

DB6 Partitioning Administrator Quick Guide



Applies to:

SAP NetWeaver 7.00 and higher, Application Server ABAP on DB2 for Linux, UNIX, and Windows. For more information, visit the [Landscape Design and Architecture homepage](#).

Summary

The DB6 Partitioning Administrator is a new ABAP based tool for SAP OLTP Systems which run on DB2 for Linux, UNIX, and Windows. It helps to perform table range partitioning on a specified set of tables which usually grow large. The partitioning takes place on an indirectly time dependent criteria (on numbers drawn from number ranges) which allows to distinguish between 'cold' and 'hot' data.

Author: Johannes Heinrich

Company: SAP AG

Created on: 17 June 2011

Author Bio

Johannes Heinrich works as a senior developer in the IBM DB2 for Linux, UNIX, and Windows development team at SAP. He holds a degree in Computer Science from the University of Kaiserslautern, Germany and is a IBM Certified DBA for DB2 9 on Linux, UNIX, and Windows.

Table of Contents

| | |
|---|----|
| Aim..... | 3 |
| Scope..... | 3 |
| Suitable Tables | 3 |
| Prerequisites | 3 |
| Limitations..... | 4 |
| Local and Global Indexes | 4 |
| Installation..... | 4 |
| Usage..... | 5 |
| Calculate Partitions | 5 |
| Performance of Report RSDB6PARTGEN | 6 |
| Results of Report RSDB6PARTGEN..... | 6 |
| Apply the Partitioning..... | 7 |
| Partition the table using DB6CONV | 7 |
| Check the Partitioning..... | 8 |
| Converting a Partitioned Table Back to a Normal Table | 9 |
| Partitioning Scheme..... | 9 |
| Partition Types | 9 |
| Partition Names..... | 10 |
| Tablespace Names | 10 |
| Feedback | 11 |
| Related Content..... | 12 |
| Copyright..... | 13 |

Aim

We expect the following benefits from range partitioned tables:

- Easier maintenance of large database objects
- Separation between operational ('hot') and aged ('cold') data
- Performance improvements

Scope

Currently the DB6 Partition Administrator offers the following functionality:

- Calculate suitable partitions for 30 of the largest SAP OLTP tables (Report RSDB6PARTGEN). Tables which are candidates for partitioning have an entry in table DB6PARTADMIN.
- Carry out the range partitioning of these tables (using DB6CONV).
- Provide detailed information about the partitions of a partitioned table (Report RSDB6PARTMON).

It is planned to provide another report which can be used for partition maintenance. This report will

- Provide new partitions for forthcoming data, if needed.
- Merge older partitions with a low filling degree, if possible.
- Require a SAP system based on the forthcoming DB2 version 10.

Suitable Tables

The following tables - if they exist and are filled with at least 100.000 records - will be considered for range partitioning:

BDCP, BDCP2, BDCPS, CDCLS, CDHDR, CKIS, CKMI1, DBERCHZ1, DBERCHZ3, DBERDL, DFKKOP, DFKKREP06, EBAN, EDI40, EDIDC, EDIDS, EKBE, EKET, EKKN, EKKO, GLPCA, IMRG, LIPS, RESB, SWWLOGHIST, SWWWIHEAD, SWW_CONT, VBAP, VBRP

The following additional tables will be considered in a forthcoming version of the DB6 Partitioning Administrator:

BSIS, COEP, RFBLG, MKPF, MSEG, RSEG

Prerequisites

To use the DB6 Partition Administrator the following prerequisites must be fulfilled:

- The SAP system must run with DB2 for Linux, UNIX, and Windows version 9.7 or higher.
- The SAP system must be based at a SAP NetWeaver 7.00 or higher.
- The required changes for range partitioned tables in the ABAP DDIC as described in SAP note [1379362](#) must be implemented in the SAP system, either through an appropriate SP level or through the correction instructions attached to the SAP note.
- The SAP system must be a Unicode system.
- The SAP user who runs the DB6 Partitioning Administrator must have the authorization S_ADMI_FCD (same as for DBACOCKPIT).
- In DBACOCKPIT, a password must be set for the DB2 instance owner so that connections via the DB6ADM connection are possible. See SAP note [1072206](#) for the necessary steps. Without this connection the DDL statements for the table (as it currently exists) and the tablespaces cannot be extracted from the database.
- The table must have at least 100.000 records.

