

BEx Monitor



ADDON.NDBWGEN

Release 30B



Copyright

© Copyright 2002 SAP AG. 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 AG. The information contained herein may be changed without prior notice.

Some software products marketed by SAP AG and its distributors contain proprietary software components of other software vendors.

Microsoft®, WINDOWS®, NT®, EXCEL®, Word®, PowerPoint® and SQL Server® are registered trademarks of Microsoft Corporation.

IBM®, DB2®, DB2 Universal Database, OS/2®, Parallel Sysplex®, MVS/ESA, AIX®, S/390®, AS/400®, OS/390®, OS/400®, iSeries, pSeries, xSeries, zSeries, z/OS, AFP, Intelligent Miner, WebSphere®, Netfinity®, Tivoli®, Informix and Informix® Dynamic Server™ are trademarks of IBM Corporation in USA and/or other countries.

ORACLE® is a registered trademark of ORACLE Corporation.

UNIX®, X/Open®, OSF/1®, and Motif® are registered trademarks of the Open Group.

Citrix®, the Citrix logo, ICA®, Program Neighborhood®, MetaFrame®, WinFrame®, VideoFrame®, MultiWin® and other Citrix product names referenced herein are trademarks of Citrix Systems, Inc.

HTML, DHTML, XML, XHTML are trademarks or registered trademarks of W3C®, World Wide Web Consortium, Massachusetts Institute of Technology.

JAVA® is a registered trademark of Sun Microsystems, Inc.

JAVASCRIPT® is a registered trademark of Sun Microsystems, Inc., used under license for technology invented and implemented by Netscape.

SAP, SAP Logo, R/2, RIVA, R/3, SAP ArchiveLink, SAP Business Workflow, WebFlow, SAP EarlyWatch, BAPI, SAPPHIRE, Management Cockpit, mySAP, mySAP.com, and other SAP products and services mentioned herein as well as their respective logos are trademarks or registered trademarks of SAP AG in Germany and in several other countries all over the world. MarketSet and Enterprise Buyer are jointly owned trademarks of SAP Markets and Commerce One. All other product and service names mentioned are the trademarks of their respective owners.

Icons

Icon	Meaning
	Caution
	Example
	Note
	Recommendation
	Syntax

Typographic Conventions

Type Style	Description
<i>Example text</i>	Words or characters that appear on the screen. These include field names, screen titles, pushbuttons as well as menu names, paths and options. Cross-references to other documentation.
Example text	Emphasized words or phrases in body text, titles of graphics and tables.
EXAMPLE TEXT	Names of elements in the system. These include report names, program names, transaction codes, table names, and individual key words of a programming language, when surrounded by body text, for example, SELECT and INCLUDE.
Example text	Screen output. This includes file and directory names and their paths, messages, source code, names of variables and parameters as well as names of installation, upgrade and database tools.
EXAMPLE TEXT	Keys on the keyboard, for example, function keys (such as F2) or the ENTER key.
Example text	Exact user entry. These are words or characters that you enter in the system exactly as they appear in the documentation.
<Example text>	Variable user entry. Pointed brackets indicate that you replace these words and characters with appropriate entries.

- BEx Monitor 5
 - OLAP Trace Tool 6
 - Setting an OLAP Trace 8
 - Query Monitor 9
 - Query Properties 11
 - Read mode 12
 - Cache Mode 14
 - Accuracy 15
 - Technical Information 16
- OLAP: Cache Monitor 19
 - Global Cache Parameters 22
 - Cache Data Removal and Swapping 24
 - Cache Structure 27
 - Logical File Names 29



BEx Monitor

Purpose

In the *Business Information Warehouse SAP Easy Access Menu*, under *Business Explorer*, you can find a range of entries that allow the administrator to check the execution of queries or query views in the backend of the BW system.

You can also change settings for InfoProviders on an individual query level in Customizing.

Implementation Considerations

The *SAP Easy Access* initial screen was introduced in BW 3.0A SP11 = BW 3.0B SP04.

Features

Query Monitor

For more information, see [Query Monitor \[Seite 9\]](#).

OLAP Trace Tool

For more information, see [OLAP Trace Tool \[Extern\]](#).

OLAP: Cache Monitor

For more information, see [OLAP Cache Monitor \[Seite 19\]](#).

ICM Monitor

For more information, see [Monitoring the ICM with the ICM Monitor \[Extern\]](#).



OLAP Trace Tool

Purpose

The OLAP Trace tool allows you to trace data transferred to the BW system as a result of actions made in the frontend. You can also trace values for variables. The OLAP trace is stored in the BW system.

You can use the OLAP trace to reproduce activities in the frontend. The OLAP trace tool is particularly useful when testing and when examining frontend activities involving SAP support.

The log itself, created when an OLAP trace is created in the frontend system, is of little use.

Implementation Considerations

Note the following:

1. Keep the trace as short as possible. It may be useful to log the problem with several traces.
2. The trace is no longer valid once the query for which the trace was set is changed. In particular, the trace loses its validity when carrying out the following actions:
 - Changing the hierarchy
 - Changing restrictions
 - Adding or removing characteristics and key figures
 - Reading data (The trace finds nothing if the data for a navigation step that was traced is all of a sudden no longer there).
3. When using the OLAP trace tool to assist SAP Support in investigating a problem, you also need to make sure that the test user you make available to SAP Support has the necessary authorizations.
 - Authorization for executing the query with all navigation steps and data requests.
 - Authorization for the OLAP trace tool (RSRTRACE or RSRCATTRACE)

Integration

As of BW 3.0A SP11 = BW 3.0B SP04, you can reach the OLAP trace tool from the *SAP Easy Access* screen, from *Business Explorer* → *BEx Monitor* → *Trace Tool*.

