02 SP16 / 7.30 SP12 / 7.31 SP12 / 7.40 SP07 and an empty result is returned.</p>

Unexpected negative values in the output	The following reasons can be responsible for unexpected negative values in the output: Historic analysis across a database restart: In this case performance counters are reset to 0 and the delta mechanism of history views calculates significant negative values for the interval containing the restart time. Inactive nodes / services: Some performance values for inactive nodes or services (e.g. standby node) may be shown with "-1" values in certain monitoring views, so you have to focus on the active nodes and services to avoid negative values. Due to a SAP HANA bug with Revision 2.00.020 the calculation of a time range (e.g. via SECONDS_BETWEEN) may return a negative rather than a positive number. This problem is fixed with Revision 2.00.021 and higher.
Wrong timestamps in output	The following reasons can be responsible for wrong timestamps in the output: Due to a SAP HANA bug with Revision 2.00.020 the calculation of a time range (e.g. via SECONDS_BETWEEN) may return a negative rather than a positive number. This problem is fixed with Revision 2.00.021 and higher.
The output sort order is wrong	If the sort order of the output is wrong, you can consider the following points: Make sure that the ORDER_BY parameter in the modification section is set correctly and based on your needs (if it exists). Analysis commands that are sorting based on the analytic ROW_NUMBER function may suffer from a SAP HANA bug that can return the records in a wrong sort order. This problem is fixed with SAP HANA >= 1.00.122.19, >= 2.00.024.05 and >= 2.00.033. In many cases you can manually add the required ORDER BY at the end of the query as a temporary workaround.
Output doesn't match the expectation	If the output of a SQL statement doesn't match the expectation (and has perhaps changed compared to a previous statement version), you have to make sure that the "Modification section" of the command is configured based on your needs. For example, different values in AGGREGATE_BY can result in completely different outputs.

Further details can be provided as part of a TPO or System Administration service for HANA.

Starting with patch levels 7. 51 SP3, 7.50 SP8, 7.40 SP18, 7.31 SP21, 7.30 SP18 and 7.02 SP20 it is possible to load SQL statements from a local directory into the SQL editor of the SAP ABAP transaction DBACOCKPIT using the "Others" menu ("Upload SQL Scripts" / "Download SQL Scripts").

With earlier patch levels you can alternatively implement the attached report Z_INSERT_INTO_SQL_EDITOR:

1. If you already implemented an earlier version of the SQL statement collection before, you can purge it via "DELETE FROM DB6SCRIPT WHERE SCRIPTNAME LIKE 'HANA%'". Be aware that all imported commands starting with 'HANA' will be deleted with this command.

2. Implement the Z_INSERT_INTO_SQL_EDITOR on ABAP side.
3. Extract the SQL statement collection into one directory of your local PC.
4. Run report Z_INSERT_INTO_SQL_EDITOR and enter the directory in the PATH field of the selection screen.

If SQL statements fail that are uploaded via Z_INSERT_INTO_SQL_EDITOR (e.g. "[257]: sql syntax error") while they work fine when they are executed directly, the existing line length limit of 256 characters can be responsible and it is required to adjust the original statement from the statement collection.

It is possible to load the SQL statement collection into SAP HANA Studio, see SAP Note [2073112](#) ("Can the SAP HANA SQL statement collection be imported into SAP HANA Studio?") for more details.

Failures during upload can be treated as described in the overview above.

Feedback and suggestions are welcome: martin.fraendorfer@sap.com

Software Components

Software Component	Release
HDB	1.00 - 1.00
HDB	2.00 - 2.00

This document refers to

SAP Note/KBA	Title
2621976	Parameter validation failed: data.body.statements[0].statement should NOT be longer than 128000 characters
2405756	Cannot iterate over result set rows: SAP DBTech JDBC: [257] (at 197): sql syntax error: incorrect syntax near "WITH": line 1 col7 (at pos 196)
2380176	FAQ: SAP HANA Database Trace
2310449	Could not retrieve attachment in notes
2222220	FAQ: SAP HANA DBACOCKPIT
2222200	FAQ: SAP HANA Network
2186744	FAQ: SAP HANA Parameters
2185556	FAQ: SAP HANA Cockpit
2177604	FAQ: SAP HANA Technical Performance Optimization Service

2154870	How-To: Understanding and defining SAP HANA Limitations
2147247	FAQ: SAP HANA Statistics Server
2127458	FAQ: SAP HANA Loads and Unloads
2101244	FAQ: SAP HANA Multitenant Database Containers (MDC)
2073112	FAQ: SAP HANA Studio
1999997	FAQ: SAP HANA Memory
1999993	How-To: Interpreting SAP HANA Mini Check Results
1999880	FAQ: SAP HANA System Replication
2453766	Selecting Service Thread Sample Data Fails With an "invalid character encoding" Error Due to Columns PASSPORT_ACTION or PASSPORT_COMPONENT_NAME
2418887	Selecting the statement_string From m_sql_plan_cache Fails due to Invalid CESU-8 Encoding