- Statistics for the table must be available and up to date.
- Number ranges in the corresponding number range object must be defined and used.

All these prerequisites are checked programmatically during the execution of report RSDB6PARTGEN.

Limitations

The DB6 Partition Administrator currently has the following limitations:

- Partitioning takes place for the records in one specified client only. All records in other clients will end up in two special partitions.
- External number ranges are not supported. If the number range objects which correspond to the table has external number ranges defined the table cannot be partitioned.
- Tables which use subobjects and/or a fiscal year in their number range definition cannot be partitioned. This affects the tables BSIS, COEP, MKPF, MSEG, RFBLG and RSEG.
- It is not possible to re-partition an already partitioned table. Before calculating and applying another partitioning schema the table must be de-partitioned first.

Local and Global Indexes

DB2 will use local (partitioned) indexes whenever possible. However, there is a restriction that the fields of the partitioning key must be a subset of the fields of every unique index. If a unique index exists which does not have all fields of the partitioning key included it will stay as a global (non-partitioned) index. Apart from customer defined unique indexes (unlikely) this restriction affects only table GLPCA where index GLPCA~2 cannot be partitioned.

Report RSDB6PARTGEN will check all defined indexes and write a warning to the application log if it finds an index which cannot be partitioned. Once the table is partitioned you cannot create unique indexes which do not contain the fields of the partitioning key.

Installation

Together with this documentation you should have received two transports, one for the DB6 Partitioning Administrator and one for a slightly improved version of the DB6CONV tool.

Just import both transports into your SAP system. The archive for DB6CONV contains a .txt file with more detailed instructions for the import if DB6CONV is already installed in your system.

Usage

Calculate Partitions

Go to SAP transaction SE38 and execute report RSDB6PARTGEN. You will be asked for the following input values.

DB6 Partition Administrator: Partition Generation

Table Information

| | |
|---------------------|---|
| Table Name | <input checked="" type="checkbox"/> ▼ |
| Client | <input checked="" type="checkbox"/> |
| Partition Size (GB) | <input style="width: 80%;" type="text" value="20"/> |
| Empty Partitons (%) | <input style="width: 80%;" type="text" value="20"/> |

Options

- Use separate tablespaces
- Prepare DB6CONV conversion

- *Table Name*: The drop-down list contains all partitioning candidates which are available in the system. Select the table you like to partition.
- *Client*: Select the client in which the table data is located. Partitions will be created only for the data in the client specified here.
- *Partition size*: Specify the size of each data partition in GBytes. The report tries to calculate partition boundaries so that each partition has approximately the same size. You can specify a partition size between 1 and 500 GB here, we recommend to use a value between 2 GB for smaller and 20 GB for larger tables.
- *Empty Partitions*: The number of spare partitions to be added in relation to the calculated partitions for the existing data. Example: The report calculates that 50 data partitions are needed to partition the existing table data. If you specified 20 percent empty partitions, the report will add 10 (20 percent of 50) spare partitions of the same size to accommodate forthcoming records inserted into the table.
- *Use separate tablespaces*: Indicates if all partitions should be located in the regular, index and long tablespace in which the table is currently located or if separate tablespaces should be used for every partition. If this box is checked (default), a DB2 command line script will be generated which can be used to create the required tablespaces.
- *Prepare DB6CONV conversion*: If checked you can use report DB6CONV immediately after report RSDB6PARTGEN finished successfully to actually convert the non-partitioned table to a partitioned table. Uncheck this box if you plan to carry out the partitioning manually on the command line, e.g. by using the DB2 UDF ADMIN_MOVE_TABLE.