Further information about setting an OLAP trace can be found under [Setting an OLAP Trace \[Seite 8\]](#).

Features

The following functions are available in the OLAP trace tool toolbar:

Activate User

The user from the *User* field is added to the list. This user's actions are logged.



All users

The system puts the entry ALLUSERS into the list. All users are now logged. Those users that were already activated retain their own configuration however.

Deactivate User

The user selected in the table is deleted from the list. This user's actions are not logged any more.

Configure User

You get to the *Set Trace Configuration* screen. Enter a standard description for the active log process. You can also choose topic areas and individual routines that you should look at in greater detail. You can also activate the *Detailed Runtime Analysis*.

User Logs

The *Processing/Deletion of Logged Callups* screen appears. The system displays the logs that belong to the current user.

CATT Logs

All Logs

The *Processing/Deletion of Logged Callups* screen appears. The system displays all OLAP traces set in the BW system.

Navigating within Logs

You are in the *Processing/Deletion of Logged Callups* screen. Each entry in the list contains the following **entries for each log**:

- System
- Log No.
- User
- Date
- Time
- Description (If necessary)



Delete the logs you no longer need. The number of logs is restricted by the system. When the maximum number is reached, the system starts to delete the oldest logs.

Double-clicking on a log takes you to the “Call up tree” Each entry in the list can contain entries for the following **characteristics for the log considered in greater detail**:

- SUB (The nesting depth is displayed with an “X” or “XX” in this column).
- Function module
- Parameter (for function module)
- InfoCube
- Report name
- CPU time (µs) (required by the trace).



You get back to the *Set Trace Configuration* screen via *Configuration*. You can enter a description of the current log here.

Double-clicking on a row takes you to the *ABAP Debugger*. The log is executed up until the function module is called up (0 nesting depth). If a module with subrc <>0 that was also described is to be exited, a breakpoint is inserted here.



The logs can also be executed via remote login.



Setting an OLAP Trace

Procedure

1. Activate the trace mode.

You have various options to activate the trace mode **for all additional actions on the frontend**:

- Call up the BEx Analyzer. Via the *Settings* icon, choose *Trace*, before logging into the BW system.
- From the *Business Explorer SAP Easy Access* screen, choose *BEx Monitor* → Trace Tool. Enter a user and choose *Edit* → *Activate user*. You can activate all users via *All Users*. For more information, see [OLAP Trace Tool \[Extern\]](#).

To activate the trace mode **for a particular query** in the Web, you need to add the parameter **&trace=X** to the query URL.

2. Log into the BW system and carry out the required steps.

The frontend actions are logged. The log file SAPBEX.trc is generated on the frontend computer.



Make sure that you keep the log as short as possible.

To restrict logging to important steps, it is useful to start the BEx Analyzer without the trace at first, then to load the workbook. Only at this point should you interrupt the connection to the OLAP server, via the *Settings* icon, and activate the trace. You can then create a new connection via  *Refresh*. The actions you subsequently carry out are now logged.

3. Deactivate the trace mode in the same way you activated it in step 1.

Result

You have logged actions in the frontend. The logs and the OLAP trace number can be found in the *OLAP Trace Tool*.

When making contact with SAP support, your customer message can contain the technical name (from the InfoCube or query) as well as the number of the OLAP Trace (*Log Number*) in addition to the user, and the date and time the log was created.



Query Monitor

Purpose

The query monitor tests, checks, and manages BW queries. By using the query monitor, you can test queries or query views, as well as regenerate, check, or change query properties.

With the query monitor you have the option of displaying technical information about queries. Furthermore, the query monitor allows the entry into the cache monitor.

The query monitor is a transaction for administrators, because previous technical knowledge is presumed.

Implementation Considerations

You can use this component from the BW system without a BW Frontend. However, a prerequisite would be that you have already defined the BW queries you want in the BEx Query Designer.

Integration

As of BW 3.0A SP11 = BW 3.0B SP04, you can reach the query monitor from *SAP Easy Access* screen from *Business Explorer* → *BEx Monitor* → *Query Monitor*.

Features

Query

Specify the technical name for the BW query you want in the entry field. Input help is available for searching in *History*, *InfoAreas*, *Roles*, and *Favorites*.

View

If you want to test a query view, specify the technical name for the query view you want in the entry field. Input help is available. (You can find additional information about query views under [Query Views \[Extern.\]](#))

The following functions are available in the query monitor toolbar:

Pushbutton	Function
 <i>Execute</i>	Execute the query or query view You arrive in a screen with the query display in an HTML control. You can test the functions with certain restrictions in the display as you would on the Web. In the upper part of the entry screen, the URL is displayed with the complete parameterization. By using the respective pushbuttons, you can set bookmarks.
 <i>Generate Report</i>	Regenerates the query code The code is stored as a generated report.  With a regeneration, the agreement is always given with the query definition.
 <i>Properties</i>	You arrive at the <i>Query Properties</i> dialog box. You can find additional information under Query Properties [Seite 11] .