This document is referenced by

SAP Note/KBA	Title
2116157	FAQ: SAP HANA Consistency Checks and Corruptions
2114710	FAQ: SAP HANA Threads and Thread Samples
2186744	FAQ: SAP HANA Parameters
2600030	Parameter Recommendations in SAP HANA Environments
1999998	FAQ: SAP HANA Lock Analysis
2399990	How-To: Analyzing ABAP Short Dumps in SAP HANA Environments
2180165	FAQ: SAP HANA Expensive Statements Trace
1999997	FAQ: SAP HANA Memory
2222200	FAQ: SAP HANA Network
1999930	FAQ: SAP HANA I/O Analysis
2400005	FAQ: SAP HANA Persistence
2380176	FAQ: SAP HANA Database Trace
2000003	FAQ: SAP HANA
2142945	FAQ: SAP HANA Hints

2000002	FAQ: SAP HANA SQL Optimization
2222218	FAQ: SAP HANA Database Server Management Console (hdbcons)
1999880	FAQ: SAP HANA System Replication
2100040	FAQ: SAP HANA CPU
2470289	FAQ: SAP HANA Non-Uniform Memory Access (NUMA)
2222250	FAQ: SAP HANA Workload Management
2169283	FAQ: SAP HANA Garbage Collection
2119087	How-To: Configuring SAP HANA Traces
2749491	Differences in EWA report for HANA Parameter recommendations when compared with output of HANA parameter check script HANA_Configuration_Parameters_1.00.90+.txt
2600076	FAQ: SAP HANA Inverted Individual Indexes
2147247	FAQ: SAP HANA Statistics Server
2400022	FAQ: SAP HANA Smart Data Integration (SDI)
2000000	FAQ: SAP HANA Performance Optimization
2127458	FAQ: SAP HANA Loads and Unloads
2699939	SAP HANA Emergency Suitcase
2593571	FAQ: SAP HANA Integrated liveCache
2160391	FAQ: SAP HANA Indexes
2535951	FAQ: SAP HANA Users and Schemas
2154870	How-To: Understanding and defining SAP HANA Limitations
2340450	FAQ: SAP HANA Table Replication
2536974	The system is running without alerts, but not all checks are being performed
2638539	Difference between Statisticsserver check status 'Inactive' and 'Disabled'
2572224	How-To: Repairing SAP HANA Tables
2731449	After Table Deletion and Successfully Delta Merge, Table Size Is Not Decreased
2200772	FAQ: SAP HANA Statement Routing and Client Distribution Mode
2729072	SAP Minicheck shows SLES15 is not supported.
2502256	FAQ: SAP HANA Caches
2477204	FAQ: SAP HANA Services and Ports
2717351	Failed to run HANA_Configuration_MiniChecks*.txt due to "Cannot iterate over result set row" error.

1977268	How to handle HANA Alert 40: 'Total memory usage of column-store tables'
2670064	FAQ: SAP HANA Multi-Dimensional Services (MDS)
2600095	FAQ: SAP HANA Sequences
2050579	How to handle HANA Alert 68: 'total memory usage of row store'
2400006	FAQ: SAP HANA Statement Hints
2081845	How to handle HANA Alert 65: 'Runtime of the log backups currently running'
2506811	FAQ: SAP HANA Dynamic Result Cache
2621976	Parameter validation failed: data.body.statements[0].statement should NOT be longer than 128000 characters
1998599	How-To: Analyzing high SAP HANA Memory Consumption due to Translation Tables
2559245	How to Export SQL Query Results on SAP HANA Studio and Database Explorer
2079396	How to handle HANA Alert 49: 'long-running blocking transaction'
2473915	SAP HANA Last Accessed Time for Tables and Columns
2439125	Basic SAP HANA Support Question Catalog
1977291	Handle HANA Alert 28: 'Most recent savepoint operation'
1977220	How to handle HANA Alert 54: Savepoint Duration
1794297	Secondary Indexes for S/4HANA and the business suite on HANA
2633077	Rowstore LOB Garbage is not Collected and the Number of Disk LOBs Keeps Increasing
2418299	SAP HANA: Partitioning Best Practices / Examples for SAP Tables
2488924	Linux: Recommended values for maximum stack size of processes
2603173	STATEMENT_HASH in M_SERVICE_THREAD_SAMPLES cut to 31 Characters
2565285	Temporary Performance Degradation After Upgrade Due to SAP Note 2413261
2301382	Increased Used Memory Size due to Pool/PersistenceManager/PersistentSpace/DefaultLPA /Page (Rev. 110 - 122.05)
2405822	SAP HANA: After a landscape redistribution, tables take up more memory than previously.
2402503	Increased memory requirement of BW objects in SAP HANA
706478	Preventing Basis tables from increasing considerably
2271684	Accelerated search does not use bind variables
2285065	DVM analysis on HANA based system shows tables with entries but without size

Attachments

File Name	File Size	Mime Type
SQLStatements_1.00.99-.zip	263	application/x-zip-compressed
Z_INSERT_INTO_SQL_EDITOR.TXT	8	text/plain
SQLStatements.zip	1708	application/x-zip-compressed