Performance of Report RSDB6PARTGEN

The performance of the report is mostly determined by a SELECT statement which uses DB2's ROWNUM OVER() function to determine the partition boundaries. This SELECT results in an index scan on the primary index of the table under investigation and it is repeated for every defined number range interval.

You can start report RSDB6PARTGEN in batch by defining a variant and creating an appropriate job in SAP transaction SM36.

Results of Report RSDB6PARTGEN

After execution the report presents the results of its calculations in the form of application log entries.

Display logs

| Overview | Nu... |
|---|-------|
| ▼ DB6 Partition Administrator V0.9 | 53 |
| • Started at 11.02.2011 11:02:20 User: HEINRICHJO | 1 |
| ▶ INIT | 14 |
| ▼ PARTITION CALCULATION | 27 |
| • Partitioning key: RCLNT , GL_SIRID | 1 |
| • Index GLPCA~0 will be partitioned. | 1 |
| • Index GLPCA~1 will be partitioned. | 1 |
| • Index GLPCA~2 will be global! | 1 |
| • Index GLPCA~3 will be partitioned. | 1 |
| • Index GLPCA~7 will be partitioned. | 1 |
| ▼ Processing number range 01 | 10 |
| • Data partition GLPCA_001_01_X_X_00001 starting at 0000000000000000 | 1 |
| • Data partition GLPCA_001_01_X_X_00002 starting at 0000000000001170 | 1 |
| • Data partition GLPCA_001_01_X_X_00003 starting at 0000000000002340 | 1 |
| • Spare partition GLPCA_001_01_X_X_00004 starting at 0000000000003510 | 1 |
| • Spare partition GLPCA_001_01_X_X_00005 starting at 0000000000004680 | 1 |

Message Text

- PARTITION CALCULATION
- Partitioning key: RCLNT , GL_SIRID
- Index GLPCA~0 will be partitioned.
- Index GLPCA~1 will be partitioned.
- Index GLPCA~2 will be global!
- Index GLPCA~3 will be partitioned.
- Index GLPCA~7 will be partitioned.
- Processing number range 01
- Data partition GLPCA_001_01_X_X_00001 starting at 00000000000000903
- Data partition GLPCA_001_01_X_X_00002 starting at 0000000000001170947
- Data partition GLPCA_001_01_X_X_00003 starting at 0000000000002340978
- Spare partition GLPCA_001_01_X_X_00004 starting at 0000000000003510978
- Spare partition GLPCA_001_01_X_X_00005 starting at 0000000000004680978
- Logical max partition GLPCA_001_01_X_X_00006_MAX starting at 0000000000005850978 ending at 09999999
- 3 data partition(s) calculated
- 2 spare partitions calculated
- 1 logical max partition calculated
- Processing number range 02

The upper half of the screen holds an hierarchical view of the log entries generated by the report. By double-clicking on one entry the log is displayed again in the lower half of the screen. Here you can use the provided sorting and filtering functionality and export the log in various formats.

Furthermore, you will find two DB2 CLP scripts in the SAP system global directory (e.g. /usr/sap/<SID>/SYS/global/).

- The script tablespace_<tablename>_<timestamp>.sql contains the CREATE TABLESPACE statements for the tablespaces of every newly generated partition. This script will be only generated if 'Use separate tablespaces' was checked at the selection screen of report RSDB6PARTGEN.

- The script partition_<tablename>_<timestamp>.sql contains the complete CREATE TABLE statement with the calculated partitioning clause. It can be used to provide the table DDL directly to the DB2 UDF ADMIN_MOVE_TABLE.

Finally, the log as shown on the screen is saved to a file named partition_<tablename>_<timestamp>.log

Apply the Partitioning

You can use the following two methods to actually partition the table:

- Use DB6CONV to partition the table. DB6CONV perform the conversion offline or online. The online conversion is carried out with the help of ADMIN_MOVE_TABLE.
- Use ADMIN_MOVE_TABLE directly.

