



User Guide | PUBLIC

Document Version: 7.40-7.56 – 2024-03-13

Technical Monitoring Cockpit (On-Premise, 7.4-7.56)

Content

1	Technical Monitoring Cockpit (On-Premise, 7.4-7.56)	12
1.1	Learning the Basics	15
1.2	Architecture: Engine and Content	16
1.3	Availability of the Technical Monitoring Cockpit	18
1.4	Video Library	19
1.5	What's New	20
2	Getting Started with the User Interface	22
2.1	Search	23
2.2	Using the Help	24
2.3	Supported UI Technologies	26
2.4	UI Components	26
	Add to Favorites	26
	Aggregation	27
	Apply	30
	Chart Settings	30
	Data Providers	31
	Filter	31
	Metrics and Characteristics	33
	Pivot	33
	Row Settings	35
	Select Characteristics (Drill Down)	37
	Select Metrics	37
	Show Identifiers Only	38
	Sort	38
	Time Slider	38
	Zoom Into Chart	40
3	Authorizations and Security	41
3.1	Authorization Concept of the Technical Monitoring Cockpit: Overview	42
3.2	Authorization Objects for the Technical Monitoring Cockpit	43
3.3	Authorizations: Examples	45
4	Setting Up the Technical Monitoring Cockpit	49
4.1	Settings for SAPUI5	49
	Browser Support	49
	Setting Up Web Services for SAPUI5	50
	Setting Up Ports for SAPUI5	50

	Checking the SAPUI5 Setup.	50
5	Managing Your System Landscape.	52
5.1	Setting Up the System Landscape for Monitoring.	52
	Prerequisites for the Host.	52
	Adding a Database or Host to the Monitoring Landscape.	53
	Usage Modes of the Managed System.	54
	Configuration of the Data Collection for the Application Server ABAP.	55
	Configuring the Data Collection Statistics for Host Data Providers.	65
5.2	Displaying and Updating Your Managed Systems.	66
	Managed Systems.	66
	System Details.	67
6	System Monitoring.	73
6.1	Coverage of Your Entire Technology Stack.	74
6.2	Monitoring Screens.	75
6.3	Selecting a System for Monitoring.	76
6.4	Dashboards.	76
	Database System Details.	77
	Application System Details.	77
6.5	System Monitoring: When to Use What.	78
6.6	Data Explorer: Create Your Own Monitoring Screens.	79
	What Is the Data Explorer?.	79
	What Are Data Providers?.	80
	What Are Technical Data Providers?.	81
	Example: Guided Tour on How to Use the Data Explorer.	82
7	Adapting the Technical Monitoring Cockpit.	104
7.1	Creating a Dashboard.	105
7.2	Changing the Tile Arrangement on the Launchpad or on a Dashboard in SAPUI5.	106
7.3	Changing the Tile Arrangement on the Launchpad or on a Dashboard in SAP GUI.	107
7.4	Adding a New Folder to the Menu.	107
7.5	Changing or Deleting a Menu Item in SAPUI5.	108
7.6	Changing or Deleting a Menu Item in SAP GUI.	109
8	Tools.	111
8.1	Content Explorer.	111
8.2	Screen Catalog.	111
9	Monitoring the SAP HANA Database.	113
9.1	Analysis.	113
	SQL Statement Analysis.	113
9.2	Information.	114

	Alerts.	115
	Backup and Recovery.	116
	CPU.	116
	Configuration.	118
	I/O.	119
	Logs.	119
	Memory.	121
	Storage.	128
	Time Spent.	129
	Workload.	130
9.3	Tools.	137
	SQL EXPLAIN.	137
9.4	Data Providers in the Data Explorer.	138
	Data Providers for the SAP HANA Database.	138
	Technical Data Providers for the SAP HANA Database.	165
10	Monitoring SAP Adaptive Server Enterprise.	167
10.1	Preconfigured Monitoring Screens for SAP ASE.	167
	Configuration.	167
	CPU.	172
	Job Scheduler.	174
	Memory.	176
	Performance.	178
	Storage.	180
	Time Spent.	183
	Workload.	186
10.2	Data Providers for SAP Adaptive Server Enterprise.	193
	Available Free Memory.	193
	Cache Hit Ratios.	193
	Cache Pool Info.	194
	CPU Utilization.	194
	Database Dump History.	194
	Database Options.	195
	Database Space Usage.	195
	Data Caches.	196
	Data Cache Objects.	196
	Data Cache Usage.	197
	Data Indexes.	197
	Device IO Statistics.	198
	Device Space Usage.	198
	Expensive SQLs by CPU Time.	199
	Expensive SQLs by Total I/O.	199

Frequently Used Tables.	200
Global Trace Flags and Switches.	200
HADR Device Usage.	201
HADR Resource Configuration.	201
HADR sap_collect_log.	202
HADR sap_get_log.	202
HADR sap_help.	203
HADR sap_send_trace.	203
HADR sap_set.	204
HADR sap_status active_path.	204
HADR sap_status path.	204
HADR sap_status_resource.	205
HADR sap_status route.	205
HADR Ticket History.	206
HADR Version.	206
IO Queue Statistics.	207
Job Scheduler - Activity.	207
Job Scheduler - History.	208
Job Scheduler - Jobs.	208
Job Scheduler - Schedules.	209
Locks per Process.	209
Logins.	210
Log Threshold.	210
Long Running Transactions.	210
Memory Usage Percent.	211
Memory Utilization.	211
Monitor Configuration.	212
Number of Open Objects.	213
Number of Open Partitions.	213
Procedure Cache Module Usage.	214
Procedure Cache Size in MB.	214
Procedure Cache Usage.	215
Resource Limits.	215
Resource Utilization.	216
Running Processes.	216
SAP ASE License.	217
SAP ASE Version.	217
SAP Configuration Check.	218
SAP NetWeaver Version.	218
Segments.	219
Server Configuration.	219

	Spinlock Contention.	220
	SQL Statements.	220
	Stored Procedure - Database Space Usage.	221
	System Utilization.	221
	Temporary Database Activity.	222
	Thread Activity	222
	Time Spent.	223
	Top 10 Modules Using ProcCache.	223
	Top 10 Tables By Growth.	224
	Top 10 Tables by Size.	224
	Top 5 User Process By CPU Time.	224
	Transaction Log Usage.	225
	Triggers.	225
	Unused Indexes.	226
	User Connections	226
	Workload Statistics.	226
10.3	Technical Data Provider for SAP Adaptive Server Enterprise.	227
11	Monitoring IBM Db2 for i.	228
11.1	Data Providers for IBM Db2 for i.	228
	SQL Statements.	228
	Table Scans.	228
	Query Sorts.	229
	Created Indexes.	230
	Advised Indexes.	231
	Used Indexes.	231
	Temporary Files.	232
	Temporary Bitmaps Creation and Merges.	233
	Temporary Hash Tables.	234
	Host Variables.	234
11.2	Technical Data Providers for IBM Db2 for i.	235
	Technical Data Provider Based on Db2 for IBM i Data Sources.	235
	Technical Data Provider Based on SAP Database Performance Collector for IBM i.	236
	Technical Data Provider Based on PLAN_CACHE Table of the SAP Database Performance Collector for IBM i.	237
12	Monitoring IBM Db2 for Linux, UNIX, and Windows.	238
12.1	Preconfigured Monitoring Screens for IBM Db2 for Linux, UNIX, and Windows.	238
	Backup and Recovery.	238
	Configuration.	238
	Memory.	240
	Network.	242

	Time Spent.	243
	Workload.	244
12.2	Data Providers for IBM Db2 for Linux, UNIX, and Windows.	246
	Backup.	246
	Costly Transactions.	247
	Database Parameters.	247
	Memory Overview.	248
	Memory Statistics.	249
	Network Statistics.	250
	Time Spent.	250
	Workload Overview.	251
	Workload Statistics.	251
12.3	Technical Data Providers for IBM Db2 for Linux, UNIX, and Windows.	252
	Technical Data Provider Based on Db2 Data Sources.	252
	Technical Data Provider Based on SAP Tables.	252
12.4	Tools.	253
	SQL EXPLAIN.	253
13	Monitoring IBM Db2 for z/OS.	255
13.1	Preconfigured Monitoring Screens for IBM Db2 for z/OS.	255
	Alerts.	255
	Bufferpools.	264
	Configuration.	270
	DB2 Subsystem Activity.	270
	DB2 Subsystem Summary.	275
	Jobs.	278
	SQL.	279
	Space.	284
	Storage.	302
	Thread Activity.	314
13.2	Data Providers for IBM Db2 for z/OS.	315
	DB2 Installation Parameters.	315
	DB2 Alerts.	315
	Deadlocks Overview.	316
	Statement Cache History.	317
	Statement Cache Overview.	318
	Timeouts Overview.	318
	Indexes Overview.	319
	Tables Overview.	320
	Top SQL Statements.	321
	Thread Activity Overview.	321
13.3	Technical Data Providers for IBM Db2 for z/OS.	322

Technical Data Provider DB2LOCKESCAL.	322
Technical Data Provider DB2QBSTT.	323
Technical Data Provider DB2QDB02_OBJECT.	323
Technical Data Provider DB2QDB02_SCHEMA.	324
Technical Data Provider DB2QDB02_SMS.	324
Technical Data Provider DB2QDBPT1.	325
Technical Data Provider DB2QDSCHIST.	325
Technical Data Provider DB2QGLOTIME.	326
Technical Data Provider DB2QLODLK.	326
Technical Data Provider DB2QLODRS.	327
Technical Data Provider DB2QLOGS.	327
Technical Data Provider DB2QLOTHW.	328
Technical Data Provider DB2QLOTRS.	328
Technical Data Provider DB2QLRUR.	329
Technical Data Provider DB2QSYSW1.	329
Technical Data Provider DB2QSYSW2.	330
Technical Data Provider DB2QTHDDSC.	330
Technical Data Provider DB2QTHDDSCP.	331
Technical Data Provider DB2QTHDDL.	331
Technical Data Provider DB2QTHDLRP.	332
Technical Data Provider DB2QTHDLST.	332
Technical Data Provider DB2QW0225.	333
Technical Data Provider SYSIBM.DB2_SYSPARM.	333
Technical Data Provider SYSIBM.SYSAUXRELS.	334
Technical Data Provider SYSIBM.SYSCOLUMNS.	334
Technical Data Provider SYSIBM.SYSINDEXES.	335
Technical Data Provider SYSIBM.SYSINDEXPART.	335
Technical Data Provider SYSIBM.SYSINDEXSPACESTATS.	336
Technical Data Provider SYSIBM.SYSTABLEPART.	336
Technical Data Provider SYSIBM.SYSTABLES.	337
Technical Data Provider SYSIBM.SYSTABLES_PROFILES.	337
Technical Data Provider SYSIBM.SYSTABLESPACE.	338
Technical Data Provider SYSIBM.SYSTABLESPACESTATS.	338
Technical Data Provider SYSIBM.SYSTABSTATS.	339
14 Monitoring Microsoft SQL Server.	340
14.1 Preconfigured Monitoring Screens for SQL Server.	340
Alerts.	340
CPU.	341
Configuration.	342
High Availability.	343
I/O.	344

	Memory.	346
	SQL.	348
	Space.	348
	Tables.	350
	Time Spent.	350
	Workload.	351
14.2	Data Providers for Microsoft SQL Server.	352
	Batch Requests/sec.	352
	CPU Utilization.	353
	Database File Size History.	353
	Database Size.	354
	File I/O Statistics.	354
	I/O Wait Statistics.	355
	Memory Statistics.	356
	Page Life Expectancy.	356
	SQL Server Setup Check.	357
	Table Size.	358
	Transactions per Second.	358
	Wait Statistics.	359
14.3	Technical Data Provider for Microsoft SQL Server.	360
15	Monitoring the Oracle Database.	361
15.1	Data Providers for Oracle.	361
	CPU Sessions.	361
	Memory SGA Allocation.	361
	Memory SGA Consumption.	362
	Space Datafiles.	362
	Space Tablespaces.	363
	Space Tempfiles.	363
	Time Spent Sessions.	364
	Workload Processes.	364
	Workload Sessions.	365
	Workload Statistics.	365
	Workload Users.	366
15.2	Technical Data Provider for the Oracle Database.	366
16	Monitoring the Application Server ABAP.	367
16.1	Analysis.	367
	ABAP Work Process Analysis.	367
	SQL Trace Analysis.	368
16.2	Information.	371
	CPU.	371

	Memory.	372
	Time Spent.	374
	Traces.	375
	Workload.	376
16.3	Data Providers for the Application Server ABAP.	380
	Application Hit List by DB Access.	380
	Application Hit List by Response Time.	381
	Application Statistics.	382
	Database Connections.	383
	Database Procedures.	384
	Enqueue Lock Table.	385
	Front-End Communication.	385
	Hit List by Database Access.	386
	Hit List by Response Time.	387
	Load from External Systems.	388
	Memory Consumption.	389
	Memory Utilization (Application Server ABAP).	390
	RFC Client Calls.	391
	RFC Client Destinations.	392
	RFC Server Calls.	393
	RFC Server Destinations.	394
	Response Time Distribution.	395
	Spool Activity.	396
	TREX Statistics.	396
	Table Access.	397
	Time Profile.	398
	SQL Trace Directory.	399
	Transaction Detail.	399
	User and Transaction Profile.	400
	User-Related Data.	401
	VMC Statistics.	402
	Web Client Calls.	403
	Web Client Destinations.	403
	Web Server Calls.	404
	Web Server Destinations.	405
	Workload Overview.	405
	Work Process Utilization.	406
16.4	Technical Data Providers for the Application Server ABAP.	407
	Technical Data Provider TA_ABAP_WP_INFO_SESSION_LIST.	408
	Technical Data Provider TA_ABAP_WP_INFO_WP_LIST.	408
	Technical Data Provider TA_ABAP_WP_INFO_WP_LOAD.	409

	Technical Data Provider TA_ABAP_WP_INFO_WP_SUMMARY.	410
17	Monitoring the Host.	411
17.1	Preconfigured Monitoring Screens for the Host.	411
17.2	Data Providers for the Host.	411
	Configuration of the Computer System.	411
	Configuration of the Host Computer System.	412
	Configuration of the Operating System.	412
	Configuration of the Virtual Computer System.	413
	Configuration of CIM Provider Capabilities.	414
	Configuration of File Systems.	414
	Configuration of the Host Agent.	415
	Configuration of Network Ports.	415
	Configuration of Processors.	416
	Disks.	416
	File Systems.	417
	Host Agent Components.	417
	Host CPU Utilization.	418
	Host Computer System.	419
	Host Memory Utilization (Linux).	420
	Host Memory Utilization (Windows).	421
	Memory Overview.	422
	Network Ports.	423
	OS Processes.	423
	Operating System.	424
	Processors.	424
	Virtual CPU Utilization.	425
	Virtual Computer System.	426
	Virtual Memory Utilization.	426
17.3	Technical Data Provider for the Host.	428

1 Technical Monitoring Cockpit (On-Premise, 7.4-7.56)


The technical monitoring cockpit is a monitoring and analysis tool for SAP systems based on SAP NetWeaver and the ABAP Platform.

📘 Note

This documentation refers to the technical monitoring cockpit that is part of SAP NetWeaver (SAP BASIS 7.4-7.56). For more information about recent versions of the technical monitoring cockpit that is part of ABAP Platform 2022 and higher, see [Technical Monitoring in the ABAP Platform](#).

You want to monitor and analyze your database and host? And the local Application Server ABAP? Virtual or physical computer systems, too? No problem. It's all in one place with the technical monitoring cockpit.

The technical monitoring cockpit has the following features:



- [Getting Started with the User Interface \[page 22\]](#)
- [System Monitoring \[page 73\]](#)
- [Authorizations and Security \[page 41\]](#)

- [Adapting the Technical Monitoring Cockpit \[page 104\]](#)
- [Coverage of Your Entire Technology Stack \[page 74\]](#)
- [Supported UI Technologies \[page 26\]](#)

Coverage of the Whole Technology Stack (Application Server ABAP, Database, Host)

Previously, you had to use layer-specific monitors for each component of the technology stack, and they all might have had a different way of querying, storing, and displaying data. With the technical monitoring cockpit, you no longer need to work with multiple monitors. Now, you can identify, locate, and analyze problems throughout your SAP technology stack, including the local Application Server ABAP, databases, and hosts (depending on your environment). All technical monitoring data can be queried and analyzed the same way.

For more information, see [Coverage of Your Entire Technology Stack \[page 74\]](#).

Uniform Operations and Consistent Layout with All Databases and Operating Systems Supported by SAP

Databases and operating systems of different vendors typically work differently and require individual training. The same is true for the utilities around those products. This isn't the case for the technical monitoring cockpit, as it works the same way for all.

For more information, see [Getting Started with the User Interface \[page 22\]](#).

Ready to Use

Thanks to preconfigured monitoring and analysis screens as well as dashboards, the technical monitoring cockpit is ready to use.

For more information, see [System Monitoring \[page 73\]](#).

SAP Fiori User Experience

The user interface (UI) of the technical monitoring cockpit is based on SAP Fiori design principles, which allows you to use the latest features available for SAP UIs. If you still need or want to use SAP GUI, you can do so as an alternative.

For more information, see [Supported UI Technologies \[page 26\]](#) and [Getting Started with the User Interface \[page 22\]](#).

Access Control

Monitoring and administration roles can vary. You might want to control who has access to which system with which authorization. With the authorization objects and usage types for managed systems, you can control visibility and access to systems.

For more information, see [Authorization Concept of the Technical Monitoring Cockpit: Overview \[page 42\]](#).

Simple Adaptation

You can adapt your launchpad with just a few clicks. Configure and save your favorite monitoring queries as menu items or as tiles on the launchpad, which makes them available to other users, too.

For more information, see [Adapting the Technical Monitoring Cockpit \[page 104\]](#).

Target Users of the Technical Monitoring Cockpit

- Operators and administrators of the Application Server ABAP
- Database operators and administrators
- Operating system administrators
- Database and operating system support engineers at SAP and SAP technology partners

Calling Up the Technical Monitoring Cockpit

To call up the technical monitoring cockpit, use transaction `STMC_UI5` for SAPUI5 and `STMC` for SAP GUI.

Note


Before you start working with the technical monitoring cockpit, make sure that you follow the instructions in SAP Note [2721488](#) .

Also check SAP Note [2818920](#)  for current updates or corrections to the documentation for the technical monitoring cockpit.

1.1 Learning the Basics

You want to start using the technical monitoring cockpit and don't know where to begin? Browse our most important help topics.

Must-Know Help Topics

Introduction  Read the introduction	Guided Tour  Read an example of how to explore monitoring data	System Landscape  Set up systems for monitoring
Security  Learn about roles and authorizations	Adaptation  Create favorites & adapt the launchpad to your needs	

- [Technical Monitoring Cockpit \(On-Premise, 7.4-7.56\) \[page 12\]](#)
- [Example: Guided Tour on How to Use the Data Explorer \[page 82\]](#)
- [Adding a Database or Host to the Monitoring Landscape \[page 53\]](#)
- [Authorizations and Security \[page 41\]](#)
- [Adapting the Technical Monitoring Cockpit \[page 104\]](#)


1.2 Architecture: Engine and Content

The technical monitoring cockpit consists of two major building blocks: an **engine**, which is database- and operating system-independent, and **content**, which is specific to databases, operating systems, or Application Server ABAP.

In this documentation, we occasionally talk about the “content” and the “engine” of the technical monitoring cockpit, and you might wonder what this is.

What Is Content? What Is the Engine?

The following figure outlines the architecture of the technical monitoring cockpit, including its separation of engine and content:



- [Adapting the Technical Monitoring Cockpit \[page 104\]](#)
- [Monitoring Screens \[page 75\]](#)
- [What Are Data Providers? \[page 80\]](#)
- [Dashboards \[page 76\]](#)
- [Example: Guided Tour on How to Use the Data Explorer \[page 82\]](#)
- [Search \[page 23\]](#)
- [System Monitoring \[page 73\]](#)
- [System Topology \[page 68\]](#)
- [Supported UI Technologies \[page 26\]](#)
- [Coverage of Your Entire Technology Stack \[page 74\]](#)
- [Availability of the Technical Monitoring Cockpit \[page 18\]](#)

Architecture of the Technical Monitoring Cockpit

The engine is the part of the technical monitoring cockpit with fundamental operations that are independent from databases, operating systems, or Application Server ABAP. The engine interprets the content and provides functions that build up the user interface and handle the user interaction. In addition, the engine provides core services such as topology detection and connection handling.

The content is the part of the technical monitoring cockpit that is platform-specific. It comprises, for example, monitoring and analysis screens, data providers, and dashboards specific to databases, operating systems, and Application Server ABAP, including charts and table overviews.

For What Systems Can You Get Content?

Content is available for various databases and operating systems supported by SAP as well as for the Application Server ABAP. For more information, also about required releases, versions, and support packages, see SAP Note [2721488](#) .

What Does the Content Contain?

The following content is available in the technical monitoring cockpit:

Launchpad

The launchpad is your landing page when you start the technical monitoring cockpit, just like a home page. It's the entry point to all important activities. With a content update, you might get additional tiles on the launchpad.

For more information about the launchpad, see [Adapting the Technical Monitoring Cockpit \[page 104\]](#).

Monitoring Screens

Monitoring screens provide you with a preconfigured view of the most important metrics in a particular monitoring situation for your database, host (including your operating system), or for the Application Server ABAP. Preconfigured monitoring screens are an easy start for administrators who want to monitor their systems and want to see the most important KPIs at a glance without having to select from the large range of available metrics and characteristics that are available in the technical monitoring cockpit. With a content update, more preconfigured monitoring screens might be available for your platform.

For more information about monitoring screens, see [Monitoring Screens \[page 75\]](#)

Data Providers

Data providers are objects in the technical monitoring cockpit that consist of a selection of characteristics and metrics for databases, hosts, and Application Server ABAP. Each data provider has its own aggregated view of important characteristics and metrics. With a content update, more data providers might be available for your platform.

For more information about data providers, see [What Are Data Providers? \[page 80\]](#).

Associated Go-To Screens

Have you found a monitoring screen or a data provider for specific metrics, but it's not exactly what you need? Then try the [Go To..](#) feature on the UI: it will lead you to other, associated screens that contain the same or

similar metrics, but with a different focus. With a content update, more or more finely tuned associated go-to screens might be available.

Dashboards

A dashboard is an arrangement of tiles that you can create yourself according to your monitoring needs and preferences. With a content update, you might find more dashboards available for your.

For more information about dashboards, see [Dashboards \[page 76\]](#).

Metrics

When you call up, for example, a data provider or preconfigured monitoring screen, you see a few selected metrics that are aggregated, as predefined by SAP. You can add or deselect metrics, and much more. After a content update, data providers or preconfigured monitoring screens might contain updated or additional metrics.

For more information about metrics, see [Example: Guided Tour on How to Use the Data Explorer \[page 82\]](#).

Analysis

The available screens in the menu under [Analysis](#) provide you with more detailed information than the monitoring screens, helping you to run a root cause analysis of system issues. With a content update, you might find additional analysis screens ready for use.

For more information about analysis screens, see [System Monitoring \[page 73\]](#).

Evaluation

For some of your monitoring and administration work, you need information whether settings deviate from the standard. For example, you might want to see whether the database parameters for SAP HANA are set according to SAP recommendations. For this purpose, some of the available monitoring screens and data providers contain evaluations. With a content update, such evaluations might be updated, or more evaluations are added to monitoring screens and data providers.

For an example of a monitoring screen with evaluations, see [Parameter Check \[page 148\]](#) for the SAP HANA database.

Search Index

The powerful search function lets you search across the whole technical monitoring cockpit, no matter if it's data providers, monitoring or analysis screens, tables, or any other object or field you're looking for. With every content update, the search index is automatically updated, too.

For more information, see [Search \[page 23\]](#).

1.3 Availability of the Technical Monitoring Cockpit

Although the technical monitoring cockpit is part of and delivered with every SAP on-premise system based on SAP NetWeaver 7.40 and higher, its content and content updates are independent of SAP NetWeaver releases

and support packages (read on to learn more). In addition, the technical monitoring cockpit is also available for SAP BTP, ABAP environment.

On-Premise Systems

Implement SAP Note [2721488](#) to get the technical monitoring cockpit engine.

New features for monitoring the application server ABAP, databases, and operating systems are shipped as part of the content for the technical monitoring cockpit, while the engine remains unchanged. Updates are shipped via SAP Note transport-based correction instruction (TCI). For more information, see SAP Note [2721488](#). In addition, you might also find updated content with support packages for the supported SAP NetWeaver releases.

SAP BTP, ABAP Environment

There's also another version of the technical monitoring cockpit with a different feature scope, which is available as part of SAP BTP, ABAP environment. For more information, see the documentation for the technical monitoring cockpit (cloud version) at https://help.sap.com/viewer/tmc_cloud.

Related Information

[Architecture: Engine and Content \[page 16\]](#)

1.4 Video Library

Watch our videos to get an overview of the technical monitoring cockpit.

- [Introduction](#)
- [Guided Tour of the Data Explorer](#)
- [How to Use the Help](#)

1.5 What's New

Get an overview of what's new with the technical monitoring cockpit.

Content Version	What's New
7.40-7.56 (technical ID: 02.1909.0000)	This content is available as part of systems based SAP NetWeaver and ABAP Platform (SAP BASIS release 7.40 to 7.56). For more information about the availability of the technical monitoring cockpit for on-premise and the relevant releases and support packages, see SAP Note 2721488. New database supported: IBM Db2 for i Corrected and enhanced data providers for IBM Db2 for z/OS
2.0.0 (02.1908.0016)	New database supported: IBM Db2 for z/OS New features and content: • Updated SAP UI5 layout • New chart types and settings • New data providers and new preconfigured monitoring screens • New dashboards <i>Database System Details</i> and <i>Application System Details</i> • Analysis of SQL statements and work processes
1.7.0	New database supported: Microsoft SQL Server New data providers for the following: • SAP ASE • SAP HANA database Note Version 1.7.0 and lower is not supported any longer.
1.6.0	New databases and operating systems supported: • SAP ASE database • Oracle database • Operating system: AIX New data providers and preconfigured monitoring screens for the following: • IBM Db2 for Linux, UNIX, and Windows • SAP HANA database
1.5.1	Patches only
1.5.2	
1.5.3	


Content Version	What's New
1.5.0	First version released to customers with the following operating systems and databases supported: • SAP HANA database • IBM Db2 (for Linux and Windows only) • Linux with VMware • Windows with Hyper-V

2 Getting Started with the User Interface

The user interface of the technical monitoring cockpit provides a consistent user experience for the entire technology stack. All technical monitoring data can be queried and analyzed the same way for databases, the Application Server ABAP, hosts and operating systems.

Consistent Layout

In the technical monitoring cockpit, the layouts of monitoring screens and data providers are structured identically, no matter whether you choose a monitoring screen or data provider for a database, the host or operating system, or the Application Server ABAP. In the table view of the screens, you will always find the same elements, such as a time slider and buttons for selecting metrics and characteristics, for filtering, for sorting, and for configuring row settings. In this way, you get a uniform overall user experience.



Embedded Help

By choosing the help button at the top of the screen, you can turn on help hot spots, which indicate where help is available. With a click on a hotspot, you can call up the embedded help of the technical monitoring cockpit.



SAPUI5 and SAP GUI

You can run the technical monitoring cockpit with SAPUI5 or with SAP GUI.

Related Information

[Using the Help \[page 24\]](#)

[Supported UI Technologies \[page 26\]](#)

2.1 Search

With the powerful search function of the technical monitoring cockpit, you can search for data providers, monitoring screens, and tables that match your search words.

Here's what you'll get in the search results:

- Dashboards or monitoring screens (defined by SAP or by you) whose labels, tooltips, or fields contain your search words
If you choose a monitoring screen in the search results, you will be directly forwarded to that screen.
- Data providers whose labels, tooltips, or fields contain your search words
If you choose a data provider in the search results, the data explorer opens with the selected data provider.
- Database objects such as tables whose labels, tooltips, or fields contain your search words

If you choose the database object in the search results, the data explorer opens with a data provider that gives you information about this database object (for example, about performance, storage, or configuration).

2.2 Using the Help


The technical monitoring cockpit has a convenient embedded help that you can switch on and off whenever you need.

Note

The display of Help on SAPUI5 screens is slightly different from SAP GUI screens.

SAPUI5

On SAPUI5 screens, use the ⓘ [Show/Hide Help Hotspots](#) button to activate the help. You can easily see where help topics are available then: The relevant UI elements are highlighted by green bubbles, so-called hotspots. If you choose a hotspot, a help window appears that contains the context-specific information about the UI element.




Help on SAPUI5 Screens

Do you want an overview of all the available help topics? Then choose [Open in New Window](#), and the complete documentation for the technical monitoring cockpit is displayed. From the table of contents on the left, you can then choose whatever topic you like, it's all there.


SAP GUI

On SAP GUI screens, choose the  *Activate Help Mode* button. You can then see  *Show Help* buttons for all UI elements for which help is available.



Help Buttons on SAP GUI Screens

If you choose one of the *Show Help* buttons, a help window appears that contains the context-specific information about the UI element.



Help Window in SAP GUI

Do you want an overview of all available help topics? Then choose the  *Open in New Browser Window* button in the help window, and the complete documentation for the technical monitoring cockpit is displayed. From the table of contents on the left, you can then choose whatever topic you like, it's all there.

Note

The display of help on SAP GUI-based screens is not fully supported and might lead to occasional problems. For example, single sign-on might not always work. If you encounter errors, use the SAPUI5-based screens.

2.3 Supported UI Technologies

With SAPUI5, you get the up-to-date SAP user interface (UI) technology, which is based on HTML5. The screens are built according to SAP Fiori design principles, with a focus on usability and responsive design. SAPUI5 is available for the technical monitoring cockpit in all SAP systems. To call up the technical monitoring cockpit with SAPUI5, use transaction `STMC_UI5`.

SAP GUI is also available for the technical monitoring cockpit in all on-premise SAP systems. With SAP GUI, you get SAP's older UI technology. The SAP GUI-based user interface has been implemented for all administrators who still work with many SAP GUI-based transactions and would like to switch to the technical monitoring cockpit without having a disruption in the UI display. To call up the technical monitoring cockpit with SAP GUI, use transaction `STMC`.

Related Information

[Settings for SAPUI5 \[page 49\]](#)

2.4 UI Components

2.4.1 Add to Favorites

Use the [Add to Favorites](#) button to save your settings, including selected filters, sorting, metrics, characteristics, and chart settings for a particular data provider in the data explorer. As a result, the data provider with its settings will be saved together with the preconfigured monitoring screens delivered by SAP in the menu of the technical monitoring cockpit.

Procedure

1. If you're satisfied with your settings and selections, choose the [Add to Favorites](#) button.
2. Enter a name, description, and a package for your favorite.

The name will appear in the menu in the folder that you assign. If you later create a tile on the launchpad for your favorite, the description will be displayed on the tile.

3. If you always want to use your favorite for the same time range and point in time, select [Keep Time](#).
4. If you want to insert your favorite in the menu in a folder other than the [Favorites](#) folder, choose [Select Parent Folder](#) and then select the folder that you want.

Result

As a result, your favorite is now available as a preconfigured monitoring screen in the menu. You can also add this favorite as a tile to the launchpad or a dashboard later. It's also possible to change the name, description, or folder of the favorite later.

Related Information

[Changing or Deleting a Menu Item in SAPUI5 \[page 108\]](#)

[Changing or Deleting a Menu Item in SAP GUI \[page 109\]](#)

[Changing the Tile Arrangement on the Launchpad or on a Dashboard in SAPUI5 \[page 106\]](#)

[Changing the Tile Arrangement on the Launchpad or on a Dashboard in SAP GUI \[page 107\]](#)

2.4.2 Aggregation

Aggregation is a process to summarize metrics, for example, by grouping values of multiple rows into a single result. Aggregations are calculated using functions such as average, count, maximum, minimum, or sum.

Context

In the technical monitoring cockpit, when you call up data providers or preconfigured monitoring screens, metrics are often shown in an aggregated form. Why is aggregation used here?

The sources from which the technical monitoring cockpit retrieves its data provide more details than are required for a health check or for identifying and locating a potential issue. Therefore, the technical monitoring cockpit reduces the complexity by aggregating data first before displaying metrics to the user. Aggregation functions such as sum or maximum are modeled on the data provider layer and cannot be changed by users.

To better understand how data displayed by the technical monitoring cockpit is aggregated, you need to know that the technical monitoring cockpit uses two types of aggregation: aggregation by characteristics and an additional aggregation, aggregation by time. Aggregation by characteristics is used to condense data with the same characteristic values, as typically done with a **GROUP BY** clause in an SQL statement. Standard aggregation functions include sum, maximum, and minimum.

Aggregation by time is used to condense data even further. In the technical monitoring cockpit, time aggregations are used to calculate the average, the minimum, or the maximum of metrics for a certain time

period. For example, the technical monitoring cockpit might sum up the memory size according to different memory heaps first and then take the average heap size of the summed-up heap sizes for a certain time interval.

Example

Let's have a closer look at an example. We assume that you are interested in the memory utilization of your database. The fine-grained time series retrieved by the technical monitoring cockpit from its data sources roughly looks like this:

Time Stamp	H1 Physical Memory (Characteristic)	H2 HANA Allocated Memory (Characteristic)	Average Memory [MB] (Metric)
2018.04.28 10:33:00	Not yet Allocated	Not yet Allocated	113,861.1
2018.04.28 10:33:00	Reserved by Operating System	Reserved by Operating System	103,198.9
2018.04.28 10:33:00	SAP HANA Memory Pool (Free)	SAP HANA Memory Pool (Free)	558,413.2
2018.04.28 10:33:00	SAP HANA Used Memory	HANA Used Code and Stack	4,171.8
2018.04.28 10:33:00	SAP HANA Used Memory	Memory Usage of Database Management	252,333.7
2018.04.28 10:34:00	Not yet Allocated	Not yet Allocated	113,861.0
2018.04.28 10:34:00	Reserved by Operating System	Reserved by Operating System	103,198.9
2018.04.28 10:34:00	SAP HANA Memory Pool (Free)	SAP HANA Memory Pool (Free)	558,264.5
2018.04.28 10:34:00	SAP HANA Used Memory	HANA Used Code and Stack	4,171.8
2018.04.28 10:34:00	SAP HANA Used Memory	Memory Usage of Database Management	252,482.5

(Actually, this example already contains pre-aggregated averages by minute, but for the sake of simplicity, let's ignore this.) In this example, for every minute of the day, you see an update of the metric [Average Memory](#), grouped by the two characteristics [H1 Physical Memory](#) and [H2 HANA Allocated Memory](#).

You might not want such a detailed overview; perhaps you are only interested in how much memory is used by SAP HANA, but no further details about how memory is allocated within SAP HANA. That's why the technical


monitoring cockpit uses aggregation and sums up the memory for HANA-allocated memory types into a single result per minute. The corresponding column *H2 HANA Allocated Memory* disappears:

Time Stamp	H1 Physical Memory	Average Memory [MB]
2018.04.28 10:33:00	Not yet Allocated	113,861.1
2018.04.28 10:33:00	Reserved by Operating System	103,198.9
2018.04.28 10:33:00	SAP HANA Memory Pool (Free)	558,413.2
2018.04.28 10:33:00	SAP HANA Used Memory	256,505.6
2018.04.28 10:34:00	Not yet Allocated	113,861.0
2018.04.28 10:34:00	Reserved by Operating System	103,198.9
2018.04.28 10:34:00	SAP HANA Memory Pool (Free)	558,264.5
2018.04.28 10:34:00	SAP HANA Used Memory	256,654.3

However, this type of aggregation alone is often not enough. You might want to see the average memory for a longer time period rather than scroll through results for every single minute, as you can see here. This is where time aggregation is helpful. In the technical monitoring cockpit, you can define a time period that you are interested in. Then the technical monitoring cockpit calculates new averages based on your selected period. For the sake of simplicity, let's assume you set the period of interest from 10:33:00 to 10:34:00 (2 minutes). Then the technical monitoring cockpit will aggregate the data further, applying an average function. The result looks like this:

H1 Physical Memory	Average Memory [MB]
Not yet Allocated	113,861.0
Reserved by Operating System	103,198.9
SAP HANA Memory Pool (Free)	558,338.9
SAP HANA Used Memory	256,579.9

This is a simplified example – in the technical monitoring cockpit, you would probably look at longer periods than just two minutes like in the example used above, but look at metrics for an hour or a day. Such an aggregation is shown by default when you call up the data provider [Memory Utilization](#):



In the data explorer, you can now drill down, that is, select further characteristics such as HANA-allocated memory, DB management memory, and memory usage of tables to see more detailed (less aggregated) results and even display data at the finest level of detail.

How to Find Aggregation Information

Now how do you know which aggregation has been used? In the technical monitoring cockpit, you can find the type of time aggregation in the column label text of a metric. If no aggregation type is specified in the label text, the underlying aggregation type is aggregation by characteristics. The aggregation by characteristics is the most obvious aggregation that you would expect, for example, sum in most cases, so there's no need to specify it on the UI. Only if a time aggregation is applied that is different from the aggregation by characteristics, is an aggregation type shown in the column label text.

2.4.3 Apply

Don't forget to choose the [Apply](#) button (or press [Enter](#)). Otherwise, you won't get any results.

Only after you have chosen [Apply](#), does the technical monitoring cockpit apply the time periods, filters, sort orders, metrics, drilldowns, and number of rows that you have specified.

2.4.4 Chart Settings

If you go to [Settings](#), you can adapt your chart in many ways for a better visualization of your data:

- Change the chart type (time- or hierarchy-based, or histogram) and data series type (for example, pie chart, donut, treemap, columns, lines, Gantt, and more).
- Select a relative and/or stacked chart view as well as a second axis.
- Choose between different color schemes, and reverse or swap colors.

Related Information

[Changing Chart Settings \[page 97\]](#)

2.4.5 Data Providers

Data providers are predefined views that contain metrics useful for monitoring databases, operating systems, or application servers.

Use data providers to display metrics or to build your own favorite monitoring screen for a particular monitoring scenario.

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[Add to Favorites \[page 26\]](#)

2.4.6 Filter

Use the [Filter](#) button to limit your results to criteria that you define.

In the [Filter](#) popup, you can use the following operators:

Operator	Description
Equal	The field's value is equal to the value specified.
Not Equal	The field's value is not equal to the value specified.
Greater Than	The field's value is greater than the value specified. Note that the system sorts the values alphanumerically.
Less Than	The field's value is less than the value specified.

Operator	Description
<i>Like</i>	The field's value is like the value specified. When you specify the value for <i>Like</i>, you can use the % character or the * character as wildcards to represent zero, one, or multiple characters. The wildcard + (0 or 1 occurrences) also works. Depending on the position of the wildcards, you get the following, for example: - With a%: All values beginning with the character <i>a</i> - With %a: All values ending with the character <i>a</i> - With %a%: All values with the character <i>a</i> in any position
<i>Not Like</i>	The field's value is different from the value specified. For example, if you enter a%, you get all values that do not begin with <i>a</i>.
<i>In List</i>	The field contains one of the values specified. For example, for the field <i>User</i>, you select <i>In List</i> and enter the values Peter, Paul, and Mary, the system limits the results to all values containing Peter, or Paul, or Mary.
<i>Not in List</i>	The field contains none of the values specified. For example, for the field <i>User</i>, you select <i>Not in List</i> and enter the values Peter, Paul, and Mary, the system limits the results to all values containing neither Peter, nor Paul, nor Mary.
<i>Is Null</i>	Field has no value; it's empty.
<i>Is Not Null</i>	Field has a value; it's not empty.

Filter by User Names

To filter for the name of the user who is currently logged on in the system, you can choose the operator *Equal* and enter the variable **@UI/USERNAME@** as value. If this filter is set for a data provider or monitoring screen, a user only sees data filtered by his or her user name.

If you want to see the traces of all users, you need to remove the filter **owner=@UI/USERNAME@**. This filter is only visible for you on the UI if you have the relevant authorizations of authorization object **S_TMC_DP** (field **S_TMC_FLD1**).

❖ Example

In the data provider *SQL Trace Directory*, by default, only your own traces are displayed. Only with the relevant authorizations can you see traces of other users.

2.4.7 Metrics and Characteristics

Data providers give you a view of relevant metrics and characteristics. Metrics are numeric values, whereas characteristics are attributes that further characterize metrics.

❧ Example

CPU Time = 0.2 s (metric), Service Name = 'webdispatcher' (characteristic)

2.4.8 Pivot

Use the [Pivot](#) button to swap rows and columns of the table and chart while keeping the time column as it is.

Here's an example:

In the first table, you can see the processing times for steps and transactions in the columns, with the hosts in the rows.

Host	Time Stamp	Processing Time/Step [ms]	Processing Time/Transaction [ms]
Host 1	2018-02-16 21:59:59	0.1	1.2
Host 2	2018-02-16 21:59:59	0	0
Host 3	2018-02-16 21:59:59	0	0
Host 4	2018-02-16 21:59:59	0.2	2
Host 1	2018-02-16 22:29:59	0.1	1.3
Host 2	2018-02-16 22:29:59	0.1	0.1
...	...	...	...

After pivoting the table, you can see processing times for four hosts in the columns, with the processing time per step or transaction in each row:

Time Stamp	Metrics	Host 1	Host 2	Host 3	Host 4
2018-02-16 21:59:59	Processing Time/Step [ms]	0.1	0	0	0.2
2018-02-16 21:59:59	Processing Time/Transaction [ms]	1.2	0	0	2
2018-02-16 22:29:59	Processing Time/Step [ms]	0.1	0.1	0.1	0.4

Time Stamp	Metrics	Host 1	Host 2	Host 3	Host 4
2018-02-16 22:29:59	Processing Time/ Transaction [ms]	1.3	0.1	0.1	3
2018-02-16 22:29:59	Processing Time/ Step [ms]	0	0	0	0
2018-02-16 22:29:59	Processing Time/ Transaction [ms]	0	0	0	0


Especially when you turn the tables into charts, pivoting can make a difference and improve the visualization of your data (depending on your use case).

Here's an example of a chart before and after pivoting:

Before, there's a pie chart for every row (= characteristic of the drilldown) with two metrics :



After pivoting, there are only two charts now, one for each selected metric:



In the pivoted chart, you get a better view of which consumer holds which share of the overall percentage of CPU time and the overall number of CPUs (simply hover over or click on the different chart areas).

Related Information

[Example: Guided Tour on How to Use the Data Explorer \[page 82\]](#)

2.4.9 Row Settings

You can use the [Rows](#) button to limit the number of rows and to create a top n analysis.

When you choose the [Rows](#) button, remember the following:

- If you choose a maximum number of n rows, the system will limit the table rows to the **first** n results that are returned by the system. These are not necessarily the **top** n results.
- If you want to create a top n analysis, you have to combine the [Sort](#) button with the [Rows](#) button:
 1. Sort the table by **metric** using the [Sort](#) button. Make sure you pick a metric as your first sorting criterion and **not** a characteristic. It's also mandatory that you do **not** select a time characteristic. For more information, see the note below.
 2. Choose the [Rows](#) button, specify the number of rows you want to display, and select the [Top N Aggregation](#) checkbox.

3. Don't forget to choose [Apply](#).

The system then shows the aggregated values of the top n entries in the row labelled with the total symbol at the end of the result list.

You can optionally select two additional rows to be displayed: One (with the symbol "[T](#)") showing the total aggregation values, the other one ("[R](#)") showing the aggregated values of the rest, that is, of the entries that don't belong to the top n entries.

Do not confuse the concept of aggregated values of the top n entries (with the sigma symbol Σ) and total aggregation values (symbol "[T](#)"). The aggregated values of the top n entries are calculated considering only the first n different characteristics combinations. The total aggregation values are calculated considering all possible characteristics combinations.

For example, if the CPU utilization of only the top 10 services is displayed, the [Total Aggregation](#) row indicates the aggregated values of all services, and the [Rest Aggregation](#) row contains the aggregated values for the CPU utilization of the remaining services that are not displayed.

Let's assume there are more than 10 services, let's say 15. For the calculation of the [Top N Aggregation](#), the system uses the first 10 services, and for the calculation of the [Total Aggregation](#) row, the 15 services are used. If there are exactly 10 services, the values in both rows are identical.

Note

You can sort by metrics and characteristics at the same time, but for the extra aggregation rows to be displayed, the following is mandatory:

- The **first** sorting criterion must be a metric.
- You did **not** select a time characteristic (such as timestamp) earlier from the drilldown options behind the [Characteristics](#) button.

When Does a Top N Analysis Make Sense?

A top n analysis only makes sense if you have selected at least one characteristic (make sure it's not a time characteristic) **and**, as your first sorting element, a metric. If you select a characteristic as your primary sorting criterion (let's say "service name"), it doesn't help you much when analyzing, for example, the CPU utilization of your database. What you rather need is numbers, that is, metrics (for example, CPU times) that are sorted in a helpful way. Conversely, if you select metrics, you need at least one characteristic as well to create a meaningful analysis.

Related Information

[Sort \[page 38\]](#)

[Example: Guided Tour on How to Use the Data Explorer \[page 82\]](#)

[Aggregation \[page 27\]](#)

2.4.10 Select Characteristics (Drill Down)

When you call up a data provider, the metrics displayed are aggregated (SUM, MAX, MIN, AVG, etc.) as predefined by SAP. You do a drilldown by choosing the [Characteristics](#) button and selecting/deselecting characteristics. By doing so, you change the grouping of the metrics.

ⓘ Note

Don't forget to choose the [Apply](#) button after you've confirmed the [Characteristics](#) popup window. Otherwise, your changes won't be applied.

After a drilldown, the result set typically comprises more than one row (metric), and you can select additional drilldown characteristics.

Key Characteristics

If you want to select the key characteristics (= deepest level of drilldown), you don't have to do it manually, but you can use the [Characteristics](#) button and then choose the [Select All Keys](#) button in the [Characteristics](#) popup window. All the key characteristics of that data provider are then automatically selected and displayed. You can still select more (non-key) characteristics manually, but this won't deepen the drilldown, it'll broaden it. This means you won't get more result rows, but it'll add more columns to your result list.

Related Information

[Metrics and Characteristics \[page 33\]](#)

[Example: Guided Tour on How to Use the Data Explorer \[page 82\]](#)

2.4.11 Select Metrics

The set of metrics that is displayed in columns is preset. You can select additional metrics or hide metrics that you don't need, and change the sequence of the displayed columns by using the [Move Up](#) and [Move Down](#) buttons and/or the [Move to Top](#) and [Move to Bottom](#) buttons respectively.

Don't forget to choose the [Apply](#) button after you've confirmed the [Metrics](#) popup window. Otherwise, your changes won't be applied.

Related Information

[Metrics and Characteristics \[page 33\]](#)

2.4.12 Show Identifiers Only

You can limit the number of shown columns to only ad-hoc computed identifiers by choosing the toggle button with the key symbol ([Show All Columns](#) / [Show Identifiers Only](#)).

2.4.13 Sort

Use the [Sort](#) button to sort the complete set of results returned for your request.

You can also sort data on the screen by directly clicking on the relevant column. Note, however, that this kind of sorting only sorts the data that is displayed on the screen. If you need a complete data refresh, you must choose the [Sort](#) button and choose [Apply](#).

If you want to run a top n analysis, make sure you refresh your data completely using [Sort](#) and [Apply](#). Then limit your result to n rows. A top n analysis where you first limit the number of results choosing [Rows](#) and then sort the entries displayed on the screen leads to meaningless results.

For more information about top n analysis, see [Row Settings](#) [page 35].

2.4.14 Time Slider

Use the time slider to change the time period for the metrics on the screen:

- Web UI only: Drag and drop the time range on the time slider to the period that you want.
- Enter start and end times in the date/time fields. In the end time field on the right, you can also enter any word, for example, [now](#) or [today](#), or any character, and the field will switch to the current date and time.
- Use the arrows to move the time range forward and backward.
- Enter a time period in the entry field in the middle of the time slider using the following shortcuts:

Shortcuts for Time Periods

Time Period	Abbreviation	Example
Year	<Number>Y	1Y
Month	<Number>M	1M
Week	<Number>W	7W
Day	<Number>D	1D
hour	<Number>h	1h
minute	<Number>m	30m
second	<Number>s	30s

Note

In SAP GUI, the time slider shows the time zone of the monitored system whereas in SAP UI5 it's your browser time zone that is shown.

If you want to use the current date and time as ending point of your time range, you can just enter **now** (or simply **n**).

You can combine the shortcuts with the terms “last” or “align” as in the following examples:

Combined Shortcut Examples for the Time Slider

Examples	Description
last 12h L12h l 12h Last 12h 1h l 1hL	Data will be displayed for the last 12 hours until now.  Note It's OK to abbreviate last and it doesn't matter if you use a blank afterwards or not. Upper or lower case is fine, and you can type last (or simply L or l) at the beginning or end, either combination works.
last 5M 14D 12h L5M, 14D, 12h	Data will be displayed for the last 5 months and 14 days and 12 hours.  Note Possible delimiters between the time units are blanks, commas, semicolons, or hyphens.
1M 1D 12h align 1M 1D 12h a 1M 1D 12h A	If you add align to your entry, the time range calculation is aligned (that is, rounded) to the smallest given unit that you enter, in this case hours.  Note It's OK to abbreviate align and it doesn't matter if you use blanks or upper or lower case. You can type align (or simply A or a) at the beginning or end, either combination works.
The time slider is on the current date and time and you enter one of the following time ranges: 1h 1h a 1h l	If today's current date and time is, for example, 4-27-2017 5:15 p.m., the following will be displayed: Data from 4:15 p.m. to 5:15 p.m. today Data from 4:00 p.m. to 5:00 p.m. today Data from 4:15 p.m. to 5:15 p.m. today

Examples	Description
You set the time slider to a certain time range completely in the past and the ending point is not now:	Let's say your time range in the past is 4-20-2017 4:15 p.m. to 5:15 p.m. In this case, the following will be displayed:
1h	Data from 4:15 p.m to 5:15 p.m. on that day in the past
1h a	Data from 4:00 p.m. to 5:00 p.m. on that day in the past
1h l	Data from 4:15 p.m. to 5:15 p.m. today (because the attribute "last" is always aligned to "now", that is, to the current date and time).

2.4.15 Zoom Into Chart

You see a second, smaller chart below your chart, right below the magnifiers? That's a zoom function.

Here's how you zoom in and out:

- In the smaller chart, keep the mouse button pressed on the desired start time and drag to the right until the end of your desired time period. The bigger chart above now displays the data of this selected time interval, which is highlighted. You can extend or decrease the interval by dragging at the beginning or the end of the highlighted section.
- To return to the original settings without zoom, drag and drop the highlighted time period to the left until the beginning of the time scale. Then drag the end of the selected time period to the end of the time scale on the right until the highlighting for the time period disappears.

Of course, to zoom in and out, you can also simply use the magnifiers below the chart.