By using the respective pushbuttons, the following functions are available for selection:

Pushbutton	Function
------------	----------

 <i>Technical Info</i>	<p>You arrive at the <i>Technical Information</i> screen.</p> <p>You can find additional information under Technical Information [Seite 16].</p>
 <i>Cache Monitor</i>	<p>You arrive at the <i>Cache Monitor</i> screen.</p> <p>For more information, see OLAP: Cache Monitor [Seite 19].</p>
 <i>Query Variants</i>	<p>If the query contains variables that are ready for entry, the following screen displays the query variants. You can change these or create new variants.</p> <p>You can find additional information about query variants under Saving Query Variants [Extern].</p>
 <i>Test IGS</i>	<p>Test Internet Graphics Server</p> <p>With the Internet Graphics Server (IGS) from SAP, you can create charts and maps in the BW system for Web-based reporting.</p> <p>By using <i>Test IGS</i>, the Internet Graphics Server is checked to see whether it is accessible. If the system displays a graphic with two sets of three bars, you have a connection. You can now connect to performance analyses.</p> <p>You can find additional information in the implementation guide (IMG) of the <i>Business Information Warehouse</i> under <i>Reporting-Relevant Settings</i> → <i>Web-Based Settings</i> → <i>IGS Settings</i>.</p>

Restrictions

The function that offers the BW Frontend is not entirely supported in the query monitor. On the other hand, there are settings that you can only make here and not in the Frontend (see  *Properties* and [Query Properties \[Seite 11\]](#)).



Query Properties

Use

From the *Query Properties* dialog box for the query monitor, you can make settings for a BW query concerning the read mode, the cache mode, and the accuracy.

Integration

Read mode and cache mode

In Customizing, you can specify the standard settings for an InfoProvider for the read mode and the cache mode. These need to be valid for those queries that will be created for the corresponding InfoProvider. Furthermore, for the InfoProvider for which data changes do not occur under the control of the BW system (for example, virtual InfoCubes and transactional ODS objects), you can specify the cache validity, meaning the data retention period in the cache.



You can find additional information in the implementation guide (IMG) of the *Business Information Warehouse* under *Reporting-Relevant Settings* → *General Reporting-Settings* → *InfoProvider Settings*.

In the *Query Monitor*, you can change the standard settings for the read mode and the cache mode for the queries that already exist. You can determine the cache validity exclusively for the InfoProvider in Customizing.

Accuracy

You can specify the accuracy with which a key figure needs to be calculated in the InfoObject maintenance.

You can only specify the accuracy with which all query key figures need to be calculated in the *Query Monitor* (and not in the BEx Query Designer from the BW Frontend). You can use this to “override” the affected settings for individual key figures in the InfoObject maintenance, if necessary.

Prerequisites

You have created the query you want in the BEx Query Designer. The settings specified for the basic InfoProvider in Customizing, or those that were already specified earlier in the *Query Monitor*, apply to the query (see above in the *Integration* section)

Features

- [Read Mode \[Seite 12\]](#)
- [Cache Mode \[Seite 14\]](#)
- [Accuracy \[Seite 15\]](#)



Read mode

Definition

The read mode determines how the OLAP processor gets data during navigation. You can set the mode in Customizing for an InfoProvider and in the Query Monitor for a query.

Use

The following types are supported:

1. *Query to be read when you navigate or expand hierarchies (H)*

The amount of data transferred from the database to the OLAP processor is the smallest in this mode. However, it has the highest number of read processes.

In the following mode *Query to Read Data during Navigation*, the data for the fully expanded hierarchy is requested for a hierarchy drilldown. In the *Query to be read when you navigate or expand hierarchies* mode, the data across the hierarchy is aggregated and transferred to the OLAP processor on the hierarchy level that is the lowest in the start list. When expanding a hierarchy node, the children of this node are then read.

You can improve the performance of queries with large presentation hierarchies by creating aggregates on a middle hierarchy level that is greater or the same as the hierarchy start level.

2. *Query to Read Data during Navigation (X)*

The OLAP processor only requests data that is needed for each navigational status of the query in the Business Explorer. The data that is needed is read for each step in the navigation.

In contrast to the *Query to be read when you navigate or expand hierarchies* mode, presentation hierarchies are always imported completely on a leaf level here.

The OLAP processor can read data from the main memory when the nodes are expanded.

When accessing the database, the best aggregate table is used and, if possible, data is aggregated in the database.

3. *Query to Read All Data At Once (A)*

There is only one read process in this mode. When you execute the query in the Business Explorer, all data in the main memory area of the OLAP processor that is needed for all possible navigational steps of this query is read. During navigation, all new navigational states are aggregated and calculated from the data from the main memory.

The read mode *Query to be read when you navigate or expand hierarchies* significantly improves performance in almost all cases compared to the other two modes. The reason for this is that only the data the user wants to see is requested in this mode.

Compared to the *Query to be read when you navigate or expand hierarchies*, the setting *Query to Read Data During Navigation* only improves performance for queries with presentation hierarchies.

Unlike the other two modes, the setting *Query to Read All Data At Once* also improves performance for queries with free characteristics. The OLAP processor aggregates on the corresponding query view. For this reason, the aggregation concept, that is, working with pre-aggregated data, is least supported in the *Query to Read All Data At Once* mode.



We recommend you choose the mode *Query to be read when you navigate or expand hierarchies*.

Only choose a different read mode in exceptional circumstances. The read mode *Query to Read All Data At Once* may be of use in the following cases:

- The InfoProvider does not support selection. The OLAP processor reads significantly more data than the query needs anyway.

- A user exit is active in a query. This prevents data from already being aggregated in the database.

Integration

You can find additional information under [Query Properties \[Seite 11\]](#).



Cache Mode

Definition

The cache mode determines whether or not the query results and navigation states that were calculated using the OLAP processor are to be stored in a cross-transactional cache and, where required, in a background store. You can set the mode in Customizing for an InfoProvider and in the *Query Monitor* for a query.

Use

The following modes are supported for the global cache:

1. *Cache is Active*

The cross-transactional cache is active (default). When the cache memory has been exhausted, excess data is removed, i.e. deleted, from it. When a query is subsequently requested, the result is read from the InfoProvider (the database) once again.

2. *Cache is Inactive*

The cross-transactional cache is deactivated.

3. *Cache is Active and Persistent*

The cross-transactional cache is active. When the caching memory is used up, the excess data is written to a background store (file or cluster-table) and can be loaded back into the cache memory when making another request.



The background store can not be used again when loading data once more.



Various parameters determine whether it is better to use the setting *Cache is Active* or *Cache is Active and Persistent*. These include:

- How often the query is requested
- The complexity of the query
- How often data is loaded

Integration

Further information about the various ways of setting the cache mode for an InfoProvider or a query can be found under [Query Properties \[Seite 11\]](#).

Further information about the global cache parameter and the amount of memory used by the runtime object can be found under [OLAP Cache Monitor \[Seite 19\]](#).



Accuracy

Definition

The determination of a specific accuracy decides which number format is used internally for calculating key figures. The degree of accuracy is determined in the InfoObject maintenance for each key figure and in the query monitor for all query key figures.

Use

When converting decimal figures with numbers following the decimal point into binary floating-point numbers, unavoidable rounding errors occur. This is because many decimals (for example, 0.1) correspond to periodic fractions, and are therefore unable to be displayed exactly.

If you choose *Calculate with Packed Numbers* in the query monitor, packed numbers (ABAP data type P, with nine decimal places) are used as the internal number format for all query key figures instead of binary decimals. Rounding errors are avoided with commercial calculations with this fixed-point number, because in doing so no conversion of numbers into the dual system takes place.

Choose this option for those application areas in which a high accuracy is required (for example, for queries in the accounting and cost accounting areas).



Because calculating with packed numbers causes a high processing status overall, we recommend not using this number format for all query key figures, but rather only for those key figures for which a very high degree of accuracy is required.

You have to make this setting in the InfoObject maintenance for the respective key figure:

- i. From the SAP Easy Access Menu, choose *Modeling* → *Object Maintenance* → *InfoObject*.
- ii. Enter the name of the key figure and choose  *Maintain*.
- iii. To calculate with packed numbers, choose from the *Additional Properties* tab page under *Diverse* with the option *Key Figure with Maximum Accuracy*.



Technical Information

Use

From the *Technical Information* screen, you get an overview of important, technical, reference numbers for a query.

Integration

In order to access the *Technical Information* screen, choose a query from the *Query Monitor* screen and then  *Technical Info*.

Prerequisites

The technical information is automatically created.

Features

The technical information is grouped by category for the selected query. Depending on the query definition, all or only a part of the categories is displayed. The first line of the HTML control contains links to the categories displayed in table form in the lower area of the page.

System-side entries are specified in the first column for some of the categories:

The symbol  means that the value of the technical reference numbers corresponds exactly to the system-side recommendation.

The symbol  means that the value of the technical reference numbers possibly does not correspond to the system-side recommendation.

The symbol  means that the value of the technical reference numbers does not correspond to the system-side recommendation. Therefore, a change can be quite meaningful. Such a calculation does not necessary mean that an error must exist.

General Data

- Data target(s)

Query Definition Data

- Technical name of the query definition
- Last change to the query definition (date and time)

OLAP-Relevant Data

- Query name (technical name and internal name)
- Name of the generated report (by using *Display Gen. Report*, you can display the program; see below).
- Query Generation Time (see under *Cache-Relevant Data*)
- Version (external/internal)
- Generation status (active/inactive)
- Read mode (You can find additional information under [Read Mode \[Seite 12\]](#).)
- Released for OLE DB for OLAP (You can find additional information under [Query Properties \[Extern\]](#).)
- Calculating before the aggregation (You can find additional information under [Calculated Key Figure Properties \[Extern\]](#).)
- Virtual characteristics/key figures
- Attributes in calculated key figures

- Currency translation in the query definition
- Query can use aggregates
- Value field groups (groups of key figure IDs or selection elements and basic sizes for the form elements, combined by specific properties, for example, their aggregation behavior or by specific steps of the process that are required, for example, calculating as packed numbers)
- Key figure IDs for display (KIDs)
- Calculated key figure IDs
- Selecting Constants (You can find additional information under [Selecting Constants \[Extern\]](#).)
- Fiscal year variant
- Query key date
- The storage table Sp cannot be permanently buffered (see under *Cache-Relevant Data*)

Non-Cumulative Flags

- There is a non-cumulative value with AV1
- There are non-cumulative values with LAST or FIRST
- There are non-cumulative values with MAX or MIN
 - You can find additional information about the types of exception aggregation for non-cumulative values under [Aggregation Behavior of Non-Cumulative Values \[Extern\]](#).
- More than a provider of non-cumulatives (see [Evaluating Non-Cumulative Values \[Extern\]](#) in the InfoCube definition section)
- NCUMTIM is not unique (see [Time Reference Characteristic \[Extern\]](#))
- Query does not interrupt the critical interval (see [Validity Period \[Extern\]](#))

Statistics

This section runs the statistics by specifying the current and maximum number:

- InfoCubes involved for the MultiProvider
- FEMS number (selection groups)
- Selection rows
- Columns in the SP storage table
- Free characteristics
- Basic key figures
- Formula components
- Hierarchies
- Hierarchy nodes
- Variables

Invisible Drilldown Characteristics

If a query has invisible drilldown characteristics, they are displayed in this section with their technical name and description, for example:

0CURRENCY Units or currency field

Cache-Relevant Data

- The storage table Sp cannot be permanently buffered (corresponds to the specification under *OLAP-Relevant Data*)

In addition, this section includes important time stamps:

- Query generation time (corresponds to the specification under *OLAP-Relevant Data*)
- Time of the last loading of data into the InfoProvider (last data change)
- Last ValueSet data change
- Last master data change for navigation attributes



This time stamp influences the invalidation of the cache. If the last time stamp for a cache entry for the selected query is older than one of the above-named time stamps, then the cache is rebuilt.

Load Size for a Generated Report

This section restricts ABAP programs. If these are not complied with, then the report is not generated, and the query is not executed.

- PXA Buffer
- Static control blocks
- Variable control blocks
- Literals
- Structure descriptions
- Includes

The following functions are available in the screen **toolbar**:

Function	What you need to know
<i>Display Gen. Report</i>	Goes to the ABAP editor You see the ABAP program that was generated for the query you selected. (You can find the name of the generated report under <i>OLAP-Relevant Data</i> .)
<i>Query Def. Designer</i>	Goes to the ABAP debugger. You see the coding for the interface between the BEx Query Designer and the OLAP processor.
<i>Query Def. Processor</i>	Goes to the ABAP debugger. You see a few lines of the coding according to the above-named interface. For this break point, the complex data structure SX_REPORT and internal control tables are displayed for the query.



OLAP: Cache Monitor

Purpose

From the *Cache Monitor* screen, you can look at the global cache parameters, the amount of memory used by the query runtime objects and the current underlying cache structure.

Integration

As of BW 3.0A SP11 = BW 3.0B SP04, you access the OLAP Cache Monitor (transaction RSRCACHE) from the *SAP Easy Access* screen, from *Business Explorer* → *BEx Monitor* → *OLAP: Cache Monitor*. From the *Query Monitor* screen, you can get to the OLAP Cache Monitor via  *Cache Monitor*.

Settings for the global cache parameters are made in Customizing (see [Global Cache Parameters \[Seite 22\]](#)).

The cache-relevant settings that were set for an InfoProvider or query can “overwrite” the global cache parameters.

- You can set the cache mode for an InfoProvider in Customizing.
- You can change the cache mode for a query in the *Query Monitor* screen (see [Query Properties \[Seite 11\]](#)).



You can get an overview of the *cache-relevant data* for a query (for example, important timestamps) from the *Query Monitor* screen via  *Technical Information* (See [Technical Information \[Seite 16\]](#)).

Features

System-side entries are specified in the first column for some of the categories. You can find additional information under [Technical Information \[Seite 16\]](#).

Cache Parameters

The global cache parameters are displayed in the upper screen area. The display is based on the settings made in Customizing.

Example:

	Cache active	Local Cache Size	100 MB	Global Cache Size	200 MB
	Cache Persistence Mode	Flat File			
	Cache Persistence Logical File Name	BW_OLAP_CACHE			

Choose  *Cache Parameter* to change any settings. Further information can be found under [Global Cache Parameters \[Seite 22\]](#).

Tab Page: Shared Memory

In the Cache Monitor screen on the Shared Memory tab page, you can see the important cache-relevant figures for query runtime objects.

The table on the left-hand side gives an overview of current memory consumption. Further information can be found from [Cache Data Removal and Swapping \[Seite 24\]](#).

The table on the right-hand side gives an overview of current cache structure. You can choose from different display modes:

 *List Display*

 *Hierarchical Display*

Further information can be found under [Cache Structure \[Seite 27\]](#).

The following functions are available in the screen **toolbar**:

Function	What you need to know
 (Refresh)	
 (Delete)	Deletes cache data in the shared memory. This function is useful when testing.
 Cache Parameter	See Global Cache Parameters [Seite 22] .
 Buffer Monitor	<p>Choose  <i>Buffer Monitor</i> or  <i>Buffer Overview</i> to check whether enough cache memory is available.</p> <p>You can use  <i>Buffer Monitor</i> to reach the <i>Buffer Monitoring Tool</i>, which gives you an overview of the buffer for your application server. In the <i>Buffer Overview</i> screen area, the shared memory rows (buffer: <i>Exp/ImpMem</i>) contain data relevant to the OLAP cache:</p> <ul style="list-style-type: none"> • <i>Size</i>: Maximum size of the cache memory in bytes • <i>Free bytes</i>: Free cache memory in bytes • <i>Directory</i>: Maximum number of directories, that is, individual packages in the cache memory • <i>Dir. Free</i>: Number of free directories • <i>Efficiency</i>: Degree to which requests can be met, in %, - corresponds to the <i>Hitratio</i> in the <i>Buffer Overview</i>.
 Buffer Overview	<p>The <i>Tune Summary</i> screen appears (transaction ST02). The shared memory row (buffer: <i>Exp/Imp SHM</i>) contains data of interest for the OLAP cache:</p> <ul style="list-style-type: none"> • <i>Hitratio</i>: [%] - corresponds to <i>Efficiency</i> in the <i>Buffer Monitor</i> • <i>Allocated</i>: [kB] • <i>Free space</i>: [kB] and [%] • <i>Dir. size</i>: Entries • <i>Free directory</i>: Entries and [%] <p></p> <p>Note that both the number of bytes and the number of directories have a limiting effect. You can find out which of these sizes may be critical from the relationship between the respective maximum value and free capacity value.</p> <p></p> <p>Further information can be found under The Buffer Monitor [Extern].</p>
 Logical File Name	See Logical File Names [Seite 29] (transaction FILE).
 Directory Overview	<p>The <i>SAP Directories</i> screen appears.</p> <p></p> <p>DIR_GLOBAL /usr/<SYSID>/global</p> <p>Double-click on the physical path to get to the file display. If you have kept to the naming conventions, you can find your physical file under CACHE_<HOST>_<SYSID>_<PARAM_1>.</p> <p>Further information about defining logical and physical file paths and –names</p>

	can be found under Logical File Names [Seite 29] .
--	--



Global Cache Parameters

Use

Global cache parameters are the global (valid system-wide) settings for the OLAP cache.