Partition the table using DB6CONV

If you decided to use separate tablespaces, create the tablespaces first by using the generated script.

1. Log on the an SAP application server with user <sid>adm.
2. Change to the SAP system global directory and list the latest scripts.
cd glo
ls -ltr
3. Review the generated tablespace script and adapt it if necessary.
4. Connect to DB2 as user db2<sid>.
db2 connect to <SID> user db2<sid>
5. Execute the script.
db2 -tvf tablespace__<tablename>_<timestamp>.sql
6. Check that all tablespaces were created successfully.

Now go to SAP transaction SE38 and execute report DB6CONV and create a new conversion job. We recommend to use the parameters as shown in the example below for the conversion.

DB6CONV 5.06 - New Conversion

Specify Target Data Class

Single Table Conversion Tablespace Conversion Long->LOB Conversion

Name of the Table to Be Converted:

Target Data Tablespace:

Target Index Tablespace:

Target Long Tablespace (Optional):

Online Conversion Offline Conversion

Use COPY WITH INDEXES

Stop Before Switch

Data Transfer Not Logged

Sort Data by Primary Key

COMMIT Size (MB)

Compress Data

Delete Source Table After Conversion

Update Statistics After Conversion

First enter the name of the table which is going to be partitioned. As *Target Data Tablespace*, you should specify the tablespace the *table is currently using* for regular data. You can do the same for the *Target Index Tablespace*. Note that the tablespace you define here will be used for non-partitioned indexes. The entries made in these two fields will determine the <tablespace-clause> of the newly created table. The partitions are placed into separate tablespaces (if requested) by using the <partition-tablespace-options> in the <partitioning-clause> of the table DDL. Finally, we recommend that you check the *Online Conversion* radio button (so that the table remains accessible during the conversion) and deselect the *Use COPY WITH INDEXES* checkbox (for performance reasons).

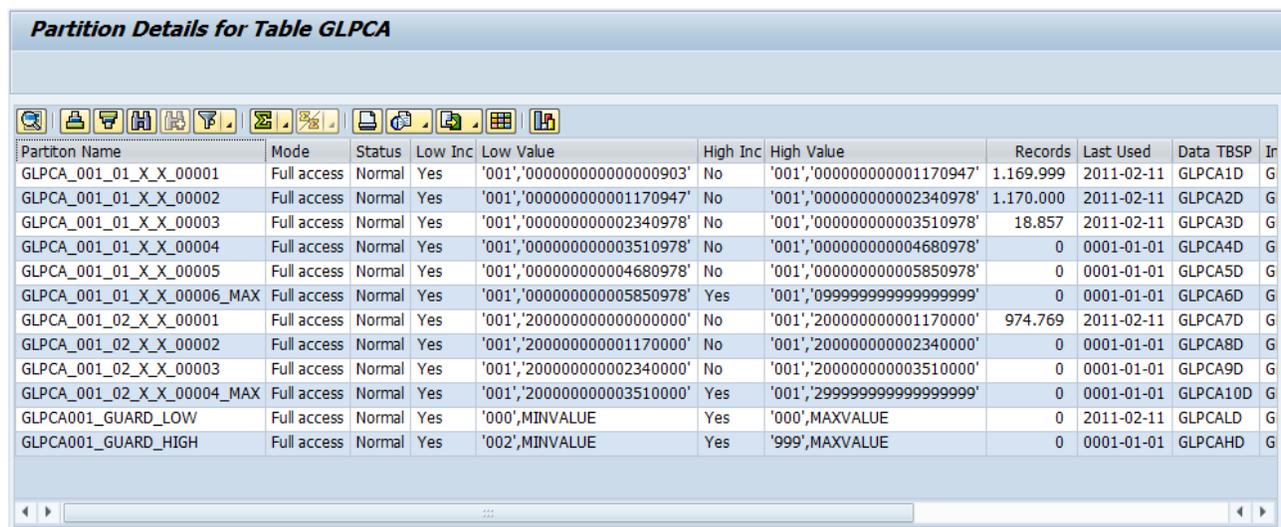
Save the input parameters and execute the conversion job. A detailed documentation for report DB6CONV is attached to SAP note [1513862](#).

Check the Partitioning

Go to SAP transaction SE38 and start the report RSDB6PARTMON. On the selection screen, enter the table name and hit F8. If the table is partitioned a grid control which contains details on the table partitions appears on the screen.

Example:

Partition Details for Table GLPCA



| Partition Name | Mode | Status | Low Inc | Low Value | High Inc | High Value | Records | Last Used | Data TBSP | Ir |
|----------------------------|-------------|--------|---------|----------------------------|----------|----------------------------|-----------|------------|-----------|----|
| GLPCA_001_01_X_X_00001 | Full access | Normal | Yes | '001','000000000000000903' | No | '001','000000000001170947' | 1.169.999 | 2011-02-11 | GLPCA1D | G |
| GLPCA_001_01_X_X_00002 | Full access | Normal | Yes | '001','000000000001170947' | No | '001','000000000002340978' | 1.170.000 | 2011-02-11 | GLPCA2D | G |
| GLPCA_001_01_X_X_00003 | Full access | Normal | Yes | '001','000000000002340978' | No | '001','000000000003510978' | 18.857 | 2011-02-11 | GLPCA3D | G |
| GLPCA_001_01_X_X_00004 | Full access | Normal | Yes | '001','000000000003510978' | No | '001','000000000004680978' | 0 | 0001-01-01 | GLPCA4D | G |
| GLPCA_001_01_X_X_00005 | Full access | Normal | Yes | '001','000000000004680978' | No | '001','000000000005850978' | 0 | 0001-01-01 | GLPCA5D | G |
| GLPCA_001_01_X_X_00006_MAX | Full access | Normal | Yes | '001','000000000005850978' | Yes | '001','099999999999999999' | 0 | 0001-01-01 | GLPCA6D | G |
| GLPCA_001_02_X_X_00001 | Full access | Normal | Yes | '001','200000000000000000' | No | '001','200000000001170000' | 974.769 | 2011-02-11 | GLPCA7D | G |
| GLPCA_001_02_X_X_00002 | Full access | Normal | Yes | '001','200000000001170000' | No | '001','200000000002340000' | 0 | 0001-01-01 | GLPCA8D | G |
| GLPCA_001_02_X_X_00003 | Full access | Normal | Yes | '001','200000000002340000' | No | '001','200000000003510000' | 0 | 0001-01-01 | GLPCA9D | G |
| GLPCA_001_02_X_X_00004_MAX | Full access | Normal | Yes | '001','200000000003510000' | Yes | '001','299999999999999999' | 0 | 0001-01-01 | GLPCA10D | G |
| GLPCA001_GUARD_LOW | Full access | Normal | Yes | '000',MINVALUE | Yes | '000',MAXVALUE | 0 | 2011-02-11 | GLPCALD | G |
| GLPCA001_GUARD_HIGH | Full access | Normal | Yes | '002',MINVALUE | Yes | '999',MAXVALUE | 0 | 0001-01-01 | GLPCAHD | G |

The columns in this table have the following meaning:

- *Partition Name*: Name of the partition. See below for an explanation of the naming schema for partitions.
- *Mode*: Current access restrictions as reported in field ACCESS_MODE of table SYSCAT.DATAPARTITIONS.
- *Status*: Current partition status as reported in field STATUS of table SYSCAT.DATAPARTITIONS.
- *Low Inc*: Low value included, YES or NO (field LOWINCLUSIVE from SYSCAT.DATAPARTITIONS).
- *Low Value*: Low key value for this partition (field LOWVALUE from SYSCAT.DATAPARTITIONS).
- *High Inc*: High value included, YES or NO (field HIGHINCLUSIVE from SYSCAT.DATAPARTITIONS).
- *High Value*: High key value for this partition (field HIGHVALUE from SYSCAT.DATAPARTITIONS).
- *Records*: Number of records in this partition as reported by field CARD in SYSCAT.DATAPARTITIONS.
- *Last Used*: Last used timestamp as reported by field LASTUSED of table SYSCAT.DATAPARTITIONS.

- *Data TBSP*: Tablespace used by this partition for regular data (as reported by field TBSPACEID in table SYSCAT.DATAPARTITIONS).
- *Index TBSP*: Tablespace used by the *partitioned (local)* indexes of this partition. Field INDEX_TBSPACEID from SYSCAT.DATAPARTITIONS.
- *Long TBSP*: Tablespace used by LONG/LOB data of this partition. Field LONG_TBSPACEID from SYSCAT.DATAPARTITIONS.
- *Data size*: Size of the regular data of this partition in MB as reported from UDF ADMIN_GET_TAB_INFO, field DATA_OBJECT_P_SIZE.
- *Index size*: Size of the partitioned (local) indexes of this partition in MB as reported from UDF ADMIN_GET_TAB_INFO, field INDEX_OBJECT_P_SIZE.
- *Long size*: Size of the long/LOB data of this partition in MB as reported from UDF ADMIN_GET_TAB_INFO, field LONG_OBJECT_P_SIZE.

Converting a Partitioned Table Back to a Normal Table

You can use the following procedure to convert a partitioned table back to a non-partitioned table:

1. Go to SAP transaction SE14 and enter the table name.
2. Click on *Storage Parameters*
3. Click on *For new Creation*
4. Select *Computed from technical Settings*
5. If necessary adapt the storage parameters, e.g. add COMPRESS for row compression.
6. Save the storage parameters.
7. Perform a normal table conversion with DB6CONV as described above.

Partitioning Scheme

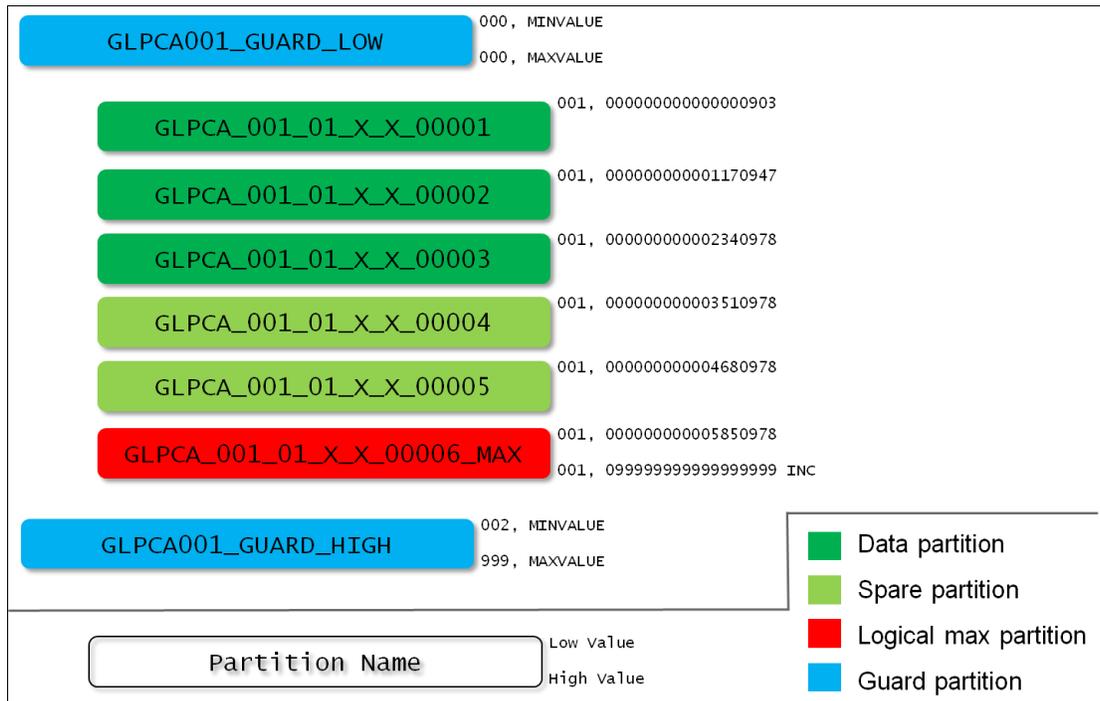
Partition Types

On a conceptual level, we differentiate between the following four partition types:

- *Data partitions* are partitions which hold the existing data after the conversion to a partitioned table took place.
- *Spare partitions* are empty partitions which are added to the partitioning schema to hold records which will be inserted into the table in the future.
- One *logical max partition* will be added from the end of the last spare partition to the end of the current number range interval. If all spare partitions are filled with data, records will end up in this logical max partition.
- Two *guard partitions* are defined to ensure that all data in other clients (e.g. test data) will end up in separate partitions.

Data partitions, spare partitions and the logical max partition are defined for every number range interval. The two guard partitions are defined once per partitioned table.

The following picture shows an example.



Partition Names

Partition names can be up to 128 characters long in DB2. Report RSDB6PARTEGN generates unique partitions which are composed of the following information:

- The name of the partitioned table
- The client
- The number range interval number
- The number range subobject (currently not used)
- The fiscal year of the number range (currently not used)
- An ascending number (local to every number range interval)
- For the logical max partition, the addition 'MAX'

Example: 'GLPCA_001_02_X_X_00004_MAX' is the logical max partition for table GLPCA in client 001 and number range interval 02. It is the fourth partition for this number range interval.

Partition names for the two guard partitions are composed of the table name, the client and this literal 'GUARD_LOW' and 'GUARD_HIGH' (e.g. 'GLPCA001_GUARD_LOW')

Tablespace Names

Tablespace names can be only up to 18 characters long in DB2 9.7, despite the fact that the system catalog already reserves 128 characters for tablespace names. To generate unique tablespace names for every partition, report RSDB6PARTGEN simply concatenates the table name, an ascending number and the literal 'D', 'I' or 'L' (for regular, index and long data).

Example: GLPCA9D is the 9th tablespace generated for regular data of the partitioned table GLPCA. You can use report RSDB6PARTMON to find out in which partition this tablespace is used.

The tablespace names for the two guard partitions are composed of the table name, the literal 'L' or 'H' (for 'low' or 'high' and the D, I or L literal as above. Example: GLPCALI is the tablespace for the local index data of the lower guard partition.

Feedback

Please send your questions, suggestions or comments to johannes.heinrich@sap.com.

Related Content

IBM DB2 9.7 Information Center topic [Partitioned tables](#)

SDN article [Using DB2 for Linux, Windows, and UNIX Table Partitioning in SAP Systems](#)

SAP Note 1379362 [DB6: Support for partitioned tables in the ABAP DDIC](#)

For more information, visit the [Landscape Design and Architecture homepage](#).

Copyright

© Copyright 2011 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, Excel, Outlook, and PowerPoint are registered trademarks of Microsoft Corporation.

IBM, DB2, DB2 Universal Database, System i, System i5, System p, System p5, System x, System z, System z10, System z9, z10, z9, iSeries, pSeries, xSeries, zSeries, eServer, z/VM, z/OS, i5/OS, S/390, OS/390, OS/400, AS/400, S/390 Parallel Enterprise Server, PowerVM, Power Architecture, POWER6+, POWER6, POWER5+, POWER5, POWER, OpenPower, PowerPC, BatchPipes, BladeCenter, System Storage, GPFS, HACMP, RETAIN, DB2 Connect, RACF, Redbooks, OS/2, Parallel Sysplex, MVS/ESA, AIX, Intelligent Miner, WebSphere, Netfinity, Tivoli and Informix are trademarks or registered trademarks of IBM Corporation.

Linux is the registered trademark of Linus Torvalds in the U.S. and other countries.

Adobe, the Adobe logo, