

Foglight for SAP HANA 5.9.4.20

Cartridge Guide

© 2019 Quest Software Inc.

ALL RIGHTS RESERVED.

This guide contains proprietary information protected by copyright. The software described in this guide is furnished under a software license or nondisclosure agreement. This software may be used or copied only in accordance with the terms of the applicable agreement. No part of this guide may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording for any purpose other than the purchaser's personal use without the written permission of Quest Software Inc.

The information in this document is provided in connection with Quest Software products. No license, express or implied, by estoppel or otherwise, to any intellectual property right is granted by this document or in connection with the sale of Quest Software products. EXCEPT AS SET FORTH IN THE TERMS AND CONDITIONS AS SPECIFIED IN THE LICENSE AGREEMENT FOR THIS PRODUCT, QUEST SOFTWARE ASSUMES NO LIABILITY WHATSOEVER AND DISCLAIMS ANY EXPRESS, IMPLIED OR STATUTORY WARRANTY RELATING TO ITS PRODUCTS INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT. IN NO EVENT SHALL QUEST SOFTWARE BE LIABLE FOR ANY DIRECT, INDIRECT, CONSEQUENTIAL, PUNITIVE, SPECIAL OR INCIDENTAL DAMAGES (INCLUDING, WITHOUT LIMITATION, DAMAGES FOR LOSS OF PROFITS, BUSINESS INTERRUPTION OR LOSS OF INFORMATION) ARISING OUT OF THE USE OR INABILITY TO USE THIS DOCUMENT, EVEN IF QUEST SOFTWARE HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. Quest Software makes no representations or warranties with respect to the accuracy or completeness of the contents of this document and reserves the right to make changes to specifications and product descriptions at any time without notice. Quest Software does not make any commitment to update the information contained in this document.

If you have any questions regarding your potential use of this material, contact:

Quest Software Inc.

Attn: LEGAL Dept

4 Polaris Way

Aliso Viejo, CA 92656

Refer to our Web site (<https://www.quest.com>) for regional and international office information.

Patents

Quest Software is proud of our advanced technology. Patents and pending patents may apply to this product. For the most current information about applicable patents for this product, please visit our website at <https://www.quest.com/legal>.

Trademarks

Quest, the Quest logo, and Join the Innovation are trademarks and registered trademarks of Quest Software Inc. For a complete list of Quest marks, visit <https://www.quest.com/legal/trademark-information.aspx>. All other trademarks and registered trademarks are property of their respective owners.

Table of Contents

Table of Contents	3
Introduction	9
Description	9
Key Features.....	9
Foglight for SAP HANA Requirements.....	11
Installing and Configuring Agents	12
SAP HANA Configuration.....	13
Database User Setup.....	13
Resource Tracking.....	13
Cartridge Installation	14
Creating and Configuring Agents	15
Agent Properties	16
System Database Connection	16
Tenant Database Connections.....	17
Collection Periods	17
SQL Options.....	17
Roles.....	18
Upgrading the Agent.....	19
Dashboards	20
SAP HANA Systems.....	20
Host Overview.....	20
Memory.....	21
Host Load	21
Services	22
SQL Plans.....	22
Expensive Statements.....	23
Blocked Transactions	23
Table Persistence	24
Native Alerts.....	24
Rules.....	25

SAP HANA Database Availability	25
SAP HANA Host CPU Used Percent	25
SAP HANA Memory Instance Allocation Used Percent.....	25
SAP HANA Memory Used Physical Memory Percent.....	25
SAP HANA Native Alert	25
SAP HANA Service Used Memory Effective Allocation Percent.....	25
SAP HANA System Availability	25
SAP HANA Volume Used Disk Percent	25
Internal statistics server problem	25
Host physical memory usage	26
Disk usage	26
Inactive services	26
Restarted services	26
Host CPU Usage.....	26
Delta merge (mergedog) configuration	26
Memory usage of name server	26
Lock wait timeout configuration.....	26
Record count of non-partitioned column-store tables.....	27
Table growth of non-partitioned column-store tables	27
Internal event.....	27
Notification of all alerts	27
Notification of medium and high priority alerts	27
Notification of high priority alerts	27
Open connections	27
Unassigned volumes	27
Record count of column-store table partitions	28
Most recent savepoint operation	28
Size of delta storage of column-store tables	28
Check internal disk full event.....	28
License expiry.....	28
Log mode LEGACY	28
Log mode OVERWRITE	28

Unavailable volumes	28
Existence of data backup	29
Status of most recent data backup	29
Age of most recent data backup	29
Status of most recent log backups	29
Long-running statements	29
Total memory usage of column-store tables	29
In-memory DataStore activation	29
Long running/idling cursors	29
Memory usage of services	29
Licensed memory usage.....	30
Memory usage of main storage of column-store tables	30
RTEdump files	30
Long-running serializable transactions	30
Long-running uncommitted write transactions.....	30
Long-running blocking situations.....	30
Number of diagnosis files	30
Size of diagnosis files	30
Crashdump files	31
Pagedump files.....	31
Savepoint duration.....	31
Columnstore unloads.....	31
Python trace activity	31
Instance secure store file system (SSFS) inaccessible.....	31
Plan cache size	31
Percentage of transactions blocked.....	31
Sync/Async read ratio	32
Sync/Async write ratio	32
Expiration of database user passwords	32
Granting of SAP_INTERNAL_HANA_SUPPORT role.....	32
Total memory usage of table-based audit log	32
Runtime of the log backups currently running	32