The following table gives an overview of the different types of OLAP cache:

OLAP Cache Type	Description
Cross-transactional Cache	<p>Cross-transactional application buffer in which the results calculated by the OLAP processor and navigation states of a query are stored as a data cluster on an application server instance. With similar query requests, the OLAP processor can access data stored in the cache.</p> <p>When the OLAP processor is able to read data from the cache, this accelerates query execution, since accesses to caches are much quicker than database accesses. At the same time, the database instance is spared.</p>
Local Cache	Local OLAP processor cache, in which the results calculated by the OLAP processor are stored in the roll area for each setting.

Integration

The global cache parameters are set centrally when implementing the BW system.



You can find additional information in the implementation guide (IMG) of the *Business Information Warehouse* under *Reporting-Relevant Settings* → *General Reporting-Settings* → *Global Cache Settings*.

If you find out when evaluating data displayed in the OLAP Cache Monitor that the global cache parameters must be adjusted to the system requirements, you can change these later on.



From the *Cache Monitor* screen, via  *Cache Parameter*, you can reach the *Change View "OLAP Cache Parameters" View: Detail* screen. Make any changes that you want to make here.

Features

Cache Inactive

Setting this parameter means that the cross-transactional cache is deactivated centrally.

Local Size MB: Local cache size in MByte

This parameter determines the memory size of the local OLAP processor cache.

Note that the local cache is always used, also in the following cases:

- When the cross-transactional cache was deactivated (see the parameter *Cache Inactive*).
- When the cache was deactivated for the InfoProvider or the query (see [Cache Mode \[Seite 14\]](#))
- When it can be seen at runtime that the data cannot be buffered (see [Technical Information \[Seite 16\]](#), under *OLAP- or Cache-relevant Data: The Storage table Sp Cannot be Permanently Buffered*)

Global Size MB: Global cache size in MByte

This parameter determines the maximum memory size that all objects can use in the cross-transactional cache.

Memory use means the amount of memory used by the runtime objects. The actual memory use is generally smaller, since the runtime objects are written to the cross-transactional application buffer in compressed form.

The actual size of the cross-transactional cache depends on the size settings of the cross-transactional application buffer: The maximum cache size is the minimum from the *Size Global MB* parameter and the memory size that is actually available in the cross-transactional application buffer.

Persistence Mode

This parameter determines what happens to the data if the maximum size for the memory is reached. A proportion of the data must either be removed or swapped. A process using the Least Recently Used (LRU) algorithm determines which data is affected by this.

You can choose from the following modes:

- **Inactive:** When the cache memory has been fully exhausted, data is removed, i.e. deleted, from it.
- **Flatfile:** When the memory space available for caching has been exhausted, the data is swapped out to a file (see *Flat File Name*).
- **Cluster Table:** When the memory space available for caching has been exhausted, the data is stored in a cluster table in the database.

Flat File Name: Cache Persistence Logical File Name

Logical file name under which flat files are stored (see [Logical File Names \[Seite 29\]](#)).



BW_OLAP_CACHE



Cache Data Removal and Swapping

The overview of current memory consumption in the *Shared Memory* tab page in the *Cache Monitor* has the following values:

Column	Meaning
Maximum Cache Size	200MB by default
Current Cache Size	Sum of the size of all cache structure elements in KB (see Cache Struktur [Seite 27] , <i>Bytes</i> column)
Current Swap Size	Size of background store (flat file or cluster table) in KB
Cache reserved	Ratio between <i>Maximum Cache Size</i> and <i>Current Cache Size</i> in %
Current Entries, Total	Sum of <i>Current Cache Entries</i> and <i>Current Swap Entries</i> .
Current Cache Entries	Number of all cache entries, that is, of all cache structur elements listed individually in the table on the right-hand side.(see Cache Structure [Seite 27])
Current Swap Entries	Number of all entries in the background store
Buffer Poll Time	Time when buffer capacity was last read
Buffer Reserved	Degree to which the cache memory is used in %. This value is the same as the minimum free bytes and/or free directory entries. If you want to look at these limiting sizes in greater detail, choose Buffer Monitor or Buffer Overview (see OLAP Cache Monitor [Seite 19]).

When the capacity of the cache (the maximum cache size) is exhausted, and more data is to be written to the cache, the following solutions are available:

- Data is removed from the cache (deleted)
- Data is stored in a background store (swapped).

The following section outlines the basic principle of both procedures using status diagrams:

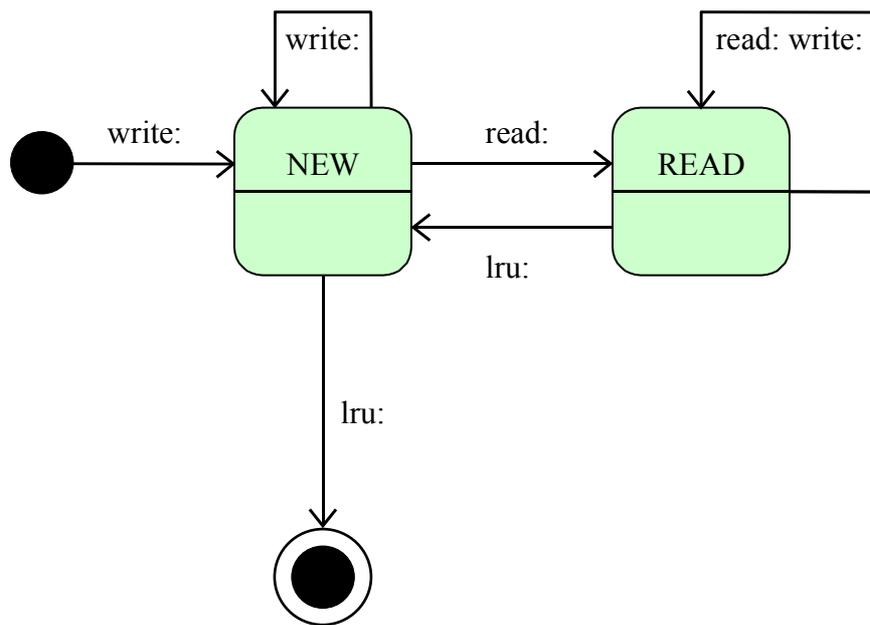
Theory: Caching with Removal

When data is written to the cache, the entry has the status NEW. (A new write receives this status).

When data is written from the cache, the entry has the status READ. (A new read and write receives this status).

The LRU replacement mechanism starts as soon as the cache memory capacity has been exhausted. This checks the status of entries and removes the entry that was last read the longest time ago.

- When the LRU algorithm comes across an entry with the status READ, it resets it to NEW.
- When the LRU algorithm comes across an entry with status NEW, it makes sure that this entry is overwritten with the new data. If you need to access the overwritten data again afterwards, this data must be read from the database again (cache miss).



Theory: Caching with Storage in a Background Store (Swap)

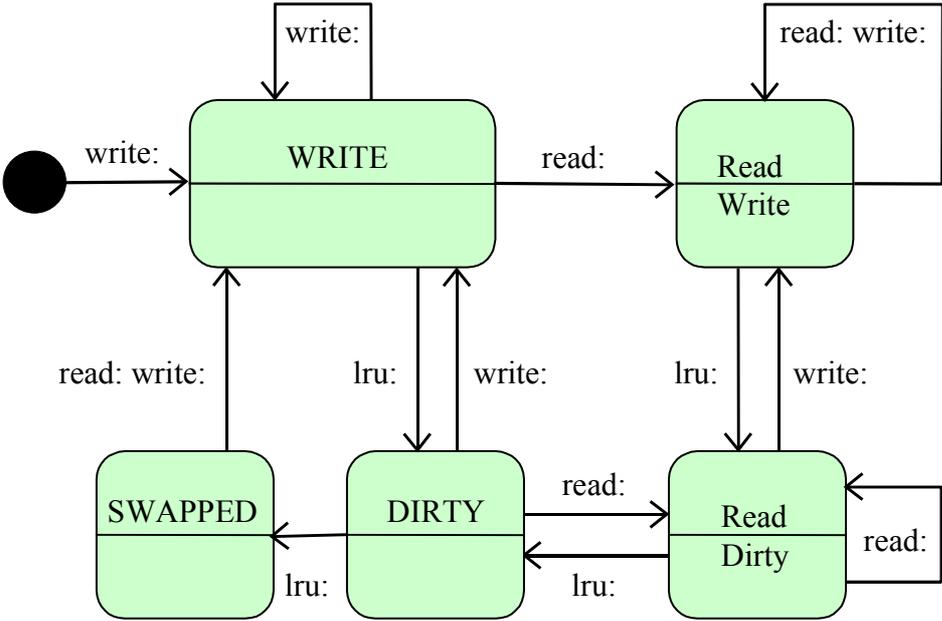
When data is written to the cache, the entry has the status WRITE. (A new write receives this status).

When data is written from the cache, the entry has the status READ WRITE. (A new read and write receives this status).

The LRU replacement mechanism starts as soon as the cache memory capacity has been exhausted. This checks the status of entries and removes the entry that was last read the longest time ago.

- When the LRU algorithm comes across an entry with the status READ WRITE, it resets it to READ DIRTY. (A new read receives this status). A new write returns the status of the entry to READ WRITE.) The DIRTY flag serves as a “lookout”, making sure that this entry is persisted.
- When the LRU algorithm comes across an entry with the status READ DIRTY, it resets it to DIRTY. The LRU algorithm does the same with entries having the status WRITE that are not read. (A new read resets the entry to the status READ DIRTY, a new write to the status WRITE).
- When the LRU algorithm comes across an entry with status DIRTY, it makes sure that this entry is stored in the background memory and indicated as SWAPPED. If the data needs to be subsequently accessed again, the data can be read from the background memory. At the same time, a new caching entry is created having the status WRITE.

After having run through the entries a maximum of two times, the LRU algorithm, by resetting the flags, has found an entry that can be overwritten.



Cache Structure

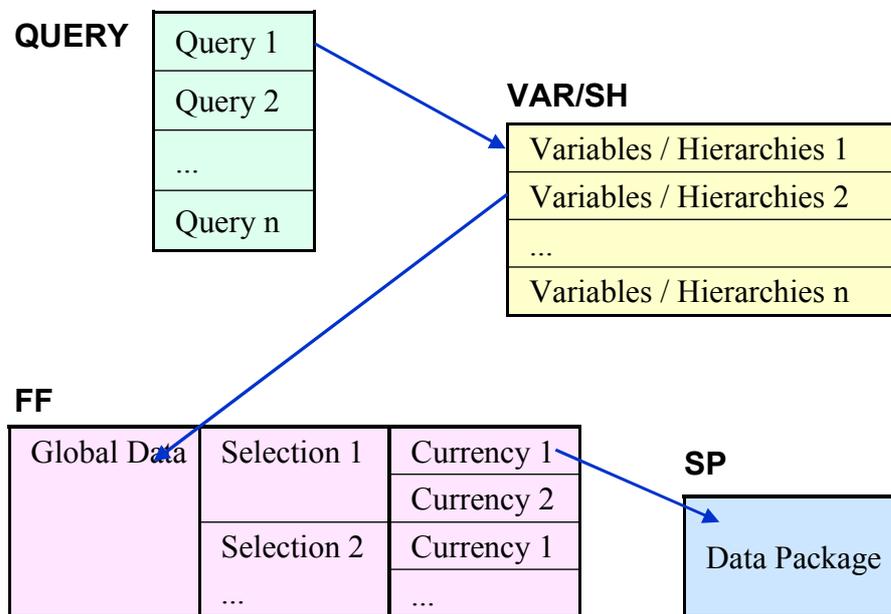
The cross-transactional OLAP cache contains the following elements per query request.

- **QUERY:** Table with queries in the cache (directory).
- **VAR/SH:** Complex structure for the variables and for the hierarchy that were requested for the query.
- **FF:** Complex structure for the entire selection for the query.
- **SP** (storage table): data package, a sequence of bytes.

Logically, the cross-transactional cache is structured as a hierarchy. This view corresponds to the  *Hierarchical Display* on the *Shared Memory* tab page.

Physically, all the elements are stored flat in the cache. Their sequence reflects the order in which they were created. This view corresponds to the  *List Display* on the *Shared Memory* tab page.

The following graphic shows the relationship between the different OLAP cache structure elements.



Display on the Shared Memory tab page.

Column	Meaning
Columns in the  <i>List Display</i> and in the  <i>Hierarchical Display</i> :	
Memory ID	Cache structure element ID
Swapped	Indicator is set when data from the cache is stored in the background store (swapped).
Read Flag	Indicator is set when data is read from the cache.
Write Flag	Indicator is set when data was written to the cache and the cache mode is actively persistent.
Dirty Flag	Indicator is set when data had to be written to the background store before data removal

	(cache mode = persistently active).
Directory	Query directory
Bytes	Cache structure element size, in bytes
Buffer ID	Cache structure element number sequence in the memory (also part of flat file name = <PARAM_1>, see Logical File Names [Seite 29])
Logical ID	Cache structure element internal logical ID
Columns that are only shown in the  <i>List Display</i> :	
Parent Memory ID	Cache structure element ID that, logically, lies on the next highest hierarchy level.
Parent Logical ID	Internal logical ID of parent element

Further information about the meaning of the individual flags can be found under [Cache Data Removal and Swapping \[Seite 24\]](#).



Logical File Names

Use

You need to have made the following settings before being able to save excess data removed from the cache in a flat file in the form of a background store.

- Defined the logical file path and assigned a platform-specific physical path.
- Defined the logical file name

From the *Cache Monitor* screen, via  *Logical File Names*, you can access the complex table maintenance for logical file paths and -names. The table maintenance consists of several sub-dialogs (views). SAP has already delivered a proportion of the required settings.

You need to check these settings and make any adjustments necessary for your file system before you can use the persistent OLAP cache with a flat file as a background store.



Note that these assignments are subject to settings made for the transport connection. (You need a request that can be transported).

Features

The following section gives an overview of the sub-dialog in the table maintenance using an example configuration. Further detailed information about the procedure can be found under [Defining Logical Path- and File Names \[Extern\]](#).

You can choose the following sub-dialogs via the Dialog Structure:

Logical File Path Definition

The (platform-independent) logical file path is evaluated at runtime to get the (platform-dependent) physical path (see *Assignment of Physical Paths to Logical Path*).



Logical file path: BW_OLAP_CACHE

Name: BW OLAP Cache

Assignment of Physical Paths to Logical Path

At runtime, the physical path assigned for the corresponding operating system platform is generated from the logical file path.

Since the physical path is platform-dependent, the customer must set it her/himself. The nearer the file is to the application server, the quicker the system can find and read it.



Logical path: BW_OLAP_CACHE

Name: BW OLAP Cache

Syntax group: UNIX (all Unix platforms)

Physical path: /usr/<SYSID>/global/<FILENAME> (You can choose the path up to <FILENAME> according to your own configuration.)



If you cannot find the syntax group for file names and paths for your operating system, you can define your own (see *Syntax Group Definition and Assignment of Operating System to Syntax Group*).

The parameter <FILENAME> must always be set for the physical file name (see *Logical File Name Definition, Cross-Client*). At runtime, the complete platform-specific file name is automatically constructed from the physical path and the physical file name.

Logical File Name Definition, Cross-Client

At runtime, a complete platform-specific file name including the path is generated for a (platform-independent) logical file name.



Logical file: BW_OLAP_CACHE (We recommend you use the delivered attributes for the logical file).

Name: BW OLAP Cache Files

Physical file: CACHE_<HOST>_<SYSID>_<PARAM_1>

Data format: BIN (binary)

Application area: BW (Business Information Warehouse)

Logical path: BW_OLAP_CACHE (see *Logical File Path Definition*). Input help is available.



The **physical file name** must have the following schema:

CACHE_<HOST>_<SYSID>_<PARAM_1>

The variables (not replaced with current values until runtime) have the following definitions:

<HOST> Computer-/server name (for SY-HOST)

<SYSID> System ID

<PARAM_1> external parameter 1 (sequence number in HEX)

Defining Variables

If you want to use variables in physical file names and –paths, you can define these via *New Entries*.

Syntax Group Definition

If required, you can define a new syntax group for operating systems with the same syntax for file names and paths via *New Entries*.



Syntax grp: UNIX

Name: all Unix platforms

Length (of file name)

Extension (determines whether the data format of the logical file name is allowed as a file name extension for this syntax group)

Active (specifies whether an operating system for this syntax group is used in the given system configuration).

Assignment of Operating System to Syntax Group

If required, you can assign your operating system to a syntax group via *New Entries*.



OP system: SunOS

Name: SUN Unix