3 Authorizations and Security

The authorization concept for the technical monitoring cockpit protects your data in two ways. First, as an administrator for on-premise systems, you can decide on the right roles and authorizations for the users of the technical monitoring cockpit itself. Second, you can decide for each system monitored by the technical monitoring cockpit whether it should be accessed with administration, monitoring, or read-only authorizations.

Get an Overview

You want to get an overview of how authorization checks work with the technical monitoring cockpit? Then this is for you: [Authorization Concept of the Technical Monitoring Cockpit: Overview \[page 42\]](#).


Learn the Technical Details

You're keen on learning all technical details of how you can control authorizations? Then this is for you:

- [Authorization Objects for the Technical Monitoring Cockpit \[page 43\]](#)
- [Authorizations: Examples \[page 45\]](#)

3.1 Authorization Concept of the Technical Monitoring Cockpit: Overview

Learn more about the two-layered authorization concept of the technical monitoring cockpit.



The technical monitoring cockpit evaluates authorizations on two layers:

- **User layer**
User authorizations are based on the authorization objects.
Depending on your preferences, you can either use S_RZL_ADM (global authorization for using the technical monitoring cockpit) or the authorization objects S_TMC_DP and S_TMC_TM.
- **System layer**
You must define a usage mode when you add a new system to the managed systems of the technical monitoring cockpit. Depending on the usage mode, authorizations might be further restricted, that is, only monitoring or read-only is allowed in the system, even if an individual user has administration rights.

📌 Note

The authorization object S_DBCON is **not** used for the technical monitoring cockpit.

Related Information

[Authorization Objects for the Technical Monitoring Cockpit \[page 43\]](#)

3.2 Authorization Objects for the Technical Monitoring Cockpit

With the technical monitoring cockpit, you have authorization objects that allow you to control what users of the technical monitoring cockpit can do, from simple monitoring with read-only authorizations up to complex administration tasks, content management, and system landscape setup.

Overview

For the technical monitoring cockpit, you can choose between the following approaches for authorizations:

- **Reuse of existing authorizations**
If you have already worked with monitoring tools such as the DBA Cockpit, you'll probably be familiar with authorization object `S_RZL_ADM`. Users with the authorizations of authorization object `S_RZL_ADM` also have all permissions to use the technical monitoring cockpit. The advantage of using `S_RZL_ADM` is that you don't need to assign any new authorizations to your users if you already use it for other monitoring tools. The disadvantage of `S_RZL_ADM` is that it does not allow you to control access to data on system, database, operating system or user level.
- **Finely-tuned authorizations for the technical monitoring cockpit and system access**
Use authorization objects `S_TMC_DP` and `S_TMC_TM` to define more detailed authorizations for different tools of the technical monitoring cockpit and for system access.

Reuse of Existing Authorizations

Use the authorization object `S_RZL_ADM` to control whether users can display or change objects in the technical monitoring cockpit. With authorizations of this authorization object, they can also use the content management and the system landscape management of the technical monitoring cockpit.

Authorization Object `S_RZL_ADM`

Activity	Description
01 (CREATE)	Users get global access to all functions in the technical monitoring cockpit, including the authorizations to create, update, and delete objects. Users can also use the content management of the technical monitoring cockpit, including uploading, activating, resetting, and deleting content. Users also need this activity for all adaptations of the technical monitoring cockpit, including creating favorites, creating dashboards, and adapting dashboards and the launchpad.

Activity	Description
03 (DISPLAY)	Users can only display data in the technical monitoring cockpit.

To view or hide sensitive data (user names) from administration users, use authorization object `S_TOOLS_EX`.

Finely Tuned Authorizations for the Technical Monitoring Cockpit and System Access

With `S_TMC_DP`, you can control whether users of the technical monitoring can create or change objects in the technical monitoring cockpit or whether they can use the content management from the monitoring landscape. In addition, you can fine-tune whether users have permissions for the application server ABAP, specific databases, or hosts. With `S_TMC_TM`, you can control whether users can add and remove systems from the monitoring landscape.

Use authorization object `S_TMC_DP` (field `S_TMC_FLDT`) to control whether users of the technical monitoring cockpit can view user names.

For more information, see the system documentation of the authorization objects and the examples in this documentation.

Related Information

[Authorizations: Examples \[page 45\]](#)

3.3 Authorizations: Examples

Here are some examples of how you can fine-tune authorizations in the technical monitoring cockpit using the authorization objects S_TMC_DP and S_TMC_TM.

Display User for SAP HANA Monitoring

To create a display user for the technical monitoring cockpit who can only access data for a particular database, use the authorization objects S_TMC_DP and S_TMC_TM with the following values:

Values for Authorization Object S_TMC_DP

Field	Value
S_TMC_SVC (<i>TMC-DP: Service ID</i>)	*
S_TMC_PRV (<i>TMC-DP: Provider ID</i>)	*
S_TMC_FLDT (<i>TMC-DP: Field Type</i>)	S (standard, user names in data providers are hidden)
S_TMC_TMST (<i>TMC-TM: System Type</i>)	DBMS
S_TMC_TMSS (<i>TMC-TM: System Subtype</i>)	HDB (for SAP HANA, for example)
S_TMC_TMAG (<i>TMC-TM: Authorization Group</i>)	
ACTVT (<i>Activity</i>)	03 (Display)

Values for Authorization Object S_TMC_TM

Field	Value
S_TMC_TMST (<i>TMC-TM: System Type</i>)	DBMS
S_TMC_TMSS (<i>TMC-TM: System Subtype</i>)	HDB (for SAP HANA, for example)
S_TMC_TMAG (<i>TMC-TM: Authorization Group</i>)	
ACTVT (<i>Activity</i>)	03 (Display)

Monitoring User

With the following values, you can create authorizations for a user in the technical monitoring cockpit who can monitor systems and also save preferred selections in the data explorer as favorites (monitoring screens) in the menu. In addition, the user can adapt the launchpad and create and change dashboards.

Values for Authorization Object S_TMC_DP

Field	Value
S_TMC_SVC (<i>TMC-DP: Service ID</i>)	CR_CORE
S_TMC_PRV (<i>TMC-DP: Provider ID</i>)	UI_*

Field	Value
S_TMC_FLDT (<i>TMC-DP: Field Type</i>)	S (standard, user names in data providers are hidden) or U (user names are shown)
S_TMC_TMST (<i>TMC-TM: System Type</i>)	*
S_TMC_TMSS (<i>TMC-TM: System Subtype</i>)	*
S_TMC_TMAG (<i>TMC-TM: Authorization Group</i>)	
ACTVT (<i>Activity</i>)	01 (Create), 02 (Change), 03 (Display), 06 (Delete), 16 (Execute)

The following values provide display authorizations for the remaining features and tools of the technical monitoring cockpit:

Values for Authorization Object S_TMC_DP: Display

Field	Value
S_TMC_SVC (<i>TMC-DP: Service ID</i>)	*
S_TMC_PRV (<i>TMC-DP: Provider ID</i>)	*
S_TMC_FLDT (<i>TMC-DP: Field Type</i>)	S
S_TMC_TMST (<i>TMC-TM: System Type</i>)	*
S_TMC_TMSS (<i>TMC-TM: System Subtype</i>)	*
S_TMC_TMAG (<i>TMC-TM: Authorization Group</i>)	
ACTVT (<i>Activity</i>)	03 (Display)

Values for Authorization Object S_TMC_TM

Field	Value
S_TMC_TMST (<i>TMC-TM: System Type</i>)	*
S_TMC_TMSS (<i>TMC-TM: System Subtype</i>)	*
S_TMC_TMAG (<i>TMC-TM: Authorization Group</i>)	
ACTVT (<i>Activity</i>)	03 (Display)

User for System Landscape Management

Do you want to create a user who adds and removes systems in the system landscape monitored by the technical monitoring cockpit? Use these values:

Values for Authorization Object S_TMC_DP: System Management

Field	Value
S_TMC_SVC (<i>TMC-DP: Service ID</i>)	CR_CORE
S_TMC_PRV (<i>TMC-DP: Provider ID</i>)	TM_*

Field	Value
S_TMC_FLDT (<i>TMC-DP: Field Type</i>)	S (standard, user names in data providers are hidden)
S_TMC_TMST (<i>TMC-TM: System Type</i>)	*
S_TMC_TMSS (<i>TMC-TM: System Subtype</i>)	*
S_TMC_TMAG (<i>TMC-TM: Authorization Group</i>)	
ACTVT (<i>Activity</i>)	01 (Create), 02 (Change), 03 (Display), 06 (Delete), 16 (Execute) With 71 (Analyze), you can grant additional permissions to see additional information about systems, such as more system properties on the System Details screen and on the Application System Details and Database System Details tiles on the launchpad.

With these values, the user can create, update, change, and delete objects in system management and execute functions such as synchronization.

The following values provide display authorizations for the remaining features and tools of the technical monitoring cockpit:

Values for Authorization Object S_TMC_DP: Display

Field	Value
S_TMC_SVC (<i>TMC-DP: Service ID</i>)	*
S_TMC_PRV (<i>TMC-DP: Provider ID</i>)	*
S_TMC_FLDT (<i>TMC-DP: Field Type</i>)	S (standard, user names in data providers are hidden)
S_TMC_TMST (<i>TMC-TM: System Type</i>)	*
S_TMC_TMSS (<i>TMC-TM: System Subtype</i>)	*
S_TMC_TMAG (<i>TMC-TM: Authorization Group</i>)	
ACTVT (<i>Activity</i>)	03 (Display)

Values for Authorization Object S_TMC_TM

Field	Value
S_TMC_TMST (<i>TMC-TM: System Type</i>)	*
S_TMC_TMSS (<i>TMC-TM: System Subtype</i>)	*
S_TMC_TMAG (<i>TMC-TM: Authorization Group</i>)	
ACTVT (<i>Activity</i>)	03 (Display)

User for System Changes

With the following values, a user can perform changes to a monitored system, such as updating parameters (for example, UPDATE DB CFG).

Values for Authorization Object S_TMC_TM

Field	Value
S_TMC_TMST (<i>TMC-TM: System Type</i>)	*
S_TMC_TMSS (<i>TMC-TM: System Subtype</i>)	*
S_TMC_TMAG (<i>TMC-TM: Authorization Group</i>)	
ACTVT (<i>Activity</i>)	01 (Create), 03 (Display), 23 (Maintain)

Values for Authorization Object S_TMC_DP: System Changes

Field	Value
S_TMC_SVC (<i>TMC-DP: Service ID</i>)	A* , D* , 0*
S_TMC_PRV (<i>TMC-DP: Provider ID</i>)	*
S_TMC_FLDT (<i>TMC-DP: Field Type</i>)	S (standard, user names in data providers are hidden)
S_TMC_TMST (<i>TMC-TM: System Type</i>)	*
S_TMC_TMSS (<i>TMC-TM: System Subtype</i>)	*
S_TMC_TMAG (<i>TMC-TM: Authorization Group</i>)	
ACTVT (<i>Activity</i>)	01 (Create), 02 (Change), 03 (Display), 06 (Delete), 16 (Execute)

The following values provide display authorizations for the remaining features and tools of the technical monitoring cockpit:

Values for Authorization Object S_TMC_DP: Display

Field	Value
S_TMC_SVC (<i>TMC-DP: Service ID</i>)	*
S_TMC_PRV (<i>TMC-DP: Provider ID</i>)	*
S_TMC_FLDT (<i>TMC-DP: Field Type</i>)	S (standard, user names in data providers are hidden)
S_TMC_TMST (<i>TMC-TM: System Type</i>)	*
S_TMC_TMSS (<i>TMC-TM: System Subtype</i>)	*
S_TMC_TMAG (<i>TMC-TM: Authorization Group</i>)	
ACTVT (<i>Activity</i>)	03 (Display)

4 Setting Up the Technical Monitoring Cockpit

The setup of the technical monitoring cockpit comprises the following basic steps: Check that you have the right minimum releases, set up SAPUI5, and add systems to your monitoring landscape.

Related Information

[Availability of the Technical Monitoring Cockpit \[page 18\]](#)

[Settings for SAPUI5 \[page 49\]](#)

[Adding a Database or Host to the Monitoring Landscape \[page 53\]](#)

4.1 Settings for SAPUI5

To run the technical monitoring cockpit with an SAPUI5 user interface, you must add settings for Web services and ports in your system.

ⓘ Note

If you run the technical monitoring cockpit with SAP GUI, no settings are required.

4.1.1 Browser Support

Check which browser versions are supported for the technical monitoring cockpit.

The Web-based user interface of the technical monitoring cockpit is based on OpenUI5. An overview of the supported browsers for OpenUI5 is available in the openUI5 documentation on SAP Help Portal: [Browser and Platform Support](#).

4.1.2 Setting Up Web Services for SAPUI5

To be able to run the technical monitoring cockpit with SAPUI5, you must add the technical monitoring cockpit to the network settings.

Procedure

1. Start transaction SICF_INST.
2. In the field *Technical Name*, enter **TECHNICAL MONITORING COCKPIT**.
3. Choose *Execute*.

4.1.3 Setting Up Ports for SAPUI5

You need SAP profile parameters to configure HTTP(S) ports, timeouts, and single sign-on tickets for running the technical monitoring cockpit with SAPUI5.

Procedure

1. Call transaction RZ10.
2. Ensure that the following profile parameters are set:
 - login/create_sso2_ticket = 2
 - login/accept_sso2_ticket = 1
 - icm/server_port_0 = PROT=HTTP,PORT=80\$\$,PROCTIMEOUT=600,TIMEOUT=600
 - icm/server_port_1 = PROT=HTTPS,PORT=443\$\$,PROCTIMEOUT=600,TIMEOUT=600

4.1.4 Checking the SAPUI5 Setup

Automatic checks help you to verify that your settings for running the technical monitoring cockpit with SAPUI5 are correct.

Procedure

1. Start the technical monitoring cockpit.
2. In the navigation pane, choose ► *Setup* ► *Web Interface (UI5)* ►.

On the following screen, you can see whether the required settings for SAPUI5 are correct.

3. If the settings for the Web services are on status *Red*, correct the settings following the instructions under *Setting Up Web Services for SAPUI5*.
4. If the settings for the ports are on status *Red*, correct them following the instructions under *Setting Up Ports for SAPUI5*.
5. If the setting for URL_ALTERNATE is on status *Yellow*, you can still run the technical monitoring cockpit.

Note

In this case, an HTTPS connection is not available, only HTTP.

If you need HTTPS, correct the settings following the instructions under *Setting Up Ports for SAPUI5*.

Related Information

[Setting Up Web Services for SAPUI5 \[page 50\]](#)

[Setting Up Ports for SAPUI5 \[page 50\]](#)

5 Managing Your System Landscape

5.1 Setting Up the System Landscape for Monitoring

Before you start using the technical monitoring cockpit, you need to set up the systems that you want to monitor in the technical monitoring cockpit.

With the technical monitoring cockpit, you can monitor the local application server ABAP as well as local and remote databases and hosts. The setup includes adding your system to the monitoring landscape, checking prerequisites, and configuration.

5.1.1 Prerequisites for the Host

You need to make sure that the monitored system is set up in a way that the SAP Host Agent can collect the data. To meet the prerequisites for the host, see the following information:

Windows

SAP Note [1409604](#)  - Virtualization on Windows: Enhanced Monitoring

AIX

SAP Note [1379340](#)  - Extended set of virtualization metrics for AIX Power Systems

Linux

- SAP Notes:

SAP Note Number	Title
1606643 	Linux: VMware vSphere host monitoring interface
962334 	Linux: SAP on SUSE Xen virtual machine

SAP Note Number	Title
1739625 	Linux: SAP on Red Hat Xen
1532458 	Monitoring of SAP on Citrix XenServer
1807851 	Linux: SAP on Huawei FusionSphere virtualization
1400911 	Linux: SAP on Red Hat KVM - Kernel-based Virtual Machine
1522993 	Linux: SAP on SUSE KVM - Kernel-based Virtual Machine
2097317 	Linux on Power: SAP on PowerKVM - Kernel-Based Virtual Machine
2025249 	Linux: SAP on Hitachi HVM LPAR virtualization
1656250 	SAP on AWS: Support prerequisites
2191498 	SAP on Linux with Azure: Enhanced Monitoring
1817065 	SAP on Oracle VM - Monitoring SAP for Oracle VM

- SAP Community:
 - <http://scn.sap.com/docs/DOC-57231> 
 - <http://scn.sap.com/docs/DOC-56601> 

5.1.2 Adding a Database or Host to the Monitoring Landscape

After you have set up the technical monitoring cockpit, set up your monitoring system landscape. With the *Managed Systems* function, you can add databases or hosts to the system landscape that you want to monitor.

Procedure

1. On the launchpad, choose the *Managed Systems* tile and then the *Add System* button. Alternatively, choose  *Setup* > *Managed Systems* > *Add System*  in the navigation pane or menu.
2. Select a database and a platform or select a host.
3. Choose *Start System Integration*.
4. Enter the parameters of your new system.

Here you can also decide how users can access the system: for monitoring only, for monitoring and administration, or in a strict read-only mode. For more information, see [Usage Modes of the Managed System \[page 54\]](#).

5. Choose [Step 2](#).

The technical monitoring cockpit now creates a system connection and adds the new system with its parameters to the managed systems of the technical monitoring cockpit.

On the [Connection Summary](#) screen, you can see the authorizations of the connection user and the synchronization messages that were created during the system integration to the managed systems.

6. Choose [Step 3](#).
7. On the [System Parameters](#) screen, you can add a description for your system.
8. Choose [Step 4](#).

The [System Details](#) screen displays a summary of all system information.

9. Choose [Finish](#) to complete the system integration to the managed system landscape.

5.1.3 Usage Modes of the Managed System

Your users have their own authorizations, but you want to control authorizations of specific managed systems in the technical monitoring cockpit for all users? Use the usage modes in the system setup.

When you add a new system to the list of managed systems in the technical monitoring cockpit, the technical monitoring cockpit prompts you to select one of the following:

Usage Modes for Managed Systems

Usage Mode	Description
Monitoring Only	The user can only monitor the managed system. Monitoring tasks can also comprise small changes in the monitored system purely related to monitoring tasks, such as starting RUNSTATS or EXPLAIN . Business data cannot be accessed or changed.
Monitoring and Administration	The user can monitor the managed system and perform administration tasks. Administration tasks can also mean that the administrator can change system objects or configuration in the managed system. Note that this is restricted to administration tasks. Business data cannot be accessed or changed.
Strict Read-Only	The access to the managed system is limited to a strict read-only mode. Note that the user can only display monitoring information. Business data cannot be accessed or displayed.

Depending on the selected usage mode, users of the technical monitoring cockpit can access a managed system for monitoring only, for monitoring and administration, or for read-only use. The usage mode might further restrict the authorizations of the user in the managed system. If a user of the technical monitoring cockpit has full administration authorizations, but the system is set to a [Strict Read-Only](#) usage mode, the user can only access the managed system in read-only mode.

5.1.4 Configuration of the Data Collection for the Application Server ABAP

To make sure that the monitoring screens and data providers for the Application Server ABAP in the technical monitoring cockpit contain data, you need to perform basic configuration steps.

General Settings

Some configuration steps are relevant for all data providers. For more information, see [Configuring the Collection of Application Server ABAP Statistics for the Technical Monitoring Cockpit \[page 55\]](#).

Screen- and Data Provider-Specific Settings

Some data providers require additional specific configuration steps. For more information, see:

- [Configuring the Collection of Application Statistics \[page 58\]](#)
- [Configuring the Collection of Database Procedure Call Statistics \[page 60\]](#)
- [Configuring the Collection of Table Access Statistics \[page 61\]](#)
- [Configuring the Collection of Transaction Detail Statistics \[page 62\]](#)
- [Configuring the Collection of Distributed Statistics Records \(DSRs\) \[page 63\]](#)
- [Configuring the Number of RFCs Logged for Each Transaction Step \[page 64\]](#)
- [Scheduling Jobs for Memory and Work Process Utilization \[page 65\]](#)

5.1.4.1 Configuring the Collection of Application Server ABAP Statistics for the Technical Monitoring Cockpit

To make sure that the Application Server ABAP data providers in the technical monitoring cockpit contain data, you need to perform basic configuration steps in transactions ST03 and ST03T.

Context

The Application Server ABAP data providers in the technical monitoring cockpit rely on the profile tables in the new transaction ST03T (see SAP Note [2274315](#)). In turn, the profile tables used in ST03T rely on the collection of statistics in the Workload Monitor (transaction ST03). To enable the collection of statistics in transactions ST03 and ST03T, you need to perform the basic configuration steps below.

Procedure

1. Enable the collection of statistics in transaction ST03. To do so, schedule the background job `SAP_COLLECTOR_FOR_PERFMONITOR` as follows:
 - a. Log on to the SAP system in client 000 as user DDIC or as a user to which the role `SAP_BC_COLLECTOR_FOR_PERFMON` is assigned.

Note

The latest version of this role is attached to SAP Note [2321829](#).

- b. Call transaction SM36 and choose *Job Wizard* in the application toolbar.
- c. In the popup window *Create a job*, choose *Continue*.
- d. In the popup window *General job information*, enter `SAP_COLLECTOR_FOR_PERFMONITOR` in the field *Job Name* and leave the default values in the other fields. Choose *Continue*.
- e. In the popup window *Job definition: Job step*, choose *ABAP program step*. Choose *Continue*.
- f. In the popup window *ABAP program step*, enter `RSCOLL00` in the field *ABAP Program Name* and leave the default values in the other fields. Choose *Continue*.
- g. On the screen *Multi-step option*, do not select *Add additional steps (optional)* and choose *Continue*.
- h. On the screen *Spool List Recipient*, choose whether you want to send the output from the step as an e-mail (*Spool List Recipient*) or if you want to send a notification after the job has terminated (*E-Mail Notification*). Choose *Continue*.
- i. In the popup window *Job definition: Start conditions*, choose *Immediately*. Choose *Continue*.
- j. In the popup window *Def. of start immediately*, select the checkbox *Period* under *Periodic Jobs* and choose *Continue*.
- k. In the popup window *Period definition*, choose *Hourly*. Choose *Continue*.
- l. In the popup window *Set job*, review the information and choose *Complete*.

For more information about the report `RSCOLL00`, see SAP Note [966309](#).

2. Enable the collection of statistics in the ST03T profile tables as follows:
 - a. Call transaction ST03.
 - b. Make sure that you are in *Expert mode*.
 - c. In the navigation pane, choose **Collector and Performance DB** > *Workload Collector* > *Configuration of Data Storage*.
 - d. On the screen *Configuration of Data Storage*, select all the profiles in the column *Profile Tables*. (Note that this column was named *Data Table* in an earlier version of the tool.)
 - e. Save your entries.

For more information, see SAP Note [2274315](#).

3. (Optional) If you want to display columns with sensitive information (user ID), make sure that you have authorization for authorization object `S_TOOLS_EX`.

To control the access to this sensitive information, an authorization check is carried out against `S_TOOLS_EX`. If an unauthorized user attempts to display a user ID column, the following error message is displayed:

Error while checking authorization; missing privileges for `S_TOOLS_EX`

Note

This step is only relevant for the following data providers that contain a *USER ID* column:

- Application Hit List by DB Access
- Application Hit List by Response Time
- Application Statistics
- Database Procedures
- Hit List by Database Access
- Hit List by Response Time
- Memory Consumption
- RFC Client Calls
- RFC Client Destinations
- RFC Server Calls
- RFC Server Destinations
- Transaction Detail
- User and Transaction
- User-Related Data
- Web Client Calls
- Web Client Destinations
- Web Server Calls
- Web Server Destinations

Results

- Data is displayed in transaction **ST03** after a maximum of a few hours.
- Data is displayed in transaction **ST03T** after a maximum of one day.
- As soon as transaction **ST03T** contains data, the Application Server ABAP data providers in the technical monitoring cockpit also contain data.

Related Information

[Application Hit List by DB Access \[page 380\]](#)

[Application Hit List by Response Time \[page 381\]](#)

[Application Statistics \[page 382\]](#)

[Database Connections \[page 383\]](#)

[Database Procedures \[page 384\]](#)

[Front-End Communication \[page 385\]](#)

[Hit List by Database Access \[page 386\]](#)

[Hit List by Response Time \[page 387\]](#)

[Load from External Systems \[page 388\]](#)

[Memory Consumption \[page 389\]](#)
[Response Time Distribution \[page 395\]](#)
[RFC Client Calls \[page 391\]](#)
[RFC Client Destinations \[page 392\]](#)
[RFC Server Calls \[page 393\]](#)
[RFC Server Destinations \[page 394\]](#)
[Spool Activity \[page 396\]](#)
[Table Access \[page 397\]](#)
[Time Profile \[page 398\]](#)
[Transaction Detail \[page 399\]](#)
[TREX Statistics \[page 396\]](#)
[User and Transaction Profile \[page 400\]](#)
[User-Related Data \[page 401\]](#)
[VMC Statistics \[page 402\]](#)
[Web Client Calls \[page 403\]](#)
[Web Client Destinations \[page 403\]](#)
[Web Server Calls \[page 404\]](#)
[Web Server Destinations \[page 405\]](#)
[Workload Overview \[page 405\]](#)

5.1.4.2 Configuring the Collection of Application Statistics

To enable the collection of application statistics, you need to set the required profile parameters and to activate the individual components of the application statistics.

Prerequisites

The SAP transaction or application has already implemented function calls to measure application statistics.

Procedure

1. Set the required profile parameters as follows:
 - a. In your SAP system, call transaction **ST03**.
 - b. Make sure that you are in *Expert mode*.
 - c. In the navigation pane, choose  *Collector and Performance DB*  *Statistics Records & File*  *Relevant Profile Parameters*  and double-click the relevant instance.
 - d. On the screen *Workload Analysis: Relevant Profile Parameters*, make sure that the following profile parameters are set:

`stat/as_level = 1`

stat/as_collect = 2

- e. If these profile parameters are not set, call transaction RZ11 and follow the standard procedure to set them accordingly.
2. Activate the individual components of the application statistics as follows:
 - a. In your SAP system, call transaction ST03.
 - b. In the navigation pane, choose ► [Collector and Performance DB](#) ► [Statistics Records & File](#) ► [Online Parameters](#) ► [Application Statistics](#) .
 - c. On the screen [Configure Application Statistics](#), select the application component that you want to activate and choose [Activate](#).

Note

You need to do this for each required application component. Make sure that the corresponding checkboxes in the [Activated](#) column are selected.

- d. Save your entries.
3. Make sure that the collector for the application statistics has been scheduled:
 - a. In your SAP System, call transaction ST03.
 - b. In the navigation pane, choose ► [Collector and Performance DB](#) ► [Workload Collector](#) ► [Instance Collector](#) ► [Control](#) .
 - c. Make sure that the checkbox for the [Application Profile](#) (AA) is selected.
 - d. Choose [Back](#).
 - e. In the navigation pane, choose ► [Collector and Performance DB](#) ► [Workload Collector](#) ► [TOTAL Collector](#) ► [Control](#) .
 - f. Make sure that the checkboxes for the [Application Profile](#) (AA) are selected ([Week.Aggr.](#), [Month.Aggr.](#), and [Gen TOTAL](#)).

Results

Application statistics are displayed in the technical monitoring cockpit after a maximum of one day.

Related Information

[Application Hit List by DB Access \[page 380\]](#)

[Application Hit List by Response Time \[page 381\]](#)

[Application Statistics \[page 382\]](#)

5.1.4.3 Configuring the Collection of Database Procedure Call Statistics

To enable the collection of database procedure call statistics, you need to enable the profile parameter `stat/dbprocrec` that specifies the number of database procedure subrecords.

Context

Some applications rely on the use of database procedures, for example, in Advanced Planning and Optimization (APO) systems. The collection of database procedure call statistics should only be enabled on such systems.

Procedure

1. In your SAP system, call transaction **ST03**.
2. Make sure that you are in *Expert mode*.
3. In the navigation pane, choose **Collector and Performance DB** > **Statistics Records & File** > **Online Parameters** > **Dialog Step Statistics**.
4. On the screen *Change Runtime Parameters of Statistics*, specify the parameter `stat/dbprocrec`.

❖ Example

For example, specify **5**. This value provides some insight into the database procedure calls without significantly affecting system performance. If you need more information about database procedures, you can increase this value and continue the analysis.

5. Choose *Activate Values*.

Results

If database procedure calls are used in the system, the statistics are displayed in the technical monitoring cockpit after a maximum of one day.

Related Information

[Database Procedures \[page 384\]](#)

5.1.4.4 Configuring the Collection of Table Access Statistics

To enable the collection of table access statistics, you need to set the required parameters.

Context

Since the collection of table access statistics can impact system performance, we recommend you to enable table statistics only to analyze specific problems and disable them as soon as they are no longer needed.

Procedure

1. In your SAP system, call transaction **ST03**.
2. Make sure that you are in *Expert mode*.
3. In the navigation pane, choose **Collector and Performance DB** > *Statistics Records & File* > *Online Parameters* > *Dialog Step Statistics*.
4. On the screen *Change Runtime Parameters of Statistics*, specify the following parameters for each SAP instance:
 - **stat/tabrec**
This parameter specifies the number N of tables that consumed the most resources during each transaction step. To avoid performance degradation, specify a value lower than 10.
 - **stat/tcode1 to stat/tcode5**
If you want to collect table access statistics for certain transactions only, you can specify a maximum of five transactions using these parameters.
5. Choose *Activate Values*.

Results

Table access statistics are displayed in the technical monitoring cockpit after a maximum of one day.

Next Steps

Turn off the collection of table access statistics after your analysis is complete. To do so, simply repeat the steps 1-5 in the procedure above. In step 4, clear the profile parameter values (stat/tabrec and stat/tcode1 to stat/tcode 5).

Related Information

[Table Access \[page 397\]](#)

5.1.4.5 Configuring the Collection of Transaction Detail Statistics

To collect transaction detail statistics, you need to activate the individual transaction codes and make sure that the collector for the transaction details has been scheduled.

Procedure

1. Set the required profile parameters as follows:
 - a. In your SAP system, call transaction **ST03**.
 - b. Make sure that you are in *Expert mode*.
 - c. In the navigation pane, choose **Collector and Performance DB** > **Workload Collector** > **Parameters**.
 - d. On the screen *Selection of Data to Cumulate*, under *Create Transaction Detail Profiles For*, enter the desired transaction codes in the *TCode* column.

These are the transaction codes that are required to start the application that has been identified as problematic.

Note

The transaction detail statistics can be recorded for up to five transactions.

- e. Save your entries by choosing *Save Values*.
2. Make sure that the collector for the transaction details has been scheduled:
 - a. In your SAP System, call transaction **ST03**.
 - b. In the navigation pane, choose **Collector and Performance DB** > **Workload Collector** > **Instance Collector** > **Control**.
 - c. Make sure that the checkbox for the *Function Code Profile* (WG) is selected.
 - d. Choose *Back*.
 - e. In the navigation pane, choose **Collector and Performance DB** > **Workload Collector** > **TOTAL Collector** > **Control**.
 - f. Make sure that the checkboxes for the *Function Code Profile* (WG) are selected (*Week.Aggr.*, *Month.Aggr.*, and *Gen TOTAL*).

Results

The detailed function codes of the enabled transactions are displayed in the [Transaction Detail](#) data provider in the technical monitoring cockpit after a maximum of one day.

Related Information

[Transaction Detail \[page 399\]](#)

5.1.4.6 Configuring the Collection of Distributed Statistics Records (DSRs)

To make sure that data is displayed in the TREX data provider in the technical monitoring cockpit, you need to perform some configuration steps to collect the distributed statistics records (DSRs).

Context

In simple terms, DSRs are an extension of the concept of statistical records to non-ABAP components (J2EE Engine, BC, and ITS).

Procedure

To configure the collection of distributed statistics records (DSRs) for TREX, proceed as described in the [documentation for the Global Workload Monitor](#) on the SAP Help Portal.

The following subtopics are the most useful:

- [Configuring the Display of DSRs of the Local System](#)
- [Configuring the Display of DSRs of Remote Systems](#)

Related Information

[TREX Statistics \[page 396\]](#)

5.1.4.7 Configuring the Number of RFCs Logged for Each Transaction Step

The default value for the number of RFCs logged for each transaction step is 5. You can change this value by configuring the profile parameter `stat/rfcrec`.

Context

This parameter applies to all RFC subrecord types (RFC Client, RFC Server, RFC Client Destination, and RFC Server Destination).

⚠ Caution

We do not recommend that you increase this value since a value higher than 5 causes a significant performance overhead in the collector.

Procedure

1. In your SAP system, call transaction `ST03`.
2. Make sure that you are in *Expert mode*.
3. In the navigation pane, choose **Collector and Performance DB** > *Statistics Records & File* > *Online Parameters* > *Dialog Step Statistics*.
4. On the screen *Change Runtime Parameters of Statistics*, specify the parameter `stat/rfcrec` for each SAP instance.

This parameters specifies the number of RFCs logged for each transaction step.

5. Choose *Activate Values*.

Related Information

[RFC Client Calls \[page 391\]](#)

[RFC Client Destinations \[page 392\]](#)

[RFC Server Calls \[page 393\]](#)

[RFC Server Destinations \[page 394\]](#)

5.1.4.8 Scheduling Jobs for Memory and Work Process Utilization

For memory and work process utilization, you must schedule jobs to get data.

For memory utilization, schedule job `SMON_MEMORY_MGMT`.

For work process utilization, schedule job `SMON_MEMORY_MGMT`.

For more information about jobs in the technical job repository, see SAP Note [2581518](#).

5.1.5 Configuring the Data Collection Statistics for Host Data Providers

To make sure that host data providers in the technical monitoring cockpit contain data, you need to enable data collection statistics. To do this, schedule the background job `SAP_COLLECTOR_FOR_PERFMONITOR`.

Context

For more information about the report `RSCOLL00`, see SAP Note [966309](#).

Procedure

1. Log on to your SAP system in client 000 as user DDIC or as a user to which the role `SAP_BC_COLLECTOR_FOR_PERFMON` is assigned.

Note

The latest version of this role is attached to SAP Note [2321829](#).

2. Call the job monitor (transaction `SM36`).
3. Choose *Standard Jobs* in the application toolbar.

As a result, SAP standard jobs are automatically scheduled, including background job `SAP_COLLECTOR_FOR_PERFMONITOR`.

4. Alternatively, if you want to schedule the job individually, choose *Job Wizard* in the application toolbar and create a job with name `SAP_COLLECTOR_FOR_PERFMONITOR` and with the following data:

Popup Window	What to Do
<i>Job definition: Job step</i>	Choose <i>ABAP program step</i> .
<i>ABAP program step</i>	In the field <i>ABAP Program Name</i> , enter RSCOLL00 . Leave the default values in the remaining fields.

Popup Window	What to Do
Multi-step option	Do not select Add additional steps (optional) .
Spool List Recipient	Choose whether you want to send the output from the step as an e-mail (E-Mail Notification) or if you want to send a notification after the job has terminated (Spool List Recipient).
Job definition: Start conditions	Choose Immediately .
Def. of start immediately	Under Periodic Jobs , select the checkbox Period .
Period definition	Choose Hourly .

5.2 Displaying and Updating Your Managed Systems

Get an overview of your system landscape using the [System Details](#) and [Managed Systems](#) screens.

Your system landscape is up and running? To display more information about your managed systems and update your system landscape, choose [System Details](#) and [Overview](#) in the menu.

[Managed Systems \[page 66\]](#)

Get an overview of your landscape of systems managed by the technical monitoring cockpit and navigate to details.

[System Details \[page 67\]](#)

Display the properties of an individual system, including its system topology.

5.2.1 Managed Systems

Get an overview of your landscape of systems managed by the technical monitoring cockpit and navigate to details.

Use the [Managed Systems](#) screen to do the following:

- Add a new system to the system landscape
- View the system details and topology of a selected managed system

You can also change the system properties or remove a system from the landscape of managed systems.

Related Information

[Adding a Database or Host to the Monitoring Landscape \[page 53\]](#)

[System Details \[page 67\]](#)

[System Topology \[page 68\]](#)


5.2.2 System Details

Display the properties of an individual system, including its system topology.

On the *System Details* screen, you can do the following:

- Switch between different systems using the selection area on the left.
- Display the topology of the selected system.
To see more details, choose one of the nodes of the topology tree.
- Change your system landscape, synchronize the display of the system topology, and remove systems from the monitoring landscape.

On the launchpad, choose the *Managed Systems* tile and go to the *System Details* screen by choosing the detail view (little arrow at the end of each row entry) for the relevant system. Alternatively, you can choose ► *Menu* ► *Setup* ► *Managed Systems* ► *System Details* ►:



Related Information

System Topology [page 68]

[Adding a Database or Host to the Monitoring Landscape \[page 53\]](#)

5.2.2.1 System Topology

Get an overview of the topology of an individual system: from database and application server to detailed information about cluster members or scale-out options such as distributed SAP HANA systems or IBM Db2 pureScale.

Automatic Reading of Topology Information


When you add a new system to your monitoring landscape in the technical monitoring cockpit, the system topology details are automatically read and added to the system landscape information in the technical monitoring cockpit. This includes information about replication scenarios, clusters, available instances, and so on. There's no need to enter this information manually. The system topology also includes information about all system connections.

Displaying the System Topology


To see the topology of a system, go to the [System Details](#) screen. For more information, see [System Details \[page 67\]](#).

Examples


The following examples illustrate what the display of system topologies might look like:



System Topology Example: Application System



System Topology Example: Replication



System Topology Example: Scale-Out with Multiple Cluster Nodes

5.2.2.2 Properties, Environments, Validity

On the [Properties](#) tab page, details of the system or connection are shown. The [Environments](#) and [Validity](#) tab pages contain technical information that is mainly intended for SAP support.

Properties

On this tab page, you can find more detailed information about your system or your connection, such as the database version, the SAP application server release, or whether your system is partitioned or running with replication. This kind of information is read from the system after you have included it into your system landscape to be monitored by the technical monitoring cockpit. The properties also include information that you entered when you added a system to the system landscape, such as the usage mode (monitoring, read only, or monitoring and administration). Technical properties such as the system hash or the technical ID might be useful information if you need to get support from SAP.

Environments and Validity

You might need the [Environments](#) and [Validity](#) tab pages if you are in contact with SAP support.

On the [Environments](#) tab page, environment IDs are shown, which are a combination of valid system types and database or operating system releases or versions for the selected system. The technical monitoring cockpit evaluates these environment IDs to ensure that system- and release-specific features are only shown if they apply to the selected system. SAP support might need this information to find out why particular features are not shown for your system.

On the [Validity](#) tab page, validity IDs are shown, which are check results whether your system has certain properties, such as database partitioning or workload type OLTP, for example. Like the environment IDs, validity IDs are evaluated by the technical monitoring cockpit to ensure that certain features are only shown if they apply to the selected system. For example, information about partitions are only shown if the relevant validity ID indicates that the database is partitioned. Again, SAP support might need this information to find out why particular features are not shown for your system.

Environments and validity are only shown for systems, for connections these tab pages are empty.

5.2.2.3 System Details: Quick Guide

5.2.2.3.1 Change System Details

Use the [Change](#) button to change the details of the selected system that are stored in the technical monitoring cockpit, such as host, port, user, password, display name, or system type (production system, test system).

5.2.2.3.2 Display Logs

Use the [Display Logs](#) button to display the logs of the system synchronization.

Related Information

[Synchronize \[page 72\]](#)

5.2.2.3.3 Synchronize

Use the [Synchronize](#) button to read the system status and topology again using the system connection.

This might be relevant when you know that the system topology has changed: For example, a new node has been added to your system cluster.

5.2.2.3.4 Delete System Connection

Use the [Delete](#) button to remove a system from the system landscape managed by the technical monitoring cockpit.

6 System Monitoring

In the technical monitoring cockpit, different types of content are available for system monitoring, depending on how much flexibility you want.

Using Preconfigured Content

The following content is available and ready for use:

Monitoring Screens in the Menu

In the menu of the technical monitoring cockpit, under [Information](#), you can find monitoring screens that cover the most important information needs of a system and database administrator, such as information about CPU utilization, memory utilization, storage, workload, and so on. Examples of such monitoring screens are [CPU Utilization by Service](#), [Top 10 Tables by Size](#), [Host Memory Utilization](#), [Disk Usage](#), [Average Transaction Response Time](#), or [Top 10 SQL Statements by Execution Time](#). The available monitoring screens depend on your chosen database, operating system, and application server.

Application System Details and Database System Details

In addition to the [Information](#) menu item, where you can find all available monitoring screens, you can also use the tiles [Application System Details](#) and [Database System Details](#) on the launchpad of the technical monitoring cockpit. The thumbnail preview on these tiles provides you with the most important key figures of your Application Server ABAP and of your database system. If you choose one of these tiles, a new dashboard opens. The dashboard shows the most important system information as tiles. When you choose a tile, a corresponding monitoring screen opens. This monitoring screen is also available under Information in the menu.

The advantage of using the [Application System Details](#) and [Database System Details](#) tiles on the launchpad is that you get a predefined selection of the most important monitoring screens for your application server and database system.

Analysis Screens in the Menu

Whereas the monitoring screens under [Information](#) purely serve to provide you with the most important key figures of your system, the screens under [Analysis](#) provide you with more detailed information so that you can run a root cause analysis of system issues. For example, for the SAP HANA database, you can use the analysis screen [SQL Statement Analysis](#) to find long-running SQL statements, display and change the full SQL statement, and observe the changes in its access plan.

Creating your Own Monitoring Screens


Instead of using what's already available, you can also make use of a collection of many views that contain a wide range of key figures for databases, hosts, operating systems, and the application server ABAP. These

views are referred to as data providers. In the data explorer of the technical monitoring cockpit, you can fine-tune the settings of each data provider (selected key figures, filtering, sorting, time periods) in such a way that it fits your needs best. You can save this as a favorite, which appears then in the menu together with the monitoring screens provided by SAP.

6.1 Coverage of Your Entire Technology Stack

The technical monitoring cockpit covers the SAP technology stack, that is, databases, operating systems, and SAP application server.

So far, you might have had to use specific monitors for each component of the technology stack, and they all might have had a different way of querying, storing, and displaying data. With the technical monitoring cockpit, you no longer need to work with multiple monitors. You can identify, locate, and analyze problems throughout the SAP technology stack, that is, the local Application Server ABAP, databases, and hosts (depending on your environment). All technical monitoring data can be queried and analyzed the same way due to a consistent user interface.



Example: CPU Utilization

For example, in an Application Server ABAP running on an SAP HANA database, you might notice issues with CPU utilization. With the technical monitoring cockpit, you have now monitoring screens available for checking the CPU utilization of the Application Server ABAP (Top 10 transactions by CPU utilization), CPU utilization by service and by workload on the SAP HANA database, and in on-premise systems, you can also check the host and the virtual CPU utilization. This is all available in the menu of the technical monitoring cockpit; there's no need to switch between different tools.

Related Information

[Technical Monitoring Cockpit \(On-Premise, 7.4-7.56\) \[page 12\]](#)

[Architecture: Engine and Content \[page 16\]](#)
[Learning the Basics \[page 15\]](#)
[Getting Started with the User Interface \[page 22\]](#)
[System Topology \[page 68\]](#)

6.2 Monitoring Screens

Monitoring screens are available under [Information](#) in the menu of the technical monitoring cockpit. You can use these screens to monitor your systems and display the characteristics and metrics for the most typical monitoring use cases.

Monitoring screens provide you with a preconfigured view of the most important metrics for a particular monitoring situation and for a particular database, for your host (including your operating system), or for the Application Server ABAP. Preconfigured monitoring screens are an easy start for administrators who want to monitor their systems and want to see the most important KPIs without having to select from the large range of available metrics and characteristics available in the technical monitoring cockpit.

If you are already familiar with the DBA Cockpit, which is an existing monitoring tool in all SAP NetWeaver-based systems, you will find many screens that you already know from the DBA Cockpit.

Where Do You Find Preconfigured Monitoring Screens?

Preconfigured monitoring screens can originate from the following sources:

- Predefined by SAP
You can find these monitoring screens in the menu under [Information](#).
- Configured by an expert user of the technical monitoring cockpit using the data explorer
You can create and save your screens in a folder of your choice, for example, under [Favorites](#).

Preconfigured monitoring screens can also appear as tiles on the launchpad.

How Are Preconfigured Monitoring Screens and Data Providers Related?

Technically speaking, a preconfigured monitoring screen is a view on a data provider that was saved with selected characteristics, metrics, filters, and sorting.

Related Information

[What Are Data Providers? \[page 80\]](#)
[What Is the Data Explorer? \[page 79\]](#)

6.3 Selecting a System for Monitoring

Learn how you select a local or remote system for monitoring in the technical monitoring cockpit.

Prerequisites

You have added the database to your monitored systems (see [Adding a Database or Host to the Monitoring Landscape \[page 53\]](#)).

Context

You can select local and remote hosts and databases as well as the local application server ABAP for monitoring.

Note

As of SAP_BASIS 7.53, SAP systems don't run on non-HANA databases anymore, but you can still select non-HANA databases for remote monitoring.

Procedure

1. To select a system, proceed as follows:
 - SAPUI5: In the menu bar to the right, choose  [More](#)  [Select a System](#) .
 - SAP GUI: Choose [Select a System](#).
2. Select a system type and name.

Depending on which system you have selected, the menu and the available monitoring screens will change.
3. You can now use the technical monitoring cockpit for monitoring the system that you've selected (see also [Monitoring Screens \[page 75\]](#)).

6.4 Dashboards

A dashboard is a defined arrangement of tiles. Create dashboards if you want to have different tile selections on your screen, depending on the work scenario you are in.

Depending on your work scenario, you can switch between different dashboards. For example, you might want to have a dedicated dashboard for monitoring scale-out solutions for your database, such as IBM Db2 pureScale, or for a distributed SAP HANA system, or for other database-specific features.

You can create your own dashboards and select the tiles that are most relevant for your company from the set of all available tiles. Available tiles include tiles provided by SAP that allow users to navigate to important configuration and administration tasks, such as content management or managed system administration. In addition, all monitoring screens from the navigation pane and all data providers with user-defined selections and filters that were saved as favorites are also available for dashboards. These tiles include thumbnails with a simplified preview of the charts, tables, or lists from the relevant data providers.

Note

The launchpad is one type of dashboard. It's the standard dashboard that appears when you start the technical monitoring cockpit.

Related Information

[Creating a Dashboard \[page 105\]](#)

6.4.1 Database System Details

The [Database System Details](#) dashboard provides you with a set of essential data to monitor your database.

The dashboard includes a set of tiles. When you choose a tile, you navigate to more detailed information on a monitoring screen.


6.4.2 Application System Details

The [Application System Details](#) dashboard provides you with a set of essential data to monitor your Application Server ABAP.

The dashboard includes a set of tiles. When you choose a tile, you navigate to more detailed information on a monitoring screen.

6.5 System Monitoring: When to Use What

Find out how you can use the technical monitoring cockpit for monitoring your systems. You'll see that casual users and experts have different options.



Users and Use Cases for Monitoring

The Casual User's Choice: Easy Monitoring with Preconfigured Monitoring Screens

You are a casual user of the technical monitoring cockpit, or you are only getting started with the technical monitoring cockpit? Then try the preconfigured monitoring screens that are shipped as part of the technical monitoring cockpit. These screens are designed to meet the most typical monitoring needs for databases, Application Server ABAP, and hosts.

You can find these preconfigured monitoring screens in the navigation pane under [Information](#).

Note

In SAP UI5, open the navigation pane by choosing the [Menu](#) button on the left.

The Super User and Expert Mode: Creating Your Own Monitoring Screens

You are already an expert in the technical monitoring cockpit? You checked the existing monitoring screens in the technical monitoring cockpit, but this is not exactly what you need? No problem, you can easily create your own monitoring screens using the data explorer. The data explorer comes with a large collection of data providers with hundreds of metrics. In the data explorer, select a data provider for your monitoring requirement, select and filter metrics, and configure a table or chart view. Then, you can save your own monitoring screen with the metrics that you want to use, without having to write a single line of code.

If you are in a super user role at your company and assist other users in the technical monitoring cockpit, you will like this: All your configured monitoring screens are also available to all other users of the technical monitoring cockpit in your system.

You can find the data explorer as a tile on the launchpad or in the navigation pane under [Tools](#).

Related Information

[Monitoring Screens \[page 75\]](#)

[Data Explorer: Create Your Own Monitoring Screens \[page 79\]](#)

6.6 Data Explorer: Create Your Own Monitoring Screens

You can easily create your own monitoring screens using the data explorer. The data explorer comes with a large collection of data providers with hundreds of metrics that you can adapt to your needs.

6.6.1 What Is the Data Explorer?

Welcome to the data explorer, the expert tool that allows you to visualize monitoring data and to configure your own monitoring and analysis screens.

The data explorer provides you with predefined views that contain monitoring data of databases, operating systems, and application servers. These predefined views are referred to as *data providers*.

In the list in the left frame of the data explorer, you can choose the data provider that you want to use. The main screen displays all metrics that are relevant for your chosen data provider.

Here's an overview of the main actions you can perform with the data explorer:

- 1 Start the data explorer (using the launchpad or search) and choose a data provider.
 
- 2 Select characteristics and metrics you want to see.

- 3 Choose a filter, sorting order, and the number of rows you want to display.

- 4 Set a time period in the time slider.
 
- 5 Choose *Apply*. Otherwise, you won't see any results.

- 6 Toggle between table and chart and go to the detail view.
 
- 7 Edit the chart (pivot, change chart type, colors, ...).
 
- 8 Don't forget to save your selections in your favorites.


1. Start the data explorer from the lauchpad or using the search function and choose a data provider from the list on the left.
2. Select a suitable drilldown (= select characteristics and metrics you want to see).
3. To limit the result set to the most important characteristics and metrics, select filters, sorting orders, and the number of rows you want to display. Optionally, you can also perform a top n analysis (for more information, see [Row Settings \[page 35\]](#)).
4. Set a time period in the time slider.
5. Choose [Apply](#). Otherwise, you won't see any results.
6. Toggle between table and chart view and/or go to the detail view, if needed.
7. Edit the chart according to your preferences (pivot, change chart type, colors, and so on).
8. If you find your selections useful for later reuse, don't forget to save them in your favorites.

Related Information

[Example: Guided Tour on How to Use the Data Explorer \[page 82\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

6.6.2 What Are Data Providers?

Data providers are objects in the technical monitoring cockpit that consist of a selection of characteristics and metrics for databases, hosts, or Application Server ABAP.

Each data provider has its own aggregated view of important characteristics and metrics. You can use the screen of the data provider to drill down to more details by selecting/deselecting more characteristics and metrics as well as filters, sorting orders, and the number of rows you want to be displayed. Each data provider also includes detail views of the data (in table format or as adaptable charts).

What's the Relation Between Data Providers and Monitoring Screens?

Each monitoring screen that you find in the menu of the technical monitoring cockpit is technically based on a data provider. As to look & feel, they are quite the same. Monitoring screens are just a little more specialized for particular monitoring scenarios than data providers.

Where Do I Find Data Providers?

You can find all data providers in the data explorer of the technical monitoring cockpit. After you have called up the data explorer, you can see a list of all available data providers in your system.

When Do I Use Data Providers?

If you are happy with the monitoring screens that you can find in the menu of the technical monitoring cockpit, there's no need to work with data providers.

Use data providers only if you prefer to create your own monitoring screens instead of using what is already preconfigured in the menu.

More Technical Background on Data Providers

Data providers are data models designed by SAP to create abstractions from the native data sources, for example, SQL or Web service interfaces provided by the database and operating system vendors. Data providers combine or aggregate data in a way that the native SQL and Web service interfaces don't provide. Hence they offer optimal data models for the use in the technical monitoring cockpit.

Why Do Some Data Providers Only Have Technical Names?

There are two types of data providers: First, there are data providers that have titles in natural language, such as *Memory Utilization* or *Blocked Transactions*. Second, depending on your system selection, you might also view data providers with technical names as titles such as *M_DISKS* in the data explorer. These are technical data providers.

A technical data provider is a data provider that is based on a platform vendor's data source.

Related Information

[What Are Technical Data Providers? \[page 81\]](#)

[System Monitoring: When to Use What \[page 78\]](#)

[Example: Guided Tour on How to Use the Data Explorer \[page 82\]](#)

6.6.3 What Are Technical Data Providers?

A technical data provider is a data provider that is based on a platform vendor's data source. A data source is any native interface from a platform vendor that can be consumed by the technical monitoring cockpit. This can be, for example, a table function or a catalog view.

Technical data providers offer a native-like access to the data source. Often there is a one-to-one relationship between the fields used in a technical data provider and the fields in the corresponding data source.

When Would You Use Them?

Technical data providers can be the technical basis for other data providers or preconfigured monitoring screens, which usually show data that has been pre-selected, transformed, and filtered. As opposed to “standard” or “modeled” data providers, technical data providers are mainly intended for database or operating system administration experts, SAP development, and SAP support who want to create their own monitoring screens based on technical data providers.

How Do You Find Them?

You can easily identify technical data providers because, as opposed to standard data providers, technical data providers have titles in upper case letters, often just the technical names or abbreviations of table functions or views, such as `M_SQL_CLIENT_NETWORK_IO`. These titles usually match the name of the consumed native interface.

How Do You Get More Information?

In this documentation for the technical monitoring cockpit, we frequently refer to the documentation of the data source of the original platform vendor for additional information. In exceptional cases, especially if there is no one-to-one relationship between a technical data provider and a data source, we also provide more information in the technical monitoring cockpit documentation itself.

6.6.4 Example: Guided Tour on How to Use the Data Explorer

Let's walk you through an example of how to use the data explorer.

Here's an overview of the main steps:

- 1 Start the data explorer (using the launchpad or search) and choose a data provider.
 
- 2 Select characteristics and metrics you want to see.

- 3 Choose a filter, sorting order, and the number of rows you want to display.

- 4 Set a time period in the time slider.
 
- 5 Choose *Apply*. Otherwise, you won't see any results.

- 6 Toggle between table and chart and go to the detail view.
 
- 7 Edit the chart (pivot, change chart type, colors, ...)
 
- 8 Don't forget to save your selections in your favorites.


Let's assume that we want to monitor the CPU utilization in an SAP HANA database. For this purpose, we'll use the data provider *CPU Utilization per Service*.

Prerequisites

In the following procedure, you will make some selections and edit some settings. If you want to save these selections and settings as a favorite, create a package in the ABAP workbench first.

Starting the Data Explorer

There are two ways to start the data explorer:

- Using the search function
Enter the keyword *CPU utilization* and you will get a list of results, including the data provider *CPU Utilization per Service*. When you choose it, the system directly navigates to the data explorer with the chosen data provider preselected.
- From the launchpad
On the launchpad, choose the *Data Explorer* tile. The list on the left contains all data providers that you can use for monitoring and analysis. Choose the data provider *CPU Utilization per Service*.

Note

To restrict the list to data providers for databases, hosts, or application servers, use the *Filter* field or the relevant toggle buttons at the top of the data provider list.

Adapting the Data Display

Selecting Characteristics (Drilldown)

When you call up a data provider, the metrics displayed are aggregated (total, maximum, minimum, average, etc.) as predefined by SAP. In our example, the metrics displayed are summed up, so you see the CPU time in percent and seconds, and the number of CPUs:

CPU Utilization per Service			
Characteristics (0) Metrics (3) Filter (0) Sort (0) Rows (1000) Apply ...			
12 PM Fri 25 12 PM Sat 26 12 PM Jan 27 12 PM Mon 28 12 PM Tue 29 Now			
2019-01-28 11:33 +01:00 1D 2019-01-29 11:33			
1 Row Show Chart Pivot ...			
CPU Time [%]		CPU Time [s]	CPUs [#]
100.0		30,398,543.8	175.7

With a drilldown (= selecting characteristics), you can change the aggregation of the metrics. After a drilldown, the result set typically comprises more than one record, and you can select additional drilldown criteria. If you want to select the key characteristics (= deepest level of drilldown), you don't have to do it manually, but you can use the *Characteristics* button and choose the *Select All Keys* button in the *Characteristics* popup window.

In our example, we choose the [Characteristics](#) button and select a drilldown by consumer and service name:

Characteristics



☒	H1 Consumers
☒	H2 Service Name
☐	Host
☐	Row Rating
☐	Timestamp (Server)
☐	Topology System ID
☐	Topology System Label

Ok


Cancel

Don't forget to choose [Go](#) or press [Enter](#). Only after you have chosen [Go](#), does the technical monitoring cockpit apply the time periods, filters, sort orders, number of rows, and metrics that you have specified. Our example data provider now looks like this:

CPU Utilization per Service					
▼ Characteristics (2) Metrics (3) Filter (0) Sort (0) Rows (1000) Apply ...					
▼ ↑↓					
12 PM Fri 25 12 PM Sat 26 12 PM Jan 27 12 PM Mon 28 12 PM Tue 29					
2019-01-28 11:33 +01:00 < 1D > 2019-01-29 11:33 🕒					
14 Rows Show Chart Pivot ...					
H1 Consumers	H2 Service Name	CPU Time [%]	CPU Time [s]	CPU's [#]	
SAP HANA Ser...	scriptserver	0.0	5,234.3	0.0	>
SAP HANA Ser...	webdispatcher	0.0	1,466.0	0.0	>
SAP HANA Ser...	xsengine	0.0	2,982.6	0.0	>
SAP HANA Ser...	dpserver	0.0	3,616.9	0.0	>
SAP HANA Ser...	indexserver	1.9	589,374.7	3.4	>
SAP HANA Ser...	daemon	0.0	0.0	0.0	>
SAP HANA Ser...	diserver	0.0	1,323.2	0.0	>
SAP HANA Ser...	nameserver	0.0	10,136.8	0.0	>
SAP HANA Ser...	compileserver	0.0	1,648.7	0.0	>
SAP HANA Ser...	preprocessor	0.1	59,573.0	0.3	>
SAP HANA Ser...	docstore	0.0	3,002.3	0.0	>
Idle	Idle	97.2	29,549,637.3	170.8	>
Other Consumers	Wait for IO	0.0	1,148.9	0.0	>
Other Consumers	Others CPU+SY...	0.5	169,398.5	0.9	>

Selecting and Deselecting Metrics

When you call up a data provider screen, you see a few selected metrics of this data provider as predefined by SAP. To add or deselect metrics, choose the [Metrics](#) button. In our example, we don't want the CPU time indicated in seconds, so we deselect it:




Don't forget to choose [Go](#) (or press `Enter`).

Specifying a Time Period

Many metrics are time-dependent, so let's get results for a concrete time period. On the time slider, we enter **7D** to get data for the last seven days:

Keep in mind that the time slider shows the time zone of the monitored system.



Don't forget to choose [Go](#).

Use the following abbreviations to enter days, weeks, or months. You can enter these abbreviations in the time-range input field in the middle under the time slider:

Shortcuts for Time Periods

Time Period	Abbreviation	Example
Year	<Number>Y	1Y
Month	<Number>M	1M
Week	<Number>W	7W
Day	<Number>D	1D
hour	<Number>h	1h
minute	<Number>m	30m
second	<Number>s	30s

You can combine these time units and also use them by adding the terms “last” or “align” as shown in the following examples:

Combined Shortcut Examples for the Time Slider

Examples	Description
last 12h L12h l 12h Last 12h 1h l 1hL	Data will be displayed for the last 12 hours until now. Note It's OK to abbreviate last and it doesn't matter if you use a blank afterwards or not. Upper or lower case is fine, and you can type last (or simply L or l) at the beginning or end, either combination works.
last 5M 14D 12h L5M, 14D, 12h	Data will be displayed for the last 5 months and 14 days and 12 hours. Note Possible delimiters between the time units are blanks, commas, semicolons, or hyphens.

Examples	Description
1M 1D 12h align	If you add align to your entry, the time range calculation is aligned (that is, rounded) to the smallest given unit that you enter, in this case hours.
1M 1D 12h a	
1M 1D 12h A	
 Note It's OK to abbreviate align and it doesn't matter if you use blanks or upper or lower case. You can type align (or simply A or a) at the beginning or end, either combination works.	
The time slider is on the current date and time and you enter one of the following time ranges:	If today's current date and time is, for example, 4-27-2017 5:15 p.m., the following will be displayed:
1h	Data from 4:15 p.m. to 5:15 p.m. today
1h a	Data from 4:00 p.m. to 5:00 p.m. today
1h l	Data from 4:15 p.m. to 5:15 p.m. today
You set the time slider to a certain time range completely in the past and the ending point is not now:	Let's say your time range in the past is 4-20-2017 4:15 p.m. to 5:15 p.m. In this case, the following will be displayed:
1h	Data from 4:15 p.m. to 5:15 p.m. on that day in the past
1h a	Data from 4:00 p.m. to 5:00 p.m. on that day in the past
1h l	Data from 4:15 p.m. to 5:15 p.m. today (because the attribute "last" is always aligned to "now", that is, to the current date and time).

You can add **last** or **align** where and as often as you want when you enter your desired time range.

By default, the time period ends with the current moment (now). You can move the time period using the arrows of the time slider or by dragging the highlighted bar on the time slider.

Filtering Data

You can filter metrics and characteristics by their values by choosing [Filter](#):

Add Filter

Select field

☒ Equal

☐ Not Equal

☐ Greater Than

☐ Less Than

☐ Like

☐ Not Like

☐ In List

☐ Not in List

☐ Is Null

☐ Is Not Null

Enter value


OK

Cancel

In our example, we don't set a filter because we want to see all values.

Sorting

You can define a sort order of the requested data by choosing [Sort](#). Let's sort the results by CPU time in percent:



With the [Move Up](#) and [Move Down](#) buttons on the left, you can set the priority of the sorting (the metric on top comes first, and so on). With the sorting buttons on the right, you indicate whether you want to sort ascending or descending.


Don't forget to choose [Go](#).

Note

If you want to sort your data, make sure that you use the [Sort](#) function in the upper toolbar. If you click on a column header to sort, only the metrics currently displayed on the screen are sorted, but not all metrics. For more information about sorting, see the [Tips and Troubleshooting](#) section at the end of this guided tour.

Limiting the Number of Rows to Be Displayed and Doing a Top N, Rest, and Total Aggregation

If your result table contains more than one entry, you can limit the number of rows to be displayed and also do a top n, rest, and total aggregation of sorted data. Choose [Rows](#) and enter your desired number of rows:

A dialog box titled "Row Settings" with a light gray background. It contains three rows of settings. The first row is "Max. Number of Rows:" followed by a text input field containing the number "1000". The second row is "Top N Aggregation (Σ):" followed by an unchecked checkbox. The third row is "Rest Aggregation (R):" followed by an unchecked checkbox. The fourth row is "Total Aggregation (T):" followed by an unchecked checkbox. At the bottom right, there are two buttons: "OK" and "Cancel". The "OK" button is highlighted with a dashed border.

Row Settings

Max. Number of Rows:

Top N Aggregation (Σ): ☐

Rest Aggregation (R): ☐


Total Aggregation (T): ☐

OK Cancel

Since the result list in our example isn't long, we leave the number of rows to be displayed as it is. If you want to find out more about doing a top n, rest, and total aggregation, see [Row Settings \[page 35\]](#).

Result

This is the result screen of our example:




For more information, see also [Aggregation \[page 27\]](#) and [Row Settings \[page 35\]](#).

Detail View and Associated Go-To Screens

By choosing the little arrow at the end of a row, you can go to the detail view of the selected row. If there are associated go-to screens or related data providers,

a popup window opens showing all available navigation options including the detail




view.

If you choose the detail view here, you get an overview of all selected characteristics and all metrics. On the left, there's a list of all selected columns of the data provider. You can limit the number of shown columns to only ad-hoc computed identifiers by choosing the toggle button with the key symbol at the top ([Show All Columns](#) / [Show Identifiers Only](#)).

Chart View


Now let's have a look how you can visualize the results. Use the [Show Chart](#) button to go to the chart view. By default, pie charts are displayed (as in our case here) unless you have selected a time characteristic in the drilldown earlier (for example, timestamp). In that case, the data is displayed in a time-based chart by default.



Some of the pie charts here are empty because the relevant metrics in our example have a zero value. The labels under the charts indicate the drilled down characteristics (here: service names). To see the values of the metrics (in our case, number of CPUs and CPU time in %), simply hover over the pie charts or click on the relevant areas of the chart.

For a different, sometimes more suitable visualisation (depending on how you want your data to be presented), you can use the [Pivot](#) button. The chart and table are transposed, that is, the rows are turned into columns and vice versa, except for time metrics, which cannot be pivoted/transposed. Hierarchy charts like these pie charts are more useful if pivoted. If we pivot the chart in our example, we get a better view of which consumer holds

which share of the overall percentage of CPU time and the overall number of CPUs:



Hover over or click on the different chart areas to see what is indicated there, or use the collapsible legend below the chart. If you want to change the chart settings, choose the [Settings](#) button in the upper right corner. To learn more about the many options you can choose in the settings, see [Changing Chart Settings \[page 97\]](#).

Add to Favorites

If you want to keep your selected characteristics, metrics, filters, sorting, row settings, and chart settings for later, you can save the screen with all these selections to your favorites. For more information, see [Add to Favorites \[page 26\]](#).

Tips and Troubleshooting

Sorting

You can also sort data on the screen by directly clicking on the relevant column header in the result list. Note, however, that this kind of sorting only sorts the data that is displayed on the screen. If you need a complete data refresh, you must choose the upper [Sort](#) button **and** choose [Go](#).

Note

If you want to run a top n analysis, this is particularly important. Make sure you refresh your data completely with [Sort](#) and [Go](#). A top n analysis without a data refresh leads to meaningless results. For more information about top n analysis, see [Row Settings \[page 35\]](#).

Related Information

[Changing Chart Settings \[page 97\]](#)


[Row Settings \[page 35\]](#)

[Pivot \[page 33\]](#)

[Add to Favorites \[page 26\]](#)

6.6.4.1 Changing Chart Settings

Apart from pivoting, there are more ways to change a chart: Choose the [Settings](#) button in the chart view and the [Chart Settings](#) popup window appears looking like this for our [Example: Guided Tour on How to Use the Data Explorer \[page 82\]](#) :



Here, you can change the chart type ([Time Series](#), [Hierarchy](#), or [Histogram](#)) as well as the data series type such as [Pie](#), [Donut](#), [Treemap](#), or [Sunburst](#). You can also choose a relative and/or stacked chart view as well as a second axis. Unless you have selected a time characteristic, the [Hierarchy](#) tab page with "pie" preselected is displayed by default.

Note

Due to technical reasons, there might be slight differences in chart rendering between SAP GUI and SAP UI5. For SAP GUI the following applies:


- If you choose a stacked chart view for one of your data series, **all** data series will be rendered as column charts.
- The maximum number of charts that can be displayed at a time is 52.
- If you have a Gantt and area chart and want to select a third or maybe fourth data series type, note that with the combination of Gantt and area chart, line and scatter charts will be rendered as area charts, too, in SAP GUI.
- If you combine a Gantt chart with other data series types, that is, if you also select, for example, a line or column chart in addition to Gantt, choosing a stacked view has no effect in SAP GUI and will only be rendered in SAP UI5.

However, you can still select and save these settings in SAP GUI for potential use in SAP UI5 where they will be rendered as intended.

As to choosing a second axis, there's a little exception: In SAP GUI, you cannot select different axes for stacked data series; all stacked data series must be on the same axis.

The color distribution in the charts depends on the first hierarchy level, that is, the first selected characteristic (here: "consumers"). Since we have three consumers, there are three different colors.

To change the color scheme, choose the color palette button so that the *Choose Color Scheme* popup window opens:



By default, color scheme *Seven* is preselected. Note that it's not a single color that you choose, but a color scheme, that is, a whole predefined array of colors. The *No. of Colors* slider gives you a preview of how the

different color schemes would look like if you changed the number of rows (= number of data series), that is, if you added, for example, characteristics and/or metrics. When you choose a color scheme, the chart settings and the chart behind the chart settings window change instantly.

If you're not satisfied and want your metrics to have a different color within the selected color scheme, choose the [Reverse Color Scheme](#) button and the colors are assigned in the opposite way.

Let's add a time characteristic to our drilldown - in our case, *Timestamp (Server)*. Since you don't want a pie chart now, but a time-based chart, go to the [Time Series](#) tab in the settings and choose, for example, a line chart:

Chart Settings

Time Series
Hierarchy
Histogram


Color		Relative	2nd Axis	Stacked	Data Series Type
		☐ NO	☐ NO	☐ NO	
	Idle -- Idle	☐ NO	☐ NO	☐ NO	
	SAP HANA Services -- indexserver	☐ NO	☐ NO	☐ NO	
	Other Consumers -- Others CPU+SYSTEM	☐ NO	☐ NO	☐ NO	
	SAP HANA Services -- preprocessor	☐ NO	☐ NO	☐ NO	
	Other Consumers -- Wait for IO	☐ NO	☐ NO	☐ NO	
	SAP HANA Services -- nameserver	☐ NO	☐ NO	☐ NO	
	SAP HANA Services -- scriptserver	☐ NO	☐ NO	☐ NO	
	SAP HANA Services -- dpserver	☐ NO	☐ NO	☐ NO	
	SAP HANA Services -- docstore	☐ NO	☐ NO	☐ NO	
	SAP HANA Services -- xsengine	☐ NO	☐ NO	☐ NO	

OK
Cancel

Our chart now looks like this after pivoting:




If a certain spot in the chart catches your attention, you can click on that point in time to have a closer look. A grey dotted line appears marking that point in time:



In the time slider, this point of time is also shown by a little red vertical line. You can then, for example, adjust the selected time period by extending the time slider's blue bar to the left and/or right around that red line to further see what was going on in your system around that point in time. Furthermore, if you go back to the table view, the relevant entry in the table is also highlighted and easy to spot.

On the *Time Series* tab page, you can choose between the following data series types: *Area*, *Column*, *Line*, *Scatter*, and *Gantt*. You can also select a relative and/or stacked chart view as well as a second axis:


Let's choose *Stacked* and *Column* to change our chart to a stacked column chart. After zooming in to get a better view with the magnifiers in the left corner of the chart, it looks like this:



You can drag the blue bar to the left and right to scroll around within your selected time range. To see what metric is indicated where, simply hover over the bars and circles or click on them to display the information. Of course, you can change and reverse the color scheme as you want just like we did with the pie charts above.


If you have a time series, there's an option that is not available for hierarchy charts: The *Swap Colors* button. This means you can assign the different colors of a color scheme to the metrics instead of the characteristics.

To do so, choose **Settings** > **Time Series** > **Color** > **Swap Colors**. Our stacked column chart then looks like this:



If you also want to see the frequency distribution of the CPU values, you can create a histogram by using the [Histogram](#) tab page in the chart settings. By default, data is distributed over 30 buckets. If you do a drilldown

(= select characteristics), for example, with "service name" and "timestamp", and set a filter for "service name equals indexserver", you get a histogram that looks like this:



To display details, simply click on or hover over the bars and legend.

Related Information

[Example: Guided Tour on How to Use the Data Explorer \[page 82\]](#)

[Zoom Into Chart \[page 40\]](#)

7 Adapting the Technical Monitoring Cockpit

Learn how to adapt the technical monitoring cockpit to your company's needs with favorites, dashboards, or with a changed launchpad.

Favorites

There's a monitoring screen in the navigation pane that you find extremely useful and want to keep? Or maybe you're already an expert and have picked the perfect selections and filters for a data provider? Then make sure to keep them by saving them as favorites so that they remain quickly available in the navigation pane under *Favorites*.

For more information about creating favorites, see [Add to Favorites \[page 26\]](#).

Changing the Launchpad

The launchpad is the home page and entry point to all important activities in the technical monitoring cockpit. The launchpad is the standard dashboard that you always get when you start the technical monitoring cockpit or when you return to the entry screen using the *Launchpad* button. You can change the launchpad like a dashboard by defining your own selection of tiles.

ⓘ Note

As opposed to the standard SAP Fiori launchpad, you cannot personalize the launchpad of the technical monitoring cockpit. When you change the launchpad, this is not a personalization, but the changes affect all users of the technical monitoring cockpit. However, you can define dashboards for specific scenarios. All users can use these dashboards.

For more information, see the following:

- [Changing the Tile Arrangement on the Launchpad or on a Dashboard in SAPUI5 \[page 106\]](#)
- [Changing the Tile Arrangement on the Launchpad or on a Dashboard in SAP GUI \[page 107\]](#)

Dashboards

A dashboard is an arrangement of tiles that you can create. Creating dashboards can be useful if you want to have different tile selections on your screen, depending on the work scenario you are in.

Note

The launchpad is one type of dashboard. It's the standard dashboard that appears when you start the technical monitoring cockpit.

Related Information

[Dashboards \[page 76\]](#)

[Creating a Dashboard \[page 105\]](#)

7.1 Creating a Dashboard

Creating dashboards can be useful if you want to have different tile selections on your screen, depending on your work situation.

Procedure

1. Open the navigation pane and proceed as follows:
 - **SAPUI5:** Choose [Edit Favorites...](#) From the buttons that appear, choose [Create Dashboard...](#)
 - **SAP GUI:** Choose [Create Dashboard...](#)
2. On the [Create New Dashboard](#) popup window, enter a name for the new dashboard.
3. If you also select the [Dashboard Is System-Specific](#) checkbox, the dashboard will always be called up with the currently selected system type.

A system type might be, for example, a specific database currently selected such as SAP HANA, an operating system, or the application server.

4. Select a parent folder for your dashboard.

You can choose any folder that is part of the navigation pane. By default, [Favorites](#) is selected.

5. Choose [Add](#).

A new empty dashboard is created under the parent folder that you have specified.

6. Follow the instructions for changing a dashboard to add tiles to your new, empty dashboard.

Related Information

[Dashboards \[page 76\]](#)

[Changing the Tile Arrangement on the Launchpad or on a Dashboard in SAPUI5 \[page 106\]](#)

7.2 Changing the Tile Arrangement on the Launchpad or on a Dashboard in SAPUI5

Are you unhappy with the selection or with the order of the tiles on the launchpad? Then change it! If you want to have different tile collections depending on your work situation, you can create multiple dashboards and add your favorite screens as tiles.

Procedure

1. Go to the launchpad or choose the dashboard that you want to edit.
2. To switch to edit mode, choose  *Edit Dashboard* in the bottom right corner of the main screen.

The technical monitoring cockpit now displays all available tiles. Tiles that are assigned to a group such as *Information* or *Tools* are displayed in their group. All other tiles are located under *Others (Unassigned)*.

3. Adapt the launchpad or the selected dashboard as follows:

- Add tiles by choosing  *Add*.

Note

If you have saved a specific selection, drilldown, and filtering of a data provider as a favorite, you will also find this favorite among the available tiles.

- Remove tiles by choosing  *Delete*.
 - To move a tile to a group, choose the *Edit* icon in the right lower corner of the tile.
You can now change the group assignment of the tile, its position, and its size.
4. To leave the edit mode, choose *Done* in the bottom right corner of the main screen.

Changes to the launchpad or to a dashboard are saved automatically.

7.3 Changing the Tile Arrangement on the Launchpad or on a Dashboard in SAP GUI

Are you unhappy with the selection or with the order of the tiles on the launchpad? Then change it! If you want to have different tile collections depending on your work situation, you can create multiple dashboards and add your favorite screens as tiles.

Procedure

1. Go to the launchpad or double-click the dashboard that you want to edit.
2. To switch to edit mode, choose  - *Edit Dashboard* on the main screen at the top.

The technical monitoring cockpit now displays all available tiles. Tiles that are assigned to a group such as *Information* or *Tools* are displayed in their group. All other tiles are located under *Others (Unassigned)*.

3. Adapt the launchpad or the selected dashboard as follows:

- To add a tile, choose the  - *Add Tile* button.

Note

If you have saved a specific selection, drilldown, and filtering of a data provider as a favorite, you will also find this favorite among the available tiles.

- To remove a tile, choose the  - *Remove Tile* button.
 - To move a tile to another group, choose the  - *Tile Settings* button.
With the tile settings, you can also change the size of the tile, the position of the tile in the group, and the time range and metrics that are displayed as preview on the thumbnail of the tile.
4. Leave the edit mode by choosing  - *Edit Dashboard* on the main screen at the top.

Changes to the launchpad or to a dashboard are saved automatically.

7.4 Adding a New Folder to the Menu

You can store your favorites in one of the folders in the menu of the technical monitoring cockpit. If the available folders don't fit your needs, you can create new folders.

Procedure

1. Open the menu and proceed as follows:

- **SAPUI5:** Choose [Edit Favorites...](#) From the buttons that appear, choose [Create Folder...](#)
 - **SAP GUI:** Choose [Show/Hide Navigation](#). From the buttons that appear at the top of the navigation pane, choose [Create Folder...](#)
2. On the [Create New Folder](#) popup window, enter a name for the new folder.
 3. If you also select the [Folder Is System-Specific](#) checkbox, the folder will always be called up with the currently selected system type.

A system type might be, for example, a specific database currently selected such as SAP HANA, an operating system, or the application server.
 4. Select a parent folder for your new folder.

You can choose any folder that is part of the navigation pane. By default, [Favorites](#) is selected.
 5. Choose [Add](#).

A new empty folder is created under the parent folder that you have specified.

Results

You can now add favorites (monitoring screens) and dashboards to the new folder.

Related Information

[Creating a Dashboard \[page 105\]](#)

[Add to Favorites \[page 26\]](#)

7.5 Changing or Deleting a Menu Item in SAPUI5

You can change the name of a menu item (dashboard or monitoring screen), its description and folder. You can also delete a menu item if you no longer need it.

Context

Note

Keep in mind that your changes and deletions are visible to all users in the system.

Deleting a dashboard will not delete the tiles selected for the dashboard.

Procedure

1. Open the menu and choose  [Edit Favorites](#) at the top of the menu.
2. Navigate to the item (monitoring screen/favorite, dashboard, or folder) that you want to change.
3. To change the name, folder, or description of a menu item, proceed as follows:
 - a. Choose  [Edit](#) next to the menu item.
 - b. Change the name, description, or folder of the menu item.
 - c. Choose [OK](#).
4. To delete the menu item, choose [Delete](#) from the buttons next to the menu item.
5. After you have completed your changes, choose  [Leave Edit Mode](#) at the top of the menu.

Related Information

[Changing the Tile Arrangement on the Launchpad or on a Dashboard in SAPUI5 \[page 106\]](#)

7.6 Changing or Deleting a Menu Item in SAP GUI

You can change the name, description, or folder of an existing dashboard or monitoring screen (favorite) in the menu (navigation pane). You can also delete a dashboard or monitoring screen if you no longer need it.

Context

Note

Keep in mind that your changes and deletions are visible to all users in the system.

Deleting a dashboard will not delete the tiles selected for the dashboard.

Procedure

- To change the name, description, or parent folder of a menu item (monitoring screen or dashboard), select the menu item and choose [Change...](#) in the navigation pane.
- To delete a menu item, select it and choose [Delete...](#) in the navigation pane.

Related Information

[Changing the Tile Arrangement on the Launchpad or on a Dashboard in SAP GUI \[page 107\]](#)

8 Tools

8.1 Content Explorer

The content explorer is a tool for browsing the folders and files of uploaded content for the technical monitoring cockpit.

Note

The content explorer is a tool designed for SAP support. If you are an SAP customer or partner, you don't need to use it.

8.2 Screen Catalog

The screen catalog provides you with a list of all screens available for use in the technical monitoring cockpit. In addition, you can also modify the tile settings of preconfigured monitoring screens.

Screens in the Screen Catalog

Screens available in the screen catalog include the following:

- Preconfigured monitoring screens created by SAP, which you can also find in the menu under [Information](#)
- Complex screens for analytical scenarios created by SAP, which you can also find in the menu under [Analysis](#)
- Data providers with their settings that you have stored as favorites in the data explorer
- Dashboards

Screens in the screen catalog have either been delivered by SAP as part of the content for the technical monitoring cockpit, or they are screens that you have created yourself using the [Create Favorites](#) function in the data explorer or using the [Create Dashboard](#) function of the technical monitoring cockpit.

Changing Tile Settings

You can add each preconfigured monitoring screen that you have created yourself or that SAP has delivered as tile on the launchpad or on a dashboard. In the screen catalog, you can change the tile settings by choosing the button [Change Tile Settings](#). The idea behind this is that you reduce the amount of information to make the thumbnail preview on the tile a bit more 'light-weight'.

You can change the following:

- Reduce the result set for a top n list by setting the maximum number of rows to a smaller value than defined for the preconfigured monitoring screen
- Reduce the number of characteristics, which will result in a more aggregated set of returned data
- Reduce the number of metrics to get a result set with fewer metrics

Related Information

[Add to Favorites \[page 26\]](#)

[Creating a Dashboard \[page 105\]](#)

9 Monitoring the SAP HANA Database

9.1 Analysis

In the menu, under [Analysis](#), you can find the available monitoring screens that allow you to analyze in detail specific issues with your system.

9.1.1 SQL Statement Analysis

Use the SQL statement analysis to analyse the performance of SQL statements within a given time range. The SQL statement analysis provides you with a chronology of performance metrics such as execution time and lock-wait time, and the database access plans.

Plan Cache

On the [Plan Cache](#) tab, the SQL statements from the plan cache are shown for a selected time range. (The default time range is the last 24 hours.) The default configuration of the tab page shows accumulated values for the execution time, number of executions, average execution time, lock-wait time, and number of lock waits.

Note

If you want to switch from a table view to a chart view, pivot the table first before you choose [Show Chart](#).

Analysis

The [Analysis](#) tab shows the corresponding time series for the statement that you selected on the [Plan Cache](#) tab: the number of result records, of executions, and of lock waits are shown.

If multiple access plans are available, you can choose between the access plans in the [Access Plans for Selected Statement](#) screen area. The time series is updated depending on the access plan chosen. If the database has created a new access plan during the selected time period, multiple access plans may exist as a result.

The left axis refers to the changes of lock waits, result records, and executions over time in percent. The right axis indicates the relative development of the execution time and lock-wait time in percent.

To display more details of the chart, do the following:

- Hover over the colored circles next to the left or right axis to display the minimum, maximum, mean, and median.

- Click on a column or line to see details, for example, the lock-wait time in microseconds for a particular time slot.

SQL Statement

On the [SQL Statement](#) tab, you can see the complete SQL statement that you selected. You can change the SQL statement. To generate an access plan for the changed SQL statement, choose [Calculate New Prepared Plan](#) or [Calculate New Executed Plan](#).

A changed SQL statement is reverted back to its original if you go to the [Plan Cache](#) tab, choose the original statement again, and then return to the [SQL Statement](#) tab.

Access Plan

Here you can see the database access plans (prepared and/or executed) of the selected SQL statement (or, if you have changed the SQL statement on the [SQL Statement](#) tab, the changed SQL statement). To see more details about an operator or a table in the access plan, choose a row in the access plan.

Related Information

[What Is the Data Explorer?](#) [page 79]

[What Are Data Providers?](#) [page 80]

9.2 Information

In the menu, under [Information](#), you can find the available monitoring screens preconfigured by SAP for the SAP HANA database.

9.2.1 Alerts

9.2.1.1 Delta Merge Errors

Get an overview of potential problems with delta merges in the last 24 hours. Only entries are displayed for which the field *Success* is not **TRUE**.

Where Does the Data Come From?

This monitoring screen is based on the monitoring view **HOST_DELTA_MERGE_STATISTICS**. For more information about the monitoring view, see the SAP HANA database documentation.

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

[Delta Merge Statistics \[page 131\]](#)

9.2.1.2 HANA Alerts

Get an overview of the current SAP HANA database alerts from the last 24 hours.

What Else Do You Need to Know?

Best practice recommendations for handling specific SAP HANA alerts are published in SAP Knowledge Base Articles (KBAs) available in SAP Service Marketplace. For most SAP HANA alerts, there is one dedicated KBA.

Where Does the Data Come From?

This monitoring screen is based on the data provider *Alerts*. For more information, see the data provider documentation.

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

[Alerts \[page 138\]](#)

9.2.2 Backup and Recovery

9.2.2.1 Backup Overview

Get an overview of the last 20 data and log backups that have been scheduled on your SAP HANA database, sorted by the end time of the backup.

The selection is restricted to the last 100 days.

Where Does the Data Come From?

This monitoring screen is based on the data provider [Backup](#). For more information, see the data provider documentation.

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

[Backup \[page 139\]](#)

9.2.3 CPU

9.2.3.1 CPU Utilization by Service

Monitor how the CPU is utilized by SAP HANA database services and by other consumers during the last 7 days.

The CPU utilization is shown as number of CPUs used and as percentage of available CPU resources.

What Else Do You Need to Know?

The resource CPU determines the computational power of the hardware. If your system has sufficient computational power at any time, you know that your software applications are running smoothly. To prevent performance issues during runtime, ensure that CPU utilization is below 100%. Be aware that idle CPUs imply costs without additional value, so you should also avoid this.

Where Does the Data Come From?

This monitoring screen is based on the data provider [CPU Utilization](#). For more information, see the data provider documentation.

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

[CPU Utilization per Service \[page 140\]](#)

9.2.3.2 CPU Utilization by Workload

Find out which workload type on the SAP HANA server consumes CPU capacity.

The workload types might be SQL statements, backups, delta merges, garbage collections, savepoints, non-HANA consumers on the same host, and so on.

What Else Do You Need to Know?

The resource CPU determines the computational power of the hardware. If your system has sufficient computational power at any time, you know that your software applications are running smoothly. To prevent performance issues during runtime, ensure that CPU utilization is below 100%. Be aware that idle CPUs imply costs without additional value, so you should also avoid this.

Where Does the Data Come From?

This monitoring screen is based on the data provider [CPU Utilization by Workload](#). For more information, see the data provider documentation.

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

[CPU Utilization per Workload \[page 140\]](#)

9.2.4 Configuration

9.2.4.1 Host Information

Get an overview of host information such as machine or the operating system configuration.

Where Does the Data Come From?

This monitoring screen is based on the monitoring view `M_HOST_INFORMATION`. For more information about the monitoring view, see the SAP HANA database documentation.

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

9.2.4.2 Parameter Check

Get an overview which HANA configuration parameters deviate from SAP default settings.

When Do You Need This Data?

In the case of problems, you might want to check which parameters of the SAP HANA database deviate from values recommended by SAP, which might be the root cause of the problems.

What Else Do You Need to Know?

Non-optimal parametrization of the SAP HANA database can have a negative impact on the performance of the database.

9.2.5 I/O

9.2.5.1 I/O Overview

View the average size and time for read and write accesses to log or data files. The average is calculated for the last 24 hours.

Where Does the Data Come From?

This monitoring screen is based on the monitoring view `HOST_VOLUME_IO_DETAILED_STATISTICS`. For more information about the monitoring view, see the SAP HANA database documentation.

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

9.2.6 Logs

9.2.6.1 Diagnosis Files

View existing trace files, including their size and date.

Where Does the Data Come From?

This monitoring screen is based on the monitoring view `M_TRACEFILES`. For more information about the monitoring view, see the SAP HANA database documentation.

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

9.2.6.2 Expensive Statement Trace

View the entries in the expensive statement trace from the last 30 days.

You can monitor all statements that have an execution time longer than a predefined threshold and that might therefore be performance bottlenecks.

Prerequisites

Whether and which information is collected and displayed in the expensive statement trace depends on the parameters in the `global.ini` file, section `expensive_statements`.

To trace expensive statements, you need to set the parameter `enable` to `true`. By default, the parameter `enable` has the value `false` and no data is stored in the expensive statement trace.

In section `expensive_statements`, you also need to set the threshold as of which runtime you consider a statement as expensive.

Where Does the Data Come From?

This monitoring screen is based on the monitoring view `M_EXPENSIVE_STATEMENTS`. For more information about the monitoring view, see the SAP HANA database documentation.

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

9.2.6.3 Service Restarts

See the number of alerts related to service restarts of your SAP HANA database and the alert timestamps from the last 7 days.

Where Does the Data Come From?

This monitoring screen is based on the data provider [Alerts](#). For more information, see the data provider documentation.

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

[Alerts \[page 138\]](#)

9.2.7 Memory

9.2.7.1 SAP HANA Memory Utilization

Get a high-level overview of the consumers using the available physical memory. The average values of the last 24 hours are displayed.

Use the drilldown to display the memory consumption of the different hosts or the memory utilization in more detail.


When Do You Need This Data?

- To find out how much memory is used or freed up by the SAP HANA database
- To identify a possible memory resource bottleneck
- To find out how much memory is allocated by processes other than the SAP HANA database

What Else Do You Need to Know?

Memory is a fundamental SAP HANA database resource. Understanding how the SAP HANA database requests, uses, and manages this resource is crucial to understanding how the system workload is processed. Only with sufficient memory at any time can you ensure that the system workload is processed optimally.

The following figure gives you an overview of the memory structure of the SAP HANA database:



Memory Structure of the SAP HANA Database

The memory segments are autonomously managed by the SAP HANA database. The memory segment size adapts to the current workload and therefore changes over time.

Where Does the Data Come From?

This monitoring screen is based on the data provider [Memory Utilization](#). For more information, see the data provider documentation.

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

[Memory Utilization \[page 144\]](#)

9.2.7.2 Top 10 Column Indexes by Size

View the top 10 column indexes sorted by size.

Where Does the Data Come From?

This monitoring screen is based on the data provider [Column Table Index Objects](#). For more information, see the data provider documentation.

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

[Column Table Index Objects \[page 141\]](#)

9.2.7.3 Top 10 Column Tables by Growth

View the top 10 fastest growing column tables.

Where Does the Data Come From?

This monitoring screen is based on the data provider [Column Table Objects](#). For more information, see the data provider documentation.

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

[Column Table Objects \[page 142\]](#)

9.2.7.4 Top 10 Column Tables by Size

View the top 10 column tables sorted by size.

Where Does the Data Come From?

This monitoring screen is based on the data provider [Column Table Objects](#). For more information, see the data provider documentation.

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

[Column Table Objects \[page 142\]](#)

9.2.7.5 Top 10 Indexes by Size

View the top 10 indexes sorted by size.

Where Does the Data Come From?

This monitoring screen is based on the data provider [Table Index Overview](#). For more information, see the data provider documentation.

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

[Table Index Overview \[page 152\]](#)

9.2.7.6 Top 10 Row Indexes by Size

View the top 10 row indexes sorted by size.

Where Does the Data Come From?

This monitoring screen is based on the data provider [Row Table Index Objects](#). For more information, see the data provider documentation.

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

[Row Table Index Objects \[page 149\]](#)

9.2.7.7 Top 10 Row Tables by Growth

View the top 10 fastest growing row tables.

Where Does the Data Come From?

This monitoring screen is based on the data provider [Row Table Objects](#). For more information, see the data provider documentation.

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

[Row Table Objects \[page 150\]](#)

9.2.7.8 Top 10 Row Tables by Size

View the top 10 row tables sorted by size.

Where Does the Data Come From?

This monitoring screen is based on the data provider [Row Table Objects](#). For more information, see the data provider documentation.

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

[Row Table Objects \[page 150\]](#)

9.2.7.9 Top 10 Tables by Size

View the top 10 tables by size.

Where Does the Data Come From?

This monitoring screen is based on the data provider [Table Overview](#). For more information, see the data provider documentation.

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

[Table Overview \[page 153\]](#)

9.2.8 Storage

9.2.8.1 Disk Usage

Get an overview of disk configuration and utilization of the SAP HANA database.

Where Does the Data Come From?

This monitoring screen is based on the monitoring view M_DISKS. For more information about the monitoring view, see the SAP HANA database documentation.

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

9.2.8.2 Volume Space Usage

Get an overview of the current volume space usage of your SAP HANA database for data and logs.

The volume space figures are grouped by usage type (data or logs) and by server name (index server, name server, XS engine, and so on).

Where Does the Data Come From?

This monitoring screen is based on the data provider [Volume Space Usage](#). For more information, see the data provider documentation.

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

[Volume Space Usage \[page 155\]](#)

9.2.9 Time Spent

9.2.9.1 Top 10 Users by Time Spent

Get an overview of which application user causes high processing times in a SAP HANA database system.

Prerequisites

To view the IDs of application users, you need the authorizations of authorization object S_TOOLS_EX.

When Do You Need This Data?

- To see who causes high system workload in the database system
- As a starting point, if you assume that you are facing resource bottlenecks in the database system

Where Does the Data Come From?

This monitoring screen is based on the data provider *Time Spent*.

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

[Time Spent \[page 153\]](#)

9.2.10 Workload

9.2.10.1 Blocked Transactions

Get an overview of blocked transactions during the last 24 hours, including the blocked table.

Where Does the Data Come From?

This monitoring screen is based on the data provider [Blocked Transactions](#). For more information, see the data provider documentation.

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

[Blocked Transactions \[page 139\]](#)

9.2.10.2 Current Activities

View the currently active threads.

Where Does the Data Come From?

This monitoring screen is based on the data provider [Workload Overview](#). For more information, see the data provider documentation.

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

[Workload Overview \[page 156\]](#)

9.2.10.3 Delta Merge Statistics

Get a chart with the number of merged delta records and the execution time for delta merges, aggregated over all hosts for the last 7 days.

Why Do You Need This Data?

Monitoring merges can be useful for the following reasons:

- The delta merge operation transfers data from the delta storage in SAP HANA to the main storage. The delta merge operation can be expensive because it temporarily duplicates the memory needed for the main storages and it might cause some disk I/O load.
For more information about how to mitigate the impact of merges on performance, search the SAP HANA administration guide for the keyword *delta merge operation*.
- The number of merged records reflects the number of database manipulation language (DML) operations on database tables. If the number of merged records increases or decreases, this indicates major changes or possible issues with the application.

Where Does the Data Come From?

This monitoring screen is based on the monitoring view `HOST_DELTA_MERGE_STATISTICS`. For more information about the monitoring view, see the SAP HANA database documentation.

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

9.2.10.4 Load History: Blocked Transactions

View the average number of blocked transactions during the last 24 hours.

Where Does the Data Come From?

This monitoring screen is based on the data provider [Load History](#). For more information, see the data provider documentation.

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

[Load History \[page 144\]](#)

9.2.10.5 Load History: CPU

Get an overview of average CPU utilization during the last 24 hours.

Where Does the Data Come From?

This monitoring screen is based on the data provider [Load History](#). For more information, see the data provider documentation.

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

[Load History \[page 144\]](#)

9.2.10.6 Load History: Threads

View the average number of all threads, all waiting threads, and all active threads for each host for the last 24 hours.

When Do You Need This Data?

In the load history, you can see how the resource consumption of the SAP HANA database develops over time and whether changes of the resource consumption correlate with changes of the SAP HANA response times.

Where Does the Data Come From?

This monitoring screen is based on the data provider [Load History](#). For more information, see the data provider documentation.

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

[Load History \[page 144\]](#)

9.2.10.7 SQL Workload

Get an overview of the SQL workload and lock waits of the SAP HANA database system and how this impacts the database response time.

When Do You Need This Data?

- To identify peaks in the SQL workload of the SAP HANA database
- To view long lock wait times of the SAP HANA database
- To see trends in the performance of the SAP HANA database

What Else Do You Need to Know?

Peaks in the SQL workload or locked objects of the SAP HANA system can lead to resource bottlenecks and therefore to a deteriorating system performance.

Where Does the Data Come From?

This monitoring screen is based on the data provider [SQL Plan Cache](#).

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

[SQL Plan Cache \[page 150\]](#)

9.2.10.8 Top 10 Applications by Execution Time

View the top 10 applications that spent time in your SAP HANA database, sorted by execution time.

When Do You Need This Data?

To find applications that are responsible for the workload of your SAP HANA database

Where Does the Data Come From?

This monitoring screen is based on the data provider *Time Spent*. For more information, see the data provider documentation.

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

[Time Spent \[page 153\]](#)

9.2.10.9 Top 10 SQL Statements by Execution Time

View the top 10 SQL statements in your database, sorted by execution time.

When Do You Need This Data?

The total execution time of SQL statements has a direct impact on the response time of the application.

A long total execution time might be caused by a high number of executions, even if the individual executions are very fast, or by a small number of executions with long single execution times.

Note

Use the monitoring screen [Top 10 SQL Statements by Execution Time](#) to improve SQL statement execution and thus the response time of the application. This might be relevant to improve the system performance as it is perceived by individual users using the SQL statements.

What Else Do You Need to Know?

You can get support for documenting and addressing expensive SQL statements from the self-service on SQL statement tuning available in SAP Solution Manager. For more information, see SAP Note [1601951](#) .

Where Does the Data Come From?

This monitoring screen is based on the data provider [SQL Plan Cache](#). For more information, see the data provider documentation.

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

[SQL Plan Cache \[page 150\]](#)

9.2.10.10 Transactional Locks

View the current lock situation in your SAP HANA database, including the locked objects and the locking transactions.

The table is empty when the system is not affected by a current lock situation.

When Do You Need This Data?

- To find out which transactions are locked or are holding locks
- To identify transactions that show or cause high lock-wait times

If the SAP HANA database has high response times, this might be caused by high lock-wait times. Transactional locks are the result of parallel system workload. The lock-wait times of transactions might increase if locks on database resources such as tables or table rows are held for a long time.

Where Does the Data Come From?

This monitoring screen is based on the data provider [Transactional Locks](#). For more information, see the data provider documentation.

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

[Transactional Locks \[page 155\]](#)

9.2.10.11 Unloads

Get the history of column unloads.

The database might actively unload tables or individual columns from memory, for example, if a query or other processes in the database require more memory than is currently available.

Note

If the field [Column Name](#) is empty, not a single column, but the whole table was unloaded.

Where Does the Data Come From?

This monitoring screen is based on the monitoring view `M_CS_UNLOADS`. For more information about the monitoring view, see the SAP HANA database documentation.

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

9.2.10.12 Workload Statistics

View the number of SQL statements executed in your database during the last 24 hours.

When Do You Need This Data?

The number of SQL statements might be a good starting point to get a first impression of the system workload. Use the drilldown, selection, and filtering options to display more workload statistics.

Where Does the Data Come From?

This monitoring screen is based on the data provider [Workload Statistics](#). For more information, see the data provider documentation.

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

[Workload Statistics \[page 156\]](#)

9.3 Tools

9.3.1 SQL EXPLAIN

Use the SQL EXPLAIN to display access plans for SQL statements.

SQL Statement

Enter an SQL statement for which you would like to view a database access plan. You can also enter the default schema to be assumed for the objects in the SQL statement.

Prepared Access Plan

Here you can see the prepared access plan of the SQL statement that you have entered. To see more details about an operator, table, index, or view in the access plan, choose a row in the relevant details section at the bottom of the screen.

Executed Access Plan

Here you can see the executed access plan of the SQL statement that you have entered. Note that for data-changing SQL statements, the executed plan is not available.

9.4 Data Providers in the Data Explorer

9.4.1 Data Providers for the SAP HANA Database

9.4.1.1 Alerts

Get the history of SAP HANA database alerts.

When Do You Need This Data?

- To get an overview of database-specific alerts and their statuses
- To resolve the alert situation

Alerts inform you about existing or suspected database issues so that you can protect the database system from data loss or unplanned outages.

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

9.4.1.2 Backup

Get a history of data backups and log backups of your SAP HANA database.

When Do You Need This Data?

- To get an overview of the backup size and duration
- To check backup-related messages.
For example, to investigate the reasons for an alert.

What Else Do You Need to Know?

Backups are necessary to recover your system if unexpected outages or data corruption occur. To store backups, you need to know how much disk space is required for data and log backups.

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

9.4.1.3 Blocked Transactions

Get a history of blocked transactions caused by concurrent processes on your SAP HANA database.

When Do You Need This Data?

- To identify long-running transactions
- To see which transactions are blocking other transactions and which transactions are waiting

What Else Do You Need to Know?

Blocked transactions are the result of concurrent processes in a system. How long system resources are blocked depends on application efficiency and system performance. If you understand lock dependencies, you can analyze and resolve existing lock wait situations more easily.

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

9.4.1.4 CPU Utilization per Service

Monitor how the CPU is utilized by SAP HANA database services and by other consumers. You can also display the idle time of your CPU.

When Do You Need This Data?

- To get an operating system view on the CPU workload of your SAP HANA database
- To find out whether your CPU is utilized optimally by SAP HANA database services
- To make sure that CPU resources in your SAP HANA database are provisioned optimally

What Else Do You Need to Know?

The resource CPU determines the computational power of the hardware. If your system has sufficient computational power at any time, you know that your software applications are running smoothly. To prevent performance issues during runtime, ensure that CPU utilization is below 100%. Be aware that idle CPUs imply costs without additional value, so you should also avoid this.

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

9.4.1.5 CPU Utilization per Workload

Monitor how the CPU is used by SAP HANA workload and by other consumers.

The SAP HANA workload comprises workload from applications and SAP HANA internal maintenance activities.

When Do You Need This Data?

- To get an operating system view of the CPU utilization of your SAP HANA system workload
- To find out whether your CPU is utilized optimally by SAP HANA workload
- To make sure that CPU resources in your SAP HANA system are provisioned optimally
- To display the idle time of your CPU

Why Is This Relevant?

The resource CPU determines the computational power of the hardware. If your system has sufficient computational power at any time, your software applications will run smoothly.

To prevent performance issues during runtime, ensure that the CPU occupancy rate is below 100%.

Be aware that idle CPUs imply costs without additional value, so you should also avoid CPU idle time as far as possible.

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

9.4.1.6 Column Table Index Objects

Get detailed ad-hoc information about the column table indexes of your SAP HANA database.

When Do You Need This Data?

- To find out, for example, the index size or the number of entries in the index
- To find out how much time it took to create the index

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

9.4.1.7 Column Table Objects

Get a detailed history of the column tables in your SAP HANA database.

When Do You Need This Data?

- To check the growth of column tables
- To see whether the table is partitioned, compressed, or replicated
- To see information about the workload on a table such as read and write access

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

9.4.1.8 Database Objects

Get an ad-hoc overview of your SAP HANA database properties.

When Do You Need This Data?

- To get general system information about your SAP HANA database
For example, the database version and database start time.
- To get aggregated information about your SAP HANA database
For example, system size, system resources, or alerts.

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

9.4.1.9 Disk Space Usage

Get an ad-hoc overview of the disk space usage of your SAP HANA database.

When Do You Need This Data?

To find out how much disk space is used by the SAP HANA database for data, logs, traces, and backups

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

9.4.1.10 I/O Performance

Get fast information about the I/O performance of the SAP HANA database. This information is based on standard KPIs such as the average time to read one 4K block of data from disk and the average time to write one 4K block of log to disk.

When Do You Need This Data?

You need this data when an SQL statement seems to be slow even though it was written with appropriate indexes, primary key access, single `INSERT`, and so on, so it should be fast. One root cause could be slow I/O reads and writes of the underlying database.

What Else Do You Need to Know?

Basic I/O information that you usually get is only information about how many GB were written or how much time was spent for reading and writing. By having standardized KPIs, you can see at first glance when something is wrong with I/O performance. When you work in a cloud environment, such KPIs might prove useful because the database might show slowdowns of its I/O subsystem.

Related Information

[What Is the Data Explorer? \[page 79\]](#)

9.4.1.11 Load History

Get a history of the system workload and resource consumption in the SAP HANA database.

When Do You Need This Data?

- To find out in what way resources relevant to the system workload are being consumed by the SAP HANA database
- To identify long system response times

What Else Do You Need to Know?

In the load history, you can see how the workload of an SAP HANA database behaves over time and, for example, whether changes in the workload correlate with the SAP HANA database response times.

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

9.4.1.12 Memory Utilization

Get a history of the memory utilization of your SAP HANA database.


When Do You Need This Data?

- To find out how much memory is used or freed up by the SAP HANA database
- To identify a possible memory resource bottleneck
- To find out how much memory is allocated by processes other than the SAP HANA database

What Else Do You Need to Know?

Memory is a fundamental SAP HANA database resource. Understanding how the SAP HANA database requests, uses, and manages this resource is crucial to understanding how the system workload is processed. Only with sufficient memory at any time can you ensure that the system workload is processed optimally.

The following figure gives you an overview of the memory structure of the SAP HANA database:



Memory Structure of the SAP HANA Database

The memory segments are autonomously managed by the SAP HANA database. The memory segment size adapts to the current workload and therefore changes over time.

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

9.4.1.13 Network Statistics

Get an overview of network communication traffic of your SAP HANA database.

Network communication could occur on one SAP HANA node (INTRA), between different SAP HANA nodes (INTER), or between SAP HANA nodes and external clients (EXTRA).

Prerequisites

Client network time (EXTRA) might significantly contribute to the overall response time of an SQL statement. To enable tracking for client network communication, the SAP HANA database configuration under

► `indexserver.ini` ► `[sql_client_network_io]` ► `enabled` ► must be set to `true`.

When Do You Need This Data?

- To find out how much network traffic is caused by your SAP HANA database
- To evaluate the network throughput of your system environment

What Else Do You Need to Know?

SAP HANA databases store information about network communication in different tables. This data provider consolidates this information into one view.

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

9.4.1.14 Parameter Check

Get an overview which HANA configuration parameters deviate from SAP default settings.

When Do You Need This Data?

In the case of problems, you might want to check which parameters of the SAP HANA database deviate from values recommended by SAP, which might be the root cause of the problems.

What Else Do You Need to Know?

Non-optimal parametrization of the SAP HANA database can have a negative impact on the performance of the database.

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

9.4.1.15 Procedure Parameters

Get an overview of the stored procedures in your SAP HANA database.

When Do You Need This Data?

- To find out which procedures exist in your SAP HANA database
- To get more detailed information about a procedure.
For example, procedure definition, input/output parameters, size, and database tables that are accessed.

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

9.4.1.16 Resource Utilization Peaks

Get an overview of the resource utilization peaks of your SAP HANA database by comparing average and maximum values of important resource metrics.

When Do You Need This Data?

- To identify a possible resource bottleneck
- To view the average and maximum memory and CPU utilization

What Else Do You Need to Know?

The calculation of average and maximum values of resource metrics is based on raw data that is collected every ten minutes. That is, if you look at the data provider with a time granularity of ten seconds, both average and maximum values are the same. With coarse time granularities, average and maximum values might differ. The maximum values might be indicators of temporary resource bottlenecks, even though the average resource metrics appear to be low.

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

9.4.1.17 Row Table Index Objects

Get detailed ad-hoc information about the row table indexes of your SAP HANA database.

When Do You Need This Data?

To find out, for example, the index size or the number of entries in the index

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

9.4.1.18 Row Table Objects

Get a detailed history of the row tables of your SAP HANA database.

When Do You Need This Data?

- To check the growth of row tables
- To find out whether the table is partitioned, compressed or replicated

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

9.4.1.19 SQL Plan Cache

Get an overview of the current and history execution plans of SAP HANA database. This helps you identify SQL statements that are expensive with regard to runtime or resource consumption.

When Do You Need This Data?

- To find out which SQL statements are running for a long time
- To identify statements that showed varying runtime in history
- To find out which SQL statements consume many system resources
- To identify transactions, programs, or users causing expensive SQL statements

What Else Do You Need to Know?

SAP HANA databases store history and current SQL workload in different tables. This data provider combines history and current SQL workload. SQL statements that are extracted from the current workload are attributed with the `CURRENT_TIMESTAMP` of the SAP HANA database system.

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

9.4.1.20 Savepoints

Get an overview about savepoint activities of the SAP HANA database by SAP HANA services.

Prerequisites

Client network time (EXTRA) might significantly contribute to the overall response time of an SQL statement. To enable tracking of client network communication, the SAP HANA database configuration under `indexserver.ini` > `[sql_client_network_io]` > `enabled` must be set to `true`.

When Do You Need this Data?

To get an overview on network communication volume and throughput.

What Else Do You Need to Know?

Savepoints are required to synchronize changes in memory with the persistency on disk level. All modified pages of row and column store are written to disk during a savepoint. The availability of a recent savepoint improves the restart time of SAP HANA because less redo logs need to be applied to make the database consistent. Each SAP HANA host and service has its own savepoints.

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

9.4.1.21 System Operations

Get an overview of SAP HANA operations that have been running for a long time.

When Do You Need This Data?

- To find out the duration of long-running system operations.
For example, backups, savepoints, delta merges, and table consistency checks.
- To find out the number of jobs in long-running system operations

What Else Do You Need to Know?

The system operations data provider gets its data from `_SYS_STATISTICS.HOST_JOB_HISTORY`.

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

9.4.1.22 Table Index Overview

Get an ad-hoc overview of the table indexes of your SAP HANA database.

When Do You Need This Data?

To see high-level information about row or column table indexes, for example, index type, index columns, size, and number of distinct records

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

9.4.1.23 Table Overview

Get an ad-hoc overview of the tables in your SAP HANA database.

When Do You Need This Data?

To see high-level information about row or column tables, for example, table type, size, and number of records

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

9.4.1.24 Time Spent

View processing and waiting times, and also to what extent specific activities contribute to the overall processing time of the SAP HANA database. This helps you identify how much time is spent on different database operations during a specific period of time.

Note

The times shown are times spent on specific activities. Parallel processing within the SAP HANA database may cause these times to be significantly higher than the number of seconds within the selected time period.

When Do You Need This Data?

- To identify the root cause of performance problems
- To analyze the workload of an SAP HANA database.
For example, by identifying how much time is spent on SQL processing, delta merge, garbage collection, backup, table load for specific applications or users.

- To help decide where to start with performance optimization

Why Is This Relevant?

If you are experiencing performance issues, it is important to find out how time is being spent in your SAP HANA database, and which workload may potentially be causing an issue.

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

9.4.1.25 Top Memory Consumers

Monitor areas in the SAP HANA database with the highest memory requirements.

When Do You Need This Data?

- To ensure that memory resources are used optimally
- To get an overview of the highest memory consumers in your SAP HANA database. This includes code, stack, heaps, tables, indexes, and fragmentation.
- To analyze the causes of high memory usage or out-of-memory dumps

Why Is This Relevant?

Memory is a fundamental resource for the SAP HANA database. Sufficient memory must be available at any time to ensure that the system workload is processed optimally.

Procedure

To filter SQL-related memory consumers, use the column *Is SQL Workload*.

To filter non-overlapping heap segments, use the column *Is Overlapping Heap Area*.

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

9.4.1.26 Transactional Locks

Get an ad-hoc overview of the transactional locks of your SAP HANA database.

When Do You Need This Data?

- To find out which transactions are locked or are holding locks
- To identify transactions that show or cause high lock-wait times

If the SAP HANA database has high response times, this might be caused by high lock-wait times. Transactional locks are the result of parallel system workload. The lock-wait times of transactions might increase if locks on database resources such as tables or table rows are held for a long time.

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

9.4.1.27 Volume Space Usage

Get an ad-hoc overview of the volume space usage of your SAP HANA database.

When Do You Need This Data?

To find out how much volume space is used by the SAP HANA database for data and logs

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

9.4.1.28 Workload Overview

Get an ad-hoc overview of the system workload of your SAP HANA database.

When Do You Need This Data?

- To find out what the threads in your SAP HANA database are currently doing, how long they have been running for, and how much CPU resources they require
- To find out which transactions are currently in state *Running* or *Locked*
- To find out which SQL statements are currently being executed on the SAP HANA database

What Else Do You Need to Know?

Checking the current workload is a good starting point for any kind of ad-hoc analysis.

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

9.4.1.29 Workload Statistics

Get a history of the system workload of your SAP HANA database.

When Do You Need This Data?

- To find out how often specific activities have been performed on your SAP HANA database, for example, read, write or merge activities
- To identify peak workloads of your system

What Else Do You Need to Know?

Balancing system workload over time allows system resources to be used optimally.

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

9.4.1.30 MiniCheck: CPU

Get an overview of CPU usage in SAP HANA.

When Do You Need This Data?

- To find out whether there are currently any issues concerning CPU consumption
- To find out if CPU consumption is stable long-term
- To identify SQL statements that are causing excessive CPU consumption

What Else Do You Need to Know?

The CPU check focuses on the following indicators:

- Minimum CPU rate
- Hosts with varying CPU rates
- Current CPU utilization
- Peak CPU utilization
- Time since CPU utilization
- Peak system CPU utilization

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

[SAP Note 2100040](#) 

9.4.1.31 MiniCheck: Delta Merges

Get information about delta merges in SAP HANA environments

When Do You Need This Data?

- To check the delta merge configuration (`mergedog`).
`mergedog` is the system process that periodically checks whether a delta merge needs to be performed.
- To check the number of records in delta storage in column-store tables
- To check the size of delta storage in column-store tables
- To check whether the delta merge of a table was executed successfully

What Else Do You Need to Know?

A delta merge is performed when SAP HANA main memory needs to be updated with the content of the delta storage. The delta storage collects changes to SAP HANA column store tables (INSERT, UPDATE, and DELETE) that are not immediately updated in main memory.

For more information, see *The Delta Merge Operation* in the *SAP HANA Database Administration Guide*

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

[SAP Note 2057046](#) 

9.4.1.32 MiniCheck: Garbage Collection

Get an overview of the garbage collection activities of your SAP HANA system.

When Do You Need This Data?

- To find out how much data that is no longer needed has been removed
- To check whether there are any potential issues concerning increased requirements for memory and disk space as a result of data cleanup

What Else Do You Need to Know?

The garbage collection check is primarily focused on:

- Row store version consolidation
- Column store version consolidation
- Memory garbage collection
- Persistence garbage collection
- Large object (LOB) garbage collection
- liveCache garbage collection
- Calculation engine garbage collection

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

[SAP Note 2169283](#) 

9.4.1.33 MiniCheck: Input/Output

Get information about I/O operations

When Do You Need This Data?

- To analyze the impact of I/O operations in your SAP HANA database
- To identify potential I/O bottlenecks

What Else Do You Need to Know?

The following operations in SAP HANA use I/O:

- Row store load
- Column store table load
- Hybrid large objects (LOBs)
- Persistence garbage collection
- Rollback
- System replication data shipping

- Savepoints
- Data snapshots
- System replication snapshots
- Data and log backups
- Database recovery
- Commits
- System replication log shipping
- Traces

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

[SAP Note 1999930](#) 📄

9.4.1.34 MiniCheck: Locks

Get information about SAP HANA connections that are hanging because they have to wait for locks

When Do You Need This Data?

- To check the configuration of lockwait times
- To identify blocked long-running transactions
- To find out what proportion of transactions are blocked

What Else Do You Need to Know?

The check distinguishes between the following lock types:

- Record locks
- Object locks
- Metadata locks
- Read / write locks
- Barrier semaphore waits
- Semaphore waits
- Mutex waits
- Barrier waits

- Speculative locks
- liveCache locks
- Nameserver topology locks

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

[SAP Note 1999998](#) 

9.4.1.35 MiniCheck: Memory

Get information about SAP HANA memory usage

When Do You Need This Data?

- To analyze the reasons for excessive memory usage
- To identify the causes of out-of-memory situations

What Else Do You Need to Know?

The check focuses on the following memory areas:

- Physical memory on the host
- Virtual memory on the host
- Resident memory
- Memory allocated to SAP HANA
- Memory used by SAP HANA
- SAP HANA shared memory
- SAP HANA heap memory
- SAP HANA code memory
- SAP HANA stack memory

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

[SAP Note 1999997](#) 

9.4.1.36 MiniCheck: Network

Get an overview of an SAP HANA network

When Do You Need This Data?

- To check the network performance
- To check the network settings in SAP HANA environments
- To get an overview of communications between different processes in a network

What Else Do You Need to Know?

You can check the following:

- Whether the network ports used by SAP HANA for internal communication are securely configured
- The average intra-node throughput
- The average inter-node throughput
- Retransmitted TCP segments
- Bad TCP segments
- The average intra-node communication volume
- The average inter-node communication volume
- Host name resolution for non-IP addresses

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

[SAP Note 2222200](#) 

9.4.1.37 MiniCheck: SQL Optimization

Get information about long-running SQL statements

When Do You Need This Data?

- To identify long-running SQL statements that are potentially impacting system performance

What Else Do You Need to Know?

You can check the following performance indicators:

- Long execution time
- High memory consumption
- High CPU consumption
- Long lock wait time
- Large number of returned records
- Large number of executions
- High intra-node and inter-node network communication

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

[SAP Note 2000002](#) 

9.4.1.38 MiniCheck: Savepoints

Get information about savepoints in your SAP HANA system

When Do You Need This Data?

- To find out when the last savepoint was defined.
That is, when data was last written to disk.
- To identify long-running savepoint operations

- To find out whether the time taken to prepare a data snapshot exceeds a given threshold

What Else Do You Need to Know?

Savepoints are needed to synchronize changes in memory with the permanent storage on disk. When a savepoint occurs, all the changed pages in the row store and the column store are written to disk.

Savepoints are set for each SAP HANA host and service.

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

[SAP Note 2100009](#) 

9.4.1.39 MiniCheck: System Replication

Get an overview of information about SAP HANA system replication

When Do You Need This Data?

- To identify internal database events
- To identify closed connections between the primary system and a secondary system.
If connections are closed, the primary system is no longer being replicated.
- Check the configuration of systems in system replication
It is important to identify configuration parameters that do not have the same value in the primary system and a secondary system.
- Check the log replay backlog for a secondary system

What Else Do You Need to Know?

SAP HANA system replication allows you to copy and continuously synchronize an SAP HANA database to a secondary location in the same data center or a remote data center. System replication supports high availability and fast disaster recovery.

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

[SAP Note 1999880](#) 

9.4.2 Technical Data Providers for the SAP HANA Database

Technical data providers are based on a specific data source in the SAP HANA database. For example, a table function or a catalog view.

Technical data providers are intended primarily for use by platform experts. For example, database administrators and SAP support.

Data Source

You need more information about the data source of this data provider?

1. Open the [Technical Information](#) section of this help window.
2. Under [Relevant for](#), you can find the database name.
3. Under [Where Does This Data Come From?](#), you can find the technical name of the data source.
4. For more information about the data source, search for the data source name in the documentation for the SAP HANA database.

9.4.2.1 Technical Data Provider M_EXPENSIVE_STATEMENTS

By default, this technical data provider has no data, as the expensive statements trace is deactivated.

To collect data, tracing for M_EXPENSIVE_STATEMENTS needs to be enabled.

Expensive statements are individual SQL queries whose execution time was above a configured threshold. Expensive statements may reduce the performance of the database.

Data Source

You need more information about the data source of this data provider?

1. Open the [Technical Information](#) section of this help window.
2. Under [Relevant for](#), you can find the database name.

3. Under [Where Does This Data Come From?](#), you can find the technical name of the data source.
4. For more information about the data source, search for the data source name in the documentation for the SAP HANA database.

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

9.4.2.2 Technical Data Provider M_SQL_CLIENT_NETWORK_IO

By default, this technical data provider has no data.

M_SQL_CLIENT_NETWORK_IO works in a similar way to an SQL trace. To collect client request data, you need to enable it.

To enable M_SQL_CLIENT_NETWORK_IO, use the following parameter setting:

```
indexserver.ini/[sql_client_network_io]/enabled=true
```

Data Source

You need more information about the data source of this data provider?

1. Open the [Technical Information](#) section of this help window.
2. Under [Relevant for](#), you can find the database name.
3. Under [Where Does This Data Come From?](#), you can find the technical name of the data source.
4. For more information about the data source, search for the data source name in the documentation for the SAP HANA database.

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

10 Monitoring SAP Adaptive Server Enterprise

10.1 Preconfigured Monitoring Screens for SAP ASE

Get an overview of the available monitoring screens preconfigured by SAP for SAP Adaptive Server Enterprise.

10.1.1 Configuration

10.1.1.1 Data Caches

Get a history of your data, index, and log pages that are currently in use, as well as pages used recently by SAP ASE database server.

When Do You Need This Data?

To administer the data caches to reconfigure them for performance

What Else Do You Need to Know?

Configuring named data caches does not divide the default cache into separate cache structures. The named data caches that you create can be used only by databases or database objects that are explicitly bound to them. All objects not explicitly bound to named data caches use the default data cache. Therefore, SAP ASE provides user-configurable data caches to improve performance, especially for multiprocessor servers.

Where Does the Data Come From?

This monitoring screen is based on the data provider [Data Caches](#). For more information, see the data provider documentation.

Related Information

[What Are Data Providers? \[page 80\]](#)

[Data Caches \[page 196\]](#)

10.1.1.2 Global Trace Flags and Switches

Get an overview about settings that are ON or OFF in your SAP ASE database.

When Do You Need This Data?

- To check the settings (ON/OFF) of global switches and query processing parameters
- To determine underlying causes of server's behavior for query processing

What Else Do You Need to Know?

This monitoring screen displays information about SAP ASE runtime environment, displaying current setting of several switches and query processing settings.

If the query performance is not meeting the expectation in SAP ASE, it may be due to unwanted settings of these switches or parameters. The behavior can be traced by looking at the settings.

Where Does the Data Come From?

This monitoring screen is based on the data provider [Global Trace Flags and Switches](#). For more information, see the data provider documentation.

Related Information

[What Are Data Providers? \[page 80\]](#)

[Global Trace Flags and Switches \[page 200\]](#)

10.1.1.3 Logins

Get a history for each valid SAP ASE user account.

When Do You Need This Data?

To identify the names, the activation time, the login time, the assigned roles and more information on the database users

Where Does the Data Come From?

This monitoring screen is based on the data provider [Logins](#). For more information, see the data provider documentation.

Related Information

[What Are Data Providers? \[page 80\]](#)

[Logins \[page 210\]](#)

10.1.1.4 Log Threshold

Get information about log thresholds of your SAP ASE database server.

When Do You Need This Data?

- A threshold has a free space amount and a stored procedure associated with it. When log segment free space falls below the free space amount, the threshold will act as a trip wire, activating the associated stored procedure.
- The procedure may alert users, dump the transaction log, increase the log space in a database, or take some other action.

Where Does the Data Come From?

This monitoring screen is based on the data provider [Log Threshold](#). For more information, see the data provider documentation.

Related Information

[What Are Data Providers? \[page 80\]](#)

[Log Threshold \[page 210\]](#)

10.1.1.5 Resource Limits

Get an overview of parameters for resource limits including the time of day to enforce the limit and the type of action to take.

When Do You Need This Data?

- To limit your queries and transactions from individual logins or applications from monopolizing server resources
- To prevent huge reports from running during critical times of the day
- To terminate a session that causes unwanted Cartesian products to be generated

What Else Do You Need to Know?

Resource limits are bound to time ranges, allowing the system administrator to define precisely when they should be enforced. Once a resource limit has been modified, all logged-in users can see the change, including the system administrator.

Where Does the Data Come From?

This monitoring screen is based on the data provider [Resource Limits](#). For more information, see the data provider documentation.

Related Information

[What Are Data Providers? \[page 80\]](#)

[Resource Limits \[page 215\]](#)

10.1.1.6 SAP Configuration Check

Get an overview of all database parameter settings from database configuration.

The most important drilldown characteristics are preselected and the results sorted by row rating, type, and technical ID.

When Do You Need This Data?

To check if your current database parameter settings are in sync with the SAP recommendations. If they are highlighted in green, they correspond to the current recommendations of the SAP parameter setting. The color red indicates that there is a deviation. The recommendation column tells you what the current SAP recommendation is.

What Else Do You Need to Know?

For standard parameter settings, see also the following SAP Notes:

- [1581695](#)  - SYB: Configuration Guide for SAP ASE 16.0
- [1749935](#)  - SYB: Configuration Guide for SAP ASE 15.7

Where Does the Data Come From?

This monitoring screen is based on the data provider [SAP Configuration Check](#). For more information, see the data provider documentation.

Related Information

[SAP Configuration Check \[page 218\]](#)

10.1.1.7 Server Configuration

Get an overview of the configuration parameters of your SAP ASE database server.

When Do You Need This Data?

- To get an overview of your configuration parameters by group, current values, non-default value settings and the amount of memory used by your settings
- To control your various aspects of SAP ASE server behavior
- To tailor your SAP ASE database server for your installation's requirements

What Else Do You Need to Know?

Configuration parameters are grouped according to the area of SAP ASE database server to simplify the identification of all parameters you need for specific server operations and for performance tuning.

Where Does the Data Come From?

This monitoring screen is based on the data provider [Server Configuration](#). For more information, see the data provider documentation.

Related Information

[What Are Data Providers? \[page 80\]](#)

[Server Configuration \[page 219\]](#)

10.1.2 CPU

10.1.2.1 CPU Utilization

Monitor how the CPU is utilized by your SAP ASE database server. The CPU utilization is shown as the number of CPUs used and as the percentage of available CPU resources.

Three system provided pools are present per default:

- `syb_default_pool`

- `syb_system_pool`
- `syb_blocking_pool`

User created pools are possible through configuration.

When Do You Need This Data?

- `syb_default_pool`: Size this pool to get multiple engines
- `syb_system_pool`: To use for I/O threads
- `syb_blocking_pool`: Threads can be blocked instead of blocking engines

What Else Do You Need to Know?

The resource CPU determines the computational power of the hardware. If your system has sufficient computational power at any time, your software applications are running smoothly. To prevent performance issues during runtime, ensure that the CPU occupancy rate is below 100%. Be aware that idle CPUs imply costs without additional value, so you must also avoid this.

Where Does the Data Come From?

This monitoring screen is based on the data provider [CPU Utilization](#). For more information, see the data provider documentation.

Related Information

[What Are Data Providers? \[page 80\]](#)

[CPU Utilization \[page 194\]](#)

10.1.3 Job Scheduler

10.1.3.1 Job Scheduler - Activity

Monitor the job executions and log or output messages of all historic jobs on your SAP ASE job scheduler.

When Do You Need This Data?

To analyze your job execution times, job states and associated schedule information and any failure or output log messages of various jobs executed on SAP ASE job scheduler.

Where Does the Data Come From?

This monitoring screen is based on the data provider [Job Scheduler - Activity](#). For more information, see the data provider documentation.

Related Information

[Job Scheduler - Activity \[page 207\]](#)

10.1.3.2 Job Scheduler - History

Monitor the job executions and log messages of all historic jobs on your SAP ASE job scheduler.

When Do You Need This Data?

To analyze job execution times, job states and associated schedule information and any failure log messages of various jobs executed on your SAP ASE job scheduler.

Where Does the Data Come From?

This monitoring screen is based on the data provider [Job Scheduler - History](#). For more information, see the data provider documentation.

Related Information

[Job Scheduler - History \[page 208\]](#)

10.1.3.3 Job Scheduler - Jobs

Get an overview of all jobs created by using the SAP ASE job scheduler.

When Do You Need This Data?

To analyze general job properties of all jobs created on your SAP ASE server.

Where Does the Data Come From?

This monitoring screen is based on the data provider [Job Scheduler - Jobs](#). For more information, see the data provider documentation.

Related Information

[Job Scheduler - Jobs \[page 208\]](#)

10.1.3.4 Job Scheduler - Schedules

Monitor the schedule properties of jobs scheduled on your SAP ASE server.

When Do You Need This Data?

To analyze schedule properties like schedule interval, schedule days etc. for any schedule which is associated with a job. This overview enables you to analyze any load situations on job scheduler and to plan re-scheduling jobs accordingly.

Where Does the Data Come From?

This monitoring screen is based on the data provider [Job Scheduler - Schedules](#). For more information, see the data provider documentation.

Related Information

[Job Scheduler - Schedules \[page 209\]](#)

10.1.4 Memory

10.1.4.1 Cache Hit Ratios

Get information about hit ratios for statement cache, procedure cache and all data caches in your SAP ASE database.

When Do You Need This Data?

To identify performance issues by monitoring the ratios of cache hits and misses. In this way, caches can be used efficiently and higher hit rates can result in better resource utilization and latency reduction.

Where Does the Data Come From?

This monitoring screen is based on the data provider [Cache Hit Ratios](#). For more information, see the data provider documentation.

Related Information

[Cache Hit Ratios \[page 193\]](#)

10.1.4.2 Memory Usage Percent

Get an overview of the memory usage in percentage for each category on your SAP ASE instance. It is a percentage of total memory configured in various categories such as data cache, memory pools, procedure cache.

When Do You Need This Data?

To get an overview of how the memory configuration looks like and how much free memory is still available for the configuration.

Where Does the Data Come From?

This monitoring screen is based on the data provider [Memory Usage Percent](#). For more information, see the data provider documentation.

Related Information

[Memory Usage Percent \[page 211\]](#)

10.1.4.3 Memory Utilization

Get an overview of the memory utilization of your SAP ASE server instance.

When Do You Need This Data?

- To find out how much memory is used and how it is allocated in your SAP ASE database
- To identify a possible memory resource bottleneck
- To find out how memory is allocated across different memory resources e.g. Caches, Memory Pools

Where Does the Data Come From?

This monitoring screen is based on the data provider [Memory Utilization](#). For more information, see the data provider documentation.

Related Information

[What Are Data Providers? \[page 80\]](#)

[Memory Utilization \[page 211\]](#)

10.1.5 Performance

10.1.5.1 Frequently Used Tables

Get an overview of frequently used tables in your SAP ASE database.

When Do You Need This Data?

To display information about tables in your database that have been used since the server was last restarted. The overview includes the table name, index ID, logical and physical reads, lock requests and waits, and contention statistics.

Where Does the Data Come From?

This monitoring screen is based on the data provider *Frequently Used Tables*. For more information, see the data provider documentation.

Related Information

[Frequently Used Tables \[page 200\]](#)

10.1.5.2 Long Running Transactions

Get an overview of long running transactions in your SAP ASE database.

When Do You Need This Data?

To get an overview of transactions like process ids, execution time, start time, failover, state etc. for transactions that are executing for over five seconds.

Where Does the Data Come From?

This monitoring screen is based on the data provider [Long Running Transactions](#). For more information, see the data provider documentation.

Related Information

[Long Running Transactions \[page 210\]](#)

10.1.5.3 Running Processes

Get an overview about running processes in your SAP ASE database.

When Do You Need This Data?

To identify issues with running processes and their process activity such as memory used, CPU time, reads, writes and get application and host information during analysis of these scenarios.

Where Does the Data Come From?

This monitoring screen is based on the data provider [Running Processes](#). For more information, see the data provider documentation.

Related Information

[Running Processes \[page 216\]](#)

10.1.6 Storage

10.1.6.1 Database Space Usage

Monitor the space utilized by databases in SAP ASE.

When Do You Need This Data?

- To get an overview on the space utilization of your SAP ASE databases
- To find out if allocated space is optimally utilized or requires additional space for SAP ASE databases

What Else Do You Need to Know?

Your SAP ASE instance has system databases master, model, tempdb, sybsystemprocs and user databases with the same name as the SAP system ID.

Where Does the Data Come From?

This monitoring screen is based on the data provider [Database Space Usage](#). For more information, see the data provider documentation.

Related Information

[What Are Data Providers? \[page 80\]](#)

[Database Space Usage \[page 195\]](#)

10.1.6.2 Device Space Usage

Get an ad-hoc overview of the space usage of all devices within your SAP ASE database.

When Do You Need This Data?

- To achieve optimal space allocation on the SAP ASE database server

- To display the logical map of your SAP ASE database to physical storage. The device will provide information about available space to be displayed.

What Else Do You Need to Know?

This monitoring screen displays information about all the devices on your SAP ASE server including device physical name, device resize date, device size, used space and device free space.

Your SAP ASE database creates a master device containing *master*, *sybsystemprocs*, *model* and *tempdb*.

Where Does the Data Come From?

This monitoring screen is based on the data provider [Device Space Usage](#). For more information, see the data provider documentation.

Related Information

[What Are Data Providers? \[page 80\]](#)

[Device Space Usage \[page 198\]](#)

10.1.6.3 Top 10 Tables by Growth

Get the list of Top 10 growing tables in your SAP ASE database.

When Do You Need This Data?

- To monitor space utilization in your database, analyze tables and keep track of tables in the SAP ASE database
- To administer and check the growth in database space

Where Does the Data Come From?

This monitoring screen is based on the data provider [Top 10 Tables By Growth](#). For more information, see the data provider documentation.

Related Information

[What Are Data Providers? \[page 80\]](#)

[Top 10 Tables By Growth \[page 224\]](#)

10.1.6.4 Top 10 Tables by Size

Get the list of top 10 tables by the total size consumed, in your SAP ASE database.

When Do You Need This Data?

- To monitor space utilization in your database
- To identify most consuming space tables
- To administer and analyze tables to find out the space they used
- To understand how space is allocated across databases and database objects in SAP ASE

Where Does the Data Come From?

This monitoring screen is based on the data provider [Top 10 Tables by Size](#). For more information, see the data provider documentation.

Related Information

[What Are Data Providers? \[page 80\]](#)

[Top 10 Tables by Size \[page 224\]](#)

10.1.6.5 Transaction Log Usage

Get an ad-hoc overview of the transaction log space in your SAP ASE database.

When Do You Need This Data?

- To get an overview on the space utilization of transaction logs in your SAP ASE database

- To check the transaction log space in your database when the system is busy
- To find out which application is responsible for filling up the log-space
- To identify performance issues with the help of Oldest Active Transaction (OAT)

What Else Do You Need to Know?

The transaction log usage gives you information about total log space allocated, active log space, and inactive log space in the database. And shows how log records are added and released from the transaction log and when your log space runs into transaction-full state.

Where Does the Data Come From?

This monitoring screen is based on the data provider [Transaction Log Usage](#). For more information, see the data provider documentation.

Related Information

[What Are Data Providers? \[page 80\]](#)

[Transaction Log Usage \[page 225\]](#)

10.1.7 Time Spent

10.1.7.1 Time Spent

View processing and waiting times and how much certain metrics contribute to the overall processing time. This helps you identify how much time is spent on different database operations during a certain time period.

When Do You Need This Data?

- To identify the root cause of performance problems
- To analyze workload situations, for example, if there are any contentions during processing or waiting times
- To determine where to start with performance optimizations of an SAP ASE database system

What Else Do You Need to Know?

You get an overview of grouping into wait classes that classify the type of event that a process is waiting for.

Where Does the Data Come From?

This monitoring screen is based on the data provider *Time Spent*. For more information, see the data provider documentation.

Related Information

[What Are Data Providers? \[page 80\]](#)

[Time Spent \[page 223\]](#)

10.1.7.2 Time Spent Graph

Get an overview of grouping into wait classes that classify the type of event that a process is waiting for or executing.

When Do You Need This Data?

- To identify the root cause of performance problems
- To analyze workload situations, for example, if there are any contentions during processing or waiting times
- To determine where to start with performance optimizations of an SAP ASE database system

What Else Do You Need to Know?

The time spent graph is a visual aspect of the data provider *Time Spent*, which displays the average processing time, and metric contributions of an event waiting or executing a process. This helps you identify how much time is spent on different database operations during a certain time period.

Where Does the Data Come From?

This monitoring screen graph is based on the data provider [Time Spent](#). For more information, see the data provider documentation.

Related Information

[What Are Data Providers? \[page 80\]](#)

[Time Spent \[page 223\]](#)

10.1.7.3 Top 10 Statements by Elapsed Time

View the top 10 SQL statements in your database, sorted by elapsed time and CPU time.

When Do You Need This Data?

The total execution time of SQL statements has a direct impact on the response time of the application.

A long total execution time might be caused by a high number of executions, even if the individual executions are very fast, or by a small number of executions with long single execution times.

Use the monitoring screen Top 10 SQL Statements by Elapsed Time to improve SQL statement execution and thus the response time of the application. This might be relevant to improve the system performance.

If you want to improve the resource consumption of SQL statements, you need to check their degree of parallelization. Therefore, to identify statements with high resource consumption, use the screen.

Where Does the Data Come From?

This monitoring screen is based on the data provider [SQL Statements](#). For more information, see the data provider documentation.

Related Information

[What Are Data Providers? \[page 80\]](#)

[SQL Statements \[page 220\]](#)

10.1.8 Workload

10.1.8.1 Expensive SQLs By CPU Time

Get a list of SQL statements which are consuming CPU time of your SAP ASE instance.

When Do You Need This Data?

- To check query performance on your SAP ASE instance
- To identify hot queries by CPU time
- To identify and tune SQL statements for better performance
- To get more information about full SQL text, query plans, and costs for investigating by using SSQLID

What Else Do You Need to Know?

Get a run-time list of SQL statements which are consuming CPU time of your SAP ASE instance. The list also provides SSQLID, lockwait time, and total IOs for those expensive SQL statements.

Where Does the Data Come From?

This monitoring screen is based on the data provider [Most Expensive SQLs By CPU time](#). For more information, see the data provider documentation.

Related Information

[What Are Data Providers? \[page 80\]](#)

[Expensive SQLs by CPU Time \[page 199\]](#)

10.1.8.2 Expensive SQLs By Total I/O

Get a list of SQL statements performing total I/O in your SAP ASE instance.

When Do You Need This Data?

- To check query performance on your SAP ASE instance
- To identify I/O intensive statements
- To identify and tune the SQL statements for better performance
- To receive more information about full SQL text, query plans, and costs for investigating by using SSQID

What Else Do You Need to Know?

The runtime list also provides SSQID, lockwait time, and total IOs for expensive SQL statements.

Where Does the Data Come From?

This monitoring screen is based on the data provider [Most Expensive SQLs By Total I/O](#). For more information, see the data provider documentation.

Related Information

[What Are Data Providers? \[page 80\]](#)

[Expensive SQLs by Total I/O \[page 199\]](#)

10.1.8.3 Resource Utilization

Monitor the cache usage statistics for indexes, objects, and database, such as the number of metadata descriptors currently in use by the server.

When Do You Need This Data?

- To provide information regarding available metadata or auxiliary scan descriptors not currently used

- To provide information regarding number of active metadata or auxiliary scan descriptors installed in cache
- To provide information regarding the percentage of cached or active metadata or auxiliary scan descriptors

What Else Do You Need to Know?

You can monitor the maximum number of descriptors used during the peak period.

Where Does the Data Come From?

This monitoring screen is based on the data provider [Resource Utilization](#). For more information, see the data provider documentation.

Related Information

[What Are Data Providers? \[page 80\]](#)

[Resource Utilization \[page 216\]](#)

10.1.8.4 Spinlock Contention

Get an overview of the spinlock activity of your SAP ASE database server.

When Do You Need This Data?

- To identify number of times an engine encountered spinlock contention on the cache
- To get an overview of waiting times, as a percentage of the total spinlock requests for that cache

What Else Do You Need to Know?

When a user task makes changes to a cache, a spinlock denies all other tasks access to the cache. Although spinlocks are held for extremely brief durations, they can slow performance in multiprocessor systems with high transaction rates. If spinlock contention is more than 10%, you may consider using named caches or adding cache partitions.

Where Does the Data Come From?

This monitoring screen is based on the data provider [Spinlock Contention](#). For more information, see the data provider documentation.

Related Information

[What Are Data Providers? \[page 80\]](#)

[Spinlock Contention \[page 220\]](#)

10.1.8.5 System Utilization

Get an overview of statistics for each engine.

When Do You Need This Data?

- To analyze statistics on an hourly, daily, weekly or monthly basis
- To get an overview about SAP ASE resources and percentage of each active resource

Where Does the Data Come From?

This monitoring screen is based on the data provider [System Utilization](#). For more information, see the data provider documentation.

Related Information

[What Are Data Providers? \[page 80\]](#)

[System Utilization \[page 221\]](#)

10.1.8.6 Temporary Database Activity

Get a history of all open local temporary databases, including global system `tempdb` when the instance is started in `tempdb` configuration mode.

Prerequisites

`monTempdbActivity` requires the enable monitoring, per object *statistics active*, and object *lockwait timing* configuration parameters to collect data.

When Do You Need This Data?

- To review temporary database activity that occurred in the past and that were captured by the related data collector
- To review temporary database activities that are currently occurring on the database server

Where Does the Data Come From?

This monitoring screen is based on the data provider [Temporary Database Activity](#). For more information, see the data provider documentation.

Related Information

[What Are Data Providers? \[page 80\]](#)

[Temporary Database Activity \[page 222\]](#)

10.1.8.7 Thread Activity

Get a history of the active threads in your SAP ASE database.

When Do You Need This Data?

- To find out how many threads have been working in your system and how much time they have spent working on specific task types

- To identify threads host engines executes user tasks, run specific jobs and process requests from a work queue
- To find out if SAP ASE contains system-defined thread pools and, if present, user-created thread pools
- To find out which SQL statements have been executed by active threads
- To locate the application source code that issues the SQL statement

What Else Do You Need to Know?

Threads are the workers of an SAP ASE database. The history of active threads gives you an overview of the system workload and tells you which application source code caused it. Thread pools are only available when SAP ASE is configured for threaded mode.

Where Does the Data Come From?

This monitoring screen is based on the data provider *Thread Activity Times*. For more information, see the data provider documentation.

Related Information

[What Are Data Providers? \[page 80\]](#)

[Thread Activity \[page 222\]](#)

10.1.8.8 Top 5 User Process By CPU Time

Get a list of the Top 5 user processes which are consuming CPU time.

When Do You Need This Data?

- To monitor the performance of your SAP ASE server based on CPU consumption
- To analyze SAP ASE CPU consumption by ASE processes

What Else Do You Need to Know?

This list also displays information about the reads, writes, and memory used details for the Top 5 processes.

Where Does the Data Come From?

This monitoring screen is based on the data provider [Top 5 User Processes By CPU Time](#). For more information, see the data provider documentation.

Related Information

[What Are Data Providers? \[page 80\]](#)

[Top 5 User Process By CPU Time \[page 224\]](#)

10.1.8.9 Workload Statistics

Get a history of the system workload of your SAP ASE database server.

When Do You Need This Data?

- To find out how often system activities have been performed on the SAP ASE database server, for example, read, write, or merge activities
- To identify peak workload of your system

What Else Do You Need to Know?

Balancing system workload over time allows system resources to be used optimally.

Where Does the Data Come From?

This monitoring screen is based on the data provider [Workload Statistics](#). For more information, see the data provider documentation.

Related Information

[What Are Data Providers? \[page 80\]](#)

[Workload Statistics \[page 226\]](#)

10.2 Data Providers for SAP Adaptive Server Enterprise

10.2.1 Available Free Memory

Get an overview of available memory of your SAP ASE server instance.

When Do You Need This Data?

- To display memory configured, memory utilized and memory free of your ASE instance
- To have an overview regarding memory related configuration parameters for example max memory and total logical memory

Related Information

[What Are Data Providers? \[page 80\]](#)

10.2.2 Cache Hit Ratios

Get information about hit ratios for statement cache, procedure cache and all data caches in your SAP ASE database.

When Do You Need This Data?

To identify performance issues by monitoring the ratios of cache hits and misses. In this way, caches can be used efficiently and higher hit rates can result in better resource utilization and latency reduction.

Related Information

[What Are Data Providers? \[page 80\]](#)

[What Are Data Providers? \[page 80\]](#)

10.2.3 Cache Pool Info

Get configuration information for the data cache.

Related Information

[What Are Data Providers? \[page 80\]](#)

10.2.4 CPU Utilization

Monitor how the CPU is utilized by SAP ASE database server in a threaded kernel. You can also display the idle time of your CPU.

When Do You Need This Data?

- To get an operating system view on the CPU workload of your SAP ASE database
- To find out whether your CPU is utilized optimally by SAP ASE database services
- To make sure that CPU resources in your SAP ASE database are provisioned optimally

What Else Do You Need to Know?

The resource CPU determines the computational power of the hardware. If your system has sufficient computational power at any time, you know that your software applications are running smoothly. To prevent performance issues during runtime, ensure that the CPU occupancy rate is below 100%. Be aware that idle CPUs imply costs without additional value, so you should also avoid this.

10.2.5 Database Dump History

Get an overview of database and transaction log dumps.

When Do You Need This Data?

To get a history of database and transaction log dumps in your SAP ASE database

Related Information

[What Are Data Providers? \[page 80\]](#)

10.2.6 Database Options

Get information about your SAP ASE database status and various database options.

When Do You Need This Data?

- To check the status of any database currently online or offline
- To check database options currently true/false for the database

Related Information

[What Are Data Providers? \[page 80\]](#)

10.2.7 Database Space Usage

Monitor the space utilized by databases in SAP ASE.

When Do You Need This Data?

- To get an overview on the space utilization of your SAP ASE databases
- To find out if allocated space is optimally utilized or requires additional space for SAP ASE databases

What Else Do You Need to Know?

This Data Provider will provide space allocated, used and free for the database. It also provides information about database properties such as encryption key, owner ID, and durability.

Devices in your SAP ASE database are created with appropriate sizes. And the devices allocates space to databases by executing create database and alter database commands.

In order to get more understanding, see also [Device Space Usage \[page 198\]](#)

10.2.8 Data Caches

Get a history of your data, index, and log pages that are currently in use, as well as pages used recently by SAP ASE database server.

When Do You Need This Data?

To administer the data caches to reconfigure them for performance

What Else Do You Need to Know?

Configuring named data caches does not divide the default cache into separate cache structures. The named data caches that you create can be used only by databases or database objects that are explicitly bound to them. All objects not explicitly bound to named data caches use the default data cache. Therefore, SAP ASE provides user-configurable data caches to improve performance, especially for multiprocessor servers.

10.2.9 Data Cache Objects

Get an overview of the content of your objects using the data caches in your SAP ASE instance.

When Do You Need This Data?

- To determine usage of data cache for the performance check
- To identify and evaluate memory usage by the objects in data caches
- To display content of your objects such as tables, indexes, databases using the data caches in your SAP ASE instance

Related Information

[What Are Data Providers? \[page 80\]](#)

10.2.10 Data Cache Usage

Get statistical information about the usage of data caches in your SAP ASE database.

When Do You Need This Data?

- To identify IO specific issues by looking at the statistics being displayed
- To determine if the configured size for the data cache is sufficient
- To display IO statistics such as physical reads, physical writes, logical reads, page read, and number of stalls for each buffer pool

Related Information

[What Are Data Providers? \[page 80\]](#)

10.2.11 Data Indexes

Monitor the space utilized by indexes in your SAP ASE database

When Do You Need This Data?

- To get an overview on the space utilized by indexes in your SAP ASE database
- To find out if the allocated space is optimally utilized
- To find out if additional space for indexes in your SAP ASE database is needed
- To get information about properties such as database, table, owner, and last modification date

Related Information

[What Are Data Providers? \[page 80\]](#)

10.2.12 Device IO Statistics

Get a history of read and write statistics of your database devices in SAP ASE.

When Do You Need This Data?

- To check the database device performance in your SAP ASE Instance
- To display the read and write statistics of database devices in your SAP ASE
- To identify any device specific read and write performance problems

Related Information

[What Are Data Providers? \[page 80\]](#)

10.2.13 Device Space Usage

Get an ad-hoc overview of the space usage of your SAP ASE database.

When Do You Need This Data?

- To get an ad-hoc overview of the space usage of your SAP ASE database
- To find out how much disk or file space is used by the SAP ASE database for data, logs, traces, and backups
- To display logical map of your SAP ASE database to physical storage. The device will provide information on available space to be displayed
- To provide general information of the volume space usage and additional metrics
- To achieve optimal space allocation on SAP ASE database server

10.2.14 Expensive SQLs by CPU Time

Get a list of SQL statements which are consuming CPU time of your SAP ASE instance.

When Do You Need This Data?

- To check query performance on your SAP ASE instance
- To identify hot queries by CPU time
- To identify and tune SQL statements for better performance
- To get more information about full SQL text, query plans, and costs for investigating by using SSQLID

What Else Do You Need to Know?

Get a run-time list of SQL statements which are consuming CPU time of your SAP ASE instance. The list also provides SSQLID, lockwait time, and total IOs for those expensive SQL statements.

10.2.15 Expensive SQLs by Total I/O

Get a list of SQL statements performing total I/O in your SAP ASE instance.

When Do You Need This Data?

- To check query performance on your SAP ASE instance
- To identify I/O intensive statements
- To identify and tune the SQL statements for better performance
- To receive more information about full SQL text, query plans, and costs for investigating by using SSQLID

What Else Do You Need to Know?

The runtime list also provides SSQLID, lockwait time, and total IOs for expensive SQL statements.

10.2.16 Frequently Used Tables

Get an overview of frequently used tables in your SAP ASE database.

When Do You Need This Data?

To display information about tables in your database that have been used since the server was last restarted. The overview includes the table name, index ID, logical and physical reads, lock requests and waits, and contention statistics.

Related Information

[What Are Data Providers? \[page 80\]](#)

10.2.17 Global Trace Flags and Switches

Get an overview about settings that are ON or OFF in your SAP ASE database.

When Do You Need This Data?

- To verify server run-time environment
- To find out which trace flags or query processing switches are set to ON in the SAP ASE server's run-time environment

What Else Do You Need to Know?

Trace flags and switches states (ON/OFF) are impacting the run-time performance of queries and your SAP ASE database. Therefore, turn the trace flags or query processing switches to ON only to verify SAP ASE behavior for troubleshooting purposes.

10.2.18 HADR Device Usage

Get an overview of the disk space used by the SAP Replication Server on your primary and secondary SAP ASE server.

When Do You Need This Data?

To identify whether the device usage is too high. In case of high usage, you must consider adding additional space to the replication paths. This reduces the risk that the primary ASE transaction log will run out of space.

What Else Do You Need to Know?

You need to use the `sap_add_device` command to add the SAP Replication Server device space.

Related Information

[What Are Data Providers? \[page 80\]](#)

10.2.19 HADR Resource Configuration

Get an overview of all the configuration details regarding the logical hosts of your replication server.

When Do You Need This Data?

To display configuration changes to the logical hosts, RS name, DB name, property name and their configured values.

Related Information

[What Are Data Providers? \[page 80\]](#)

10.2.20 HADR sap_collect_log

Use this stored procedure to collect the SAP ASE, Replication Server (RS), and Replication Management Agent (RMA) logs from different hosts in a HADR system in one local directory.

When Do You Need This Data?

To collect logs from HADR hosts, specify the number of days. The monitor parameters can be filtered by the number of days.

Related Information

[What Are Data Providers? \[page 80\]](#)

10.2.21 HADR sap_get_log

Monitor the information provided for the Replication Management Agent (RMA) and Replication Server logs.

When Do You Need This Data?

- To retrieve the log entries from the Replication Server log and the RMA log after a given timestamp
- To use the sap_get_log for RMA command to retrieve the Replication Management log entries even if the Replication Server is stopped

Related Information

[What Are Data Providers? \[page 80\]](#)

10.2.22 HADR sap_help

Display a list of available commands, or detailed information for the specified command.

When Do You Need This Data?

- To lookup available HADR commands

Related Information

[What Are Data Providers? \[page 80\]](#)

10.2.23 HADR sap_send_trace

Use this Replication Management Agent (RMA) stored procedure to get the latency calculation time. This command inserts an *rs_ticket* into the source database(s). This command sends trace to all participating databases for this host.

When Do You Need This Data?

To send the trace to a specific database, specify the database name by applying the filter to the name of the monitoring parameter.

Related Information

[What Are Data Providers? \[page 80\]](#)

10.2.24 HADR sap_set

Use this stored procedure of the Replication Management Agent (RMA) to read the values of the configuration parameters required to set up the replication system for SAP ASE HADR.

Related Information

[What Are Data Providers? \[page 80\]](#)

10.2.25 HADR sap_status active_path

Get an overview of the replication paths that are active from the primary site to the standby site.

When Do You Need This Data?

To receive information about latency time, latency and other replication server status

Related Information

[What Are Data Providers? \[page 80\]](#)

10.2.26 HADR sap_status path

Monitor the health and latency status of the replication paths.

When Do You Need This Data?

- To check the health and latency status of the replication paths
- To monitor information on the replication modes that were configured
- To monitor the current replication states in the custom HADR with DR node environment, distribution mode and path, replication server and latency

Related Information

[What Are Data Providers? \[page 80\]](#)

10.2.27 HADR sap_status_resource

Use this Replication Management Agent (RMA) stored procedure to monitor estimated failover time, device size, simple persistent queue (SPQ) size, usage, backlog, truncation backlog, replication route queue truncation backlog, SAP ASE transaction log size, and backlog.

Related Information

[What Are Data Providers? \[page 80\]](#)

10.2.28 HADR sap_status route

Monitor the sequence of queues, threads, and servers that the data is transacting in the replication path.

When Do You Need This Data?

To monitor information about the queues, threads, and servers

Related Information

[What Are Data Providers? \[page 80\]](#)

10.2.29 HADR Ticket History

Get a history of latency information. This is the amount of time taken for a transaction committed on the primary server to be applied on the standby server.

When Do You Need This Data?

- To monitor latency of the transaction after trace is sent from the primary host to the standby host
- To replicate from site1 to site2 and initiate the trace at site1(primary)
- To monitor the latency after executing the *sap_send_trace* command

Related Information

[What Are Data Providers? \[page 80\]](#)

10.2.30 HADR Version

Get an overview of your host, site name and version information of HADR and your SAP ASE.

When Do You Need This Data?

To check the version of your SAP ASE and HADR system

Related Information

[What Are Data Providers? \[page 80\]](#)

10.2.31 IO Queue Statistics

Get a history of the type of I/O and I/O time for your database devices in SAP ASE.

When Do You Need This Data?

- To display I/O time for database devices in SAP ASE
- To identify the type of I/O on the device
- To examine the I/O operations on database devices

Related Information

[What Are Data Providers? \[page 80\]](#)

10.2.32 Job Scheduler - Activity

Monitor the job executions and log or output messages of all historic jobs on your SAP ASE job scheduler.

When Do You Need This Data?

To analyze your job execution times, job states and associated schedule information and any failure or output log messages of various jobs executed on SAP ASE job scheduler.

Related Information

[What Are Data Providers? \[page 80\]](#)

10.2.33 Job Scheduler - History

Monitor the job executions and log messages of all historic jobs on your SAP ASE job scheduler.

When Do You Need This Data?

To analyze job execution times, job states and associated schedule information and any failure log messages of various jobs executed on your SAP ASE job scheduler.

Related Information

[What Are Data Providers? \[page 80\]](#)

10.2.34 Job Scheduler - Jobs

Get an overview of all jobs created by using the SAP ASE job scheduler.

When Do You Need This Data?

To analyze general job properties of all jobs created on your SAP ASE server

Related Information

[What Are Data Providers? \[page 80\]](#)

10.2.35 Job Scheduler - Schedules

Monitor the schedule properties of jobs scheduled on your SAP ASE server.

When Do You Need This Data?

To analyze schedule properties like schedule interval, schedule days etc. for any schedule which is associated with a job. This overview enables you to analyze any load situations on job scheduler and to plan re-scheduling jobs accordingly.

Related Information

[What Are Data Providers? \[page 80\]](#)

10.2.36 Locks per Process

Get an overview of the number of locks held per process in your SAP ASE database.

When Do You Need This Data?

- To identify the total number of locks held in your SAP ASE database
- To get an overview of locks held per process ID
- To receive database information about where the locks held exist

Related Information

[What Are Data Providers? \[page 80\]](#)

10.2.37 Logins

Get a history for each valid SAP ASE user account.

When Do You Need This Data?

To identify the names, the activation time, the login time, the assigned roles and more information about the database users

10.2.38 Log Threshold

Get information about log thresholds of your SAP ASE database server.

When Do You Need This Data?

- A threshold has a free space amount and a stored procedure associated with it. When log segment free space falls below the free space amount, the threshold will act as a trip wire, activating the associated stored procedure.
- The procedure may alert users, dump the transaction log, increase the log space in a database, or take some other action.

10.2.39 Long Running Transactions

Get an overview of long running transactions in your SAP ASE database.

When Do You Need This Data?

To get an overview of transactions like process ids, execution time, start time, failover, state etc. for transactions that are executing for over five seconds

Related Information

[What Are Data Providers? \[page 80\]](#)

10.2.40 Memory Usage Percent

Get an overview of the memory usage in percentage for each category on your SAP ASE instance. It is a percentage of total memory configured in various categories such as data cache, memory pools, procedure cache.

When Do You Need This Data?

To get an overview of how the memory configuration looks like and how much free memory is still available for the configuration.

Related Information

[What Are Data Providers? \[page 80\]](#)

10.2.41 Memory Utilization

Get an overview of the memory utilization of your SAP ASE server instance.

When Do You Need This Data?

- To find out how much memory is used and how it is allocated in your SAP ASE database
- To identify a possible memory resource bottleneck
- To find out how memory is allocated across different memory resources e.g. caches, memory pools


What Else Do You Need to Know?

Memory exists in SAP ASE as total logical or physical memory.

- *Total logical memory* is the sum of the memory required for all `sp_configure` parameters. The total logical memory must remain available, but can be in use at a given moment. The total logical memory value can change due to changes in the configuration parameter values.
- *Total physical memory* is the sum of all shared memory segments in SAP ASE. The total physical memory is the amount of memory SAP ASE uses at a given moment. You can verify this value with the read-only configuration parameter *total physical memory*. The value can only increase, because SAP ASE does not shrink memory pools once they are allocated. You can decrease the amount of total physical memory by changing the configuration parameters and restarting SAP ASE.

When SAP ASE starts, it allocates:

- Memory used by SAP ASE for nonconfigurable data structures
- Memory for all user-configurable parameters, including the data cache, the procedure cache, kernel resource memory e.g. memory pools, and the default data cache
- This figure illustrates how SAP ASE allocates memory as you change some of the memory configuration parameters:



10.2.42 Monitor Configuration

Get an overview of run-time cache usage statistics for metadata descriptors for indexes, objects, databases, and the kernel resource memory pool.

When Do You Need This Data?

- To determine the run-time utilization of various resources configured on your SAP ASE instance
- To fine-tune the configured values of the parameters based on their utilization status

What Else Do You Need to Know?

The stored procedure is executed with *all* options to list out the usage statistics of all the configuration parameters. But it can also be used for various configuration parameters. Therefore, you select the individual parameter for example *Open Objects* to get the utilization statistics.

Related Information

[What Are Data Providers? \[page 80\]](#)

10.2.43 Number of Open Objects

Get an overview of usage statistics for open objects currently configured in your SAP ASE database.

When Do You Need This Data?

To identify the number of free, active and maximum open objects used at peak periods for your SAP ASE database

Related Information

[What Are Data Providers? \[page 80\]](#)

10.2.44 Number of Open Partitions

Get an overview of usage statistics for open partitions currently configured in your SAP ASE database.

When Do You Need This Data?

To identify the number of free, active and maximum open partitions used at peak periods for your SAP ASE database

Related Information

[What Are Data Providers? \[page 80\]](#)

10.2.45 Procedure Cache Module Usage

Get statistical information about the usage of procedure cache modules in your SAP ASE database.

When Do You Need This Data?

- To identify issues with the procedure cache, you can look at the statistics displayed
- To display statistics for procedure cache module usage, such as number of active pages, number of reused pages, and high watermark pages

Related Information

[What Are Data Providers? \[page 80\]](#)

10.2.46 Procedure Cache Size in MB

Get the size of the procedure cache in your SAP ASE database.

When Do You Need This Data?

To determine whether the configured size for the procedure cache is sufficient

Related Information

[What Are Data Providers? \[page 80\]](#)

10.2.47 Procedure Cache Usage

Get statistical information about the usage of procedure cache in your SAP ASE database.

When Do You Need This Data?

- To identify issues with the cache usage, you can view the cache size displayed for each object
- To determine whether the configured size is sufficient for the procedure cache

Related Information

[What Are Data Providers? \[page 80\]](#)

10.2.48 Resource Limits

Get an overview of parameters for resource limits including the time of day to enforce the limit and the type of action to take.

When Do You Need This Data?

- To limit your queries and transactions from individual logins or applications from monopolizing server resources
- To prevent huge reports from running during critical times of the day
- To terminate a session that causes unwanted Cartesian products to be generated

What Else Do You Need to Know?

Resource limits are bound to time ranges, allowing the system administrator to define precisely when they should be enforced. Once a resource limit has been modified, all logged-in users can see the change, including the system administrator.

10.2.49 Resource Utilization

Monitor the cache usage statistics for indexes, objects, and database, such as the number of metadata descriptors currently in use by the server.

When Do You Need This Data?

- To provide information regarding available metadata or auxiliary scan descriptors not currently used
- To provide information regarding number of active metadata or auxiliary scan descriptors installed in cache
- To provide information regarding the percentage of cached or active metadata or auxiliary scan descriptors

What Else Do You Need to Know?

You can monitor the maximum number of descriptors used during the peak period.

10.2.50 Running Processes

Get an overview about running processes in your SAP ASE database.

When Do You Need This Data?

To identify issues with running processes and their process activity such as memory used, CPU time, reads, writes and get application and host information during analysis of these scenarios.

Related Information

[What Are Data Providers? \[page 80\]](#)

10.2.51 SAP ASE License

Get information about your SAP ASE License.

When Do You Need This Data?

To check the status of your license on your SAP ASE database server

What Else Do You Need to Know?

Your SAP ASE server cannot start, if your license is expired.

Related Information

[What Are Data Providers? \[page 80\]](#)

10.2.52 SAP ASE Version

Get information about the version of your SAP ASE database server

When Do You Need This Data?

To check the data and version string for the current release of your SAP ASE database server

Related Information

[What Are Data Providers? \[page 80\]](#)

10.2.53 SAP Configuration Check

Get an overview of all database parameter settings from database configuration

When Do You Need This Data?

To check if your current database parameter settings are in sync with the SAP recommendations. If they are highlighted in green, they correspond to the current recommendations of the SAP parameter setting. The color red indicates that there is a deviation. The recommendation column tells you what the current SAP recommendation is.

Related Information

[What Are Data Providers? \[page 80\]](#)

10.2.54 SAP NetWeaver Version

Get a list of SAP NetWeaver components installed in your SAP system and the patch level of each.

When Do You Need This Data?

- To implement the corrections attached to SAP Note [2372351](#) 
- To upgrade your SAP NetWeaver component
- SAP support requires this information to investigate in issues

Related Information

[What Are Data Providers? \[page 80\]](#)

10.2.55 Segments

Get an overview of segments in your SAP ASE database.

When Do You Need This Data?

To get an overview of segment usage in your SAP ASE database

Related Information

[What Are Data Providers? \[page 80\]](#)

10.2.56 Server Configuration

Get an overview of the configuration parameters of your SAP ASE database server.

When Do You Need This Data?

- To get an overview of your configuration parameters by group, current values, non-default value settings and the amount of memory used by your settings
- To control your various aspects of SAP ASE server behavior
- To tailor your SAP ASE database server for your installation's requirements

What Else Do You Need to Know?

Configuration parameters are grouped according to the area of SAP ASE database server to simplify the identification of all parameters you need for specific server operations and for performance tuning.

10.2.57 Spinlock Contention

Get an overview of the spinlock activity of your SAP ASE database server.

When Do You Need This Data?

- To identify number of times an engine encountered spinlock contention on the cache
- To get an overview of waiting times, as a percentage of the total spinlock requests for that cache

What Else Do You Need to Know?

When a user task makes changes to a cache, a spinlock denies all other tasks access to the cache. Although spinlocks are held for extremely brief durations, they can slow performance in multiprocessor systems with high transaction rates. If spinlock contention is more than 10%, you may consider using named caches or adding cache partitions.

10.2.58 SQL Statements

Get a history of the SQL statements in your SAP ASE database.

When Do You Need This Data?

- To find out how many SQL statements have been executed in your system and how much time they have spent running
- To find out which SQL statements have been running for a long time
- To find out which SQL statements are executed frequently
- To identify users executing specific SQL statements
- To locate the application source code that issues the SQL statement
- To monitor the capability for all prepared SQL Statements which are in the global SQL Statement Cache

What Else Do You Need to Know?

SQL statements are the main contributors to system workload. For system sizing or optimization, checking costly SQL statements is a good starting point.

Related Information

[What Are Data Providers? \[page 80\]](#)

10.2.59 Stored Procedure - Database Space Usage

Get an ad-hoc overview of the database space usage in your SAP ASE database.

When Do You Need This Data?

To find out how much space is used by the SAP ASE database for data and logs

Related Information

[What Are Data Providers? \[page 80\]](#)

10.2.60 System Utilization

Get an overview of statistics for each engine.

When Do You Need This Data?

- To compute averages that provides an indication of overall trends
- To analyze statistics on an hourly, daily, weekly or monthly basis

What Else Do You Need to Know?

This overview provides one row for each engine and statistic.

10.2.61 Temporary Database Activity

Get a history of all open local temporary databases, including global system `tempdb` when the instance is started in `tempdb` configuration mode.

Prerequisites

`monTempdbActivity` requires the enable monitoring, per object *statistics active*, and object *lockwait timing* configuration parameters to collect data.

When Do You Need This Data?

- To review temporary database activity that occurred in the past and that were captured by the related data collector
- To review temporary database activities that are currently occurring on the database server

10.2.62 Thread Activity

Get a history of the active threads in your SAP ASE database.

When Do You Need This Data?

- To find out how many threads have been working in your system and how much time they have spent working on specific task types
- To identify threads host engines executes user tasks, run specific jobs and process requests from a work queue
- To find out if SAP ASE contains system-defined thread pools and, if present, user-created thread pools
- To find out which SQL statements have been executed by active threads
- To locate the application source code that issues the SQL statement

What Else Do You Need to Know?

Threads are the workers of an SAP ASE database. The history of active threads gives you an overview of the system workload and tells you which application source code caused it. Thread pools are only available when SAP ASE is configured for threaded mode.

10.2.63 Time Spent

View processing and waiting times and how much certain metrics contribute to the overall processing time. This helps you identify how much time is spent on different database operations during a certain time period.

When Do You Need This Data?

- To identify the root cause of performance problems
- To optimize your database system
- To analyze workload situations, for example, if there are any contentions during processing or waiting times

10.2.64 Top 10 Modules Using ProcCache

Get the list of top 10 modules that allocates procedure cache memory.

When Do You Need This Data?

- To determine how the procedure cache memory is being allocated among the main modules
- To understand usage and possible tuning at the server level
- To get the number of pages currently allocated (active)
- To get the maximum number of memory pages allocated since the server was started for the module (HWM abbreviation for high-water mark)
- To identify the number of pages reused in the module

Related Information

[What Are Data Providers? \[page 80\]](#)

10.2.65 Top 10 Tables By Growth

Get the list of Top 10 growing tables in your SAP ASE database.

When Do You Need This Data?

- To monitor space utilization in your database, analyze tables and keep track of tables in the SAP ASE database
- To administer and check the growth in database space

10.2.66 Top 10 Tables by Size

Get the list of top 10 tables by the total size consumed, in your SAP ASE database.

When Do You Need This Data?

- To monitor space utilization in your database
- To identify most consuming space tables
- To administer and analyze tables to find out the space they used
- To understand how space is allocated across databases and database objects in SAP ASE

10.2.67 Top 5 User Process By CPU Time

Get a list of the Top 5 user processes which are consuming CPU time.

When Do You Need This Data?

- To monitor the performance of your SAP ASE server based on CPU consumption
- To analyze SAP ASE CPU consumption by ASE processes

What Else Do You Need to Know?

This list also displays information about the reads, writes, and memory used for the Top 5 processes.

10.2.68 Transaction Log Usage

Get an ad-hoc overview of the transaction log space in your SAP ASE database.

When Do You Need This Data?

- To get an overview on the space utilization of transaction logs in your SAP ASE database
- To check the transaction log space in your database when the system is busy
- To find out which application is responsible for filling up the log-space

What Else Do You Need to Know?

The transaction log usage gives you information about total log space allocated, active log space, and inactive log space in the database. Log records are added and freed from the transaction log and it shows when your database runs into log-full state.

10.2.69 Triggers

Get an overview of triggers in your SAP ASE system or user database.

When Do You Need This Data?

To get an overview of triggers and their related tables on system or user database. Currently the data provider shows only trigger information for the user database.

Related Information

[What Are Data Providers? \[page 80\]](#)

10.2.70 Unused Indexes

Get an overview of unused indexes in your SAP ASE database.

When Do You Need This Data?

To list indexes in your database that have not been used since the server was last restarted

Related Information

[What Are Data Providers? \[page 80\]](#)

10.2.71 User Connections

Provides the number of connections for each database user.

When Do You Need This Data ?

To review the number of connections for the database users.

Related Information

[What Are Data Providers? \[page 80\]](#)

10.2.72 Workload Statistics

Get a history of the system workload of your SAP ASE database server.

When Do You Need This Data?

- To find out how often specific activities have been performed on your SAP ASE database, for example, read, write or merge activities

- To identify peak workload of your system

What Else Do You Need to Know?

Balancing system workload over time allows system resources to be used optimally.

10.3 Technical Data Provider for SAP Adaptive Server Enterprise

This is a technical data provider that is based directly on a data source of the SAP ASE database, for example, a table function or a catalog view.

Technical data providers are mainly intended for platform experts, for example, database administrators and SAP support.

Data Source

You need more information about the data source of this data provider?

1. Open the *Technical Information* section of this help window.
2. Under *Relevant for*, you can find the database name.
3. Under *Where Does This Data Come From?*, you can find the technical name of the data source.
4. For more information about the data source, search for the data source name in the documentation for the SAP ASE database. You can find the SAP ASE online documentation at https://help.sap.com/viewer/p/SAP_ASE.
5. For additional information regarding `mon`, `j s`, and `sys` tables, refer to the SAP ASE online documentation *Reference Manual: Tables* at https://help.sap.com/viewer/p/SAP_ASE.

11 Monitoring IBM Db2 for i

11.1 Data Providers for IBM Db2 for i

11.1.1 SQL Statements

You can use this data provider to display the SQL statements executed within the database.

Prerequisites

The SAP Database Performance Collector for IBM i must be active on the monitored system.

When Do You Need This Data

To identify the most expensive SQL statements

What Else Do You Need to Know?

You can set a filter on the field [Execution Context](#) to display results only applicable for the monitored SAP system.

Related Information

[What Are Data Providers? \[page 80\]](#)

[What Is the Data Explorer? \[page 79\]](#)

11.1.2 Table Scans

You can use this data provider to identify the tables within the database for which table scans have been performed.

Prerequisites

The SAP Database Performance Collector for IBM i must be active on the monitored system.

When Do You Need This Data

- To identify the tables accessed via a table scan and the reason why this kind of access has been chosen by the database optimizer
- To identify the SQL statements that cause a table scan

To identify the most expensive SQL statements

What Else Do You Need to Know?

You can set a filter on the field *Execution Context* to display results only applicable for the monitored SAP system.

Related Information

[What Are Data Providers? \[page 80\]](#)

[What Is the Data Explorer? \[page 79\]](#)

11.1.3 Query Sorts

You can use this data provider to analyze the sort operations within the database.

Prerequisites

The SAP Database Performance Collector for IBM i must be active on the monitored system.

When Do You Need This Data

To identify the SQL statements that needed a temporary sorted list to satisfy the requested order

What Else Do You Need to Know?

You can set a filter on the field [Execution Context](#) to display results only applicable for the monitored SAP system.

Related Information

[What Are Data Providers? \[page 80\]](#)

[What Is the Data Explorer? \[page 79\]](#)

11.1.4 Created Indexes

You can use this data provider to identify the tables for which temporary indexes have been created during the SQL statement execution.

Prerequisites

The SAP Database Performance Collector for IBM i must be active on the monitored system.

When Do You Need This Data

To identify the SQL statements that needed the creation of a temporary index for an optimal execution. • To identify the SQL statements that needed the creation of a temporary index for an optimal execution.

What Else Do You Need to Know?

You can set a filter on the field [Execution Context](#) to display results only applicable for the monitored SAP system.

Related Information

[What Are Data Providers? \[page 80\]](#)

[What Is the Data Explorer? \[page 79\]](#)

11.1.5 Advised Indexes

You can use this data provider to identify the tables for which the database optimizer advised the creation of permanent indexes.

Prerequisites

The SAP Database Performance Collector for IBM i must be active on the monitored system.

When Do You Need This Data

- To identify whether the database optimizer advised the creation of a permanent index for the future execution of a SQL statement
- To identify the reason why the database optimizer advised this creation, what kind of index should be created and the columns to be used

What Else Do You Need to Know?

You can set a filter on the field [Execution Context](#) to display results only applicable for the monitored SAP system.

Related Information

[What Are Data Providers? \[page 80\]](#)

[What Is the Data Explorer? \[page 79\]](#)

11.1.6 Used Indexes

You can use this data provider to identify the indexes used by the database optimizer to access the tables from within the executed SQL statements.

Prerequisites

The SAP Database Performance Collector for IBM i must be active on the monitored system.

When Do You Need This Data

To identify how the index has been used during the SQL statement execution

What Else Do You Need to Know?

You can set a filter on the field [Execution Context](#) to display results only applicable for the monitored SAP system.

Related Information

[What Are Data Providers? \[page 80\]](#)

[What Is the Data Explorer? \[page 79\]](#)

11.1.7 Temporary Files

You can use this data provider to identify the indexes used by the database optimizer to access the tables from within the executed SQL statements.

Prerequisites

The SAP Database Performance Collector for IBM i must be active on the monitored system.

When Do You Need This Data

- To identify how the how much space has been consumed for this kind of objects during the SQL statement execution
- To identify how the how much space has been consumed for this kind of objects during the SQL statement execution
- To identify the reason why the database optimizer chose this kind of access

What Else Do You Need to Know?

You can set a filter on the field [Execution Context](#) to display results only applicable for the monitored SAP system.

Related Information

[What Are Data Providers? \[page 80\]](#)

[What Is the Data Explorer? \[page 79\]](#)

11.1.8 Temporary Bitmaps Creation and Merges

You can use this data provider to identify the SQL statements that required the creation or merge of temporary bitmaps during its execution.

Prerequisites

The SAP Database Performance Collector for IBM i must be active on the monitored system.

When Do You Need This Data

- To identify how the how much space has been consumed for this kind of objects during the SQL statement execution.
- To identify how the how much space has been consumed for this kind of objects during the SQL statement execution.
- To identify the reason why the database optimizer chose this kind of access

What Else Do You Need to Know?

You can set a filter on the field [Execution Context](#) to display results only applicable for the monitored SAP system.

Related Information

[What Are Data Providers? \[page 80\]](#)

[What Is the Data Explorer? \[page 79\]](#)

11.1.9 Temporary Hash Tables

You can use this data provider to identify the SQL statements that required the creation of a temporary hash table during its execution. You can use this data provider to identify the SQL statements that required the creation of a temporary hash table during its execution.

Prerequisites

The SAP Database Performance Collector for IBM i must be active on the monitored system.

When Do You Need This Data

- To identify how much space has been consumed for this kind of objects during the SQL statement execution
- To identify the reason why the database optimizer chose this kind of access

What Else Do You Need to Know?

You can set a filter on the field *Execution Context* to display results only applicable for the monitored SAP system.

Related Information

[What Are Data Providers? \[page 80\]](#)

[What Is the Data Explorer? \[page 79\]](#)

11.1.10 Host Variables

You can use this data provider to identify the different host variables combinations used during the SQL statements execution.

Prerequisites

The SAP Database Performance Collector for IBM i must be active on the monitored system.

When Do You Need This Data

To identify the host variable values used for the SQL statement execution

What Else Do You Need to Know?

You can set a filter on the field [Execution Context](#) to display results only applicable for the monitored SAP system.

Related Information

[What Are Data Providers? \[page 80\]](#)

[What Is the Data Explorer? \[page 79\]](#)

11.2 Technical Data Providers for IBM Db2 for i

11.2.1 Technical Data Provider Based on Db2 for IBM i Data Sources

This technical data provider is based on a Db2 for IBM i data source, for example, an IBM i catalog table or view, Db2 for i or IBM i service.

Technical data providers are mainly intended for platform experts, for example, database administrators and SAP support.

Data Source

You need more information about the data source of this data provider?

1. Open the [Technical Information](#) section of this help window.
2. Under [Relevant for](#), you can find the name of the platform provider.

3. Under [Where Does This Data Come From?](#), you can find the technical name of the data source.
4. For more information about the data source, search for its technical name on the IBM website.

Related Information

[What Are Technical Data Providers? \[page 81\]](#)

11.2.2 Technical Data Provider Based on SAP Database Performance Collector for IBM i

This technical data provider is based on a table in the `SAPDB4M` schema for SAP systems running on non-switchable ASPs or `SAPDB4M<nnn>` schema for SAP systems running on independent ASPs, where `<nnn>` is the associated ASP number, or this technical data provider is based on an user-defined table function in the `R3SAP400` library.

Technical data providers are mainly intended for platform experts, for example, database administrators and SAP support.

Data Source

You need more information about the data source for this data provider?

1. Open the [Technical Information](#) section of this help window.
2. Under [Relevant for](#), you can find the name of the platform provider.
3. Under [Where Does This Data Come From?](#), you can find the technical name of the data source.
4. For more information about the data source based on `SAPDB4M` schema, search for its technical name in the document attached to SAP Note [1622665](#) – IBM i: SAP Database Performance Collector.

Related Information

[What Are Technical Data Providers? \[page 81\]](#)

11.2.3 Technical Data Provider Based on PLAN_CACHE Table of the SAP Database Performance Collector for IBM i

This technical data provider is based on the PLAN_CACHE table in the SAPDB4M schema for SAP systems running on non-switchable ASPs or SAPDB4M<nnn> schema for SAP systems running on independent ASPs, where <nnn> is the associated ASP number.

Technical data providers are mainly intended for platform experts, for example, database administrators and SAP support.

The data provider name identifies the database monitor information being shown. For example, QQQ1000 SQL Information, QQQ3000 Table Scan, and so on.

Data Source

You need more information about the data source for this data provider?

1. Open the [Technical Information](#) section of this help window.
2. Under [Relevant for](#), you can find the name of the platform provider.
3. Under [Where Does This Data Come From?](#), you can find the technical name of the data source.
4. The PLAN_CACHE table in the SAPDB4M schema is based on the structure of the QSYS/QAQQDBMN performance statistics table, that is shipped with the system.
The value in the row identification number (QQRID) field lets you identify the kind of information that is shown.
For more information about the performance statistics table, search for it on the IBM website.

Related Information

[What Are Technical Data Providers? \[page 81\]](#)

12 Monitoring IBM Db2 for Linux, UNIX, and Windows

12.1 Preconfigured Monitoring Screens for IBM Db2 for Linux, UNIX, and Windows

12.1.1 Backup and Recovery

12.1.1.1 Backup Overview

Get an overview of the backups that ran in the last seven days, sorted by start time.

Where Does the Data Come From?

This monitoring screen is based on the data provider *Backup*. For more information, see the data provider documentation.

Related Information

Data provider [Backup](#) [page 246]

[What Are Data Providers?](#) [page 80]

[What Is the Data Explorer?](#) [page 79]

12.1.2 Configuration

12.1.2.1 Database Parameters

At a glance, see all your database parameter settings from the following different sources:

- Database configuration
- Database manager configuration
- Database registry

- Automatic maintenance

The most important drilldown characteristics are preselected and the results are sorted by row rating, type, and technical ID.

When Do You Need This Data?

To check if your current database parameter settings are in sync with the current SAP recommendations. If indicated in green, they correspond to the current SAP parameter setting recommendations, the color red indicates that there's a deviation. The *Recommendation* column tells you what the current SAP recommendation is.

Where Does the Data Come From?

This monitoring screen is based on the data provider *Database Parameters*.

What Else Do You Need to Know?

For standard parameter settings, see also the following SAP Notes:

Master SAP Note: [1334923](#) 

Db2 11.1: [2303771](#) 

Db2 10.5: [1851832](#) 

Db2 10.1: [1692571](#) 

Db2 9.7: [1329179](#) 

Related Information

Data provider [Database Parameters](#) [page 247]

[What Are Data Providers?](#) [page 80]

[What Is the Data Explorer?](#) [page 79]

12.1.3 Memory

12.1.3.1 Current Memory Utilization

Get an overview of the current memory utilization of your Db2 database in relation to the resources utilized by the operating system.

When Do You Need This Data?

- To find out how much memory is used by the database
- To identify a potential memory resource bottleneck
- To find out how much memory is allocated by processes other than the database

What Else Do You Need to Know?

The memory data is organized in four hierarchy layers, starting with the highest layer [H1 Physical Memory](#) where memory used by the database and the operating system is displayed separately. The lower layers [H2 Db2 Memory Level](#), [H3 Db2 Memory Group](#), and [H4 Db2 Memory Pool](#) further subdivide the memory usage of the database.

Where Does the Data Come From?

This monitoring screen is based on the data provider *Memory Overview*.

Related Information

Data provider [Memory Overview](#) [page 248]

[What Are Data Providers?](#) [page 80]

[What Is the Data Explorer?](#) [page 79]

12.1.3.2 Memory Statistics

Get a history of the memory utilization of your Db2 database.

When Do You Need This Data?

- To see the size of Db2 memory components develop over time, for example, the memory allocated to bufferpools
- To find out about memory bottleneck situations and which memory component is responsible

What Else Do You Need to Know?

The memory data is organized in four hierarchy layers, starting with the highest layer [H1 Physical Memory](#) where memory used by the database and the operating system is displayed separately. The lower layers [H2 Db2 Memory Level](#), [H3 Db2 Memory Group](#), and [H4 Db2 Memory Pool](#) further subdivide the memory usage of the database.

Where Does the Data Come From?

This monitoring screen is based on the data provider *Memory Statistics*.

Related Information

Data provider [Memory Statistics](#) [page 249]

[What Are Data Providers?](#) [page 80]

[What Is the Data Explorer?](#) [page 79]

12.1.4 Network

12.1.4.1 Network Statistics

Analyze network performance between application servers and database server to find out if there's a database or network issue when you have unsatisfactory response times.

The charts of this monitoring screen show data aggregated over all application servers. The first chart shows the average network time per roundtrip together with the underlying network load (data sent and data received). The length of the network time compared to the server time is displayed in the second chart.

When Do You Need This Data?

To get answers to the following questions:

- Are long database response times due to long processing times on the database server or to slow communication between the database server and the SAP application servers?
- How long is the network time compared to the server time?
- How long is the average network time per roundtrip and how does this metric react to changing network load?

Where Does the Data Come From?

This monitoring screen is based on the data provider *Network Statistics*. For more information, see the data provider documentation.

Related Information

Data provider [Network Statistics \[page 250\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Is the Data Explorer? \[page 79\]](#)

12.1.5 Time Spent

12.1.5.1 Elapsed vs. Processing Times

Compare the elapsed time spent in a logical component of the database with the actual processing time spent in that component.

When Do You Need This Data?

To find out if the elapsed time spent in a logical component of the database is dominated by the actual processing time or by other wait situations.

Where Does the Data Come From?

This monitoring screen is based on the data provider *Time Spent*. For more information, see the data provider documentation.

Related Information

Data provider [Time Spent \[page 250\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Is the Data Explorer? \[page 79\]](#)

12.1.5.2 Time Spent by Activity

See a breakdown of the time spent in the database engine by activity. The listed activities are Db2 monitoring elements that let you divide the overall time spent into processing times of logical engine components and wait times.

When Do You Need This Data?

To identify which database activities contribute relevantly to the overall time spent in the database engine in a given time range.

Where Does the Data Come From?

This monitoring screen is based on the data provider *Time Spent*. For more information, see the data provider documentation.

Related Information

Data provider [Time Spent \[page 250\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Is the Data Explorer? \[page 79\]](#)

12.1.6 Workload

12.1.6.1 Costly Transactions

Identify units of work (UOWs) that potentially impact system performance, and obtain information such as the program context that is responsible for the UOW.

When Do You Need This Data?

To identify database transactions that are costly in terms of:

- Lock wait time
- Number of locks held
- Runtime while holding log space
- Amount of log space used

Where Does the Data Come From?

This monitoring screen is based on the data provider *Costly Transactions*.

Related Information

Data provider [Costly Transactions \[page 247\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Is the Data Explorer? \[page 79\]](#)

12.1.6.2 Current Activities

Get an overview of the current workload on your Db2 database.

When Do You Need This Data?

- To monitor active and inactive applications that are connected to your database
- To find out what the Db2 agents in your database are currently doing (processing, waiting, and so on) and how long database units of work (transactions) are running
- To identify SQL statements and system or user routines that are currently being executed on the database

Where Does the Data Come From?

This monitoring screen is based on the data provider *Workload Overview*.

Related Information

Data provider [Workload Overview \[page 251\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Is the Data Explorer? \[page 79\]](#)

12.1.6.3 Workload Statistics

Get a history of database performance metrics indicating the time dependence of the workload of your Db2 database.

When Do You Need This Data?

- To find out how often database activities were performed, for example, reads, writes, and updates of table rows, package cache or catalog cache activities, and sorts
- To identify times of high and low workloads in the database

Where Does the Data Come From?

This monitoring screen is based on the data provider *Workload Statistics*.

Related Information

Data provider [Workload Statistics \[page 251\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Is the Data Explorer? \[page 79\]](#)

12.2 Data Providers for IBM Db2 for Linux, UNIX, and Windows

12.2.1 Backup

Get a list with information about backups that ran on your database.

When Do You Need This Data?

- To get an overview of backup size and runtime
- To see whether a backup finished successfully or failed
- To get further information, such as the backup type

What Else Do You Need to Know?

Backups are necessary to recover your system if unexpected outages or data corruption occurred.

Note

Backup size information is readily available as of Db2 10.5. For Db2 9.7 and Db2 10.1, the registry variable DB2_BAR_STATS must be set to [ON](#) to enable the recording of backup size information.

Related Information

[What Are Data Providers? \[page 80\]](#)

[What Is the Data Explorer? \[page 79\]](#)

12.2.2 Costly Transactions

Identify units of work (UOWs) that potentially impact system performance, and obtain information such as the program context that is responsible for the UOW.

When Do You Need This Data?

To identify database transactions that are costly in terms of:

- Lock wait time
- Number of locks held
- Runtime while holding log space
- Amount of log space used

Related Information

[What Are Data Providers? \[page 80\]](#)

[What Is the Data Explorer? \[page 79\]](#)

12.2.3 Database Parameters

At a glance, see all your database parameter settings from the following different sources:

- Database configuration
- Database manager configuration
- Database registry
- Automatic maintenance

When Do You Need This Data?

To check if your current database parameter settings are in sync with the current SAP recommendations. If indicated in green, they correspond to the current SAP parameter setting recommendations, the color

red indicates that there's a deviation. The [Recommendation](#) column tells you what the current SAP recommendation is.

What Else Do You Need to Know?

For standard parameter settings, see also the following SAP Notes:

Master SAP Note: [1334923](#) 

Db2 11.1: [2303771](#) 

Db2 10.5: [1851832](#) 

Db2 10.1: [1692571](#) 

Db2 9.7: [1329179](#) 

Related Information

[What Are Data Providers?](#) [page 80]

[What Is the Data Explorer?](#) [page 79]

12.2.4 Memory Overview

Get an overview of the current memory utilization of your Db2 database in relation to the resources utilized by the operating system.

When Do You Need This Data?

- To find out how much memory is used by the database
- To identify a potential memory resource bottleneck
- To find out how much memory is allocated by processes other than the database

What Else Do You Need to Know?

The memory data is organized in four hierarchy layers, starting with the highest layer [H1 Physical Memory](#) where memory used by the database and the operating system is displayed separately. The lower layers [H2 Db2 Memory Level](#), [H3 Db2 Memory Group](#), and [H4 Db2 Memory Pool](#) further subdivide the memory usage of the database.

Related Information

[What Are Data Providers? \[page 80\]](#)

[What Is the Data Explorer? \[page 79\]](#)

12.2.5 Memory Statistics

Get a history of the memory utilization of your Db2 database.

When Do You Need This Data?

- To see the size of Db2 memory components develop over time, for example, the memory allocated to bufferpools
- To find out about memory bottleneck situations and which memory component is responsible

What Else Do You Need to Know?

The memory data is organized in four hierarchy layers, starting with the highest layer *H1 Physical Memory* where memory used by the database and the operating system is displayed separately. The lower layers *H2 Db2 Memory Level*, *H3 Db2 Memory Group*, and *H4 Db2 Memory Pool* further subdivide the memory usage of the database.

Related Information

[What Are Data Providers? \[page 80\]](#)

[What Is the Data Explorer? \[page 79\]](#)

12.2.6 Network Statistics

Analyze network performance between application server and database server to find out if there is a database or a network issue when you have unsatisfactory response times.

Prerequisites

You need at least Db2 Version 9.7 Fix Pack 4 or higher (see also SAP Note [1625157](#) ).

When Do You Need This Data?

To get answers to the following questions:

- Are the long database response times due to long processing times on the database server or to slow communication between the database server and the SAP application server?
- How are the network statistics distributed among the members or the application servers of your system?
- How long is the network time compared to the server time?
- How long is the average network time per roundtrip and how does this metric react to changing network load?

Related Information

[What Are Data Providers? \[page 80\]](#)

[What Is the Data Explorer? \[page 79\]](#)

12.2.7 Time Spent

View processing and waiting times and how much certain metrics contribute to the overall processing time. This helps you identify how much time is spent on different database operations during a certain time period.

When Do You Need This Data?

- To identify the root cause of performance problems
- To analyze workload situations, for example, by identifying what the most active service classes and members are and if there are any outliers
- To determine where to start with performance optimizations of your database system

Related Information

[What Are Data Providers? \[page 80\]](#)

[What Is the Data Explorer? \[page 79\]](#)

12.2.8 Workload Overview

Get an overview of the current workload on your Db2 database.

When Do You Need This Data?

- To monitor active and inactive applications that are connected to your database
- To find out what the Db2 agents in your database are currently doing (processing, waiting, and so on) and how long database units of work (transactions) are running
- To identify SQL statements and system or user routines that are currently being executed on the database

Related Information

[What Are Data Providers? \[page 80\]](#)

[What Is the Data Explorer? \[page 79\]](#)

12.2.9 Workload Statistics

Get a history of database performance metrics indicating the time dependence of the workload of your Db2 database.

When Do You Need This Data?

- To find out how often database activities were performed, for example, reads, writes, and updates of table rows, package cache or catalog cache activities, and sorts.
- To identify times of high and low workloads in the database.

Related Information

[What Are Data Providers? \[page 80\]](#)

[What Is the Data Explorer? \[page 79\]](#)

12.3 Technical Data Providers for IBM Db2 for Linux, UNIX, and Windows

12.3.1 Technical Data Provider Based on Db2 Data Sources

This is a technical data provider that is based directly on an IBM Db2 data source, for example, a table function or a catalog view.

Technical data providers are mainly intended for platform experts, for example, database administrators and SAP support.

Data Source

You need more information about the data source of this data provider?

1. Open the [Technical Information](#) section of this help window.
2. Under [Relevant for](#), you can find the name of the platform provider.
3. Under [Where Does This Data Come From?](#), you can find the technical name of the data source.
4. For more information about the data source, search for its technical name on the IBM website.

Related Information

[What Are Technical Data Providers? \[page 81\]](#)

12.3.2 Technical Data Provider Based on SAP Tables

This data provider is based on an SAP table in the SAPTOOLS schema. Tables in the SAPTOOLS schema are used to store collected monitoring data.

Technical data providers are mainly intended for platform experts, for example, database administrators and SAP support.

Related Information

[What Are Technical Data Providers? \[page 81\]](#)

12.4 Tools

12.4.1 SQL EXPLAIN

Use the EXPLAIN tool to display access plans for SQL statements.

SQL Statement

On the [SQL Statement](#) tab page, enter an SQL statement for which you'd like to see a database access plan. In the [Parameters](#) section, you can also enter the default schema to be assumed for the objects in the SQL statement, or values for the query degree and optimization level. If you call the EXPLAIN without specifying these parameters, the default values are used, which will be shown after the EXPLAIN run.

In the [Optimization Guideline](#) section, choose one of the following options that will be applied when the EXPLAIN call is performed:

- [No](#): Use no optimization guideline
- [Own](#): Use the optimization guideline that you enter in the text input field.
- [From Stmt](#): Use the optimization guideline attached to the original SQL statement in the comment. The button is only enabled if the input SQL statement has an optimization guideline comment attached.
- [From Profile](#): Use the optimization guideline provided by an optimizer profile that matches the SQL statement. The button is only enabled if there's a match in the optimizer profile table DB6_OPTPROFILE of the database for the input SQL statement.

For more information about optimizer profiles, see SAP Note [1818503](#) . For the syntax of optimization guidelines, see SAP Note [868888](#) .

Overview

On the [Overview](#) tab page, you can see information about the source of the SQL statement, the database environment, the SQL compiler, and summary metrics.

For more information about the input fields, look up *EXPLAIN_INSTANCE* table in the IBM Knowledge Center.

Prepared Access Plan

On this tab page, you'll find the prepared access plan of the SQL statement that you've entered. For more details about an operator, table, index, or view in the access plan, choose a row in the relevant detail section at the bottom of the screen.

Optimized Statement

This tab page shows the SQL statement that was re-written by the Db2 optimizer.

EXPLAIN Snapshot

On this tab page, you'll find the output of an EXPLAIN snapshot. For more information about what's shown here, look up *EXPLAIN snapshot* in the IBM Knowledge Center.

13 Monitoring IBM Db2 for z/OS

13.1 Preconfigured Monitoring Screens for IBM Db2 for z/OS

13.1.1 Alerts

13.1.1.1 Active Log Shortages

Get an overview of active log shortages for the selected period in the time selection.

When Do You Need This Data?

To monitor the cases when the active log size was insufficient and identify the most extended unit of work in the application.

Where Does the Data Come From?

This monitoring screen is based on the technical data provider *DB2QLOGS*. For more information, see the technical data provider documentation.

Related Information

[Technical Data Provider DB2QLOGS \[page 327\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

[What Is the Data Explorer? \[page 79\]](#)

13.1.1.2 DB2 Alerts

Get an overview of database-specific alerts.

Where Does the Data Come From?

The stored procedure SAPCL starts a database alert router process that connects to Db2 via RRS and catches Db2 alerts via the IFI interface of Db2. This process exclusively listens to Db2 alerts. It is a separate process because capturing database events is an awaiting task.

This monitoring screen is based on the multiple monitoring views that are based on the technical data providers as follows:

- *Active Log Shortages* – DB2QLOGS
- *Lock Escalations* – DB2LOCKESCAL
- *Long Running Transactions* – DB2QLRUR
- *Deadlocks* – DB2QLODLK and DB2QLODRS
- *Timeouts* – DB2QLOTHW and DB2QLOTRS

For more information about alerts, see the *Database Administration Guide for IBM Db2 for z/OS* at https://help.sap.com/viewer/db2_administration_guide.

Related Information

[Technical Data Provider DB2LOCKESCAL \[page 322\]](#)

[Technical Data Provider DB2QLODLK \[page 326\]](#)

[Technical Data Provider DB2QLODRS \[page 327\]](#)

[Technical Data Provider DB2QLOGS \[page 327\]](#)

[Technical Data Provider DB2QLOTHW \[page 328\]](#)

[Technical Data Provider DB2QLOTRS \[page 328\]](#)

[Technical Data Provider DB2QLRUR \[page 329\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

[What Is the Data Explorer? \[page 79\]](#)

13.1.1.3 Lock Escalations

Get an overview of potential problems with lock escalations for the selected period in the time selection.

When Do You Need This Data?

When DB2's IRLM (Inter System Resource Lock Manager), prevent using uncommitted data from the ongoing transaction and get information about multiple locking events on rows, pages, or tables.

Where Does the Data Come From?

This monitoring screen is based on the technical data provider *DB2LOCKESCAL*. For more information, see the technical data provider documentation.

Related Information

[Technical Data Provider DB2LOCKESCAL \[page 322\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

[What Is the Data Explorer? \[page 79\]](#)

13.1.1.4 Long Running Transactions

Get an overview of potential problems with long running transactions for the selected period in the time selection.

Where Does the Data Come From?

This monitoring screen is based on the technical data provider *DB2QLRUR*. For more information, see the technical data provider documentation.

What Else Do You Need to Know?

Long-running transactions are either long-running units of recovery (UR), which change at least one row or long-running read-only transactions.

For more information, see the *Ensure Optimal Db2 Settings* section in the *Database Administration Guide for IBM Db2 for z/OS* at https://help.sap.com/viewer/db2_administration_guide.

Related Information

[Technical Data Provider DB2QLRUR \[page 329\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

[What Is the Data Explorer? \[page 79\]](#)

13.1.1.5 Deadlocks

13.1.1.5.1 Deadlocks Overview

Get an overview of the current lock situation in your Db2 for z/OS database during a specific period, including the locked objects and the locking transactions. The table is empty when the system is not affected by a current lock situation.

When Do You Need This Data?

- To find out which transactions are locked or are blocking the others
- To identify transactions that show or cause high lock-wait times

Where Does the Data Come From?

This monitoring screen is based on the technical data providers *DB2QLODLK*, *DB2QLODRS*, and data provider *Deadlocks*. For more information, see the technical data provider and data provider documentation.

What Else Do You Need to Know?

If the SAP Db2 for z/OS database has high response times, this might be caused by high lock-wait times. Deadlocks are the result of parallel system workload. Deadlocks' wait times might increase if locks on database resources such as tables or table rows are held long.

Related Information

Data Provider [Deadlocks Overview \[page 316\]](#)

Technical Data Provider DB2QLODLK [\[page 326\]](#)

Technical Data Provider DB2QLODRS [\[page 327\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

[What Is the Data Explorer? \[page 79\]](#)

13.1.1.5.2 Holders

Get an overview of holders for the selected period in the time selection.

When Do You Need This Data?

To identify statements that are not yet committed and find a reason why.

Where Does the Data Come From?

This monitoring screen is based on the data provider *Deadlocks*. For more information, see the data provider documentation.

Related Information

Data Provider [Deadlocks Overview \[page 316\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

[What Is the Data Explorer? \[page 79\]](#)

13.1.1.5.3 Resource

Get an overview of the resources for the selected period in the time selection.

When Do You Need This Data?

To identify all statements involved in the deadlock and on which resource.

Where Does the Data Come From?

This monitoring screen is based on the data provider *Deadlocks*. For more information, see the data provider documentation.

Related Information

Data Provider [Deadlocks Overview \[page 316\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

[What Is the Data Explorer? \[page 79\]](#)

13.1.1.5.4 Victims

Get an overview of victims for the selected period in the time selection.

When Do You Need This Data?

To identify all the statements waiting to be executed in the deadlock having the lowest priority on execution.

Where Does the Data Come From?

This monitoring screen is based on the data provider *Deadlocks*. For more information, see the data provider documentation.

Related Information

Data Provider [Deadlocks Overview \[page 316\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

[What Is the Data Explorer? \[page 79\]](#)

13.1.1.5.5 Waiters

Get an overview of waiters for the selected period in the time selection.

When Do You Need This Data?

To identify all statements which are waiting to be executed in the deadlock.

Where Does the Data Come From?

This monitoring screen is based on the data provider *Deadlocks*. For more information, see the data provider documentation.

Related Information

Data Provider [Deadlocks Overview \[page 316\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

[What Is the Data Explorer? \[page 79\]](#)

13.1.1.6 Timeouts

13.1.1.6.1 Blockers

Get an overview of blockers for the selected period in the time selection.

When Do You Need This Data?

To find SQL statement which was a cause of timeout in the system.

Where Does the Data Come From?

This monitoring screen is based on the data provider *Timeouts*. For more information, see the data provider documentation.

Related Information

Data Provider [Timeouts Overview](#) [page 318]

[What Are Data Providers?](#) [page 80]

[What Are Technical Data Providers?](#) [page 81]

[What Is the Data Explorer?](#) [page 79]

13.1.1.6.2 Resource

Get an overview of the resources for the selected period in the time selection.

When Do You Need This Data?

To identify resources involved in SQL timeout.

Where Does the Data Come From?

This monitoring screen is based on the data provider *Timeouts*. For more information, see the data provider documentation.

Related Information

Data Provider [Timeouts Overview \[page 318\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

[What Is the Data Explorer? \[page 79\]](#)

13.1.1.6.3 Timeouts Overview

Get an overview of timeouts for the selected period in the time selection.

When Do You Need This Data?

- To identify the root cause of performance problems
- To optimize your database system
- To analyze workload situations, for example, if there are any contentions during processing or waiting times.
- To identify the starting point with performance optimization.

Where Does the Data Come From?

This monitoring screen is based on the technical data providers *DB2QLOTHW* and *DB2QLOTRS*. For more information, see the technical data provider documentation.

Related Information

Technical Data Provider [DB2QLOTHW \[page 328\]](#)

Technical Data Provider [DB2QLOTRS \[page 328\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

[What Is the Data Explorer? \[page 79\]](#)

13.1.1.6.4 Victims

Get an overview of victims for the selected period in the time selection.

When Do You Need This Data?

To find SQL statements that were not executed because of the system's timeout.

Where Does the Data Come From?

This monitoring screen is based on the data provider *Timeouts*. For more information, see the data provider documentation.

Related Information

Data Provider [Timeouts Overview \[page 318\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

[What Is the Data Explorer? \[page 79\]](#)

13.1.2 Bufferpools

13.1.2.1 Activity

Provides information about the main memory that the database manager has allocated.

Where Does the Data Come From?

The monitoring screen data is based on the technical data provider *DB2QBSTT*. For more information, see the technical data provider documentation.

What Else Do You Need to Know?

Data on the monitoring screen use the record trace block IFCID 002 - Buffer Pool Activity from the Db2 Instrumentation Facility Component Identifier (IFCID).

For more information, search for *IFCID 002* in the IBM Documentation at <https://www.ibm.com/docs/en> .

Related Information

[Technical Data Provider DB2QBSTT \[page 323\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

[What Is the Data Explorer? \[page 79\]](#)

13.1.2.2 Configuration

Provides information about the main memory that the database manager has allocated.

Where Does the Data Come From?

The monitoring screen data is based on the technical data provider *DB2QDBPT1*. For more information, see the technical data provider documentation.

What Else Do You Need to Know?

Data on the monitoring screen use the record trace block IFCID 002 - Buffer Pool Activity from the Db2 Instrumentation Facility Component Identifier (IFCID).

For more information, search for *IFCID 202* in the IBM Documentation at <https://www.ibm.com/docs/en> .

Related Information

[Technical Data Provider DB2QDBPT1 \[page 325\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

[What Is the Data Explorer? \[page 79\]](#)

13.1.2.3 Defaults

Provides information about the main memory that the database manager has allocated and provides all the necessary defaults values of buffer pools.

Where Does the Data Come From?

The monitoring screen data is based on the technical data provider *SYSIBM.DB2_SYSPARM* and use the *SYSPROC.ADMIN_INFO_SYSPARM* store procedure.

For more information, see the technical data provider documentation.

Related Information

[Technical Data Provider SYSIBM.DB2_SYSPARM \[page 333\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

[What Is the Data Explorer? \[page 79\]](#)

13.1.2.4 Pool Details

13.1.2.4.1 General

In Bufferpools Activity, provides operation-specific details for each buffer pool name.

Where Does the Data Come From?

This monitoring screen is based on the technical data providers *DB2QDBPT1* and *DB2QBSTT*. For more information, see the technical data provider documentation.

What Else Do You Need to Know?

Data on the monitoring screen use the record trace block IFCID 002 - Buffer Pool Activity from the Db2 Instrumentation Facility Component Identifier (IFCID).

For more information, search for *IFCID 002* in the IBM Documentation at <https://www.ibm.com/docs/en> .

Related Information

[Technical Data Provider DB2QDBPT1 \[page 325\]](#)

[Technical Data Provider DB2QBSTT \[page 323\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

[What Is the Data Explorer? \[page 79\]](#)

13.1.2.4.2 Query

In Bufferpools Activity, provides operation-specific details for each buffer pool name.

Where Does the Data Come From?