Storage snapshot is prepared	32
Table growth of rowstore tables	33
Total memory usage of row store.....	33
Enablement of automatic log backup.....	33
Consistency of internal system components after system upgrade.....	33
Row store fragmentation.....	33
Number of log segments.....	33
Overflow of rowstore version space.....	33
Overflow of metadata version space.....	33
Rowstore version space skew.....	33
Discrepancy between host server times.....	34
Database disk usage.....	34
Connection between systems in system replication setup	34
Configuration consistency of systems in system replication setup.....	34
Availability of table replication	34
Cached view size	34
Timezone conversion	34
Table consistency.....	34
Insecure instance SSFS encryption configuration.....	35
Insecure systemPKI SSFS encryption configuration.....	35
Internal communication is configured too openly	35
Granting of SAP HANA DI support privileges	35
Auto merge for column-store tables	35
Missing volume files.....	35
Status of HANA platform lifecycle management configuration	35
Plan cache hit ratio	36
Root keys of persistent services are not properly synchronized.....	36
Streaming License expiry	36
Log replay backlog for system replication secondary.....	36
Availability of Data Quality reference data (directory files).....	36
Long-running tasks.....	36
Granting of SAP HANA DI container import privileges	36

LOB garbage collection activity.....	36
HANA version	36
Unsupported operating system in use.....	37
SQL access for SAP HANA DI technical users	37
Existence of system database backup	37
Usage of deprecated features.....	37
Log shipping backlog for system replication secondary	37
Total Open Transactions Check.....	37
ASYNCR replication in-memory buffer overflow.....	37
Fallback snapshot consistency	37
Fallback snapshot age	38
Backup history broken	38
Catalog Consistency	38
Replication status of Replication Log.....	38
Record count of non-partitioned column-store tables (include).....	38
Record count of column-store table partitions (include)	38
LDAP Enabled Users without SSL	38
Dbospace usage	38
Dbospace status.....	39
Dbospace file status	39
Inactive Streaming applications.....	39
Inactive Streaming project managed adapters.....	39
Streaming project physical memory usage.....	39
Streaming project CPU usage	39
Number of publishers of streaming project.....	39
Number of subscribers of streaming project.....	39
Row throughput of subscriber of streaming project	39
Transaction throughput of subscriber of streaming project	40
Row throughput of publisher of streaming project.....	40
Transaction throughput of publisher of streaming project.....	40
Bad rows of project managed adapter	40
High latency of project managed adapter	40

Large queue of stream of streaming project	40
Large store of stream of streaming project	40
Agent availability.....	40
Agent memory usage	40
Remote Subscription exception.....	41
Reports.....	42
SAP HANA Services Report.....	42
SAP HANA SQL Report.....	42
SAP HANA Storage Report	42
SAP HANA System Summary.....	42

Introduction

Description

SAP HANA is an in-memory, relational database management system. It operates primarily as a column-oriented store, but row-oriented tables are also supported. Both transactional and analytical workloads are supported. SAP HANA is fully ACID compliant. Various features include a JSON store, spatial data processing, text analytics and search, time-series analytics, streaming data processing, and graph data processing.

An SAP HANA deployment requires a certified hardware appliance installed with one of several supported operating systems. Hardware requirements include specification of the CPU architecture, core count, minimum RAM, and minimum storage. The supported operating systems are Red Hat Enterprise Linux for SAP Solutions, Red Hat Enterprise Linux for SAP HANA, SUSE Linux Enterprise Server for SAP Applications, and SUSE Linux Enterprise Server.

SAP HANA offers an optional, embedded web application server XS Advanced (XSA) based on Cloud Foundry. XSA offers native support for Node.js and JavaEE as well as extensibility through custom runtimes.

Transactional updates to the database are performed using multi-version concurrency control to enable read consistency and throughput. Writes are performed using row-level write locks with blocked transaction monitoring.

An SAP HANA system contains a single system database and zero or more isolated tenant databases. Multiple servers working together comprise the database system: name server, index server, preprocessor server, compile server, script server, XS Advanced runtime, and others. A single system can be distributed across multiple hosts to improve scalability and can utilize replication for high availability.

Key Features

The Foglight for SAP HANA Agent monitors the current and historical status of SAP HANA database systems. Database configuration and performance statistics are collected and presented in real-time dashboards and also available in packaged reports. Alerts are generated on respective components of the system when indicators of current or potential future performance or availability issues arise. Setup involves creating a monitoring agent for each target database system. Such agents are fully remote, meaning that they are typically configured on a separate host. Installing the agent on the Foglight Agent Manager ensures that nothing needs to be installed on the monitored system. A database user with appropriate permissions is used by the agent to collect data. An agent may also be configured to collect operating system metrics purpose of monitoring the database host and logs.

Data is collected from each configured SAP HANA system via the official JDBC driver. Collections take place at regular intervals, called periods, which are configurable by collection type in the Foglight Agent Properties. Critical metrics, such as system availability, are collected frequently, while data expected to change less frequently, such as configuration settings or individual table metrics, are collected at longer

intervals. The Foglight Agent queries Hana system tables as well as the embedded statistics service for the latest values of key parameters and metrics. These metrics are collected and compared against thresholds and historical performance to identify system possible performance and availability issues. Database hosts are monitored at the operating system level by the Infrastructure Agent which can also be used to monitor SAP HANA native logs.

Foglight for SAP HANA Requirements

Foglight for SAP HANA is compatible with **SAP HANA 2.0 SPS 02 REV 20** and later. The multitenant database container (MDC) architecture introduced in SAP HANA 1.0 SPS 09 (Revision 90) and made mandatory in SAP HANA 2.0 SPS 01 (Revision 10) is the only architecture supported.

Foglight for SAP HANA may be installed on an **FMS** of version **5.9.2** or later. Agents require a **Foglight Agent Manager** of version **5.8.5.2** or later.

Installing and Configuring Agents

Installation of Foglight for SAP HANA is detailed in the following sections. It is recommended that the installation be performed in the order below:

- [SAP HANA Configuration](#)
- [Cartridge Installation](#)
- [Creating and Configuring Agents](#)

SAP HANA Configuration

Database User Setup

To configure the monitoring of an SAP HANA system, the Foglight agent requires a database user with sufficient privileges on the SYSTEMDB and each tenant database for which monitoring is desired.

Create an ordinary database user on each database in the SAP HANA system as follows:

```
CREATE USER <username> PASSWORD <password> NO FORCE_FIRST_PASSWORD_CHANGE;  
GRANT CATALOG READ TO <username>;  
GRANT SELECT ON SCHEMA _SYS_STATISTICS TO <username>;
```

The username (and password) may be chosen independently on each database.

Resource Tracking

Some features of the Foglight for SAP Hana depend on resource tracking within SAP HANA. These include SQL Query Plan memory monitoring and Expensive Statement tracking. The following steps are necessary to enable this functionality.

Set the global [resource_tracking] properties 'enable_tracking' and 'memory_tracking' to 'on.'


CPU consumption estimation is on by default (for all SAP HANA versions greater than SAP HANA 1.0 SPS 11). Check that the global [resource_tracking] property 'cpu_time_measurement_mode' is still set to 'on'.

To enable expensive statement tracking, set the global [expensive_statement] property 'enable' to 'true.' Optionally modify the following parameters from their defaults:

- 'threshold_duration' in microseconds,
- 'threshold_memory' in bytes (only effective if resource and memory tracking are also enabled),
- 'threshold_cpu_time' in microseconds (only effective if resource and CPU time tracking are also enabled).

Cartridge Installation

1. Open Foglight Management Console.
2. A separate license may be required for the SAP HANA cartridge. From the navigation pane, select **Dashboards > Administration**. Then select **Licenses**. Click the “Install” button and select the .license file.
3. From the navigation pane, select: **Dashboards > Administration > Cartridges > Cartridge Inventory**. The Cartridge Inventory screen appears.
4. Load the SAPHANA Agent into the FMS by clicking on “Install Cartridge” and browsing to the location of the agent CAR file. Leave the “Enable on Install” check box checked.
5. Once the installation has completed on the Foglight Management Server, the SAP HANA cartridge will appear in the table of installed cartridges.



6. Deploy the cartridge to the Foglight Agent Manager that will be used to host the agent. This step must also be repeated after every cartridge upgrade. Use the “Deploy Agent Package” button on the Agent Status or Agent Managers page to perform this step.

Creating and Configuring Agents

The **Agent Status** page can be used to create new agents and configure and manage existing agents. To access the page from the navigation panel, select: **Dashboards > Administration > Agents > Agent Status**.

Use the following steps to create a new agent instance:

1. Click the Create Agent button and follow the instructions for the cards:
 - a. **Host Selector** - Choose the Agent Manager on which the agent should run. Considerations for this may include physical or virtual locality to the monitored instance, allocated resources, or grouping with other agents of the same type or monitored environment.
 - b. **Agent Type and Instance Name** – Select the SAPHANA Agent type. Then, select the “Specify Name” radio button and provide a name for the agent created. This is not canonical and should be representative of the database system that this agent will monitor.
 - c. **Summary** – Click Finish.
2. Once the agent is created, click the checkbox next to the newly created agent.
3. Click the “Edit Properties” button.
4. Select “Modify the default properties for this agent.”
5. Edit the agent properties for the SAP HANA agent instance – see [Agent Properties](#) below.
6. Click “Save,” return to the **Agent Status** screen, reselect the agent as necessary and click the “Activate” button.