This monitoring screen is based on the technical data provider *DB2QBSTT*. For more information, see the technical data provider documentation.

What Else Do You Need to Know?

Data on the monitoring screen use the record trace block IFCID 002 - Buffer Pool Activity from the Db2 Instrumentation Facility Component Identifier (IFCID).

For more information, search for *IFCID 002* in the IBM Documentation at <https://www.ibm.com/docs/en> .

Related Information

[Technical Data Provider DB2QBSTT \[page 323\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

[What Is the Data Explorer? \[page 79\]](#)

13.1.2.4.3 Read

In Bufferpools Activity, provides operation-specific details for each buffer pool name.

Where Does the Data Come From?

This monitoring screen is based on the technical data provider *DB2QBSTT*. For more information, see the technical data provider documentation.

What Else Do You Need to Know?

Data on the monitoring screen use the record trace block IFCID 002 - Buffer Pool Activity from the Db2 Instrumentation Facility Component Identifier (IFCID).

For more information, search for *IFCID 002* in the IBM Documentation at <https://www.ibm.com/docs/en> .

Related Information

[Technical Data Provider DB2QBSTT \[page 323\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

[What Is the Data Explorer? \[page 79\]](#)

13.1.2.4.4 Sort

In Bufferpools Activity, provides operation-specific details for each buffer pool name.

Where Does the Data Come From?

This monitoring screen is based on the technical data provider *DB2QBSTT*. For more information, see the technical data provider documentation.

What Else Do You Need to Know?

Data on the monitoring screen use the record trace block IFCID 002 - Buffer Pool Activity from the Db2 Instrumentation Facility Component Identifier (IFCID).

For more information, search for *IFCID 002* in the IBM Documentation at <https://www.ibm.com/docs/en> .

Related Information

[Technical Data Provider DB2QBSTT \[page 323\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

[What Is the Data Explorer? \[page 79\]](#)

13.1.2.4.5 Write

In Bufferpools Activity, provides operation-specific details for each buffer pool name.

Where Does the Data Come From?

This monitoring screen is based on the technical data provider *DB2QBSTT*. For more information, see the technical data provider documentation.

What Else Do You Need to Know?

Data on the monitoring screen use the record trace block IFCID 002 - Buffer Pool Activity from the Db2 Instrumentation Facility Component Identifier (IFCID).

For more information, search for *IFCID 002* in the IBM Documentation at <https://www.ibm.com/docs/en> .

Related Information

[Technical Data Provider DB2QBSTT \[page 323\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

[What Is the Data Explorer? \[page 79\]](#)

13.1.3 Configuration

13.1.3.1 Installation Parameters

Get an overview of the current installation parameters settings in your Db2 database.

Where Does the Data Come From?

This monitoring screen is based on the data provider *DB2 Installation Parameters*. For more information, see the data provider documentation.

Related Information

Data Provider [DB2 Installation Parameters \[page 315\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

[What Is the Data Explorer? \[page 79\]](#)

13.1.4 DB2 Subsystem Activity

13.1.4.1 Data Sharing Locking

Get an overview of the overall Db2 database performance.

Where Does the Data Come From?

The monitoring screen data is based on the technical data provider *DB2QSYSW1*. For more information, see the technical data provider documentation.

What Else Do You Need to Know?

Data on the monitoring screen the record trace block IFCID 002 - Data Sharing Locking Data from the Db2 Instrumentation Facility Component Identifier (IFCID).

For more information, search for *IFCID 002* in the IBM Documentation at <https://www.ibm.com/docs/en> .

Related Information

[Technical Data Provider DB2QSYSW1 \[page 329\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

[What Is the Data Explorer? \[page 79\]](#)

13.1.4.2 Locking Activity

Get an overview of the overall Db2 database performance.

Where Does the Data Come From?

The monitoring screen data is based on the technical data provider *DB2QSYSW2*. For more information, see the technical data provider documentation.

What Else Do You Need to Know?

Data on the monitoring screen use the record trace block IFCID 002 - Service Controller Data from the Db2 Instrumentation Facility Component Identifier (IFCID).

For more information, search for *IFCID 002* in the IBM Documentation at <https://www.ibm.com/docs/en> .

Related Information

[Technical Data Provider DB2QSYSW2 \[page 330\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

[What Is the Data Explorer? \[page 79\]](#)

13.1.4.3 Logging Activity

Get an overview of the overall Db2 database performance.

Where Does the Data Come From?

The monitoring screen data is based on the technical data provider *DB2QSYSW1*. For more information, see the technical data provider documentation.

What Else Do You Need to Know?

Data on the monitoring screen use the record trace block IFCID 001 - Log Manager Data from the Db2 Instrumentation Facility Component Identifier (IFCID).

For more information, search for *IFCID 001* in the IBM Documentation at <https://www.ibm.com/docs/en> .

Related Information

[Technical Data Provider DB2QSYSW1 \[page 329\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

[What Is the Data Explorer? \[page 79\]](#)

13.1.4.4 Open/Close Activity

Get an overview of the overall Db2 database performance.

Where Does the Data Come From?

The monitoring screen data is based on the technical data provider *DB2QSYSW2*. For more information, see the technical data provider documentation.

What Else Do You Need to Know?

Data on the monitoring screen use the record trace block IFCID 002 - Service Controller Data from the Db2 Instrumentation Facility Component Identifier (IFCID).

For more information, search for *IFCID 002* in the IBM Documentation at <https://www.ibm.com/docs/en> .

Related Information

[Technical Data Provider DB2QSYW2 \[page 330\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

[What Is the Data Explorer? \[page 79\]](#)

13.1.4.5 RID List Processing

Get an overview of the overall Db2 database performance.

Where Does the Data Come From?

The monitoring screen data is based on the technical data provider *DB2QSYW2*. For more information, see the technical data provider documentation.

What Else Do You Need to Know?

Data on the monitoring screen use the record trace block IFCID 002 - DB2 Subsystem Services from the Db2 Instrumentation Facility Component Identifier (IFCID).

For more information, search for *IFCID 002* in the IBM Documentation at <https://www.ibm.com/docs/en> .

Related Information

[Technical Data Provider DB2QSYW2 \[page 330\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

[What Is the Data Explorer? \[page 79\]](#)

13.1.4.6 Subsystem Services

Get an overview of the overall Db2 database performance.

Where Does the Data Come From?

The monitoring screen data is based on the technical data provider *DB2QSYSW2*. For more information, see the technical data provider documentation.

What Else Do You Need to Know?

Data on the monitoring screen use the record trace block IFCID 002 - DB2 Subsystem Services from the Db2 Instrumentation Facility Component Identifier (IFCID).

For more information, search for *IFCID 002* in the IBM Documentation at <https://www.ibm.com/docs/en> .

Related Information

[Technical Data Provider DB2QSYSW2 \[page 330\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

[What Is the Data Explorer? \[page 79\]](#)

13.1.5 DB2 Subsystem Summary

13.1.5.1 DB2 Resources

Get an overview of the overall Db2 database performance.

Where Does the Data Come From?

The monitoring screen data is based on the technical data providers *DB2QBSTT*, *DB2QSYSW1*, and *DB2QSYSW2*. For more information, see the technical data provider documentation.

What Else Do You Need to Know?

Data on the monitoring screen use the record trace blocks IFCID 001 - Log Manager Data, IFCID 002 - Buffer Pool Activity, and IFCID 002 – EDM Pool Data from the Db2 Instrumentation Facility Component Identifier (IFCID).

For more information, search for *IFCID 001* and *IFCID 002* in the IBM Documentation at <https://www.ibm.com/docs/en> .

Related Information

[Technical Data Provider DB2QBSTT \[page 323\]](#)

[Technical Data Provider DB2QSYSW1 \[page 329\]](#)

[Technical Data Provider DB2QSYSW2 \[page 330\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

[What Is the Data Explorer? \[page 79\]](#)

13.1.5.2 Global Times

Get an overview of the overall Db2 database performance.

Where Does the Data Come From?

The monitoring screen data is based on the technical data provider *DB2QGLOTIME*. For more information, see the technical data provider documentation.

What Else Do You Need to Know?

Data on the monitoring screen use the record trace blocks IFCID 003 - Instrumentation Accounting Data and IFCID 147 - Instrumentation Accounting Data Overflow from the Db2 Instrumentation Facility Component Identifier (IFCID).

For more information, search for *IFCID 003* and *IFCID 147* in the IBM Documentation at <https://www.ibm.com/docs/en> .

Related Information

[Technical Data Provider DB2QGLOTIME \[page 326\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

[What Is the Data Explorer? \[page 79\]](#)

13.1.5.3 Hardware Resources

Get an overview of the overall Db2 database performance.

Where Does the Data Come From?

The monitoring screen data is based on the technical data providers *DB2QGLOTIME* and *DB2QW0225*. For more information, see the technical data provider documentation.

What Else Do You Need to Know?

Data on the monitoring screen use the record trace blocks IFCID 003 - Instrumentation Accounting Data and IFCID 225 - Address Space Summary and from the Db2 Instrumentation Facility Component Identifier (IFCID).

For more information, search for *IFCID 003* and *IFCID 225* in the IBM Documentation at <https://www.ibm.com/docs/en> .

Related Information

[Technical Data Provider DB2QGLOTIME \[page 326\]](#)

[Technical Data Provider DB2QW0225 \[page 333\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

[What Is the Data Explorer? \[page 79\]](#)

13.1.5.4 Top SQL Statements

Get an overview of the overall Db2 database performance.

Where Does the Data Come From?

The monitoring screen data is based on the technical data provider *DB2QDSCHIST*. For more information, see the technical data provider documentation.

What Else Do You Need to Know?

Data on the monitoring screen use the record trace blocks IFCID 316 - SQL Statement Statistics and IFCID 317 - SQL Statement String from the Db2 Instrumentation Facility Component Identifier (IFCID).

For more information, search for *IFCID 316* and *IFCID 317* in the IBM Documentation at <https://www.ibm.com/docs/en> .

Related Information

[Technical Data Provider DB2QDSCHIST \[page 325\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

[What Is the Data Explorer? \[page 79\]](#)

13.1.6 Jobs

13.1.6.1 DB2 Table Profiles

Provides additional statistics for a specific table necessary for selecting the optimal access path.

When Do You Need This Data?

Each table has default profiles, which collect statistics by default with SAP automated Runstats. You can define a custom RUNSTATS profile in the Db2 catalog table `SYSIBM.SYSTABLES_PROFILES` for the specific table, if needed.

Where Does the Data Come From?

The monitoring screen data is based on the technical data provider `SYSIBM.SYSTABLES` and `SYSIBM.SYSTABLES_PROFILES`. For more information, see the technical data provider documentation.

What Else Do You Need to Know?

For more information, check out the `RUNSTATS` command in the IBM Documentation at <https://www.ibm.com/docs/en> .

Related Information

[Technical Data Provider SYSIBM.SYSTABLES \[page 337\]](#)

[Technical Data Provider SYSIBM.SYSTABLES_PROFILES \[page 337\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

[What Is the Data Explorer? \[page 79\]](#)

13.1.7 SQL

13.1.7.1 Statement Cache Current View

13.1.7.1.1 Overview

Get a current overview of the SQL statement cache in your Db2 for z/OS database.

When Do You Need This Data?

- To find out how many SQL statements have been executed in your system and how much time they have spent running
- To find out which SQL statements have been running
- To find out which SQL statements are executed
- To identify users executing specific SQL statements
- To locate the application source code that issues the SQL statement

Where Does the Data Come From?

The monitoring screen data is based on the technical data provider *DB2QTHDDSC*. For more information, see the technical data provider documentation.

What Else Do You Need to Know?

Data on the monitoring screen use the record trace blocks IFCID 316 - SQL Statement Statistics and IFCID 317 – SQL Statement String from the Db2 Instrumentation Facility Component Identifier (IFCID).

For more information, search for *IFCID 316* and *IFCID 317* in the IBM Documentation at <https://www.ibm.com/docs/en> .

Related Information

[Technical Data Provider DB2QTHDDSC \[page 330\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

[What Is the Data Explorer? \[page 79\]](#)

13.1.7.2 Statement Cache History

13.1.7.2.1 History Average Statistic

Provides data about the average time spent on specific SQL statements

Where Does the Data Come From?

The monitoring screen data is based on the technical data provider *DB2QDSCHIST* and *Statement Cache History* data provider. For more information, see the technical data provider and data provider documentation.

Related Information

Data Provider [Statement Cache History \[page 317\]](#)

Technical Data Provider [DB2QDSCHIST \[page 325\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

[What Is the Data Explorer? \[page 79\]](#)

13.1.7.2.2 History Bind Parameter

Shows date for dynamic parameters of SQL statements.

Where Does the Data Come From?

The monitoring screen data is based on the technical data provider *DB2QDSCHIST* and *Statement Cache History* data provider. For more information, see the technical data provider and data provider documentation.

Related Information

Data Provider [Statement Cache History \[page 317\]](#)

Technical Data Provider [DB2QDSCHIST \[page 325\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

[What Is the Data Explorer? \[page 79\]](#)

13.1.7.2.3 History Cache Parameter

Shows date for dynamic parameters of SQL statements.

Where Does the Data Come From?

The monitoring screen data is based on the technical data provider *DB2QDSCHIST* and *Statement Cache History* data provider. For more information, see the technical data provider and data provider documentation.

Related Information

Data Provider [Statement Cache History \[page 317\]](#)

Technical Data Provider [DB2QDSCHIST \[page 325\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

[What Is the Data Explorer? \[page 79\]](#)

13.1.7.2.4 History Elapsed Times

Provides the data about the total time spent on the specific SQL statement

Where Does the Data Come From?

The monitoring screen data is based on the technical data provider *DB2QDSCHIST*. For more information, see the technical data provider documentation.

What Else Do You Need to Know?

Data on the monitoring screen use the record trace block IFCID 316 - SQL Statement Statistics from the Db2 Instrumentation Facility Component Identifier (IFCID).

For more information, search for *IFCID 316* in the IBM Documentation at <https://www.ibm.com/docs/en>  .

Related Information

[Technical Data Provider DB2QDSCHIST \[page 325\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

[What Is the Data Explorer? \[page 79\]](#)

13.1.7.2.5 History Overview

Get a current overview of the SQL statement cache in your Db2 for z/OS database.

When Do You Need This Data?

Get a current overview of the SQL statement cache in your Db2 for z/OS database.

- To find out how many SQL statements have been executed in your system and how much time they have spent running
- To find out which SQL statements have been running for a long time
- To find out which SQL statements are executed frequently
- To identify users executing specific SQL statements
- To locate the application source code that issues the SQL statement
- To monitor the capability for all prepared SQL Statements which are in the global SQL Statement Cache

Where Does the Data Come From?

The monitoring screen data is based on the technical data provider *DB2QDSCHIST*. For more information, see the technical data provider documentation.

What Else Do You Need to Know?

Data on the monitoring screen use the record trace blocks *IFCID 316* - SQL Statement Statistics and *IFCID 317* – SQL Statement String from the Db2 Instrumentation Facility Component Identifier (IFCID).

For more information, search for *IFCID 316* and *IFCID 317* in the IBM Documentation at <https://www.ibm.com/docs/en>  .

Related Information

[Technical Data Provider DB2QDSCHIST \[page 325\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

[What Is the Data Explorer? \[page 79\]](#)

13.1.7.2.6 History Total Statistic

Provide a snapshot for a selected period of time of the system statement cache.

Where Does the Data Come From?

The monitoring screen data is based on the technical data provider *DB2QDSCHIST* and *Statement Cache History* data provider. For more information, see the technical data provider and data provider documentation.

Related Information

Data Provider [Statement Cache History \[page 317\]](#)

[Technical Data Provider DB2QDSCHIST \[page 325\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

[What Is the Data Explorer? \[page 79\]](#)

13.1.8 Space

13.1.8.1 Indexes

13.1.8.1.1 Details

Various properties and detailed information about the index

Where Does the Data Come From?

The monitoring screen data is based on the technical data providers *DB2QDB02_OBJECT* and *SYSIBM.SYSINDEXES*. For more information, see the technical data provider documentation.

Related Information

[Technical Data Provider DB2QDB02_OBJECT \[page 323\]](#)

[Technical Data Provider SYSIBM.SYSINDEXES \[page 335\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

[What Is the Data Explorer? \[page 79\]](#)

13.1.8.1.2 History

Show the history of an index for the selected period on the time slider.

Where Does the Data Come From?

The monitoring screen data is based on the technical data providers *DB2QDB02_OBJECT* and *SYSIBM.SYSINDEXES*. For more information, see the technical data provider documentation.

Related Information

[Technical Data Provider DB2QDB02_OBJECT \[page 323\]](#)

[Technical Data Provider SYSIBM.SYSINDEXES \[page 335\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

[What Is the Data Explorer? \[page 79\]](#)

13.1.8.1.3 Indexes Overview

The monitor provides data about space allocated for specified indexes, plus other details that might help analyze.

When Do You Need This Data?

When you need to find a specific index from all presented in the system

Where Does the Data Come From?

The monitoring screen data is based on the technical data providers *DB2QDB02_OBJECT* and *SYSIBM.SYSINDEXES* and *SYSIBM.SYSINDEXAPART*. For more information, see the technical data provider documentation.

What Else Do You Need to Know?

If you want to analyze the index in detail, select it and follow additional statistics, such as Details, History, and Structure.

Related Information

[Technical Data Provider DB2QDB02_OBJECT \[page 323\]](#)

[Technical Data Provider SYSIBM.SYSINDEXES \[page 335\]](#)

[Technical Data Provider SYSIBM.SYSINDEXPART \[page 335\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

[What Is the Data Explorer? \[page 79\]](#)

13.1.8.1.4 Structure

The monitor provides the fields of an index and the database specifics of each area.

Where Does the Data Come From?

The monitoring screen data is based on the technical data provider *SYSIBM.SYSINDEXES*. For more information, see the technical data provider documentation.

Related Information

[Technical Data Provider SYSIBM.SYSINDEXES \[page 335\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

[What Is the Data Explorer? \[page 79\]](#)

13.1.8.2 KPI

13.1.8.2.1 Fastest Growing Tables

Shows average growth over daily history for the configured time range in the time slider

When Do You Need This Data?

To check the database growth.

Where Does the Data Come From?

The monitoring screen data is based on the technical data provider *DB2QDB02_OBJECT*. For more information, see the technical data provider documentation.

Related Information

[Technical Data Provider DB2QDB02_OBJECT \[page 323\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

[What Is the Data Explorer? \[page 79\]](#)

13.1.8.2.2 Index Limits

Shows the limits of database indexes

When Do You Need This Data?

To check the database growth.

Where Does the Data Come From?

The monitoring screen data is based on the technical data providers *DB2QDB02_OBJECT*, *SYSIBM.SYSINDEXES*, *SYSIBM.SYSINDEXPART*, *SYSIBM.SYSTABLES*, and *SYSIBM.SYSTABLESPACE*. For more information, see the technical data provider documentation.

What Else Do You Need to Know?

The last column (Used Space) indicates the percentage of space that has been required per index partition.

Related Information

[Technical Data Provider DB2QDB02_OBJECT \[page 323\]](#)

[Technical Data Provider SYSIBM.SYSINDEXES \[page 335\]](#)

[Technical Data Provider SYSIBM.SYSINDEXPART \[page 335\]](#)

[Technical Data Provider SYSIBM.SYSTABLES \[page 337\]](#)

[Technical Data Provider SYSIBM.SYSTABLESPACE \[page 338\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

[What Is the Data Explorer? \[page 79\]](#)

13.1.8.2.3 Largest Freespace

Shows the remaining free space in the database

When Do You Need This Data?

To check the database growth.

Where Does the Data Come From?

The monitoring screen data is based on the technical data provider *DB2QDB02_OBJECT*. For more information, see the technical data provider documentation.

What Else Do You Need to Know?

The column (% Free) indicates the percentage of space that has not been used.

Related Information

[Technical Data Provider DB2QDB02_OBJECT \[page 323\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

[What Is the Data Explorer? \[page 79\]](#)

13.1.8.2.4 Largest Indexes

Shows the largest indexes in the database

When Do You Need This Data?

To check the database growth.

Where Does the Data Come From?

The monitoring screen data is based on the technical data provider *DB2QDB02_OBJECT*. For more information, see the technical data provider documentation.

Related Information

[Technical Data Provider DB2QDB02_OBJECT \[page 323\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

[What Is the Data Explorer? \[page 79\]](#)

13.1.8.2.5 Largest Multi-Table TS

Shows the largest tablespaces in the database

When Do You Need This Data?

To check the database growth.

Where Does the Data Come From?

The monitoring screen data is based on the technical data provider *DB2QDB02_OBJECT*. For more information, see the technical data provider documentation.

Related Information

[Technical Data Provider DB2QDB02_OBJECT \[page 323\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

[What Is the Data Explorer? \[page 79\]](#)

13.1.8.2.6 Largest Tables

Shows the largest tables in the database

When Do You Need This Data?

To check the database growth.

Where Does the Data Come From?

The monitoring screen data is based on the technical data provider *DB2QDB02_OBJECT*. For more information, see the technical data provider documentation.

Related Information

[Technical Data Provider DB2QDB02_OBJECT \[page 323\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

[What Is the Data Explorer? \[page 79\]](#)

13.1.8.2.7 Partition Limits

Shows the limits for the partitions in the database

When Do You Need This Data?

To check the database growth.

Where Does the Data Come From?

The monitoring screen data is based on the technical data provider *SYSIBM.SYSTABLESPACE*. For more information, see the technical data provider documentation.

Related Information

[Technical Data Provider SYSIBM.SYSTABLESPACE \[page 338\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

[What Is the Data Explorer? \[page 79\]](#)

13.1.8.2.8 Range Partitioned Tables

When Do You Need This Data?

To check the database growth.

Where Does the Data Come From?

The monitoring screen data is based on the technical data providers *SYSIBM.SYSTABLEPART*, *SYSIBM.SYSTABLES*, *SYSIBM.SYSTABLESPACE*, *SYSIBM.SYSTABLESPACESTATS*, and *DB2QDBO2_OBJECT*. For more information, see the technical data provider documentation.

Related Information

[Technical Data Provider DB2QDB02_OBJECT \[page 323\]](#)
[Technical Data Provider SYSIBM.SYSTABLES \[page 337\]](#)
[Technical Data Provider SYSIBM.SYSTABLEPART \[page 336\]](#)
[Technical Data Provider SYSIBM.SYSTABLESPACE \[page 338\]](#)
[Technical Data Provider SYSIBM.SYSTABLESPACESTATS \[page 338\]](#)
[What Are Data Providers? \[page 80\]](#)
[What Are Technical Data Providers? \[page 81\]](#)
[What Is the Data Explorer? \[page 79\]](#)

13.1.8.2.9 Tablespace Limits

Shows the tablespace limits in the database

When Do You Need This Data?

To check the database growth.

Where Does the Data Come From?

The monitoring screen data is based on the technical data providers *SYSIBM.SYSTABLEPART*, *SYSIBM.SYSTABLESPACE*, and *DB2QDB02_OBJECT*. For more information, see the technical data provider documentation.

What Else Do You Need to Know?

The last column (% Used) indicates the percentage of space that has been required per tablespace partition.

Related Information

[Technical Data Provider DB2QDB02_OBJECT \[page 323\]](#)
[Technical Data Provider SYSIBM.SYSTABLEPART \[page 336\]](#)

[Technical Data Provider SYSIBM.SYSTABLESPACE \[page 338\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

[What Is the Data Explorer? \[page 79\]](#)

13.1.8.3 SAP System Overview

13.1.8.3.1 DB2 Subsystem

Provides a complete cross-system overview of the database systems being monitored

Where Does the Data Come From?

The monitoring screen data is based on the technical data provider *DB2QDB02_SCHEMA*. For more information, see the technical data provider documentation.

What Else Do You Need to Know?

Shows technical DB2 limits and object sizes in the database

Related Information

[Technical Data Provider DB2QDB02_SCHEMA \[page 324\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

[What Is the Data Explorer? \[page 79\]](#)

13.1.8.3.2 Details

Overview information for tables, tablespaces, and indexes.

Where Does the Data Come From?

The monitoring screen data is based on the technical data provider *DB2QDB02_SCHEMA*. For more information, see the technical data provider documentation.

Related Information

[Technical Data Provider DB2QDB02_SCHEMA \[page 324\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

[What Is the Data Explorer? \[page 79\]](#)

13.1.8.3.3 Disk Space

This monitoring screen provides information about the space used by tables in the database.

Where Does the Data Come From?

The monitoring screen data is based on the technical data provider *DB2QDB02_SMS*. For more information, see the technical data provider documentation.

Related Information

[Technical Data Provider DB2QDB02_SMS \[page 324\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

[What Is the Data Explorer? \[page 79\]](#)

13.1.8.3.4 History

Shows the history of database space information for the selected period of time on the time slider

Where Does the Data Come From?

The monitoring screen data is based on the technical data provider *DB2QDB02_SCHEMA*. For more information, see the technical data provider documentation.

Related Information

[Technical Data Provider DB2QDB02_SCHEMA \[page 324\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

[What Is the Data Explorer? \[page 79\]](#)

13.1.8.4 Tables

13.1.8.4.1 Base Table

Mark the required partition in the list of partitions for any additional analysis.

Where Does the Data Come From?

The monitoring screen data is based on the technical data provider *SYSIBM.SYSAUXRELS*. For more information, see the technical data provider documentation.

Related Information

[Technical Data Provider SYSIBM.SYSAUXRELS \[page 334\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

[What Is the Data Explorer? \[page 79\]](#)

13.1.8.4.2 Indexes

Provides information about all indexes that exist for this table in the database.

When Do You Need This Data?

The monitoring screen data is based on the technical data provider *SYSIBM.SYSINDEXES*. For more information, see the technical data provider documentation.

Related Information

[Technical Data Provider SYSIBM.SYSINDEXES \[page 335\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

[What Is the Data Explorer? \[page 79\]](#)

13.1.8.4.3 Structure

Provide information about the fields of a table and the database specifics of each area.

Where Does the Data Come From?

The monitoring screen data is based on the technical data provider *SYSIBM.SYSCOLUMNS*. For more information, see the technical data provider documentation.

Related Information

[Technical Data Provider SYSIBM.SYSCOLUMNS \[page 334\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

[What Is the Data Explorer? \[page 79\]](#)

13.1.8.4.4 Tables Overview

The monitor provides data about space allocated for specified tables and other details that might help analyze.

When Do You Need This Data?

When you need to find a specific table from all presented in the system, it gives a basic overview of the parameters.

Where Does the Data Come From?

The monitoring screen data is based on the technical data providers *DB2QDB02_OBJECT*, *SYSIBM.SYSTABLES*, and *SYSIBM.SYSTABLESPACE*. For more information, see the technical data provider documentation.

What Else Do You Need to Know?

If you want to analyze the table in detail, select it and follow additional statistics, such as Details, History, Indexes, and Structure.

Related Information

[Technical Data Provider DB2QDB02_OBJECT \[page 323\]](#)

[Technical Data Provider SYSIBM.SYSTABLES \[page 337\]](#)

[Technical Data Provider SYSIBM.SYSTABLESPACE \[page 338\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

[What Is the Data Explorer? \[page 79\]](#)

13.1.8.4.5 Tables Details

Various properties for tables are shown, such as the time of creation, the last change, allocated space in the database, etc.

Where Does the Data Come From?

The monitoring screen data is based on the technical data providers *DB2QDB02_OBJECT* and *SYSIBM.SYSTABLES*. For more information, see the technical data provider documentation.

Related Information

[Technical Data Provider DB2QDB02_OBJECT \[page 323\]](#)

[Technical Data Provider SYSIBM.SYSTABLES \[page 337\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

[What Is the Data Explorer? \[page 79\]](#)

13.1.8.4.6 Tablespace History

Provide information about the history of the tablespace for the selected period on the time slider.

Where Does the Data Come From?

The monitoring screen data is based on the technical data provider *DB2QDB02_OBJECT*. For more information, see the technical data provider documentation.

Related Information

[Technical Data Provider DB2QDB02_OBJECT \[page 323\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

[What Is the Data Explorer? \[page 79\]](#)

13.1.8.5 Tablespaces

13.1.8.5.1 Details

History of a tablespace for the selected period of time on time slider

Where Does the Data Come From?

The monitoring screen data is based on the technical data providers *DB2QDB02_OBJECT* and *SYSIBM.SYSTABLESPACE*. For more information, see the technical data provider documentation.

Related Information

[Technical Data Provider DB2QDB02_OBJECT \[page 323\]](#)

[Technical Data Provider SYSIBM.SYSTABLESPACE \[page 338\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

[What Is the Data Explorer? \[page 79\]](#)

13.1.8.5.2 History

History of a tablespace for the selected period of time on time slider

Where Does the Data Come From?

The monitoring screen data is based on the technical data providers *DB2QDB02_OBJECT* and *SYSIBM.SYSTABLESPACE*. For more information, see the technical data provider documentation.

Related Information

[Technical Data Provider DB2QDB02_OBJECT \[page 323\]](#)

[Technical Data Provider SYSIBM.SYSTABLESPACE \[page 338\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

[What Is the Data Explorer? \[page 79\]](#)

13.1.8.5.3 Partition Details

Mark the required partition in the list of partitions for any additional analysis.

Where Does the Data Come From?

The monitoring screen data is based on the technical data providers *DB2QDB02_OBJECT* and *SYSIBM.SYSTABLESPACE*. For more information, see the technical data provider documentation.

Related Information

[Technical Data Provider DB2QDB02_OBJECT \[page 323\]](#)

[Technical Data Provider SYSIBM.SYSTABLESPACE \[page 338\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

[What Is the Data Explorer? \[page 79\]](#)

13.1.8.5.4 Partitions

Mark the required partition in the list of partitions for any additional analysis.

Where Does the Data Come From?

The monitoring screen data is based on the technical data providers *DB2QDB02_OBJECT* and *SYSIBM.SYSTABLESPACE*. For more information, see the technical data provider documentation.

Related Information

[Technical Data Provider DB2QDB02_OBJECT \[page 323\]](#)

[Technical Data Provider SYSIBM.SYSTABLESPACE \[page 338\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

[What Is the Data Explorer? \[page 79\]](#)

13.1.8.5.5 Tables

Tables of a tablespace and the table creator of each table

Where Does the Data Come From?

The monitoring screen data is based on the technical data providers *DB2QDB02_OBJECT* and *SYSIBM.SYSTABLESPACE*. For more information, see the technical data provider documentation.

Related Information

[Technical Data Provider DB2QDB02_OBJECT \[page 323\]](#)

[Technical Data Provider SYSIBM.SYSTABLESPACE \[page 338\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

[What Is the Data Explorer? \[page 79\]](#)

13.1.8.5.6 Tablespaces overview

The monitor provides data about space allocated for specified tablespaces, plus other details that might help analyze.

When Do You Need This Data?

When you need to find specific tablespace from all presented in the system

Where Does the Data Come From?

The monitoring screen data is based on the technical data providers *DB2QDB02_OBJECT* and *SYSIBM.SYSTABLESPACE*. For more information, see the technical data provider documentation.

What Else Do You Need to Know?

If you want to analyze tablespace in detail, select it and follow additional statistics, such as Details, History, Partitions, and Tables.

Related Information

[Technical Data Provider DB2QDB02_OBJECT \[page 323\]](#)

[Technical Data Provider SYSIBM.SYSTABLESPACE \[page 338\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

[What Is the Data Explorer? \[page 79\]](#)

13.1.9 Storage

13.1.9.1 EDM Pool

Get an overview of the Environmental Descriptor Manager (EDM) pool activity connected with Db2.

Where Does the Data Come From?

The monitoring screen data is based on the technical data provider *DB2QSYSW2*. For more information, see the technical data provider documentation.

What Else Do You Need to Know?

You can view this information on the EDM Pool monitoring screen:

- EDM statistics lists total counts during the last sample period and the “per second” count for various tasks associated with the EDM Pool.

- DBD Pool lists information about database descriptors (DBDs).
- SKEL Pool utilization shows the EDM skeleton pool information and lists information about package search.
- Statement Pool lists information about the statements in the global cache and EDM Pool.

For more information, see the *Calculating EDM pool sizes* IBM documentation at <https://www.ibm.com/docs/en/db2-for-zos/12?topic=storage-calculating-edm-pool-sizes> .

Related Information

[Technical Data Provider DB2QSYW2 \[page 330\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

[What Is the Data Explorer? \[page 79\]](#)

13.1.9.2 Virtual and Real Storage

13.1.9.2.1 Details

13.1.9.2.1.1 Pool Details

Get an overview of the Environmental Descriptor Manager (EDM) pool activity connected with Db2.

Where Does the Data Come From?

The monitoring screen data is based on the technical data provider *DB2QW0225*. For more information, see the technical data provider documentation.

Related Information

[Technical Data Provider DB2QW0225 \[page 333\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

[What Is the Data Explorer? \[page 79\]](#)

13.1.9.2.1.2 Shared&Common Storage

Get information about the Db2 statistics.

Where Does the Data Come From?

The monitoring screen data is based on the technical data provider *DB2QW0225*. For more information, see the technical data provider documentation.

What Else Do You Need to Know?

Data on the monitoring screen use the record trace blocks IFCID 316 – SQL Statement Statistics and IFCID 317 – SQL Statement String from the Db2 Instrumentation Facility Component Identifier (IFCID).

For more information, search for *IFCID 316* and *IFCID 317* in the IBM Documentation at <https://www.ibm.com/docs/en> .

Related Information

[Technical Data Provider DB2QW0225 \[page 333\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

[What Is the Data Explorer? \[page 79\]](#)

13.1.9.2.1.3 Stat. Cache and xPROC Details

Get information about the Db2 statistics.

Where Does the Data Come From?

The monitoring screen data is based on the technical data provider *DB2QW0225*. For more information, see the technical data provider documentation.

What Else Do You Need to Know?

Data on the monitoring screen use the record trace blocks IFCID 316 – SQL Statement Statistics and IFCID 317 – SQL Statement String from the Db2 Instrumentation Facility Component Identifier (IFCID).

For more information, search for *IFCID 316* and *IFCID 317* in the IBM Documentation at <https://www.ibm.com/docs/en> .

Related Information

[Technical Data Provider DB2QW0225 \[page 333\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

[What Is the Data Explorer? \[page 79\]](#)

13.1.9.2.1.4 Thread Information

Get information about the Db2 statistics.

Where Does the Data Come From?

The monitoring screen data is based on the technical data provider *DB2QW0225*. For more information, see the technical data provider documentation.

What Else Do You Need to Know?

Data on the monitoring screen use the record trace blocks IFCID 316 – SQL Statement Statistics and IFCID 317 – SQL Statement String from the Db2 Instrumentation Facility Component Identifier (IFCID).

For more information, search for *IFCID 316* and *IFCID 317* in the IBM Documentation at <https://www.ibm.com/docs/en> .

Related Information

[Technical Data Provider DB2QW0225 \[page 333\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

[What Is the Data Explorer? \[page 79\]](#)

13.1.9.2.2 Summary

13.1.9.2.2.1 IRLM Pool Statistics

Get information about the Db2 statistics.

Where Does the Data Come From?

The monitoring screen data is based on the technical data provider *DB2QW0225*. For more information, see the technical data provider documentation.

What Else Do You Need to Know?

Data on the monitoring screen use the record trace blocks IFCID 316 – SQL Statement Statistics and IFCID 317 – SQL Statement String from the Db2 Instrumentation Facility Component Identifier (IFCID).

For more information, search for *IFCID 316* and *IFCID 317* in the IBM Documentation at <https://www.ibm.com/docs/en> .

Related Information

[Technical Data Provider DB2QW0225 \[page 333\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

[What Is the Data Explorer? \[page 79\]](#)

13.1.9.2.2.2 DBM1

13.1.9.2.2.2.1 DBM1-24bit Storage

Provide the information about DB2 memory for the database address space DBM1 below the 16 MB line.

Where Does the Data Come From?

The monitoring screen data is based on the technical data provider *DB2QW0225*. For more information, see the technical data provider documentation.

What Else Do You Need to Know?

Data on the monitoring screen use the record trace blocks IFCID 316 – SQL Statement Statistics and IFCID 317 – SQL Statement String from the Db2 Instrumentation Facility Component Identifier (IFCID).

For more information, search for *IFCID 316* and *IFCID 317* in the IBM Documentation at <https://www.ibm.com/docs/en> .

Related Information

[Technical Data Provider DB2QW0225 \[page 333\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

[What Is the Data Explorer? \[page 79\]](#)

13.1.9.2.2.2.2 DBM1-31bit Storage

Provide the information about DB2 memory for the database address space DBM1 above the 16 MB line.

Where Does the Data Come From?

The monitoring screen data is based on the technical data provider *DB2QW0225*. For more information, see the technical data provider documentation.

What Else Do You Need to Know?

Data on the monitoring screen use the record trace blocks IFCID 316 – SQL Statement Statistics and IFCID 317 – SQL Statement String from the Db2 Instrumentation Facility Component Identifier (IFCID).

For more information, search for *IFCID 316* and *IFCID 317* in the IBM Documentation at <https://www.ibm.com/docs/en> .

Related Information

[Technical Data Provider DB2QW0225 \[page 333\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

[What Is the Data Explorer? \[page 79\]](#)

13.1.9.2.2.3 DBM1-64bit Storage

Provide the information about DB2 memory for the database address space DBM1 in 64-bit private pools.

Where Does the Data Come From?

The monitoring screen data is based on the technical data provider *DB2QW0225*. For more information, see the technical data provider documentation.

What Else Do You Need to Know?

Data on the monitoring screen use the record trace blocks IFCID 316 – SQL Statement Statistics and IFCID 317 – SQL Statement String from the Db2 Instrumentation Facility Component Identifier (IFCID).

For more information, search for *IFCID 316* and *IFCID 317* in the IBM Documentation at <https://www.ibm.com/docs/en> .

Related Information

[Technical Data Provider DB2QW0225 \[page 333\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

[What Is the Data Explorer? \[page 79\]](#)

13.1.9.2.2.4 DBM1-Real&Aux. Storage

Provide the information about DB2 memory for the database address space DBM1 real and auxiliary slots in use.

Where Does the Data Come From?

The monitoring screen data is based on the technical data provider *DB2QW0225*. For more information, see the technical data provider documentation.

What Else Do You Need to Know?

Data on the monitoring screen use the record trace blocks IFCID 316 – SQL Statement Statistics and IFCID 317 – SQL Statement String from the Db2 Instrumentation Facility Component Identifier (IFCID).

For more information, search for *IFCID 316* and *IFCID 317* in the IBM Documentation at <https://www.ibm.com/docs/en> .

Related Information

[Technical Data Provider DB2QW0225 \[page 333\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

[What Is the Data Explorer? \[page 79\]](#)

13.1.9.2.2.3 DIST

13.1.9.2.2.3.1 DIST-24bit Storage

Provide the information about DB2 memory for the database address space DIST below the 16 MB line.

Where Does the Data Come From?

The monitoring screen data is based on the technical data provider *DB2QW0225*. For more information, see the technical data provider documentation.

What Else Do You Need to Know?

Data on the monitoring screen use the record trace blocks IFCID 316 – SQL Statement Statistics and IFCID 317 – SQL Statement String from the Db2 Instrumentation Facility Component Identifier (IFCID).

For more information, search for *IFCID 316* and *IFCID 317* in the IBM Documentation at <https://www.ibm.com/docs/en> .

Related Information

[Technical Data Provider DB2QW0225 \[page 333\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

[What Is the Data Explorer? \[page 79\]](#)

13.1.9.2.2.3.2 DIST-31bit Storage

Provide the information about DB2 memory for the database address space DIST above the 16 MB line.

Where Does the Data Come From?

The monitoring screen data is based on the technical data provider *DB2QW0225*. For more information, see the technical data provider documentation.

What Else Do You Need to Know?

Data on the monitoring screen use the record trace blocks IFCID 316 - SQL Statement Statistics and IFCID 317 - SQL Statement String from the Db2 Instrumentation Facility Component Identifier (IFCID).

For more information, search for *IFCID 316* and *IFCID 317* in the IBM Documentation at <https://www.ibm.com/docs/en> .

Related Information

[Technical Data Provider DB2QW0225 \[page 333\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

[What Is the Data Explorer? \[page 79\]](#)

13.1.9.2.2.3.3 DIST-64bit Storage

Provide the information about DB2 memory for the database address space DIST in 64-bit private pools.

Where Does the Data Come From?

The monitoring screen data is based on the technical data provider *DB2QW0225*. For more information, see the technical data provider documentation.

What Else Do You Need to Know?

Data on the monitoring screen use the record trace blocks IFCID 316 - SQL Statement Statistics and IFCID 317 - SQL Statement String from the Db2 Instrumentation Facility Component Identifier (IFCID).

For more information, search for *IFCID 316* and *IFCID 317* in the IBM Documentation at <https://www.ibm.com/docs/en> .

Related Information

[Technical Data Provider DB2QW0225 \[page 333\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

[What Is the Data Explorer? \[page 79\]](#)

13.1.9.2.2.3.4 DIST-Real&Aux. Storage

Provide the information about DB2 memory for the database address space DIST real and auxiliary slots in use.

Where Does the Data Come From?

The monitoring screen data is based on the technical data provider *DB2QW0225*. For more information, see the technical data provider documentation.

What Else Do You Need to Know?

Data on the monitoring screen use the record trace blocks IFCID 316 - SQL Statement Statistics and IFCID 317 - SQL Statement String from the Db2 Instrumentation Facility Component Identifier (IFCID).

For more information, search for *IFCID 316* and *IFCID 317* in the IBM Documentation at <https://www.ibm.com/docs/en> .

Related Information

[Technical Data Provider DB2QW0225 \[page 333\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

[What Is the Data Explorer? \[page 79\]](#)

13.1.9.3 Workfiles

13.1.9.3.1 In Memory

Provide information about Db2 work files for processing SQL statements stored in memory.

Where Does the Data Come From?

The monitoring screen data is based on the technical data provider *DB2QW0225*. For more information, see the technical data provider documentation.

Related Information

[Technical Data Provider DB2QW0225 \[page 333\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

[What Is the Data Explorer? \[page 79\]](#)

13.1.9.3.2 Temporary Database

Provide information about Db2 work files for processing SQL statements stored in a temporary database.

Where Does the Data Come From?

The monitoring screen data is based on the technical data provider *DB2QW0225*. For more information, see the technical data provider documentation.

Related Information

[Technical Data Provider DB2QW0225 \[page 333\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

[What Is the Data Explorer? \[page 79\]](#)

13.1.10 Thread Activity

13.1.10.1 Overview

Get a global view of the activity of a thread for a specific Db2 subsystem.

When Do You Need This Data?

- To find out how many threads have been working in your system and how much time they have spent working on specific task types
- To identify threads, host engines execute user tasks, run specific jobs, and process requests from a work queue
- To find out if SAP Db2 contains system-defined thread pools and, if present, user-created thread pools
- To find out which active threads have executed SQL statements
- To locate the application source code that issues the SQL statement

Where Does the Data Come From?

The monitoring screen data is based on the technical data providers *DB2QTHDDL* and *DB2QTHDLST*. For more information, see the technical data provider documentation.

What Else Do You Need to Know?

This monitoring screen is connected by reference with detailed monitors, providing additional information. Stepping backward to Overview from details can result in an empty data list because the thread can be already closed at the time.

Related Information

[Technical Data Provider DB2QTHDDL \[page 331\]](#)

[Technical Data Provider DB2QTHDLST \[page 332\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

13.2 Data Providers for IBM Db2 for z/OS

13.2.1 DB2 Installation Parameters

Get an overview of the current installation parameter settings in your Db2 database and compare it to the SAP recommendations.

When Do You Need This Data?

- To find information about Parameter Name, Installation Panel, and Macro of all relevant Db2 installation parameters
- To find out the current value of the Db2 configuration parameters
- To get a quick overview, if the system conforms to SAP's recommendations

Where Does the Data Come From?

The data for this provider comes from the `SYSPROC.ADMIN_INFO_SYSPARM` procedure. This stored procedure returns the system parameters, application default module, and internal resource lock manager (IRLM) parameters of a connected Db2 subsystem. The data about the SAP recommendations are part of the content of the Technical Monitoring Cockpit.

Related Information

[What Are Data Providers? \[page 80\]](#)

[What Is the Data Explorer? \[page 79\]](#)

13.2.2 DB2 Alerts

Get the history of Db2 database alerts.

Database alerts collected by the Db2 monitor are timeouts, deadlocks, lock escalations, long-running units of recovery (long-running transactions), and active log shortages. This data provider lists the number of events for the respective category. For every category, there are more specific data providers available.

When Do You Need This Data?

- To get an overview of database-specific alerts and their statuses
- To resolve the alert situation

Alerts inform you about existing or potential database issues so that you can protect the database system from data loss or unplanned outages.

Where Does the Data Come From?

This data provider is based on the technical data providers *DB2QLOTHW* and *DB2QLOTRS*. For more information, see the technical data provider documentation.

Related Information

[Technical Data Provider DB2QLOTHW \[page 328\]](#)

[Technical Data Provider DB2QLOTRS \[page 328\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

[What Is the Data Explorer? \[page 79\]](#)

13.2.3 Deadlocks Overview

Get an overview of the deadlock events in your Db2 database.

When Do You Need This Data?

- To find out which transactions are locked or are blocking the others
- To investigate application errors caused by deadlocks on the database
- To get information about the resources involved in deadlock situations

Where Does the Data Come From?

This data provider is a product of INNER JOIN of the technical providers *DB2QLODLK* and *DB2QLODRS*. For more information, see the technical data provider documentation.

Related Information

[Technical Data Provider DB2QLODLK \[page 326\]](#)

[Technical Data Provider DB2QLODRS \[page 327\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

[What Is the Data Explorer? \[page 79\]](#)

13.2.4 Statement Cache History

Get a history of the most expensive SQL statements in your Db2 database during a specific period.

When Do You Need This Data?

- To find out how many SQL statements have been executed in your system and how long they have been running
- To find out which SQL statements have been running for a long time
- To find out which SQL statements are executed frequently
- To identify users executing specific SQL statements
- To locate the application source code that issues the SQL statement
- To monitor the capability for all prepared SQL Statements that are in the global SQL Statement Cache

Where Does the Data Come From?

This data provider is based on the technical data provider *DB2QDSCHIST*. For more information, see the technical data provider documentation.

Related Information

[Technical Data Provider DB2QDSCHIST \[page 325\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

[What Is the Data Explorer? \[page 79\]](#)

13.2.5 Statement Cache Overview

Get a current overview of the SQL statement cache in your Db2 database.

When Do You Need This Data?

- To find out how many SQL statements have been executed in your system and how long they have been running
- To find out which SQL statements have been running
- To find out which SQL statements are executed
- To identify users executing specific SQL statements
- To locate the application source code that issues the SQL statement

Where Does the Data Come From?

This data provider is based on the technical data provider *DB2QTHDDSC*. For more information, see the technical data provider documentation.

Dynamic statement cache has limitations on the number of returned data sizes. Filtering by value is recommended.

Related Information

[Technical Data Provider DB2QTHDDSC \[page 330\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

[What Is the Data Explorer? \[page 79\]](#)

13.2.6 Timeouts Overview

Get information about timeout events in your Db2 database.

When Do You Need This Data?

- To investigate application errors caused by timeouts in your database

- To get information about which threads are blocking others
- To get information about the resources involved in timeout events

Where Does the Data Come From?

This data provider is a product of `INNER JOIN` of the technical providers `DB2QLOTRS` and `DB2QLOTHW`. For more information, see the technical data provider documentation.

Related Information

[Technical Data Provider DB2QLOTHW \[page 328\]](#)

[Technical Data Provider DB2QLOTRS \[page 328\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

[What Is the Data Explorer? \[page 79\]](#)

13.2.7 Indexes Overview

Get an ad-hoc overview of the table indexes of your Db2 database.

When Do You Need This Data?

To get an overview of the index columns for the indexes belonging to the database tables.

Where Does the Data Come From?

This data provider is based on the technical data providers `DB2QDB02_OBJECT`, `SYSIBM.SYSINDEXES`, and `SYSIBM.SYSINDEXPART`. For more information, see the technical data provider documentation.

Related Information

[Technical Data Provider DB2QDB02_OBJECT \[page 323\]](#)

[Technical Data Provider SYSIBM.SYSINDEXES \[page 335\]](#)

[Technical Data Provider SYSIBM.SYSINDEXPART \[page 335\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

[What Is the Data Explorer? \[page 79\]](#)

13.2.8 Tables Overview

Get an ad-hoc overview of the tables in your Db2 database.

When Do You Need This Data?

To see information about the tables in your system, including the number of partitions, size, and growth.

Where Does the Data Come From?

This data provider is based on the technical data providers `DB2QDB02_OBJECT`, `SYSIBM.SYSTABLEPART`, `SYSIBM.SYSTABLES`, and `SYSIBM.SYSTABLESPACE`. For more information, see the technical data provider documentation.

Related Information

[Technical Data Provider DB2QDB02_OBJECT \[page 323\]](#)

[Technical Data Provider SYSIBM.SYSTABLEPART \[page 336\]](#)

[Technical Data Provider SYSIBM.SYSTABLES \[page 337\]](#)

[Technical Data Provider SYSIBM.SYSTABLESPACE \[page 338\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

[What Is the Data Explorer? \[page 79\]](#)

13.2.9 Top SQL Statements

Get a history of the top SQL statements in your Db2 database.

When Do You Need This Data?

- To find out which statements consume the most time (accumulated elapsed time) in your database.
- To locate the application source code that issues the SQL statement.

Where Does the Data Come From?

This data provider is based on the technical data provider *DB2QDSCHIST*. For more information, see the technical data provider documentation.

Related Information

[Technical Data Provider DB2QDSCHIST \[page 325\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

[What Is the Data Explorer? \[page 79\]](#)

13.2.10 Thread Activity Overview

Get an overview of the active threads in your Db2 database.

When Do You Need This Data?

- To find out which threads are currently executing in your database
- To find information about the end-user, transaction, and SAP program, currently run in this thread

Where Does the Data Come From?

This data provider is based on the technical data provider *DB2QTHDLST*. For more information, see the technical data provider documentation.

What Else Do You Need to Know?

Threads are the workers of a Db2 database. The display of active threads gives you an overview of the system workload and the workload caused by specific application source code.

Related Information

[Technical Data Provider DB2QTHDLST \[page 332\]](#)

[What Are Data Providers? \[page 80\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

[What Is the Data Explorer? \[page 79\]](#)

13.3 Technical Data Providers for IBM Db2 for z/OS

13.3.1 Technical Data Provider DB2LOCKESCAL

This technical data provider is based on the record trace block IFCID 337 - Lock Escalation Occurrences from the Db2 Instrumentation Facility Component Identifier (IFCID).

The record trace block IFCID 337 corresponds to the table *DB2LOCKESCAL* in an SAP system.

More About the Data Source

For more information, search for *IFCID 337* in the IBM Documentation at <https://www.ibm.com/docs/en> .

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

13.3.2 Technical Data Provider DB2QBSTT

This technical data provider is based on the record trace block IFCID 172 - Deadlock Data from the Db2 Instrumentation Facility Component Identifier (IFCID).

More About the Data Source

For more information, search for *IFCID 172* in the IBM Documentation at <https://www.ibm.com/docs/en> .

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

13.3.3 Technical Data Provider DB2QDB02_OBJECT

This technical data provider is based on the SAP Collector Table DB2QDB02_OBJECT.

This table stores table or index space level information, especially information about used and allocated sizes. The SAP collector builds a history that allows reviewing that development over time.

More About the Data Source

SAPCL alert router uses the stored procedure ADMIN_DS_LIST. The result is stored in table DB2QDB02_OBJECT.

For more information, see the stored procedure ADMIN_DS_LIST in the IBM Documentation at <https://www.ibm.com/docs/en> .

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

13.3.4 Technical Data Provider DB2QDB02_SCHEMA

This technical data provider is based on the SAP collector table DB2QDB02_SCHEMA.

This table contains information on the Schema level for the objects in the database, especially about the number of Tables, Indexes, and Tablespaces in the Schema and the respective size information. The SAP collector builds a history that makes it possible to review these numbers' development over time.

More About the Data Source

SAPCL alert router uses the stored procedure ADMIN_DS_LIST. The result is stored in table DB2QDB02_SCHEMA.

For more information, see the stored procedure ADMIN_DS_LIST in the IBM Documentation at <https://www.ibm.com/docs/en> .

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

13.3.5 Technical Data Provider DB2QDB02_SMS

This technical data provider is based on the SAP collector table DB2QDB02_SMS.

This table contains information about the Storage pools associated with the monitored database. The SAP collector builds a history of these values to monitor the change over time.

More About the Data Source

SAPCL alert router uses the stored procedure ADMIN_INFO_SMS to determine the size of a volume, storage group, or copy pool. The result is stored in table DB2QDB02_SMS for the copy pools used by the backup system.

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

13.3.6 Technical Data Provider DB2QDBPT1

This technical data provider is based on the record trace block IFCID 202 - Buffer Pool Attributes from the Db2 Instrumentation Facility Component Identifier (IFCID).

More About the Data Source

For more information, search for *IFCID 202* in the IBM Documentation at <https://www.ibm.com/docs/en>  .

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

13.3.7 Technical Data Provider DB2QDSCHIST

This technical data provider is based on the record trace blocks IFCID 316 - SQL Statement Statistics and IFCID 317 - SQL Statement String from the Db2 Instrumentation Facility Component Identifier (IFCID).

The SAP collector builds a history for the TOP statements regarding runtime and stores it in this table.

More About the Data Source

For more information, search for *IFCID 316* and *IFCID 317* in the IBM Documentation at <https://www.ibm.com/docs/en>  .

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

13.3.8 Technical Data Provider DB2QGLOTIME

This technical data provider is based on the record trace block IFCID 369 - Aggregated Accounting Statistics from the Db2 Instrumentation Facility Component Identifier (IFCID).

The record trace block IFCID 369 corresponds to the table **DB2QGLOTIME** in an SAP system.

More About the Data Source

For more information, search for *IFCID 369* in the IBM Documentation at <https://www.ibm.com/docs/en> .

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

13.3.9 Technical Data Provider DB2QLODLK

This technical data provider is based on the record trace block IFCID 172 – Deadlock Data from the Db2 Instrumentation Facility Component Identifier (IFCID).

The record trace block IFCID 172 contains the information in table **DB2QLODLK** in an SAP system. It includes information on the deadlock event and the victim.

More About the Data Source

For more information, search for *IFCID 172* in the IBM Documentation at <https://www.ibm.com/docs/en> .

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

13.3.10 Technical Data Provider DB2QLODRS

This technical data provider is based on the record trace block IFCID 172 – Deadlock Data from the Db2 Instrumentation Facility Component Identifier (IFCID).

The record trace block IFCID 172 contains the information in table **DB2QLODRS** in an SAP system. It includes information on the blocked resources and which thread was holding or waiting for the resource.

More About the Data Source

For more information, search for *IFCID 172* in the IBM Documentation at <https://www.ibm.com/docs/en> .

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

13.3.11 Technical Data Provider DB2QLOGS

This technical data provider is based on the record trace block IFCID 330 - Active Log Space Shortage from the Db2 Instrumentation Facility Component Identifier (IFCID).

More About the Data Source

For more information, search for *IFCID 330* in the IBM Documentation at <https://www.ibm.com/docs/en> .

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

13.3.12 Technical Data Provider DB2QL0THW

This technical data provider is based on the record trace blocks IFCID 196 - Holder and IFCID 196 - Timeout Data from the Db2 Instrumentation Facility Component Identifier (IFCID).

The record trace block IFCID 196 - Holder corresponds to the table **DB2QL0THW** in an SAP system.

More About the Data Source

For more information, search for *IFCID 196* in the IBM Documentation at <https://www.ibm.com/docs/en> .

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

13.3.13 Technical Data Provider DB2QL0TRS

This technical data provider is based on the record trace blocks IFCID 196 - Timeout Header and IFCID 196 - Timeout Data from the Db2 Instrumentation Facility Component Identifier (IFCID).

The record trace block IFCID 196 - Timeout Header corresponds to the table **DB2QL0TRS** in an SAP system.

More About the Data Source

For more information, search for *IFCID 196* in the IBM Documentation at <https://www.ibm.com/docs/en> .

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

13.3.14 Technical Data Provider DB2QLRUR

This technical data provider is based on the record trace block IFCID 313 - Uncommitted Unit of Recovery from the Db2 Instrumentation Facility Component Identifier (IFCID).

More About the Data Source

For more information, search for *IFCID 313* in the IBM Documentation at <https://www.ibm.com/docs/en> .

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

13.3.15 Technical Data Provider DB2QSYW1

This technical data provider is based on the record trace blocks IFCID 001 - IFCID Data, IFCID 225 - System Storage Usage, and IFCID 226 - Page Latch Contention Start from the Db2 Instrumentation Facility Component Identifier (IFCID).

More About the Data Source

Search for *IFCID 001*, *IFCID 225*, and *IFCID 226* in the IBM Documentation at <https://www.ibm.com/docs/en> .

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

13.3.16 Technical Data Provider DB2QSYW2

This technical data provider is based on the record trace block IFCID 002 - DB2® Statistics from the Db2 Instrumentation Facility Component Identifier (IFCID).

The record trace block IFCID 002 corresponds to the table **DB2QSYW2** in an SAP system.

More About the Data Source

For more information, search for *IFCID 002* in the IBM Documentation at <https://www.ibm.com/docs/en> .

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

13.3.17 Technical Data Provider DB2QTHDDSC

This technical data provider is based on the record trace blocks IFCID 316 - SQL Statement Statistics and IFCID 317 - SQL Statement String from the Db2 Instrumentation Facility Component Identifier (IFCID).

The record trace block IFCID 316 - SQL Statement Statistics corresponds to the table **DB2QTHDDSC** in an SAP system.

More About the Data Source

For more information, search for *IFCID 316* and *IFCID 317* in the IBM Documentation at <https://www.ibm.com/docs/en> .

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

13.3.18 Technical Data Provider DB2QTHDDSCP

This technical data provider is based on the record trace block IFCID 316 - SQL Statement Statistics and IFCID 317 - SQL Statement String from the Db2 Instrumentation Facility Component Identifier (IFCID).

This table contains Statement Cache entries associated with Alerts (Timeouts, Deadlocks, or Lock escalations) happening in the system.

More About the Data Source

For more information, search for *IFCID 316* and *IFCID 317* in the IBM Documentation at <https://www.ibm.com/docs/en> .

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

13.3.19 Technical Data Provider DB2QTHDDTL

This technical data provider is based on the record trace block IFCID 124 - SQL Statement Record from the Db2 Instrumentation Facility Component Identifier (IFCID).

More About the Data Source

The record trace block IFCID 124 corresponds to the table DB2QTHDDTL in an SAP system.

For more information, search for *IFCID 124* in the IBM Documentation at <https://www.ibm.com/docs/en> .

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

13.3.20 Technical Data Provider DB2QTHDLRP

This technical data provider is based on the record trace block IFCID 150 - Thread Locking from the Db2 Instrumentation Facility Component Identifier (IFCID).

The record trace block IFCID 150 corresponds to the table DB2QTHDLRP in an SAP system.

DB2QTHDLRP returns data from the snapshots triggered by the *Thread Activity Overview* [page 314].

More About the Data Source

For more information, search for *IFCID 150* in the IBM Documentation at <https://www.ibm.com/docs/en> .

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

13.3.21 Technical Data Provider DB2QTHDLST

This technical data provider is based on the record trace block IFCID 148 - Thread Statistics from the Db2 Instrumentation Facility Component Identifier (IFCID).

The record trace block IFCID 148 corresponds to the table DB2QTHDLST in an SAP system.

More About the Data Source

For more information, search for *IFCID 148* in the IBM Documentation at <https://www.ibm.com/docs/en> .

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

13.3.22 Technical Data Provider DB2QW0225

This technical data provider is based on the record trace block IFCD 000 from the Db2 Instrumentation Facility Component Identifier (IFCID).

More About the Data Source

For more information, search for *IFCD 000* in the IBM Documentation at <https://www.ibm.com/docs/en> .

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

13.3.23 Technical Data Provider SYSIBM.DB2_SYSPARM

This technical data provider is based on the Db2 catalog table `SYSIBM.DB2_SYSPARM`.

Catalog tables contain tablespaces, columns, indexes, privileges, application plans, and packages. `SYSIBM` qualifies all catalog tables in Db2.

The `SYSIBM.DB2_SYSPARM` table stores the results from the `SYSPROC.ADMIN_INFO_SYSPARM` procedure.

This stored procedure returns the system parameters, application default module, and internal resource lock manager (IRLM) parameters of a connected Db2 subsystem or member of its data-sharing group.

More About the Data Source

For more information, search for the `SYSIBM.DB2_SYSPARM` table in the IBM Documentation at <https://www.ibm.com/docs/en> .

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

13.3.24 Technical Data Provider SYSIBM.SYSAUXRELS

This technical data provider is based on the Db2 catalog table `SYSIBM.SYSAUXRELS`. Catalog tables contain tablespaces, columns, indexes, privileges, application plans, and packages. `SYSIBM` qualifies all catalog tables in Db2. The `SYSIBM.SYSAUXRELS` table contains one row for each auxiliary table created for a LOB column.

More About the Data Source

For more information, search for the `SYSIBM.SYSAUXRELS` table in the IBM Documentation at <https://www.ibm.com/docs/en> .

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

13.3.25 Technical Data Provider SYSIBM.SYSCOLUMNS

This technical data provider is based on the Db2 catalog table `SYSIBM.SYSCOLUMNS`. Catalog tables contain tablespaces, columns, indexes, privileges, application plans, and packages. `SYSIBM` qualifies all catalog tables in Db2. The `SYSIBM.SYSCOLUMNS` table contains one row for every column of each table and view.

More About the Data Source

For more information, search for the `SYSIBM.SYSCOLUMNS` table in the IBM Documentation at <https://www.ibm.com/docs/en> .

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

13.3.26 Technical Data Provider SYSIBM.SYSINDEXES

This technical data provider is based on the Db2 catalog table `SYSIBM.SYSINDEXES`.

Catalog tables contain tablespaces, columns, indexes, privileges, application plans, and packages. `SYSIBM` qualifies all catalog tables in Db2.

The `SYSIBM.SYSINDEXES` table contains one row for each index.

More About the Data Source

For more information, search for the `SYSIBM.SYSINDEXES` table in the IBM Documentation at <https://www.ibm.com/docs/en> .

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

13.3.27 Technical Data Provider SYSIBM.SYSINDEXPART

This technical data provider is based on the Db2 catalog table `SYSIBM.SYSINDEXPART`.

Catalog tables contain tablespaces, columns, indexes, privileges, application plans, and packages. `SYSIBM` qualifies all catalog tables in Db2.

The `SYSIBM.SYSINDEXPART` table contains one row for each nonpartitioned secondary index (NPSI) and one row for each partition of a partitioning index or a data-partitioned secondary index.

More About the Data Source

For more information, search for the `SYSIBM.SYSINDEXPART` table in the IBM Documentation at <https://www.ibm.com/docs/en> .

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

13.3.28 Technical Data Provider SYSIBM.SYSINDEXSPACESTATS

This technical data provider is based on the Db2 catalog table `SYSIBM.SYSINDEXSPACESTATS`.

Catalog tables contain tablespaces, columns, indexes, privileges, application plans, and packages. `SYSIBM` qualifies all catalog tables in Db2.

The `SYSIBM.SYSINDEXSPACESTATS` table contains real-time statistics for index spaces.

More About the Data Source

For more information, search for the `SYSIBM.SYSINDEXSPACESTATS` table in the IBM Documentation at <https://www.ibm.com/docs/en> .

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

13.3.29 Technical Data Provider SYSIBM.SYSTABLEPART

This technical data provider is based on the Db2 catalog table `SYSIBM.SYSTABLEPART`.

Catalog tables contain tablespaces, columns, indexes, privileges, application plans, and packages. `SYSIBM` qualifies all catalog tables in Db2.

The `SYSIBM.SYSTABLEPART` table contains one row for each nonpartitioned table space and one row for each partition of the partitioned tablespace.

More About the Data Source

For more information, search for the `SYSIBM.SYSTABLEPART` table in the IBM Documentation at <https://www.ibm.com/docs/en> .

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

13.3.30 Technical Data Provider SYSIBM.SYSTABLES

This technical data provider is based on the Db2 catalog table SYSIBM.SYSTABLES.

Catalog tables contain tablespaces, columns, indexes, privileges, application plans, and packages. SYSIBM qualifies all catalog tables in Db2.

The SYSIBM.SYSTABLES table contains one row for each table, view, or alias.

More About the Data Source

For more information, search for SYSIBM.SYSTABLES table in the IBM Documentation at <https://www.ibm.com/docs/en> .

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

13.3.31 Technical Data Provider SYSIBM.SYSTABLES_PROFILES

This technical data provider is based on the Db2 catalog table SYSIBM.SYSTABLES_PROFILES.

Catalog tables contain tablespaces, columns, indexes, privileges, application plans, and packages. SYSIBM qualifies all catalog tables in Db2.

The SYSIBM.SYSTABLES_PROFILES table contains one row for each profile associated with a table in SYSIBM.SYSTABLES.

More About the Data Source

For more information, search for the SYSIBM.SYSTABLES_PROFILES table in the IBM Documentation at <https://www.ibm.com/docs/en> .

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

13.3.32 Technical Data Provider SYSIBM.SYSTABLESPACE

This technical data provider is based on the Db2 catalog table SYSIBM.SYSTABLESPACE.

Catalog tables contain tablespaces, columns, indexes, privileges, application plans, and packages. SYSIBM qualifies all catalog tables in Db2.

The SYSIBM.SYSTABLESPACE table contains one row for each tablespace.

More About the Data Source

For more information, search for the SYSIBM.SYSTABLESPACE table in the IBM Documentation at <https://www.ibm.com/docs/en> .

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

13.3.33 Technical Data Provider SYSIBM.SYSTABLESPACESTATS

This technical data provider is based on the Db2 catalog table SYSIBM.SYSTABLESPACESTATS.

Catalog tables contain tablespaces, columns, indexes, privileges, application plans, and packages. SYSIBM qualifies all catalog tables in Db2.

The SYSIBM.SYSTABLESPACESTATS table contains real-time statistics for tablespaces.

More About the Data Source

For more information, search for the SYSIBM.SYSTABLESPACESTATS table in the IBM Documentation at <https://www.ibm.com/docs/en> .

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

13.3.34 Technical Data Provider SYSIBM.SYSTABSTATS

This technical data provider is based on the Db2 catalog table SYSIBM.SYSTABSTATS.

Catalog tables contain tablespaces, columns, indexes, privileges, application plans, and packages. SYSIBM qualifies all catalog tables in Db2.

The SYSIBM.SYSTABSTATS table contains one row for each partition of the partitioned tablespace.

More About the Data Source

For more information, search for the SYSIBM.SYSTABSTATS table in the IBM Documentation at <https://www.ibm.com/docs/en> .

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

14 Monitoring Microsoft SQL Server

14.1 Preconfigured Monitoring Screens for SQL Server

Get an overview of the available monitoring screens preconfigured by SAP for Microsoft SQL Server.

14.1.1 Alerts

14.1.1.1 SQL Server Setup Check

This screen shows the results of various checks that are executed on SQL Server. The checks are designed to make sure the server and the database are configured properly for use by SAP NetWeaver.

Prerequisites

The checks are executed once per day by the SQL Agent job `<SID>_TMC_Collector`. This job must be enabled and schema `SAPT00LS` must exist in the database.

Where Does the Data Come From?

This monitoring screen is based on the data provider [SQL Server Setup Check](#). For more information, see the data provider documentation.

Related Information

Data provider [SQL Server Setup Check \[page 357\]](#)

[What Are Data Providers? \[page 80\]](#)

14.1.2 CPU

14.1.2.1 Process CPU Utilization

Get an overview of the CPU utilization by SQL Server compared to all other processes on the server.

Prerequisites

The screen shows historical data that is collected by the SQL Agent job `<SID>_TMC_Collector`. The job must be enabled and schema `SAPT00LS` must exist in the database.

When Do You Need This Data?

- To see the overall workload of the server during a designated period
- To determine if the SQL Server is using too many CPU resources, or if other processes are taking CPU resources away from SQL Server
- To get a quick overview of the CPU resource consumption on the server, which also helps determine at what time of day certain jobs or activities are causing a high load

Where Does the Data Come From?

This monitoring screen is based on the data provider [CPU Utilization](#). For more information, see the data provider documentation.

Related Information

Data provider [CPU Utilization](#) [page 353]

[What Are Data Providers?](#) [page 80]

14.1.3 Configuration

14.1.3.1 Configuration Parameters

This screen shows a list of configuration parameters of SQL Server.

When Do You Need This Data?

To get an overview of configurable parameters for SQL Server, such as memory settings and many other important settings.

Where Does the Data Come From?

This monitoring screen is based on the SQL Server monitoring view *sys.configurations*.

For more information about the monitoring view, see the Microsoft SQL Server database documentation at <https://docs.microsoft.com/en-us/sql/relational-databases/system-catalog-views/sys-configurations-transact-sql> .

14.1.3.2 Database Properties

Get a detailed list of the properties and their values for the database being monitored.

When Do You Need This Data?

If problems related to data collation, update statistics, or logging occur, you can use this screen to determine the configuration of related parameters.

Where Does the Data Come From?

This monitoring screen is based on the SQL Server monitoring view *sys.databases*.

For more information about the monitoring view, see the Microsoft SQL Server database documentation at <https://docs.microsoft.com/en-us/sql/relational-databases/system-catalog-views/sys-databases-transact-sql> .

14.1.3.3 Server Properties

Get a detailed list of the properties and their values for the database server being monitored.

When Do You Need This Data?

To verify that the server settings are correct. The screen provides detailed version and setting information that helps you determine if updates are needed or if configuration changes are necessary.

Where Does the Data Come From?

This monitoring screen is based on the SQL Server function *SERVERPROPERTY*.

For more information about the function, see the Microsoft SQL Server database documentation at <https://docs.microsoft.com/en-us/sql/t-sql/functions/serverproperty-transact-sql> .

14.1.4 High Availability

14.1.4.1 Availability Groups

This screen shows a list of Always On availability groups maintained on the server.

When Do You Need This Data?

To view the configuration of Always On, which availability groups are maintained, and the databases and replicas used in each group.

Where Does the Data Come From?

This monitoring screen is based on the data provider *SYS.AVAILABILITY_GROUPS* that uses the SQL Server monitoring view *sys.availability_groups*.

For more information about the monitoring view, see the Microsoft SQL Server database documentation at <https://docs.microsoft.com/en-us/sql/relational-databases/system-catalog-views/sys-availability-groups-transact-sql> .

Related Information

[What Are Data Providers? \[page 80\]](#)

14.1.5 I/O

14.1.5.1 File I/O Performance

This screen shows the average time of I/O operations in milliseconds over the last 24 hours for all data files and the average number of bytes transferred per 20-minute period.

To examine the metrics for individual files or to check the performance of log files, you can toggle between the chart and table view and change the filter conditions.

Prerequisites

The screen shows historical data that is collected by the SQL Agent job <SID>_TMC_Collector. The job must be enabled and schema SAPTOOLS must exist in the database.

When Do You Need This Data?

To get a quick overview of the I/O performance of the system. This helps you determine whether I/O bottlenecks are causing performance issues on the system.

Where Does the Data Come From?

This monitoring screen is based on the data provider *File I/O Statistics*. For more information, see the data provider documentation.

Related Information

Data provider [File I/O Statistics \[page 354\]](#)

[What Are Data Providers? \[page 80\]](#)

14.1.5.2 I/O Wait Statistics

This screen shows a graph of two SQL Server wait statistics:

- **PAGEIOLATCH_SH** wait time per request shows the time SQL Server waits on average for I/O requests on the data files.
- **WRITELOG** wait time per request, which is the average wait time on completion of I/O requests on the log files.

Prerequisites

The screen shows historical data that is collected by the SQL Agent job `<SID>_TMC_Collector`. The job must be enabled and schema `SAPT00LS` must exist in the database.

When Do You Need This Data?

If the `PAGEIOLATCH_SH` or `WRITELOG` wait time per request is high, the I/O subsystem must be checked. There could be a misconfiguration or malfunction at the I/O system level, a high network latency, or other workload could be causing contention.

Where Does the Data Come From?

This monitoring screen is based on the data provider *I/O Wait Statistics* that uses the general data provider *Wait Statistics*. The data is collected every 20 minutes from system table `sys.dm_os_wait_stats`. For more information, see the data provider documentation.

Related Information

Data provider [I/O Wait Statistics \[page 355\]](#)

Data provider [Wait Statistics \[page 359\]](#)

[What Are Data Providers? \[page 80\]](#)

14.1.6 Memory

14.1.6.1 Available Memory for SQL

Get an overview of the system memory state and the available physical memory in KB.

There are four possible values for memory state:

- Available physical memory is high
- Available physical memory is low
- Physical memory usage is steady
- Physical memory state is transitioning. This means the system flags for high and low memory state are both set. This can only occur temporarily while the memory state is transitioning between high and low.

When Do You Need This Data?

To determine if there is memory pressure at operating system level.

Where Does the Data Come From?

This monitoring screen is based on the technical data provider `SYS.DM_OS_SYS_MEMORY`.

Related Information

[What Are Data Providers? \[page 80\]](#)

14.1.6.2 Hit Ratios

Monitor the quality of the data buffer and the procedure buffer. The graph shows the cache hit ratios over the last 24 hours.

Prerequisites

The screen shows historical data that is collected by the SQL Agent job `<SID>_TMC_Collector`. The job must be enabled and schema `SAPT00LS` must exist in the database.

When Do You Need This Data?

It's important to make sure the hit ratios of both buffers stay well above 90%. A lower value indicates a need to allocate more memory to SQL Server.

Where Does the Data Come From?

This monitoring screen is based on the data provider [Memory Statistics](#). For more information, see the data provider documentation.

Related Information

Data provider [Memory Statistics](#) [page 356]

[What Are Data Providers?](#) [page 80]

14.1.6.3 Page Life Expectancy

Get an overview of the page life expectancy (PLE) on the server over the last 24 hours.

Prerequisites

The screen shows historical data that is collected by the SQL Agent job `<SID>_TMC_Collector`. The job must be enabled and schema `SAPT00LS` must exist in the database.

When Do You Need This Data?

To determine if it's necessary to allocate more memory for SQL Server to increase the data buffer size because the PLE should stay above 300 seconds.

Where Does the Data Come From?

This monitoring screen is based on the data provider [Page Life Expectancy](#). For more information, see the data provider documentation.

Related Information

Data provider [Page Life Expectancy \[page 356\]](#)

[What Are Data Providers? \[page 80\]](#)

14.1.7 SQL

14.1.7.1 Highest Logical Reads

This screen shows a list of statements that have the highest number of logical reads as determined by the current state of the statement cache in SQL Server. You can choose different metrics and a different sorting order if you need other criteria for expensive statements.

When Do You Need This Data?

To view the most expensive SQL statements currently in use by the server. If statements are poorly optimized, they will very often appear near the top of this list.

Where Does the Data Come From?

This monitoring screen is based on the technical data provider `SYS.DM_EXEC_QUERY_STATS`.

Related Information

[What Are Data Providers? \[page 80\]](#)

14.1.8 Space

14.1.8.1 Database File Size History

View the history of the average allocated size in MB for the last 365 days for all data files.

To examine the metrics for individual files, you can toggle between the chart and table view and change the filter conditions.

Prerequisites

The screen shows historical data that is collected by the SQL Agent job <SID>_TMC_Collector. The job must be enabled and schema SAPTOOLS must exist in the database. The data for this screen is collected once per day, and the data is retained for one year.

Where Does the Data Come From?

This monitoring screen is based on the data provider [Database File Size History](#). For more information, see the data provider documentation.

Related Information

Data provider [Database File Size History \[page 353\]](#)

[What Are Data Providers? \[page 80\]](#)

14.1.8.2 Database Size History

View the history of the total database size and the allocated space within the database for the last 365 days.

Only data for the data files is displayed, not the log files. To view the log file history, change the filter on *TYPE_DESC* from *ROWS* to *LOG*.

Prerequisites

The screen shows historical data that is collected by the SQL Agent job <SID>_TMC_Collector. The job must be enabled and schema SAPTOOLS must exist in the database. The data for this screen is collected once per day and retained for one year.

Where Does the Data Come From?

This monitoring screen is based on the data provider [Database File Size History](#). For more information, see the data provider documentation.

Related Information

Data provider [Database File Size History \[page 353\]](#)

[What Are Data Providers? \[page 80\]](#)

14.1.9 Tables

14.1.9.1 Tables By Size

Get an overview of the largest tables in the database.

The list is sorted by total space in MB and shows the used and unused space and the number of rows in the tables.

Where Does the Data Come From?

This monitoring screen is based on the data provider [Table Size](#).

Related Information

Data provider [Table Size \[page 358\]](#)

[What Are Data Providers? \[page 80\]](#)

14.1.10 Time Spent

14.1.10.1 Wait Time

Get an overview of wait statistics on SQL Server for the last 10 days. The list is sorted by wait time per request. It shows all wait categories except *Others* and *Idle* and includes only categories that have a non-zero wait time per request.

Prerequisites

The screen shows historical data that is collected by the SQL Agent job <SID>_TMC_Collector. The job must be enabled and schema SAPTOOLS must exist in the database.

When Do You Need This Data?

To find out if the server spends too much time waiting for I/O, network, locking, or other resources.

Where Does the Data Come From?

This monitoring screen is based on the data provider [Wait Statistics](#). For more information, see the data provider documentation.

Related Information

Data provider [Wait Statistics](#) [page 359]

[What Are Data Providers?](#) [page 80]

14.1.11 Workload

14.1.11.1 Transactions Per Second

View the transaction rate over the last 24 hours.

Prerequisites

The screen shows historical data that is collected by the SQL Agent job <SID>_TMC_Collector. The job must be enabled and schema SAPTOOLS must exist in the database.

Where Does the Data Come From?

This monitoring screen is based on the data provider [Transactions per Second](#). For more information, see the data provider documentation.

Related Information

Data provider [Transactions per Second](#) [page 358]

14.2 Data Providers for Microsoft SQL Server

14.2.1 Batch Requests/sec

Get a history of the number of batch requests per second executed on the server. This counter is a good indicator of how much activity is being processed by SQL Server.

Prerequisites

The SQL Agent job <SID>_TMC_Collector must be enabled and schema SAPTOOLS must exist.

When Do You Need This Data?

If there is an uneven workload on the SQL Server over time, this data provider shows the time of day when most activity occurs. Peaks can be correlated with jobs or user activity that may be occurring at that time.

What Else Do You Need to Know?

This data provider shows data that is also available as the SQL Server performance counter *Batch Requests/sec*. The counter can be useful when batches consist of single statements instead of calls to stored procedures containing multiple queries and calls to other procedures. This is the case in most SAP systems.

Related Information

[What Are Data Providers? \[page 80\]](#)

[What Is the Data Explorer? \[page 79\]](#)

14.2.2 CPU Utilization

Monitor how the CPU is utilized by SQL Server and by other processes on the server. You can also display the idle time of your CPU.

Prerequisites

The SQL Agent job <SID>_TMC_Collector must be enabled and schema SAPTOOLS must exist.

When Do You Need This Data?

- To find out when CPU load has been high in the SQL Server process
- To see how much CPU load came from other processes besides SQL Server
- To see how much idling time is spent by the CPU

Related Information

[What Are Data Providers? \[page 80\]](#)

[What Is the Data Explorer? \[page 79\]](#)

14.2.3 Database File Size History

Get a history of the size of each individual file used by the database.

Prerequisites

The SQL Agent job <SID>_TMC_Collector must be enabled and schema SAPTOOLS must exist.

When Do You Need This Data?

To see the history of the allocated and free space within each file of the database. This includes both data files (TYPE_DESC = 'ROWS') and log files (TYPE_DESC = 'LOG').

To see the database size over time, you can select only the data files by adding a filter. To do so, select the field *TYPE_DESC* and enter the value **ROWS** in the dialog window of the *Filter* button.

What Else Do You Need to Know?

By default, the data is sampled once per day and kept for one year.

Related Information

[What Are Data Providers? \[page 80\]](#)

[What Is the Data Explorer? \[page 79\]](#)

14.2.4 Database Size

Get the current size of the database.

When Do You Need This Data?

This data provider shows the current total allocated and free space in the data files and log files.

Related Information

[What Are Data Providers? \[page 80\]](#)

[What Is the Data Explorer? \[page 79\]](#)

14.2.5 File I/O Statistics

Get an overview of the performance over time of the storage system that contains the database files.

Prerequisites

The SQL Agent job <SID>_TMC_Collector must be enabled and schema SAPTOOLS must exist.

When Do You Need This Data?

To evaluate the performance of I/O operations, which is an important part of overall system performance analysis.

Where Does the Data Come From?

This data provider collects data from the SQL Server dynamic management view `sys.dm_io_virtual_file_stats`.

Related Information

[What Are Data Providers? \[page 80\]](#)

[What Is the Data Explorer? \[page 79\]](#)

14.2.6 I/O Wait Statistics

This data provider shows two key wait statistics counters related to I/O performance of the database data and log files: `PAGEIOLATCH_SH` time per request and `WRITELOG` time per request.

Prerequisites

The SQL Agent job `<SID>_TMC_Collector` must be enabled and schema `SAPT00LS` must exist.

When Do You Need This Data?

To see how long SQL Server waits for the completion of I/O requests for data and log respectively. Peaks indicate periods of an overloaded I/O subsystem. The values are color-coded in green, yellow, or red.

What Else Do You Need to Know?

High values for `PAGEIOLATCH_SH` indicate some pressure on the I/O subsystem for the data files of the database. This can be caused by a number of factors such as a misconfigured I/O system, an uneven distribution of the data files to drive spindles, or any kind of problem with drivers or hardware.

High values for `WRITELOG` are particularly problematic because they show that there's a bottleneck when data is being modified in the database. The log files should normally be located on a disk drive with no other use to allow for the best possible throughput when writing the log.

Related Information

[What Are Data Providers? \[page 80\]](#)

[What Is the Data Explorer? \[page 79\]](#)

14.2.7 Memory Statistics

Get a history of memory usage by SQL Server.

This data provider shows the buffer cache hit ratio, the procedure cache hit ratio and the size of the data cache, lock cache, and free memory.

Prerequisites

The SQL Agent job `<SID>_TMC_Collector` must be enabled and schema `SAPTTOOLS` must exist.

When Do You Need This Data?

To configure memory parameters for SQL Server and to determine if more memory is required.

Related Information

[What Are Data Providers? \[page 80\]](#)

[What Is the Data Explorer? \[page 79\]](#)

14.2.8 Page Life Expectancy

Get a history of the page life expectancy in seconds on SQL Server.

Page life expectancy (PLE) is the number of seconds a page stays in the buffer pool without being referenced. If a page stays longer in the buffer pool, the PLE is higher, which leads to better performance because every time a page is requested, it is likely that the server will find the data in the cache instead of reading it from the disk.

Prerequisites

The SQL Agent job <SID>_TMC_Collector must be enabled and schema SAPTOOLS must exist.

When Do You Need This Data?

If the PLE drops below 300 for long periods, it's an indication that the amount of memory allocated to the server buffer pool may have to be increased.

What Else Do You Need to Know?

If the SQL Server uses multiple NUMA nodes, use the *Instance name* characteristic to examine the PLE per NUMA node.

Related Information

[What Are Data Providers? \[page 80\]](#)

[What Is the Data Explorer? \[page 79\]](#)

14.2.9 SQL Server Setup Check

Get a list of checks that have been executed on SQL Server.

Prerequisites

The SQL Agent job <SID>_TMC_Collector must be enabled and schema SAPTOOLS must exist.

When Do You Need This Data?

To determine if SQL Server is configured according to SAP specifications and recommendations.

What Else Do You Need to Know?

The setup checks are executed once per day. The required data collection job runs every 20 minutes, but only the first run after midnight executes the setup checks and saves the results that the data provider shows.

Related Information

[What Are Data Providers? \[page 80\]](#)

[What Is the Data Explorer? \[page 79\]](#)

14.2.10 Table Size

Get an overview of the current size of all tables in the database.

When Do You Need This Data?

To determine which tables take up the most space in the database.

Related Information

[What Are Data Providers? \[page 80\]](#)

[What Is the Data Explorer? \[page 79\]](#)

14.2.11 Transactions per Second

Get the total number of transactions per second across all databases on the monitored SQL Server.

Prerequisites

The SQL Agent job <SID>_TMC_Collector must be enabled and schema SAPTOOLS must exist.

When Do You Need This Data?

Periods of high values for this counter indicate a higher workload and help identify jobs or activities that use too many resources.

Related Information

[What Are Data Providers? \[page 80\]](#)

[What Is the Data Explorer? \[page 79\]](#)

14.2.12 Wait Statistics

View the history of all wait statistics in SQL Server.

To display the different types of wait statistics such as Network I/O, Lock, or Memory, select [Wait Category](#) in the dialog window of the [Characteristics](#) button.

Prerequisites

The SQL Agent job <SID>_TMC_Collector must be enabled and schema SAPTOOLS must exist.

When Do You Need This Data?

When performance problems occur, the wait statistics help you detect bottlenecks in SQL Server. Except for the "Idle" category, high wait times could indicate a problem when SQL Server needs to acquire a resource such as disk I/O, logging, locking, network I/O, and so on.

Where Does the Data Come From?

This data provider is based on the SQL Server system view `sys.dm_os_wait_stats`. By default, the data is collected every 20 minutes to provide a history.

Related Information

[What Are Data Providers? \[page 80\]](#)

14.3 Technical Data Provider for Microsoft SQL Server

This is a technical data provider that is based directly on a data source of Microsoft SQL Server, for example, a table function or a catalog view.

Technical data providers are mainly intended for platform experts, for example, database administrators and SAP support.

15 Monitoring the Oracle Database

15.1 Data Providers for Oracle

15.1.1 CPU Sessions

Monitor the CPU consumption of your Oracle database.

When Do You Need This Data?

- To monitor database-related load on the CPU
- To track down performance bottle necks caused by high computing times of individual database sessions or related groups of sessions

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

15.1.2 Memory SGA Allocation

Get an overview of the System Global Area (SGA) of your Oracle database.

When Do You Need This Data?

To check out the memory size of the SGA and its components

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

15.1.3 Memory SGA Consumption

View a hierarchical breakdown of the System Global Area (SGA) memory consumption in your Oracle database.

When Do You Need This Data?

- To monitor SGA consumption over time
- To ensure adequate sizing of the individual SGA pools and buffers

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

15.1.4 Space Datafiles

Get an overview of the datafiles in your Oracle database.

When Do You Need This Data?

- To know which datafiles exist in the database and what size they are
- To monitor space consumption within each datafile
- To identify in good time which datafiles are running out of space

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

15.1.5 Space Tablespaces

Get an overview of the tablespaces in your Oracle database.

When Do You Need This Data?

- To know which tablespaces exist in the database and what size they are
- To monitor space consumption within each tablespace
- To identify in good time which tablespaces are running out of space

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

15.1.6 Space Tempfiles

Get an overview of the datafiles for temporary data in your Oracle database.

When Do You Need This Data?

- To know which tempfiles exist in the database and what size they are
- To monitor space consumption within each tempfile
- To identify in good time which tempfiles are running out of space

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

15.1.7 Time Spent Sessions

Monitor session activity on the database by the amount of time the sessions spend on CPU or are actively waiting.

When Do You Need This Data?

- To monitor load on the database
- To track down performance bottlenecks caused by high wait or high computing times of individual sessions or related groups of sessions.

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

15.1.8 Workload Processes

Get an overview of the number of database processes running in your Oracle database.

When Do You Need This Data?

- To monitor load on the database
- To prevent functional limitation caused by the database running out of processes

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

15.1.9 Workload Sessions

Get a historical overview of the number of database sessions running in your Oracle database.

When Do You Need This Data?

To monitor load on the database over time

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

15.1.10 Workload Statistics

Get a historical overview of selected database workload statistics for your Oracle database.

When Do You Need This Data?

To monitor load on the database over time

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

15.1.11 Workload Users

Get a historical overview of user activity on your Oracle database.

When Do You Need This Data?

To monitor user activity over time

What Else Do You Need to Know?

This view is similar to [Workload Sessions \[page 365\]](#) in that it only takes into account foreground user sessions, ignoring the background system sessions.

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

15.2 Technical Data Provider for the Oracle Database

This is a technical data provider that is based directly on a data source of the Oracle database, for example, a table function or a catalog view.

Technical data providers are mainly intended for platform experts, for example, database administrators and SAP support.

16 Monitoring the Application Server ABAP

16.1 Analysis

16.1.1 ABAP Work Process Analysis

Use the ABAP work process analysis to check the currently running work processes. You can drill down from an overview of the application server instances to the details of a single work process.

Work processes do the majority of the processing work of an SAP system. They execute dialog steps in user transactions, updates, lock administration, and so on. Checking the currently running work processes might be necessary in the case of a high resource consumption or if you are facing performance issues.

Note

This analysis screen is available as of SAP NetWeaver 7.4. To use it, you need the same authorization as for transaction SM50 in SAP GUI.

In this analysis scenario, you can display a snapshot of the status of the work processes of all application servers of your SAP system. The analysis scenario consists of one main screen with the tabs [Application Server Instances](#), [Work Processes](#), and [Work Process Details](#). The tabs are associated to each other. Selecting a row in the selection list of a tab updates the content of subsequent tabs. The default selection for each tab is the first entry of the selection lists. You must refresh the display to get updated information.

Application Server Instances

This tab gives you an aggregation of the work process activities per application server. Under [Work Process Summary](#), you find high level information about every application server instance. If you select a line, the more detailed [ABAP Work Process Load Summary for Server](#) will be updated. The load information here is grouped by work process.

Work Processes

This tab shows the list of active work processes for the application server instance that you have selected on the tab [Application Server Instances](#). A selection of available metrics and characteristics is displayed. If you click on the arrow at the end of a table row, you can navigate to the details of a single work process.

Work Process Details

This tab provides more details about the work process that you have selected on the [Work Processes](#) tab, including call stack information. Since the detail information is collected from different data sources at slightly different points in time, it may not match perfectly in the case of transient workload.

Note

The information about current SQL data and call stack information are only available for the local application server instance on which the technical monitoring cockpit is running.

More Information

For more information, see the documentation about displaying and controlling work processes in the SAP Help Portal for your SAP NetWeaver release, for example, <https://help.sap.com/viewer/d0739d980ecf42ae9f3b4c19e21a4b6e/7.31.23/en-US/46fb763b6d4c5515e10000000a1553f6.html>.

Related Information

[What Is the Data Explorer?](#) [page 79]

[What Are Data Providers?](#) [page 80]

16.1.2 SQL Trace Analysis

If you're using the SQL trace function of the ABAP performance trace in your ABAP Development Tools (ADT) or ABAP Workbench, you can access and analyze these SQL trace records in the technical monitoring cockpit.

Procedure

1. In ABAP Development Tools, right-click on your ABAP project on the [Project Explorer](#) tab and choose [SQL Trace](#) from the context menu. A popup window appears showing the application server instance, the state of the SQL Trace (on/off), the user who last changed the trace state, and the timestamp of the last change. Make sure the SQL trace is switched off.
2. Activate the trace and then perform the relevant operations you want to trace.
3. Deactivate the trace.
4. To see the results, choose the [View Trace Directory](#) button. This will open the [SQL Trace Directory](#) screen of the technical monitoring cockpit in a separate browser window. If you're using the ABAP Workbench, you need to manually start the technical monitoring cockpit and choose [Information](#) [Traces](#) [SQL Trace Directory](#) in the menu.

The [SQL Trace Directory](#) screen displays all SQL traces that are available under your user name. The SQL trace that ran last is on top of the list.

- Click on the arrow at the end of a row entry and choose the [Trace Analysis](#) screen from the list of suggested go-to screens to display the trace records and to analyze individual SQL statements that have been recorded.

The [Trace Analysis](#) screen has three tab pages:

Timestamp	Statement With Names	Program	Statement Offset	Trace Type	Duration [µs]	Accessed Records [K]
0000-00-00 00:00:00.0000000 +00:00	OTHER (SQL)		0		15,488	0
2019-05-16 14:48:41.2430000 +00:00	SET CLIENT INFO (EPP_COUNTER=9, SAP_PASSPORT={2A 54 48 2A 03 00 E6 00 00 20 20 20 20...	CL_ST05_TRACE_MAIN_M=====CP	2461	SQL	32	0
2019-05-16 14:48:41.2460000 +00:00	SELECT * FROM "PTC_DIRECTORY" WHERE "MANDT" = ? AND "TYPE" = "N"ST05 AND "INSTANCE_N...	CL_ST05_TRACE_MAIN_M=====CP	2461	SQL	3,542	0
2019-05-16 14:48:41.2460000 +00:00	SESSION (APPLICATIONUSER="BREIN", CLIENT="000", CDS_CLIENT="000", LOCALE_SAP="E", SAP...	CL_ST05_TRACE_MAIN_M=====CP	2461	SQL	14	0
2019-05-16 14:48:41.2530000 +00:00	INSERT INTO "PTC_DIRECTORY" VALUES(?, ?, ?, ?, ?, ?, ?, ?) WITH RANGE_RESTRICTIO...	CL_ST05_TRACE_MAIN_M=====CP	2523	SQL	4,671	1
2019-05-16 14:48:41.2560000 +00:00	COMMIT WORK ON CONNECTION 0	CL_ADT_ST05_KERNEL_INTERFACE==CP	920	SQL	897	0
2019-05-16 14:48:41.2650000 +00:00	SELECT * FROM "ABDQ_EXTDBPS" WHERE "CLIENT" = ? AND "USERNAME" = ? AND "RQ_TERMID"...	CL_ADT_ST05_KERNEL_INTERFACE==CP	920	SQL	2,820	0
2019-05-16 14:48:41.2650000 +00:00	SESSION (APPLICATIONUSER="BREIN", CLIENT="000", CDS_CLIENT="000", LOCALE_SAP="E", SAP...	CL_ADT_ST05_KERNEL_INTERFACE==CP	920	SQL	16	0
2019-05-16 14:48:41.3370000 +00:00	SELECT * FROM "ABDQ_EXTDBPS" WHERE "CLIENT" = ? AND "USERNAME" = ? AND "RQ_TERMID"...	CL_ADT_ST05_KERNEL_INTERFACE==CP	920	SQL	1,191	0
2019-05-16 14:48:41.3370000 +00:00	SESSION (APPLICATIONUSER="BREIN", CLIENT="000", CDS_CLIENT="000", LOCALE_SAP="E", SAP...	CL_ADT_ST05_KERNEL_INTERFACE==CP	920	SQL	20	0
2019-05-16 14:48:41.4360000 +00:00	SELECT * FROM "ABDQ_EXTDBPS" WHERE "CLIENT" = ? AND "USERNAME" = ? AND "RQ_TERMID"...	CL_ADT_ST05_KERNEL_INTERFACE==CP	920	SQL	1,135	0
2019-05-16 14:48:41.4360000 +00:00	SESSION (APPLICATIONUSER="BREIN", CLIENT="000", CDS_CLIENT="000", LOCALE_SAP="E", SAP...	CL_ADT_ST05_KERNEL_INTERFACE==CP	920	SQL	18	0
2019-05-16 14:48:41.4390000 +00:00	SELECT * FROM "USOB_AUTHVALTRC" WHERE "NAME" = ? AND "TYPE" = ? AND "OBJECT" = ? AND...	CL_PTC_UTILITIES=====CP	67	SQL	922	1
2019-05-16 14:48:41.4440000 +00:00	SET CLIENT INFO (EPP_COUNTER=11, SAP_PASSPORT={2A 54 48 2A 03 00 E6 00 00 20 20 20 2...	CL_ADT_ST05_KERNEL_INTERFACE==CP	421	SQL	35	0
2019-05-16 14:48:41.4570000 +00:00	SELECT * FROM "ABDQ_EXTDBPS" WHERE "CLIENT" = ? AND "USERNAME" = ? AND "RQ_TERMID"...	CL_ADT_ST05_KERNEL_INTERFACE==CP	572	SQL	1,179	0

The [Trace Records](#) tab page shows the SQL trace entries in chronological order together with information about the ABAP source code location, duration of the execution, and the number of accessed records. You can adapt the display of the trace records by applying, for example, filters or different sort orders, and by changing the selection of characteristics and metrics.

If you click the detail view arrow at the end of a row, the [Go to](#) popup window opens. From here, you can choose to go to the ABAP source code location, either in your ABAP Development Tools (ADT) or in a web browser.


Note

For displaying the ABAP location with ABAP Development Tools, ADT link handling must be enabled. To do so, in ABAP Development Tools, choose **Window** > **Preferences** > **ABAP Development** and select the checkbox *Open ADT Links from external documents in this installation of ABAP Development Tools*.

The screenshot shows the 'SQL Statement' tab in the Technical Monitoring Cockpit. The left sidebar contains a 'Go to' menu with the following options: 'Detail View', 'ABAP Location (ADT)', 'ABAP Location (HTML)', 'Access Plan', and 'SQL Statement'. The main area displays a table with the following columns: 'Statement Offset', 'Trace Type', 'Duration [µs]', and 'Accessed Records [#]'. The table contains 15 rows of data. A red box highlights the 'Accessed Records [#]' column.

Statement Offset	Trace Type	Duration [µs]	Accessed Records [#]
0		132,020	0
2461	SQL	64	0
2461	SQL	75,427	0
2461	SQL	16	0
2523	SQL	10,661	1
920	SQL	925	0
920	SQL	946	1
920	SQL	19	0
920	SQL	1,150	1
920	SQL	936	1
920	SQL	903	1
920	SQL	856	1
920	SQL	1,202	1
920	SQL	947	1

The [SQL Statement](#) tab page shows the currently selected SQL statement. Here, you can recalculate the prepared and - if possible - the executed access plan for the displayed SQL statement. For data-changing SQL statements, the executed plan is not available.



The [Access Plan](#) tab page shows the prepared and executed access plans of the SQL statement with details. (The executed plan is only shown if you have triggered its calculation on the previous tab page.)

Related Information

[What Is the Data Explorer?](#) [page 79]

[What Are Data Providers?](#) [page 80]

16.2 Information

16.2.1 CPU

16.2.1.1 Top 10 Transactions by CPU Utilization

Get an overview of the top 10 transactions that consumed the largest amount of CPU time during the last 24 hours.

Prerequisites

To make sure that the Application Server ABAP data providers in the technical monitoring cockpit contain data, you need to perform basic configuration steps for the data collection for the Application Server ABAP.

When Do You Need This Data?

- To find out which programs or transactions consume the largest amount of CPU resources on the applications servers
- To identify a starting point for a detailed analysis of CPU resource issues on the applications servers

What Else Do You Need to Know?

The resource CPU determines the computational power of the hardware. If your system has sufficient computational power at any time, you know that your software applications are running smoothly. To prevent performance issues during runtime, ensure that the CPU occupancy rate is below 100%. Be aware that idle CPUs imply costs without additional value.

Where Does the Data Come From?

This monitoring screen is based on the data provider [User and Transaction Profile](#). For more information, see the data provider documentation.

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

[User and Transaction Profile \[page 400\]](#)

[Configuration of the Data Collection for the Application Server ABAP \[page 55\]](#)

16.2.2 Memory

16.2.2.1 Memory Utilization (Application Server ABAP)

Get a history of the memory usage of the ABAP application server, grouped by consumer.

A consumer might be the application server, a user session, or a work process. The server has an upper limit that might be consumed up to 100%. The main part of the memory is consumed by user sessions or work processes, but it must not be consumed completely by them.

The values are calculated every five minutes as an average of more frequently captured server statistics.


Prerequisites

For memory utilization, schedule job SMON_MEMORY_MGMT.

For more information about jobs in the technical job repository, see SAP Note [2581518](#).

When Do You Need This Data?

You want to monitor the capacity of the application server memory, or you analyze a potential memory leak issue.



About 20% of the memory is allocated for the application server-wide buffers such as table buffer and program buffer. If the memory consumers of the application server such as program buffer and table buffer reach 100%, they replace objects in the buffer. To a certain extent, this is the expected behavior.

The rest of the memory is available for user sessions and the work processes and thus represents the flexible work load on the application server. If this memory is consumed 100%, memory leaks may occur. A high usage can always be considered as problematic.

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

16.2.3 Time Spent

16.2.3.1 Average Transaction Response Time

Get an overview of the average time that is spent by transactions in different steps of ABAP work processes during the last 24 hours.

The response times are drilled down and stacked by processing time, database time, roll wait time, and wait time.

Prerequisites

To make sure that the Application Server ABAP data providers in the technical monitoring cockpit contain data, you need to perform basic configuration steps for the data collection for the Application Server ABAP.

When Do You Need This Data?

- To see changes or trends in the average transaction response time
- To find out to which extent different processing steps and wait times contribute to the average transaction response time
- To identify a starting point for a detailed performance analysis

What Else Do You Need to Know?

The statistical records comprise detailed information about all processing steps of the Application Server ABAP. They are a starting point for a more detailed analysis of an application performance issue.

Where Does the Data Come From?

This monitoring screen is based on the data provider [User and Transaction Profile](#). For more information, see the data provider documentation.

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

[User and Transaction Profile \[page 400\]](#)

[Configuration of the Data Collection for the Application Server ABAP \[page 55\]](#)

16.2.4 Traces

16.2.4.1 SQL Trace Directory

Get an overview of persisted trace records of the Application Server ABAP.

Prerequisites

Traces of the ABAP application were recorded from within the ABAP Development Tools (ADT Tools) in Eclipse or from transaction ST05.

When Do You Need This Data?

- To analyze the runtime of an SAP ABAP operation, for example, the execution of SQL statements
- As an ABAP developer, to test the performance of new ABAP code

What Else Do You Need to Know?

If you click on the arrow at the end of a row, you can go to the [SQL Trace Analysis](#) screen.

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

[SQL Trace Analysis \[page 368\]](#)

16.2.5 Workload

16.2.5.1 Active Work Processes

Get an overview of the status and the activity of the work processes on the application server instances of your system.

Prerequisites

You need the same authorization as for transaction **SM50** in SAP GUI.

When Do You Need This Data?

- To analyze the workload on the application server instance in detail
- To avoid an instance overload
- To tune your application server instance

Where Does the Data Come From?

This monitoring screen is based on the technical data provider [TA_ABAP_WP_INFO_WP_LIST](#). For more information, see the technical data provider documentation.

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

[Technical Data Provider TA_ABAP_WP_INFO_WP_LIST \[page 408\]](#)

16.2.5.2 User Sessions

Get an overview of the current user sessions.

Prerequisites

You need the same authorization as for transaction SM04 in SAP GUI.

When Do You Need This Data?

- To perform further analysis regarding user monitoring as an administration task
- To log individual users off, if required

Where Does the Data Come From?

This monitoring screen is based on the technical data provider [TA_ABAP_WP_INFO_SESSION_LIST](#). For more information, see the technical data provider documentation.

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

[Technical Data Provider TA_ABAP_WP_INFO_SESSION_LIST \[page 408\]](#)

16.2.5.3 Work Process Load Summary

View summary information regarding the workload within certain process types across all application server instances of your system.

Prerequisites

You need the same authorization as for transaction SM50 in SAP GUI.

When Do You Need This Data?

To see the utilization of work processes during the last minute, the last 5 minutes and the last 15 minutes

Where Does the Data Come From?

This monitoring screen is based on the technical data provider [TA_ABAP_WP_INFO_WP_LOAD](#). For more information, see the technical data provider documentation.

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

[Technical Data Provider TA_ABAP_WP_INFO_WP_LOAD \[page 409\]](#)

16.2.5.4 Work Process Summary

Get an overview of the number and state of work processes, the work process types and more across all application server instances of your system.

Prerequisites

You need the same authorization as for transaction SM50 in SAP GUI.

When Do You Need This Data?

- To see the number of currently configured work processes for each configured work process type
- To see the number of free work processes for each configured work process type

Where Does the Data Come From?

This monitoring screen is based on the technical data provider [TA_ABAP_WP_INFO_WP_SUMMARY](#). For more information, see the technical data provider documentation.

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

[Technical Data Provider TA_ABAP_WP_INFO_WP_SUMMARY \[page 410\]](#)

16.2.5.5 Work Process Utilization

Get a history of the usage of work processes of the application server, grouped by work process type (dialog, batch, spool, update).

The figures for work processes are calculated every five minutes. Each value is an average from the last five minutes of the server statistics that were continuously captured.

Prerequisites

For work process utilization, schedule job `SMON_MEMORY_MGMT`.

For more information about jobs in the technical job repository, see SAP Note [2581518](#).


When Do You Need This Data?

Your system has a bad response time. This data can show whether the work processes were overloaded.

You can also use this data to get an overview of the capacity. A high load on the batch work processes is not critical if it is only for a short duration. The dialog work processes must not be used by more than 90%, even if it is only for a short period of time.

Example

The figure below shows a high load on batch work processes (in blue) at night:



Only rarely does the load exceed 90%. The dialog work (in red) processes in this example are always below 20%, so the system has a relatively high capacity.

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

16.3 Data Providers for the Application Server ABAP

16.3.1 Application Hit List by DB Access

Display the 40 application statistics steps with the longest database times for each task type. There are no aggregations for hit lists.

Prerequisites

- Configure the collection of Application Server ABAP statistics for the technical monitoring cockpit.
- Configure the collection of application statistics.

When Do You Need This Data?

- To identify the application statistics steps that took the longest to execute
- To find out whether it is always the same program that has the longest database times
- To find out whether this program is running in several sessions at the same time (locking work processes)
- To check whether the system workload can be distributed
- To check how many kilobytes of information are transferred by this program

What Else Do You Need to Know?

Every day, the top 40 transactions are automatically collected by the data collector for each task type (40 longest database times for **BACKGROUND** tasks, 40 longest database times for **DIALOG** tasks, etc.).

❖ Example

- If you set a filter for today (one day) and the task type **BACKGROUND**, you get the top 40 transactions from today for this task type.
- If you set a filter for today (one day) and the task types **BACKGROUND** and **DIALOG**, you get 80 records (the top 40 transactions for the task type **BACKGROUND** and the top 40 transactions for the task type **DIALOG**).
- If you set a filter for today and yesterday (two days) and the task type **BACKGROUND**, you get 80 records for this task type (the top 40 transactions from today and the top 40 transactions from yesterday).

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

[Configuring the Collection of Application Server ABAP Statistics for the Technical Monitoring Cockpit \[page 55\]](#)

[Configuring the Collection of Application Statistics \[page 58\]](#)

16.3.2 Application Hit List by Response Time

Display the 40 application statistics steps with the longest response times for each task type. There are no aggregations for hit lists.

Prerequisites

- Configure the collection of Application Server ABAP statistics for the technical monitoring cockpit.
- Configure the collection of application statistics.

When Do You Need This Data?

- To identify the application statistics steps that took the longest to execute
- To find out whether it is always the same program that has the longest response times
- To find out whether this program is running in several sessions at the same time (locking work processes)
- To check whether the system workload can be distributed
- To check how many kilobytes of information are transferred by this program

What Else Do You Need to Know?

Every day, the top 40 transactions are automatically collected by the data collector for each task type (40 longest database times for **BACKGROUND** tasks, 40 longest database times for **DIALOG** tasks, etc.).

❖ Example

- If you set a filter for today (one day) and the task type **BACKGROUND**, you get the top 40 transactions from today for this task type.
- If you set a filter for today (one day) and the task types **BACKGROUND** and **DIALOG**, you get 80 records (the top 40 transactions for the task type **BACKGROUND** and the top 40 transactions for the task type **DIALOG**).
- If you set a filter for today and yesterday (two days) and the task type **BACKGROUND**, you get 80 records for this task type (the top 40 transactions from today and the top 40 transactions from yesterday).

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

[Configuring the Collection of Application Server ABAP Statistics for the Technical Monitoring Cockpit \[page 55\]](#)

[Configuring the Collection of Application Statistics \[page 58\]](#)

16.3.3 Application Statistics

Analyze the consumption of system resources in more detail than with the analysis of statistical records aggregated in the [Workload Overview](#) data provider. While the workload statistics always analyze a complete dialog step, you can use the application statistics to analyze the resources used by an individual function within a dialog step (such as price determination).

Prerequisites

- Configure the collection of Application Server ABAP statistics for the technical monitoring cockpit.
- Configure the collection of application statistics.

When Do You Need This Data?

To obtain more details about the resource consumption than is usually possible with the collection of dialog step statistics

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

[Configuring the Collection of Application Server ABAP Statistics for the Technical Monitoring Cockpit \[page 55\]](#)

[Configuring the Collection of Application Statistics \[page 58\]](#)

16.3.4 Database Connections

See how the workload generated by database accesses is distributed across the database connections, transactions, and jobs. Usually, the highest workload is caused by using the *DEFAULT* database connection.

Prerequisites

Configure the collection of Application Server ABAP statistics for the technical monitoring cockpit.

When Do You Need This Data?

- To get more detailed information about the database times as seen from the SAP system
- To identify database connections with high workload or long response times

What Else Do You Need to Know?

Identifying database connections with high workload or long response times can help you to narrow down the performance analysis and provide guidance for monitoring data providers at database level.

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

[Configuring the Collection of Application Server ABAP Statistics for the Technical Monitoring Cockpit \[page 55\]](#)

16.3.5 Database Procedures

Display aggregated data about how much load is generated by database procedure calls. You only need these statistics in systems that use many database procedures, such as an Advanced Planning and Optimization (APO) system.

Prerequisites

- Configure the collection of Application Server ABAP statistics for the technical monitoring cockpit.
- Configure the collection of the database procedure call statistics.

When Do You Need This Data?

If you have identified long database procedure call times in other database providers, such as [Workload Overview](#), and you want to know more details about the called database procedures, for example:

- Name of the database procedures
- Name of the transaction, report, and job names that triggered the database procedure call
- User executing the database procedure calls
- Number of times the database procedures were called
- Average and total times for the database procedure calls
- Name of the database connection used

What Else Do You Need to Know?

Detailed information about long-running database procedure calls can help you to narrow down the performance analysis and provide guidance for monitoring data providers at database level.

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

[Configuring the Collection of Application Server ABAP Statistics for the Technical Monitoring Cockpit \[page 55\]](#)

[Configuring the Collection of Database Procedure Call Statistics \[page 60\]](#)

16.3.6 Enqueue Lock Table

The SAP system is equipped with a lock mechanism that synchronizes access to data on the database. The purpose of the lock mechanism is to prevent two transactions from changing the same data on the database simultaneously.

When Do You Need This Data?

- To check for entries, for example, lock objects, lock mode, and lock holder
- To find lock entries that are not cleaned up by the dispatcher, for example, if the dispatcher, operating system, or network connection fails

What Else Do You Need to Know?

Lock entries are usually set and deleted automatically when user programs access a data object and release it again.

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

16.3.7 Front-End Communication

Display information about how the workload is distributed among the presentation servers and instances.

Prerequisites

Configure the collection of Application Server ABAP statistics for the technical monitoring cockpit.

When Do You Need This Data?

- To investigate problems with communication to the presentation servers (front end)

- To investigate high front-end network, GUI, and roll wait times that were identified in the [Workload Overview](#) or other data providers
- To learn about the volume of data that is exchanged between the individual presentation servers and SAP instances
- To learn about the response times arising at the front end
- To identify how much workload each presentation server is causing in a particular instance

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

[Configuring the Collection of Application Server ABAP Statistics for the Technical Monitoring Cockpit \[page 55\]](#)

16.3.8 Hit List by Database Access

Display the 40 transactions with the longest database times for each task type. There are no aggregations for hit lists.

Prerequisites

Configure the collection of Application Server ABAP statistics for the technical monitoring cockpit.

When Do You Need This Data?

- To identify transaction steps that took the longest to execute
- To find out whether it is always the same program that has the longest database times
- To find out whether this program is running in several sessions at the same time (locking work processes)
- To check whether the system workload can be distributed
- To check how many kilobytes of information are transferred by this program

What Else Do You Need to Know?

Every day, the top 40 transactions are automatically collected by the data collector for each task type (40 longest database times for **BACKGROUND** tasks, 40 longest database times for **DIALOG** tasks, etc.).

❖ Example

- If you set a filter for today (one day) and the task type **BACKGROUND**, you get the top 40 transactions from today for this task type.
- If you set a filter for today (one day) and the task types **BACKGROUND** and **DIALOG**, you get 80 records (the top 40 transactions for the task type **BACKGROUND** and the top 40 transactions for the task type **DIALOG**).
- If you set a filter for today and yesterday (two days) and the task type **BACKGROUND**, you get 80 records for this task type (the top 40 transactions from today and the top 40 transactions from yesterday).

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

[Configuring the Collection of Application Server ABAP Statistics for the Technical Monitoring Cockpit \[page 55\]](#)

16.3.9 Hit List by Response Time

Display the 40 transactions with the longest response times for each task type. There are no aggregations for hit lists.

Prerequisites

Configure the collection of Application Server ABAP statistics for the technical monitoring cockpit.

When Do You Need This Data?

- To identify transaction steps that took the longest to execute
- To find out whether it is always the same program that has the longest response times
- To find out whether this program is running in several sessions at the same time (locking work processes)
- To check whether the system workload can be distributed
- To check how many kilobytes of information are transferred by this program

What Else Do You Need to Know?

Every day, the top 40 transactions are automatically collected by the data collector for each task type (40 longest database times for **BACKGROUND** tasks, 40 longest database times for **DIALOG** tasks, etc.).

❖ Example

- If you set a filter for today (one day) and the task type **BACKGROUND**, you get the top 40 transactions from today for this task type.
- If you set a filter for today (one day) and the task types **BACKGROUND** and **DIALOG**, you get 80 records (the top 40 transactions for the task type **BACKGROUND** and the top 40 transactions for the task type **DIALOG**).
- If you set a filter for today and yesterday (two days) and the task type **BACKGROUND**, you get 80 records for this task type (the top 40 transactions from today and the top 40 transactions from yesterday).

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

[Configuring the Collection of Application Server ABAP Statistics for the Technical Monitoring Cockpit \[page 55\]](#)

16.3.10 Load from External Systems

Assess how much workload in the system is caused by external SAP systems in the landscape.

Prerequisites

- Configure the collection of Application Server ABAP statistics for the technical monitoring cockpit.
- This data provider only contains data if there is a connection to another SAP system in the landscape.

When Do You Need This Data?

- To understand the amount of workload generated by each external system
- To find out how much database time and CPU time is consumed due to requests from external systems

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

[Configuring the Collection of Application Server ABAP Statistics for the Technical Monitoring Cockpit \[page 55\]](#)

16.3.11 Memory Consumption

Display aggregated data about total, extended, and private memory usages and about the number of work process reservations and restarts for an SAP application server.

Prerequisites

Configure the collection of Application Server ABAP statistics for the technical monitoring cockpit.

When Do You Need This Data?

- To verify memory consumption after identifying high roll times in the [Workload Overview](#) data provider
- To identify a problem with SAP Memory Management
- To learn about the memory usage per transaction step for a particular transaction or user

What Else Do You Need to Know?

When you perform a top-down system performance analysis, the memory usage statistics help you to identify a hardware bottleneck in the main memory or an opportunity to improve the configuration of SAP Memory Management. You can then use the technical monitoring cockpit to analyze the data providers on operating system level or to verify the SAP Memory Management configuration. In addition, the [Memory Consumption](#) data provider helps you to assess the impact of recent configuration tuning or hardware upgrades.

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

[Configuring the Collection of Application Server ABAP Statistics for the Technical Monitoring Cockpit \[page 55\]](#)

16.3.12 Memory Utilization (Application Server ABAP)

Get a history of the memory usage of the ABAP application server, grouped by consumer.

A consumer might be the application server, a user session, or a work process. The server has an upper limit that might be consumed up to 100%. The main part of the memory is consumed by user sessions or work processes, but it must not be consumed completely by them.

The values are calculated every five minutes as an average of more frequently captured server statistics.


Prerequisites

For memory utilization, schedule job SMON_MEMORY_MGMT.

For more information about jobs in the technical job repository, see SAP Note [2581518](#).

When Do You Need This Data?

You want to monitor the capacity of the application server memory, or you analyze a potential memory leak issue.



About 20% of the memory is allocated for the application server-wide buffers such table buffer and program buffer. If the memory consumers of the application server such as program buffer and table buffer reach 100%, they replace objects in the buffer. To a certain extend, this is the expected behavior.

The rest of the memory is available for user sessions and the work processes and thus represents the flexible work load on the application server. If this memory is consumed 100%, memory leaks may occur. A high usage can always be considered as problematic.

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

16.3.13 RFC Client Calls

Analyze the workload caused by outbound Remote Function Calls (RFC) on the sender (client) system. Statistical information is collected for the five most time-consuming outbound RFCs in transaction steps. These statistics include the function module name, user, destination, and some performance metrics.

Prerequisites

Configure the collection of Application Server ABAP statistics for the technical monitoring cockpit.

When Do You Need This Data?

- To check whether there is a problem caused by remote systems if high RFC times or high roll wait times are observed in the [Workload Overview](#) or other data providers
- To find out which RFC-enabled function modules are the most time-consuming as the basis for investigating an application performance problem
- To find out which users are calling which RFCs and to identify the corresponding workload and performance metrics related to these calls
- To find out which transactions are triggering RFC calls
- To find out how much data is transferred during expensive RFC calls

What Else Do You Need to Know?

The default value for the number of RFCs logged for each transaction step is 5.

⚠ Caution

We do not recommend that you increase this value since a value higher than 5 causes a significant performance overhead in the collector. Only do this if you are an experienced user and it is absolutely necessary.

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

[Configuring the Collection of Application Server ABAP Statistics for the Technical Monitoring Cockpit \[page 55\]](#)

[Configuring the Number of RFCs Logged for Each Transaction Step \[page 64\]](#)

16.3.14 RFC Client Destinations

Display aggregated information about all outbound Remote Function Calls (RFC) by destination. Statistical information for the five most time-consuming outbound RFCs in transaction steps is collected by destination. These statistics include the user, destination, and some performance metrics but not the function module names because they are lost when the destinations are aggregated.

Prerequisites

Configure the collection of Application Server ABAP statistics for the technical monitoring cockpit.

When Do You Need This Data?

- To obtain aggregated information about the destinations used in outbound RFCs
- To check whether there is a problem caused by remote servers if high RFC times and high roll wait times are observed in the [Workload Overview](#) or other data providers
- To see which transactions or users are involved in the outbound RFCs per destination

What Else Do You Need to Know?

The default value for the number of RFCs logged for each transaction step is 5.

Caution

We do not recommend that you increase this value since a value higher than 5 causes a significant performance overhead in the collector. Only do this if you are an experienced user and it is absolutely necessary.

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

[Configuring the Collection of Application Server ABAP Statistics for the Technical Monitoring Cockpit \[page 55\]](#)

[Configuring the Number of RFCs Logged for Each Transaction Step \[page 64\]](#)

16.3.15 RFC Server Calls

Analyze the workload caused by inbound Remote Function Calls (RFC). In this case, the system is considered as the recipient or server system. Statistical information is collected for the five most time-consuming inbound RFCs in transaction steps on the server side.

Prerequisites

Configure the collection of Application Server ABAP statistics for the technical monitoring cockpit.

When Do You Need This Data?

- To learn about the most time-consuming inbound RFCs processed in the system and the function module names
- To identify the users, transactions, and remote servers that are triggering the most time-consuming RFCs processed in the recipient system
- To obtain performance metrics that describe the workload caused by processing inbound RFCs

What Else Do You Need to Know?

The default value for the number of RFCs logged for each transaction step is 5.

⚠ Caution

We do not recommend that you increase this value since a value higher than 5 causes a significant performance overhead in the collector. Only do this if you are an experienced user and it is absolutely necessary.

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

[Configuring the Collection of Application Server ABAP Statistics for the Technical Monitoring Cockpit \[page 55\]](#)

[Configuring the Number of RFCs Logged for Each Transaction Step \[page 64\]](#)

16.3.16 RFC Server Destinations

Display aggregated information about all inbound RFC calls sorted by destination. Statistical information for the five most time-consuming inbound RFC calls in transaction steps is collected by destination. These statistics include the user, the destination, and some performance metrics but not the function module names because they are lost when the destinations are aggregated.

Prerequisites

Configure the collection of Application Server ABAP statistics for the technical monitoring cockpit.

When Do You Need This Data?

- To obtain aggregated information about the destinations used in inbound RFCs
- To find out which transactions or users are involved in the inbound RFCs for each destination

What Else Do You Need to Know?

The default value for the number of RFCs logged for each transaction step is 5.

⚠ Caution

We do not recommend that you increase this value since a value higher than 5 causes a significant performance overhead in the collector. Only do this if you are an experienced user and it is absolutely necessary.

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

[Configuring the Collection of Application Server ABAP Statistics for the Technical Monitoring Cockpit \[page 55\]](#)

[Configuring the Number of RFCs Logged for Each Transaction Step \[page 64\]](#)

16.3.17 Response Time Distribution

Quickly identify the response time distribution of the transaction steps.

Prerequisites

Configure the collection of Application Server ABAP statistics for the technical monitoring cockpit.

When Do You Need This Data?

- To identify the number of transaction steps (absolute or percentage) that are in a particular response time category (for example, response time between 0.4 s and 0.5 s)
- To find the percentage of transaction steps with response time shorter than the upper limit of a particular response time category (for example, response time between 0 s and 0.5 s)

What Else Do You Need to Know?

When you perform a top-down system performance analysis, it is important that you understand the proportion of poor response times with respect to the overall system usage.

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

[Configuring the Collection of Application Server ABAP Statistics for the Technical Monitoring Cockpit \[page 55\]](#)

16.3.18 Spool Activity

View activity statistics related to spool requests. The spool statistics are useful to check the output in each application server.

Prerequisites

Configure the collection of Application Server ABAP statistics for the technical monitoring cockpit.

When Do You Need This Data?

- To learn about the number and the properties of the spool requests that were executed on the selected instance or the entire system
- To find out the average response time and the average wait time for spool requests
- To understand which and how many logical transaction steps there were in the spool work processes

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

[Configuring the Collection of Application Server ABAP Statistics for the Technical Monitoring Cockpit \[page 55\]](#)

16.3.19 TREX Statistics

Display workload information related to a TREX instance.

Prerequisites

- Configure the collection of Application Server ABAP statistics for the technical monitoring cockpit.
- Configure the collection of Distributed Statistics Records (DSRs).

When Do You Need This Data?

- To learn about the workload generated in the TREX instance
- To learn about performance metrics related to the TREX instance and aggregated on protocol, destinations, index IDs, queries, GUID and function names

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

[Configuring the Collection of Application Server ABAP Statistics for the Technical Monitoring Cockpit \[page 55\]](#)

[Configuring the Collection of Distributed Statistics Records \(DSRs\) \[page 63\]](#)

16.3.20 Table Access

Find out which SAP transactions have accessed which database tables, and how often.

Prerequisites

- Configure the collection of Application Server ABAP statistics for the technical monitoring cockpit.
- Configure the generation of table access statistics.

When Do You Need This Data?

- To view the tables that were most frequently accessed by transactions sorted by database times
- To obtain information such as the number of records accessed through direct or sequential reads and the number of record changes for the most frequently accessed tables
- To identify tables that could benefit from the use of indexing, buffering, or more detailed tracing

What Else Do You Need to Know?

Generating table access statistics causes a high workload. It is therefore disabled by default. Only activate it temporarily for individual analyses and transactions.

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

[Configuring the Collection of Application Server ABAP Statistics for the Technical Monitoring Cockpit \[page 55\]](#)

[Configuring the Collection of Table Access Statistics \[page 61\]](#)

16.3.21 Time Profile

Display the distribution of transaction steps over a day, grouped by the hours of the day.

Prerequisites

Configure the collection of Application Server ABAP statistics for the technical monitoring cockpit.

When Do You Need This Data?

- To identify periods of high and low workload as well as long and short response times
- To find out whether the performance problem occurs at peak hours or regularly
- To find out whether the problem occurs only for certain task types (like **BACKGROUND** or **DIALOG** processes) and on specific SAP instances

What Else Do You Need to Know?

This information can help you to decide whether some jobs can be moved to a quiet period, or, if a general problem occurred, whether to monitor their behavior in more detail.

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

[Configuring the Collection of Application Server ABAP Statistics for the Technical Monitoring Cockpit \[page 55\]](#)

16.3.22 SQL Trace Directory

Get an overview of persisted trace records of the Application Server ABAP.

Prerequisites

Traces of the ABAP application were recorded from within the ABAP Development Tools (ADT Tools) in Eclipse or from transaction ST05.

When Do You Need This Data?

- To analyze the runtime of an SAP ABAP operation, for example, the execution of SQL statements
- As an ABAP developer, to test the performance of new ABAP code

What Else Do You Need to Know?

If you click on the arrow at the end of a row, you can go to the [SQL Trace Analysis](#) screen.

Related Information

[What Is the Data Explorer?](#) [page 79]

[What Are Data Providers?](#) [page 80]

[SQL Trace Analysis](#) [page 368]

16.3.23 Transaction Detail

Analyze individual transaction steps of a transaction by their function codes.

Prerequisites

- Configure the collection of Application Server ABAP statistics for the technical monitoring cockpit.
- Configure the collection of transaction detail statistics.

When Do You Need This Data?

- To get more details about a specific transaction that you identified as a candidate for performance optimization
- To help application experts to narrow down a performance problem to specific aspects of a transaction

What Else Do You Need to Know?

The collection of transaction detail statistics is disabled by default. The interpretation of the function codes requires a deep knowledge of the application logic and coding. Therefore, enable the collection only temporarily if a specific transaction was identified as problematic and if it is requested by support engineers or by the developers responsible for the application.

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

[Configuring the Collection of Application Server ABAP Statistics for the Technical Monitoring Cockpit \[page 55\]](#)

[Configuring the Collection of Transaction Detail Statistics \[page 62\]](#)

16.3.24 User and Transaction Profile

See which transactions are causing a high workload in the system as well as the users triggering these transactions. It is a starting point for a more detailed analysis of an application performance issue.

Prerequisites

Configure the collection of Application Server ABAP statistics for the technical monitoring cockpit.

When Do You Need This Data?

- To identify specific transactions with performance problems
- To find out which transactions cause the highest workload
- To find out which transactions have the longest response times
- To identify transactions with expensive SQL statements that you can optimize

- To view the distribution of transactions and their users
- To identify which transactions or users to trace

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

[Configuring the Collection of Application Server ABAP Statistics for the Technical Monitoring Cockpit \[page 55\]](#)

16.3.25 User-Related Data

Understand the distribution of the system load among the different users and the workload profile of a specific user.

Prerequisites

Configure the collection of Application Server ABAP statistics for the technical monitoring cockpit.

When Do You Need This Data?

- To identify the users who are generating the highest workload in the system
- To find out how long a user must wait and how this response time is distributed
- To see the database accesses and their associated workload that a particular user has caused

What Else Do You Need to Know?

When you perform a top-down system performance analysis, it is important to identify in good time the users who are placing heavy loads on the system, or whether there are general performance issues affecting all users. You can then trace specific users in more detail.

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

16.3.26 VMC Statistics

See how the available resources are used at specific times when the Virtual Machine Container (VMC) is operating.

From the [Workload Overview](#), you can navigate to the [VMC Statistics](#) data provider whenever you observe or suspect VMC-related performance problems and you need more detailed information about VMC-specific statistics.

Prerequisites

Configure the collection of Application Server ABAP statistics for the technical monitoring cockpit.

When Do You Need This Data?

- To check how often the VMC has been called
- To check how often the garbage collector has been called for the various memory areas (old generation, young generation, shared objects)
- To find out how much time the system has spent on the garbage collections
- To understand how many just-in-time (JIT) compilations have occurred

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

[Configuring the Collection of Application Server ABAP Statistics for the Technical Monitoring Cockpit \[page 55\]](#)

16.3.27 Web Client Calls

Display the amount of workload that is generated due to web requests for which the system acts as a web client.

Prerequisites

Configure the collection of Application Server ABAP statistics for the technical monitoring cockpit.

When Do You Need This Data?

- To learn how much data has been sent to which URL, how often this data has been sent, and how long the data transfer took
- To analyze combinations of hosts and ports that function as web servers, and how much workload is generated for each combination
- To view web request information and statistics (such as the amount of data transferred and the response times) aggregated on URL, user, transactions or destination
- To obtain the aggregation of the workload by URLs, users, reports or transaction, or host names and port numbers that act as web servers or web clients

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

[Configuring the Collection of Application Server ABAP Statistics for the Technical Monitoring Cockpit \[page 55\]](#)

16.3.28 Web Client Destinations

Display the amount of workload that is generated due to web requests for which the system acts as a web client. The information in this data provider is aggregated by different client destinations.

Prerequisites

Configure the collection of Application Server ABAP statistics for the technical monitoring cockpit.

When Do You Need This Data?

- To learn how much data has been sent to which URL, how often the data has been sent, and how long the data transfer took
- To analyze combinations of hosts and ports that function as web servers, and how much workload is generated for each combination
- To view web request information and statistics (such as the amount of transferred data and response times) aggregated on URL, user, transactions or destination
- To obtain the aggregation of the workload by URLs, users, reports or transaction, or host names and port numbers that act as web servers or web clients

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

[Configuring the Collection of Application Server ABAP Statistics for the Technical Monitoring Cockpit \[page 55\]](#)

16.3.29 Web Server Calls

Display the amount of workload that is generated by web requests for which the system acts as a web server.

Prerequisites

Configure the collection of Application Server ABAP statistics for the technical monitoring cockpit.

When Do You Need This Data?

- To investigate a performance problem of a web application in the remote system
- To learn how much data has been sent to which URL, how often this data has been sent, and how long the data transfer took
- To analyze combinations of hosts and ports that function as web servers, and how much workload is generated for each combination
- To view web request information and statistics (such as the amount of transferred data and response times) aggregated by URL, user, transactions, or destination
- To obtain the aggregation breakdown of the workload by URLs, users, reports or transaction, or host names and port numbers that act as web servers or web clients

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

[Configuring the Collection of Application Server ABAP Statistics for the Technical Monitoring Cockpit \[page 55\]](#)

16.3.30 Web Server Destinations

Display the amount of workload that is generated by web requests for which the system acts as a web server. The information in this data provider is aggregated by different server destinations.

Prerequisites

Configure the collection of Application Server ABAP statistics for the technical monitoring cockpit.

When Do You Need This Data?

- To investigate a performance problem of a web application in the remote system
- To learn how much data has been sent and to which URL, how often the data has been sent, and how long the data transfer took
- To analyze combinations of hosts and ports that function as web servers, and how much workload is generated for each combination
- To view web request information and statistics (such as the amount of transferred data and response times) aggregated by URL, user, transactions, or destination
- To obtain the aggregation of the workload by URLs, users, reports or transaction, or host names and port numbers that act as web servers or web clients

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

[Configuring the Collection of Application Server ABAP Statistics for the Technical Monitoring Cockpit \[page 55\]](#)

16.3.31 Workload Overview

Display aggregated data describing the workload that is written by the SAP kernel. The [Workload Overview](#) provides information about the most important workload data, such as CPU time, number of database

changes, and response times. It also provides a history of the workload and lets you compare the workload of multiple SAP instances of the same system.

Prerequisites

Configure the collection of Application Server ABAP statistics for the technical monitoring cockpit.

When Do You Need This Data?

With this data provider, you can investigate general system performance problems, monitor workload trends and quickly assess overall system health. You can use it for the following:

- To check the average response times in the system
- To obtain the distribution of the total response time in terms of CPU, database, or processing times
- To identify the SAP instance with the highest workload and longest response time
- To see how much workload is being generated on each instance
- To identify the task types that generate the highest workload

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

[Configuring the Collection of Application Server ABAP Statistics for the Technical Monitoring Cockpit \[page 55\]](#)

16.3.32 Work Process Utilization

Get a history of the usage of work processes of the application server, grouped by work process type (dialog, batch, spool, update).

The figures for work processes are calculated every five minutes. Each value is an average from the last five minutes of the server statistics that were continuously captured.

Prerequisites

For work process utilization, schedule job `SMON_MEMORY_MGMT`.

For more information about jobs in the technical job repository, see SAP Note [2581518](#) .


When Do You Need This Data?

Your system has a bad response time. This data can show whether the work processes were overloaded.

You can also use this data to get an overview of the capacity. A high load on the batch work processes is not critical if it is only for a short duration. The dialog work processes must not be used by more than 90%, even if it is only for a short period of time.

Example

The figure below shows a high load on batch work processes (in blue) at night:



Only rarely does the load exceed 90%. The dialog work (in red) processes in this example are always below 20%, so the system has a relatively high capacity.

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

16.4 Technical Data Providers for the Application Server ABAP

Technical data providers are based on a specific data source in the Application Server ABAP, for example, an ABAP interface that retrieves the work process load.

Technical data providers are mainly intended for platform experts, for example, database administrators and SAP support.

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

16.4.1 Technical Data Provider TA_ABAP_WP_INFO_SESSION_LIST

This technical data provider is based on an ABAP kernel method that retrieves a list of all user sessions of an Application Server ABAP instance.

For each session, information according to the ABAP Dictionary structure SSI_SESSION_INFO is provided. Basically, this information corresponds to the information that is available in the overview of transaction SM04.

Prerequisites

You need the same authorization as for transaction SM04.

More About the Data Source

For more information, see the documentation for transaction SM04 in the SAP Help Portal for your SAP NetWeaver release, for example, https://help.sap.com/doc/saphelp_nw75/7.5.5/en-US/02/3b3ad97b0b4526839377f4c2112f33/frameset.htm.

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

16.4.2 Technical Data Provider TA_ABAP_WP_INFO_WP_LIST

This technical data provider is based on an ABAP kernel method that retrieves a list of all work processes of an Application Server ABAP instance.

For each work process, information according to the ABAP Dictionary structure SSI_WORKER_INFO is provided. Basically, this information corresponds to the information that is available in the process overview of transaction SM50.

Prerequisites

You need the same authorization as for transaction **SM50** in SAP GUI.

More About the Data Source

For more information, see the documentation for transaction **SM50** in the SAP Help Portal for your SAP NetWeaver release, for example, <https://help.sap.com/viewer/d0739d980ecf42ae9f3b4c19e21a4b6e/7.52.4/en-US/5540dc7e46f949f581ba02a0d928391c.html>.

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

16.4.3 Technical Data Provider TA_ABAP_WP_INFO_WP_LOAD

This technical data provider is based on an ABAP kernel method that retrieves summary information regarding the workload within certain process types.

The same type of information is used to build parts of the info area in transaction **SM50**. In particular, the utilization of work processes of each type is available in variants: the last 1, 5, and 15 minutes in absolute values and as a percentage.

Prerequisites

You need the same authorization as for transaction **SM50** in SAP GUI.

More About the Data Source

For more information, see the documentation for the info area in transaction **SM50** in the SAP Help Portal for your SAP NetWeaver release, for example, <https://help.sap.com/viewer/d0739d980ecf42ae9f3b4c19e21a4b6e/7.52.4/en-US/3afb089178b146f0880d53e6e020938d.html>.

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

16.4.4 Technical Data Provider TA_ABAP_WP_INFO_WP_SUMMARY

This technical data provider is based on an ABAP kernel method that retrieves summary information regarding the state of the server, the number and state of work processes, the work process types and more.

The same type of information is used to build parts of the enhanced information area in transaction **SM50**.

Prerequisites

You need the same authorization as for transaction **SM50** in SAP GUI.

More About the Data Source

For more information, see the documentation for the info area in transaction **SM50** in the SAP Help Portal for your SAP NetWeaver release, for example, <https://help.sap.com/viewer/d0739d980ecf42ae9f3b4c19e21a4b6e/7.52.4/en-US/3afb089178b146f0880d53e6e020938d.html>.

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

17 Monitoring the Host

17.1 Preconfigured Monitoring Screens for the Host

The preconfigured monitoring screens that are available for the host are based on the data providers of the same name. For more information about a monitoring screen, see the documentation of the corresponding data provider.

Related Information

[Data Providers for the Host \[page 411\]](#)

17.2 Data Providers for the Host

The data providers for the host provide information about the hardware, the operating system, and virtualization.

If one of these data providers refers to a specific scenario only, for example, a specific operating system, this is mentioned in brackets in the respective title. All other data providers are generic.

17.2.1 Configuration of the Computer System

Get an overview of the configuration of your computer system.

When Do You Need This Data?

To get information such as underlying hardware, host names, and IP addresses

What Else Do You Need to Know?

In addition to the information that is part of this data provider, you might want to check the data provider [Configuration of the Operating System](#). Both data providers provide complementary information.

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

[Configuration of the Operating System \[page 412\]](#)

17.2.2 Configuration of the Host Computer System

Get an overview of the configuration settings of your physical hardware or, if your computer system is running in a virtualized environment, of the hypervisor level hosting your virtual machine.

When Do You Need This Data?

To get information about power policies assigned, activation of hyperthreading, and the underlying hardware

What Else Do You Need to Know?

For the configuration settings specified on virtual machine level, you can find complementary information in the data provider [Configuration of the Virtual Computer System](#).

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

[Configuration of the Virtual Computer System \[page 413\]](#)

17.2.3 Configuration of the Operating System

Get an overview of the configuration of your operating system.

When Do You Need This Data?

To get information such as operating system and kernel versions, time zones, and memory settings

What Else Do You Need to Know?

In addition to the information that is part of this data provider, you might want to check the data provider [Configuration of the Computer System](#) for host names and IP addresses used. Both providers provide complementary information.

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

[Configuration of the Computer System \[page 411\]](#)

17.2.4 Configuration of the Virtual Computer System

Get an overview of the configuration settings of your virtual machine if your computer system is running in a virtualized environment.

When Do You Need This Data?

To get information about configuration settings that can be specified on virtual machine level, such as limits or guarantees

What Else Do You Need to Know?

For the configuration settings specified on hypervisor or BIOS level, such as power policies assigned, activation of hyperthreading and for information about the underlying hardware used, you can find complementary information in the data provider [Configuration of the Host Computer System](#).

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

[Configuration of the Host Computer System \[page 412\]](#)

17.2.5 Configuration of CIM Provider Capabilities

Check the capabilities of the SAP CIM providers that are part of SAP Host Agent.

The information provided by SAP Host Agent is collected by several SAP CIM providers. Every SAP CIM provider delivers information related to the configuration of the underlying resources and their environment. Whether an SAP CIM provider can expose the full list of modelled properties depends on the configuration of the monitored system and the usage of certain system features.

When Do You Need This Data?

- If you expect properties that are modelled in the data provider but are not exposed by the data provider
- If you are uncertain about the impact of your configuration on the availability of metrics or properties

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

17.2.6 Configuration of File Systems

You can use this data provider to get an overview of your file systems.

When Do You Need This Data Provider?

To get information such as the type and usage of your file systems

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

17.2.7 Configuration of the Host Agent

Get built information regarding the SAP Host Agent. The SAP Host Agent consists of a variety of executables and libraries that are linked dynamically or during built time.

When Do You Need This Data?

To get built information regarding the SAP Host Agent, such as kernel version, patch number, or compilation mode to check, for example, version prerequisites for a certain provider or feature

What Else Do You Need to Know?

In addition to the information that is part of this data provider, you might want to check the data provider [Host Agent Components](#).

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

[Host Agent Components \[page 417\]](#)

17.2.8 Configuration of Network Ports

Get an overview of the network ports configured in your computer system.

When Do You Need This Data?

To get information about the network ports configured in your computer system and their device IDs

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

17.2.9 Configuration of Processors

Get an overview of the processors configured in your computer system.

When Do You Need This Data?

To get information about the underlying frequencies and processor type for every processor

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

17.2.10 Disks

Get an overview of the disks configured in your computer system.

When Do You Need This Data?

- To get an overview of the resource utilization and throughput of your disks
- To check the performance of your logical disks

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

17.2.11 File Systems

Get detailed information about the file systems configured in your computer system.

When Do You Need This Data?

To get details on size and utilization of your file systems

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

17.2.12 Host Agent Components

Get an overview of the components of your SAP Host Agent. The SAP Host Agent consists of a variety of executables and libraries that are linked dynamically or during built time.

When Do You Need This Data?

To get information such as the components and component versions of your SAP Host Agent

What Else Do You Need to Know?

In addition to the information that is part of this data provider, you might want to check the data provider [Configuration of the Host Agent](#). This data provider comprises further built information of the SAP Host Agent, such as kernel version, patch version, und compilation mode.

Related Information


[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

[Configuration of the Host Agent \[page 415\]](#)

17.2.13 Host CPU Utilization

Get structured information about how physical processors (CPUs) of the host computer system (HCS) are managed by the hypervisor. This helps you to get an insight into the CPU utilization of your virtual machine (VM) and the aggregated CPU utilization of all other VMs on the host.



The host computer system (HCS) hosts $n+1$ different virtual computer systems (VCS). The monitored system is running on the virtual computer system VCS $n+1$.

The CPU information from the host computer system and VCS $n+1$ can be aggregated on four different hierarchy levels, which are referred to by the prefixes H1, H2, H3, and H4:

- Hierarchy H1 shows the overall number of CPUs available on the host computer system.
- Hierarchy H2 shows how many CPUs of the host computer system are in use or still free.
- Hierarchy H3 drills into the CPU usage of VCS $n+1$.
- Hierarchy H4 adds information whether the CPUs used by VCS $n+1$ are guaranteed.

When Do You Need This Data?

- To find out how many physical CPUs are available on the host computer system
- To find out how many physical CPUs on the host computer system are in use by the monitored system or other virtual computer systems

- To find out whether CPUs of the monitored system are guaranteed by the hypervisor
- To find out about the CPU resource limitations of VCSn+1

What Else Do You Need to Know?

If you want to see how many virtual CPUs are used by the operating system, you can check the data provider [Virtual CPUs](#).

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

[Virtual CPU Utilization \[page 425\]](#)

17.2.14 Host Computer System

Get an overview of the resource utilization of your physical hardware or, if your computer system is running in a virtualized environment, of the hypervisor level hosting your virtual machine.

When Do You Need This Data?

To get information about the resource utilization of your host computer system on CPU and memory level

What Else Do You Need to Know?

For the resource utilization on virtual machine level, you can find complementary information in the data provider [Virtual Computer System](#).

Related Information


[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

[Virtual Computer System \[page 426\]](#)

17.2.15 Host Memory Utilization (Linux)

Get structured information about how the physical memory of the host computer system (HCS) is managed by the hypervisor. This helps you to get an insight into the memory consumption of your VM and the aggregated memory commutation of all other VMs on the host.



Model for the Utilization of Physical Memory of the Host Computer from Virtual Computer System Perspective

The memory information from the host computer system can be aggregated on four different hierarchy levels, which are referred to by the prefixes H1, H2, H3, and H4:

- Hierarchy H1 shows how much memory is available on the host computer system.
- Hierarchy H2 shows how much memory of the host computer system is in use or still free.
- Hierarchy H3 drills into the memory usage of VCSn+1.
- Hierarchy H4 adds information whether the memory used by VCSn+1 is guaranteed.

When Do You Need This Data?

- To find out how much physical memory is available on the host computer system
- To find out how much physical memory on the host computer system is in use by the monitored system or other virtual computer systems
- To find out whether the physical memory of the monitored system is guaranteed by the hypervisor

What Else Do You Need to Know?

If you want to see how much virtual memory is used by the operating system, you can check the data provider [Virtual Memory](#).

Related Information


[What Is the Data Explorer?](#) [page 79]

[What Are Data Providers?](#) [page 80]

[Virtual Memory Utilization](#) [page 426]

17.2.16 Host Memory Utilization (Windows)

Get structured information about how the physical memory of the host computer system (HCS) is managed by the hypervisor. This helps you to get an insight into the memory consumption of your VM and the aggregated memory commutation of all other VMs on the host.



Model for the Utilization of Physical Memory of the Host Computer System from Virtual Computer System Perspective

Hyper-V uses the concept of assigned memory. The hypervisor assigns a specific amount of physical memory to a virtual computer system. This assigned memory appears to be consumed at host computer system level while it is not necessarily in use by the operating system and applications at virtual computer system level.

The memory information from the host computer system can be aggregated on four different hierarchy levels, which are referred to by the prefixes H1, H2, H3, and H4:

- Hierarchy H1 shows how much memory is available on the host computer system.
- Hierarchy H2 shows how much memory of the host computer system is in use or still free.
- Hierarchy H3 drills into the memory usage of VCSn+1.
- Hierarchy H4 adds information whether the memory used by VCSn+1 is guaranteed.

When Do You Need This Data?

- To find out how much physical memory is available on the host computer system
- To find out how much physical memory on the host computer system is in use by the monitored system or other virtual computer systems
- To find out whether the physical memory of the monitored system is guaranteed by the hypervisor
- To find out about the memory resource limitations of VCSn+1

What Else Do You Need to Know?

If you want to see how much virtual memory is used by the operating system, you can check the data provider [Virtual Memory](#).

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

[Virtual Memory Utilization \[page 426\]](#)

17.2.17 Memory Overview

Get an overview of the memory utilization in a virtualized environment on operating system level as well as on physical hardware level and on hypervisor level hosting your virtual machine.

When Do You Need This Data?

To get information about the memory consumption on host, hypervisor and operating system level

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

[Host Computer System \[page 419\]](#)

[Operating System \[page 424\]](#)

[Virtual Computer System \[page 426\]](#)

17.2.18 Network Ports

Get a history of the network utilization of the network interfaces attached to your computer system.

When Do You Need This Data?

- To get an overview of the network traffic in your computer system
- To find out if your computer system has an odd number of network transmit or receive errors

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

17.2.19 OS Processes

Get an overview of the processes that are currently running in your operating system. Additionally, you can find information like times spent, priorities, and handles for every process running in the operating system.

When Do You Need This Data?

- To identify the most time-consuming processes running in your operating system (processes with the highest CPU utilization)
- To check for resource utilization and times spent for specified processes

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

17.2.20 Operating System

Get an overview of the memory and CPU resources of your operating system.

When Do You Need This Data?

- To get an overview of how the operating system manages the resources of your computer system
- To understand the computer system workload as managed by the operating system

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

17.2.21 Processors

Get an overview of the resource utilization of the processors configured in your computer system.

When Do You Need This Data?

To find out how many CPUs that are assigned to your operating system are executing on user or system level and how many CPUs are waiting or idle

Related Information


[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

17.2.22 Virtual CPU Utilization

Get structured information about how virtual processors (CPUs) are used by the operating system.

As the operating system is running in a virtual machine (VM), the visible CPUs of the operating system (OS) might or might not be backed by physical CPUs of the host system.



Model for the Utilization of Virtual Processors as Utilized by the Operating System

The CPU information from the operating system can be aggregated on three different hierarchy levels, which are referred to by the prefixes H1, H2, and H3:

- Hierarchy H1 shows the overall number of virtual CPUs as seen by the operating system.
- Hierarchy H2 shows how many virtual CPUs of the operating system are in use or still free.
- Hierarchy H3 drills into the usage of virtual CPUs.

When Do You Need This Data?

To find out how virtual CPUs are utilized by the operating system

What Else Do You Need to Know?

If you want to see how physical CPUs of the host are used, you can check the data provider [Physical Processors](#).

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

[Host CPU Utilization \[page 418\]](#)

17.2.23 Virtual Computer System

Get an overview of the CPU and memory resource utilization of your virtual machine if your computer system is running in a virtualized environment.

When Do You Need This Data?

- To get information about the resource utilization of your virtual computer system
- To find out whether certain guarantees of your virtual computer system have been changed over time

What Else Do You Need to Know?

For the CPU and memory resource utilization on hypervisor level, you can find complementary information in the data provider [Host Computer System](#).

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)


[Host Computer System \[page 419\]](#)

17.2.24 Virtual Memory Utilization

Get structured information about how virtual memory is used by the operating system.

As the operating system is running in a virtual machine (VM), the visible memory of the operating system might or might not be backed by physical memory of the host system.

Virtual Memory



Model for the Utilization of Virtual Memory as Used by the Operating System

The operating system uses memory that appears to be free from an end-user perspective to temporarily cache data, for example, from a file system. The cache is freed up immediately once the memory is requested by processes.

The data can be aggregated on three hierarchy levels, which are referred to by the prefixes H1, H2, and H3:

- Hierarchy H1 shows the total virtual memory as seen by the operating system.
- Hierarchy H2 shows how much virtual memory of the operating system is memory or swap space.
- Hierarchy H3 drills into the usage of the virtual memory.

When Do You Need This Data?

- To find out how virtual memory is used by the operating system
- To find out whether the operating system is paging in memory

What Else Do You Need to Know?

If you want to see how the physical memory of the host is managed by the hypervisor, you can check the data provider [Physical Memory](#).

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Data Providers? \[page 80\]](#)

[Host Memory Utilization \(Linux\) \[page 420\]](#)

[Host Memory Utilization \(Windows\) \[page 421\]](#)

17.3 Technical Data Provider for the Host

This is a technical data provider that is based directly on a platform provider's data source, for example, a platform SDK or any interface accessing host or hypervisor data from inside the virtual system.

Technical data providers are mainly intended for platform experts, for example, database administrators and SAP support.

Related Information

[What Is the Data Explorer? \[page 79\]](#)

[What Are Technical Data Providers? \[page 81\]](#)

Important Disclaimers and Legal Information

Hyperlinks

Some links are classified by an icon and/or a mouseover text. These links provide additional information.

About the icons:

- Links with the icon  : You are entering a Web site that is not hosted by SAP. By using such links, you agree (unless expressly stated otherwise in your agreements with SAP) to this:
 - The content of the linked-to site is not SAP documentation. You may not infer any product claims against SAP based on this information.
 - SAP does not agree or disagree with the content on the linked-to site, nor does SAP warrant the availability and correctness. SAP shall not be liable for any damages caused by the use of such content unless damages have been caused by SAP's gross negligence or willful misconduct.
- Links with the icon  : You are leaving the documentation for that particular SAP product or service and are entering an SAP-hosted Web site. By using such links, you agree that (unless expressly stated otherwise in your agreements with SAP) you may not infer any product claims against SAP based on this information.

Videos Hosted on External Platforms

Some videos may point to third-party video hosting platforms. SAP cannot guarantee the future availability of videos stored on these platforms. Furthermore, any advertisements or other content hosted on these platforms (for example, suggested videos or by navigating to other videos hosted on the same site), are not within the control or responsibility of SAP.

Beta and Other Experimental Features

Experimental features are not part of the officially delivered scope that SAP guarantees for future releases. This means that experimental features may be changed by SAP at any time for any reason without notice. Experimental features are not for productive use. You may not demonstrate, test, examine, evaluate or otherwise use the experimental features in a live operating environment or with data that has not been sufficiently backed up.

The purpose of experimental features is to get feedback early on, allowing customers and partners to influence the future product accordingly. By providing your feedback (e.g. in the SAP Community), you accept that intellectual property rights of the contributions or derivative works shall remain the exclusive property of SAP.

Example Code

Any software coding and/or code snippets are examples. They are not for productive use. The example code is only intended to better explain and visualize the syntax and phrasing rules. SAP does not warrant the correctness and completeness of the example code. SAP shall not be liable for errors or damages caused by the use of example code unless damages have been caused by SAP's gross negligence or willful misconduct.

Bias-Free Language

SAP supports a culture of diversity and inclusion. Whenever possible, we use unbiased language in our documentation to refer to people of all cultures, ethnicities, genders, and abilities.

© 2024 SAP SE or an SAP affiliate company. All rights reserved.

No part of this publication may be reproduced or transmitted in any form or for any purpose without the express permission of SAP SE or an SAP affiliate company. The information contained herein may be changed without prior notice.

Some software products marketed by SAP SE and its distributors contain proprietary software components of other software vendors. National product specifications may vary.

These materials are provided by SAP SE or an SAP affiliate company for informational purposes only, without representation or warranty of any kind, and SAP or its affiliated companies shall not be liable for errors or omissions with respect to the materials. The only warranties for SAP or SAP affiliate company products and services are those that are set forth in the express warranty statements accompanying such products and services, if any. Nothing herein should be construed as constituting an additional warranty.

SAP and other SAP products and services mentioned herein as well as their respective logos are trademarks or registered trademarks of SAP SE (or an SAP affiliate company) in Germany and other countries. All other product and service names mentioned are the trademarks of their respective companies.

Please see <https://www.sap.com/about/legal/trademark.html> for additional trademark information and notices.