To modify the properties for an existing agent, skip to steps 3-6 and subsequently select “Deactivate” followed by “Activate” to restart the agent.

Agent Properties

Below is a complete list of the configurable properties of the Foglight for SAP HANA agent and an explanation of each property.

System Database Connection	
IP or Hostname	<input type="text" value="xxx.xxx.xxx.xxx"/>
Port	<input type="text" value="30013"/>
Username	<input type="text" value="FOGLIGHTAGENT"/>
Password	<input type="password" value="••••••"/>
Use Compression?	<input type="checkbox"/>

Tenant Database Connections	
Tenants	TenantsList <input type="button" value="Edit"/> <input type="button" value="Clone"/> <input type="button" value="Delete"/>

Periods (sec)	
Availability	<input type="text" value="30"/>
Databases and Tables	<input type="text" value="900"/>
Services, Volumes, and Load	<input type="text" value="300"/>
Query Plans	<input type="text" value="300"/>
Transactions	<input type="text" value="300"/>
Alerts	<input type="text" value="60"/>
Configuration	<input type="text" value="900"/>

SQL	
Max SQL Plans Per Period	<input type="text" value="200"/>
Select Top SQL By	<input type="text" value="Execution Count"/>
Max Expensive Statements Per Period	<input type="text" value="200"/>
Expensive Statement Threshold Duration (sec)	<input type="text" value="1"/>
Max Blocked Transactions Per Period	<input type="text" value="200"/>

Options	
Max Top Tables Per Period	<input type="text" value="200"/>
Min Table Size (KiB)	<input type="text" value="1048576"/>

System Database Connection

The agent requires a connection to the SYSTEMDB database to gather many system properties and metrics.

- **IP or Hostname** – Host on which the SAP HANA system node is running.
- **Port** – The SQL port for the system database.
- **Username** – The system database username.
- **Password** – The system database user's password.
- **Use Compression?** – Whether the database connection should be compressed.

Tenant Database Connections

The same connection properties required for the system database are mandatory for every monitored tenant (though monitoring any individual tenant is optional). See above for definitions of **IP or Hostname, Port, Username, Password, and Use Compression**.

Collection Periods

The collection period fields in the agent properties set the sample frequencies and are in seconds. Any collection can be turned off by setting its period to zero. The default collection periods vary based on the type and volatility of data collected.

- **Availability** – Controls the query interval for the SAP HANA system and database availability collections.
- **Databases and Tables** – Controls the query interval for collection of data pertaining to databases present on the system and table memory and persistence metrics.
- **Services, Volumes, and Load** – Controls the query interval for collection of the service, volume, and host load metrics.
- **Query Plans** – Controls the query interval for collection of the SQL query plans and expensive statements. Note that resource tracking must be enabled in the SAP HANA system for expensive statement data to be available.
- **Transactions** – Controls the query interval for blocked transaction monitoring.
- **Alerts** – Controls the query interval for native alert collection.
- **Configuration** – Controls the query interval for system configuration tracking.

SQL Options

Max SQL Plans Per Period – The maximum number of SQL plans to retrieve per collection period.

Select Top SQL By – The metric used to determine which SQL plans are collected in the usual case when the SQL plan cache contains more plans than can be collected based on “Max SQL Plans Per Period”.

Max Expensive Statements Per Period – The maximum number of expensive statements to retrieve per collection period. Expensive Statement tracking is only possible if [Resource Tracking](#) has been configured in SAP HANA.

Expensive Statement Threshold Duration – The minimum duration of an expensive statement to qualify for retrieval. Note that only “Max Expensive Statements Per Period” expensive statements will be retrieved per period, regardless. Hence it is possible that additional expensive statements above this duration exist but are not collected by Foglight.

Roles

Two roles, SAP HANA User and SAP HANA Administrator, are installed with the cartridge. Viewing SAP HANA dashboards requires that a user be assigned one of these roles or have the core Administrator role. In the current release, the SAP HANA Administrator role is interchangeable with the SAP HANA User role. In future releases, the SAP HANA Administrator role will confer additional privileges where system modification options are available.

Upgrading the Agent

1. Go to Dashboards > Administration > Cartridges > Cartridge Inventory and click the Install Cartridge button.
2. Locate the .car file on your system and install it with auto-enable selected. If you get a message that a bundled cartridge is of an older version than the one currently enabled on your FMS and will not be enabled, ignore it and continue.
3. Once the cartridge is installed and enabled, go to Dashboards > Administration > Agents > Agent Managers. Agent Managers that can be upgraded with newer agent packages will show “yes” in the Upgradable | Agents column. Select all Agent Managers you wish to upgrade and click the Upgrade button.

Note: If an Agent Manager is not upgradable, check that the Agent Manager version is compatible with the newer agent version. If it is not, the Agent Manager will need to be upgraded first.

Dashboards

SAP HANA Systems

The main entry point to the Foglight for SAP HANA cartridge dashboards. It provides a list containing each SAP HANA system monitored by the Foglight Management Server. Each system can be identified by its SAP ID and its system code. The database hosts that comprise the system are selectable to enable drilling down to the Host Overview page (see below). The dashboard displays system properties, availability, and current alarm state.

Health	SAP ID	SAP System	Hosts	Availability	Status	F	C	W	Start Time	Version	Build
✘	HDB	01	imdbmaster	100 %	OK	1	23	0	6/10/19 8:14 AM	2.00.033.00.1535711040	fa/hana2sp03
✘	H15	00	hana-15-master	100 %	OK	1	18	3	6/7/19 11:31 AM	2.00.020.00.1500920972	fa/hana2sp02

Host Overview

The Host Overview screen provides a comprehensive overview and health status for the selected database host. Availability and monitoring status, as well as database and service up/down status, are displayed along with hardware and software installation properties. Infrastructure and storage utilization, top SQL plans, and top current alarms give a top-level overview of the performance of the database host.

Overview | Memory | Host Load | Services | SQL Plans | Expensive Statements | Blocked Transactions | Table Persistence | Native Alerts

System: HDB | Hostname: imdbmaster | Active Status: yes | Host Status: OK

Availability

Instance: **Up** 1d 1h 1m
Availability during the last 24 hours **100%**

Monitoring Agents: DB OS

Status

Services Active: 12 / 12 | Databases Active: 3 / 3

Properties

Hardware: Xen
OS: SUSE Linux Enterprise Server 12.3
Kernel: 4.4.156-94.57-default
Cores (Threads): 4 (8)
CPU: Intel(R) Xeon(R) CPU E5-2686 v4 @ 2.30GHz
OS User: hdbadm (1001)
Timezone (Offset): EDT (-14400)

Infrastructure

CPU Utilization: 3%

Memory Utilization: 85%

Performance

Workload:

Volumes

Data: 585 GB (63 GB Utilization, 11%)
Log: 325 GB (9.5 GB Utilization, 3%)
Trace: 300 GB (122 GB Utilization, 41%)

Summary

Top 10 Alarms | Top 10 SQLs

Database	ID	Latest Severity	Recent Instances	Description
T001	35	✘	1	Determines whether or not a data backup exists. Without a data backup, your database cannot be recovered.
T001	50	⚠	1	Determines the number of diagnosis files written by the system (excluding zip-files). An unusually large number of files can indicate a problem with the database (for example, problem with trace file rotation or a high number of crashes).
HDB	43	⚠	1	Determines what percentage of its effective allocation limit a service is using.
SYSTEMDB	50	⚠	1	Determines the number of diagnosis files written by the system (excluding zip-files). An unusually large number of files can indicate a problem with the database (for example, problem with trace file rotation or a high number of crashes).
T001	24	i	1	Determines whether or not there have been any high priority alerts since the last check and if so, sends a summary e-mail to specified recipients.
T001	23	i	1	Determines whether or not there have been any medium and high priority alerts since the last check and if so, sends a summary e-mail to specified recipients.
T001	22	i	1	Determines whether or not there have been any alerts since the last check and if so, sends a summary e-mail to specified recipients.
SYSTEMDB	24	i	1	Determines whether or not there have been any high priority alerts since the last check and if so, sends a summary

Memory

Memory from the host and current Utilization are displayed alongside peak and current memory usage for the SAP HANA database.



Host Load

CPU, Network Load, and Swap Usage for the database host are displayed in parallel graphs.




Services

Database service views for both the system database and each monitored tenant database show estimated CPU usage along with the service start time, process ID, and other service properties.




SQL Plans

SQL statement, hash, and a variety of performance metrics are tabulated for the collected SQL plans. Top SQL plans are collected and are sorted by criteria defined in the agent properties. Sort criteria include average execution time, total execution time, execution count, average execution memory size.



Expensive Statements

The Expensive Statements dashboard displays individual statement executions that exceed CPU, memory, or time thresholds. Resource monitoring must be configured in the SAP HANA database to enable expensive statement monitoring. See the agent configuration section above.




Blocked Transactions

Blocked transactions metrics are displayed along with the time of occurrence, lock details, lock owner, and blocked transaction hashes.




Table Persistence

Metrics related to the largest tables are displayed. The Table Persistence dashboard shows availability charts for the system database and each monitored tenant database. Tables can be filtered by database and schema. The minimum size threshold for tables as well as a maximum total number of tables to collect is configurable in the agent properties.



Native Alerts

Native SAP HANA database alerts are collected by the SAP HANA agent. Each alert is then assigned a Foglight severity (informational, warning, critical, or fatal) corresponding to their alert severity. These alerts can be seen on the main Foglight alarms dashboard along with alarms generated by all other cartridges.



Rules

SAP HANA Database Availability

Alerts if the SAP HANA database becomes unreachable.

SAP HANA Host CPU Used Percent

Alerts if CPU utilization on the host is high.

SAP HANA Memory Instance Allocation Used Percent

Alerts if the memory used by an instance as a percentage of its allocation limit is too high.

SAP HANA Memory Used Physical Memory Percent

Alerts if the physical memory usage is too high.

SAP HANA Native Alert

Publishes SAP HANA native alerts to Foglight.

SAP HANA Service Used Memory Effective Allocation Percent

Alerts if the service's usage of its effective memory allocation is too high.

SAP HANA System Availability

Alerts if the SAP HANA System becomes unreachable through the system database.

SAP HANA Volume Used Disk Percent

Alerts if the percentage of disk used on a volume is too high.

Internal statistics server problem

Identifies internal statistics server problem.

Host physical memory usage

Determines what percentage of total physical memory available on the host is used. All processes consuming memory are considered, including non-SAP HANA processes.

Disk usage

Determines what percentage of each disk containing data, log, and trace files is used. This includes space used by non-SAP HANA files.

Inactive services

Identifies inactive services.

Restarted services

Identifies services that have restarted since the last time the check was performed.

Host CPU Usage

Determines the percentage CPU idle time on the host and whether or not CPU resources are running low.

Delta merge (mergedog) configuration

Determines whether or not the 'active' parameter in the 'mergedog' section of the system configuration file(s) is 'yes.' mergedog is the system process that periodically checks column tables to determine whether or not a delta merge operation should be executed.

Memory usage of name server

Determines what percentage of allocated shared memory is being used by the name server on a host.

Lock wait timeout configuration

Determines whether the 'lock_wait_timeout' parameter in the 'transaction' section of the indexserver.ini file is between 100,000 and 7,200,000.

Record count of non-partitioned column-store tables

Determines the number of records in non-partitioned column-store tables. Current table size may not be critical. Partitioning need only be considered if tables are expected to grow rapidly (a non-partitioned table cannot contain more than 2,147,483,648 (2 billion) rows).

Table growth of non-partitioned column-store tables

Determines the growth rate of non-partitioned columns tables.

Internal event

Identifies internal database events.

Notification of all alerts

Determines whether or not there have been any alerts since the last check and if so, sends a summary e-mail to specified recipients.

Notification of medium and high priority alerts

Determines whether or not there have been any medium and high priority alerts since the last check and if so, sends a summary e-mail to specified recipients.

Notification of high priority alerts

Determines whether or not there have been any high priority alerts since the last check and if so, sends a summary e-mail to specified recipients.

Open connections

Determines what percentage of the maximum number of permitted SQL connections are open. The maximum number of permitted connections is configured in the "session" section of the indexserver.ini file.

Unassigned volumes

Identifies volumes that are not assigned a service.

Record count of column-store table partitions

Determines the number of records in the partitions of column-store tables. A table partition cannot contain more than 2,147,483,648 (2 billion) rows.

Most recent savepoint operation

Determines how long ago the last savepoint was defined, that is, how long ago a complete, consistent image of the database was persisted to disk.

Size of delta storage of column-store tables

Determines the size of the delta storage of column tables.

Check internal disk full event

Determines whether or not the disks to which data and log files are written are full. A disk-full event causes your database to stop and must be resolved.

License expiry

Determines how many days until your license expires. Once your license expires, you can no longer use the system, except to install a new license.

Log mode LEGACY

Determines whether or not the database is running in log mode "legacy." Log mode "legacy" does not support point-in-recovery and is not recommended for production systems.

Log mode OVERWRITE

Determines whether or not the database is running in log mode "overwrite". Log mode "overwrite" does not support point-in-recovery (only recovery to a data backup) and is not recommended for production systems.

Unavailable volumes

Determines whether or not all volumes are available.

Existence of data backup

Determines whether or not a data backup exists. Without a data backup, your database cannot be recovered.

Status of most recent data backup

Determines whether or not the most recent data backup was successful.

Age of most recent data backup

Determines the age of the most recent successful data backup.

Status of most recent log backups

Determines whether or not the most recent log backups for services and volumes were successful.

Long-running statements

Identifies long-running SQL statements.

Total memory usage of column-store tables

Determines what percentage of the effective allocation limit is being consumed by individual column-store tables as a whole (that is, the cumulative size of all of a table's columns and internal structures)

In-memory DataStore activation

Determines whether or not there is a problem with the activation of an in-memory DataStore object.

Long running/idling cursors

Identifies long-running/idling cursors. The threshold is based on M_INIFILE_CONTENTS.VALUE where KEY = "idle_cursor_lifetime".

Memory usage of services

Determines what percentage of its effective allocation limit a service is using.

Licensed memory usage

Determines what percentage of licensed memory is used.

Memory usage of main storage of column-store tables

Determines what percentage of the effective allocation limit is being consumed by the main storage of individual column-store tables.

RTEdump files

Identifies new runtime dump files (*rtedump*) have been generated in the trace directory of the system. These contain information about, for example, build, loaded modules, running threads, CPU, and so on.

Long-running serializable transactions

Identifies long-running serializable transactions.

Long-running uncommitted write transactions

Identifies long-running uncommitted write transactions.

Long-running blocking situations

Identifies long-running blocking situations.

Number of diagnosis files

Determines the number of diagnosis files written by the system (excluding zip-files). An unusually large number of files can indicate a problem with the database (for example, a problem with trace file rotation or a high number of crashes).

Size of diagnosis files

Identifies large diagnosis files. Unusually large files can indicate a problem with the database.

Crashdump files

Identifies new crashdump files that have been generated in the trace directory of the system.

Pagedump files

Identifies new pagedump files that have been generated in the trace directory of the system.

Savepoint duration

Identifies long-running savepoint operations.

Columnstore unloads

Determines how many columns in columnstore tables have been unloaded from memory. This can indicate performance issues.

Python trace activity

Determines whether or not the python trace is active and for how long. The python trace affects system performance.

Instance secure store file system (SSFS) inaccessible

Determines if the instance secure store in the file system (SSFS) of your SAP HANA system is accessible to the database.

Plan cache size

Determines whether or not the plan cache is too small.

Percentage of transactions blocked

Determines the percentage of transactions that are blocked.

Sync/Async read ratio

Identifies a bad trigger asynchronous read ratio. This means that asynchronous reads are blocking and behave almost like synchronous reads. This might have negative impact on SAP HANA I/O performance in certain scenarios.

Sync/Async write ratio

Identifies a bad trigger asynchronous write ratio. This means that asynchronous writes are blocking and behave almost like synchronous writes. This might have negative impact on SAP HANA I/O performance in certain scenarios.

Expiration of database user passwords

Identifies database users whose password is due to expire in line with the configured password policy. If the password expires, the user will be locked. If the user in question is a technical user, this may impact application availability. It is recommended that you disable the password lifetime check of technical users so that their password never expires (`ALTER USER <username> DISABLE PASSWORD LIFETIME`).

Granting of SAP_INTERNAL_HANA_SUPPORT role

Determines if the internal support role (`SAP_INTERNAL_HANA_SUPPORT`) is currently granted to any database users.

Total memory usage of table-based audit log

Determines what percentage of the effective memory allocation limit is being consumed by the database table used for table-based audit logging. If this table grows too large, the availability of the database could be impacted.

Runtime of the log backups currently running

Determines whether or not the most recent log backup terminates in the given time.

Storage snapshot is prepared

Determines whether or not the period, during which the database is prepared for a storage snapshot, exceeds a given threshold.

Table growth of rowstore tables

Determines the growth rate of rowstore tables

Total memory usage of row store

Determines the current memory size of a row store used by a service

Enablement of automatic log backup

Determines whether automatic log backup is enabled.

Consistency of internal system components after system upgrade

Verifies the consistency of schemas and tables in internal system components (for example, the repository) after a system upgrade.

Row store fragmentation

Check for fragmentation of row store.

Number of log segments

Determines the number of log segments in the log volume of each service. Check for number of log segments.

Overflow of rowstore version space

Determines the overflow ratio of the rowstore version space.

Overflow of metadata version space

Determines the overflow ratio of the metadata version space.

Rowstore version space skew

Determines whether the rowstore version chain is too long.

Discrepancy between host server times

Identifies discrepancies between the server times of hosts in a scale-out system.

Database disk usage

Determines the total used disk space of the database. All data, logs, traces and backups are considered.

Connection between systems in system replication setup

Identifies closed connections between the primary system and a secondary system. If connections are closed, the primary system is no longer being replicated.

Configuration consistency of systems in system replication setup

Identifies configuration parameters that do not have the same value on the primary system and a secondary system. Most configuration parameters should have the same value on both systems because the secondary system has to take over in the event of a disaster.

Availability of table replication

Monitors error messages related to table replication.

Cached view size

Determines how much memory is occupied by cached view

Timezone conversion

Compares SAP HANA internal timezone conversion with Operating System timezone conversion.

Table consistency

Identifies the number of errors and affected tables detected by `_SYS_STATISTICS.Collector_Global_Table_Consistency`.

Insecure instance SSFS encryption configuration

Determines whether the master key of the instance secure store in the file system (SSFS) of your SAP HANA system has been changed. If the SSFS master key is not changed after installation, it cannot be guaranteed that the initial key is unique.

Insecure systemPKI SSFS encryption configuration

Determines whether the master key of the secure store in the file system (SSFS) of your system's internal public key infrastructure (system PKI) has been changed. If the SSFS master key is not changed after installation, it cannot be guaranteed that the initial key is unique.

Internal communication is configured too openly

Determines whether the ports used by SAP HANA for internal communication are securely configured. If the "listeninterface" property in the "communication" section of the global.ini file does not have the value ".local" for single-host systems and ".all" or ".global" for multiple-host systems, internal communication channels are externally exploitable.

Granting of SAP HANA DI support privileges

Determines if support privileges for the SAP HANA Deployment Infrastructure (DI) are currently granted to any database users or roles.

Auto merge for column-store tables

Determines if the delta merge of a table was executed successfully or not.

Missing volume files

Determines if there is any volume file missing.

Status of HANA platform lifecycle management configuration

Determines if the system was not installed/updated with the SAP HANA Database Lifecycle Manager (HDBLCM).

Plan cache hit ratio

Determines whether or not the plan cache hit ratio is too low.

Root keys of persistent services are not properly synchronized

Not all services that persist data could be reached the last time the root key change of the data volume encryption service was changed. As a result, at least one service is running with an old root key.

Streaming License expiry

Determines how many days until your streaming license expires. Once your license expires, you can no longer start streaming projects.

Log replay backlog for system replication secondary

System Replication secondary site has a higher log replay backlog than expected.

Availability of Data Quality reference data (directory files)

Determine the Data Quality reference data expiration dates.

Long-running tasks

Identifies all long-running tasks.

Granting of SAP HANA DI container import privileges

Determines if the container import feature of the SAP HANA Deployment Infrastructure (DI) is enabled and if import privileges for SAP HANA DI containers are currently granted to any database users or roles.

LOB garbage collection activity

Determines whether or not the lob garbage collection is activated.

HANA version

Checks the installed HANA version.

Unsupported operating system in use

Determines if the operating system of the SAP HANA Database hosts is supported.

SQL access for SAP HANA DI technical users

Determines if SQL access has been enabled for any SAP HANA DI technical users. SAP HANA DI technical users are either users whose names start with '_SYS_DI' or SAP HANA DI container technical users (<container name>, <container name>#DI, <container name>#OO).

Existence of system database backup

Determines whether or not a system database backup exists. Without a system database backup, your system cannot be recovered.

Usage of deprecated features

Determines if any deprecated features were used in the last interval.

Log shipping backlog for system replication secondary

System Replication secondary site has a higher log shipping backlog than expected.

Total Open Transactions Check

The check monitors the number of open transactions per service

ASync replication in-memory buffer overflow

Checks if local in-memory buffer in ASync replication mode runs full

Fallback snapshot consistency

Determines if an inconsistent fallback snapshot exist.

Fallback snapshot age

Determines if a snapshot exists for an extended period of time.

Backup history broken

If the backup history is broken, the log_mode is internally set to overwrite; it is not ensured that the service is fully recoverable via backup.

Catalog Consistency

Identifies the number of errors and affected objects detected by `_SYS_STATISTICS.Collector_Global_Catalog_Consistency`.

Replication status of Replication Log

Check whether the status of replication log is disabled.

Record count of non-partitioned column-store tables (include)

Determines the number of records in non-partitioned column-store tables (only include tables are checked). Current table size may not be critical. Partitioning need only be considered if tables are expected to grow rapidly (a non-partitioned table cannot contain more than 2,147,483,648 (2 billion) rows).

Record count of column-store table partitions (include)

Determines the number of records in the partitions of column-store tables (only include tables are checked). A table partition cannot contain more than 2,147,483,648 (2 billion) rows.

LDAP Enabled Users without SSL

Checks if any user is enabled for LDAP Authentication and SSL is off.

Dbpace usage

Checks for the dbspace size usage.

Dbpace status

Determines whether or not all dbspaces are available.

Dbpace file status

Determines whether or not all dbspace files are available.

Inactive Streaming applications

Identifies inactive Streaming applications.

Inactive Streaming project managed adapters

Identifies inactive Streaming project managed adapters.

Streaming project physical memory usage

Determines what percentage of total physical memory available on the host is used for the streaming project.

Streaming project CPU usage

Determines the percentage CPU usage for a streaming project on the host and therefore whether or not CPU resources are running out.

Number of publishers of streaming project

Identify the large publishers of streaming project. Make sure that they will not break the streaming project.

Number of subscribers of streaming project

Identify the large subscribers of streaming project. Make sure that they will not break the streaming project.

Row throughput of subscriber of streaming project

Identify which subscriber of streaming project has low throughput measured in rows per second.

Transaction throughput of subscriber of streaming project

Identify which subscriber of streaming project has transaction throughput measured in transactions per second.

Row throughput of publisher of streaming project

Identify which publisher of streaming project has low throughput measured in rows per second.

Transaction throughput of publisher of streaming project

Identify which publisher of streaming project has transaction throughput measured in transactions per second.

Bad rows of project managed adapter

Identify which project managed adapter has much rows with error.

High latency of project managed adapter

Identify which project managed adapter has high latency.

Large queue of stream of streaming project

Identify which stream of streaming project has large queue.

Large store of stream of streaming project

Identify which stream of streaming project has large store.

Agent availability

Determines how many minutes the agent has been inactive.

Agent memory usage

Determines what percentage of total memory available to agent is used.

Remote Subscription exception

Checks for recent exceptions in remote subscriptions and remote sources.

Reports

SAP HANA Services Report

Service Report of SAP HANA displays system overview, service details including service properties, CPU and memory utilization of a service.

SAP HANA SQL Report

SQL Report of SAP HANA includes System Overview along with top 10 SQL statements sorted by provided criterion, top 10 expensive statements and top 10 blocked transactions a system.

SAP HANA Storage Report

Storage report of SAP HANA includes System Overview along with Volume Utilization, Databases Overview and important Table metrics.

SAP HANA System Summary

System summary of SAP HANA System including availability, workload, CPU utilization, memory utilization along with different volume utilization. It includes top SQL statements and top native alerts.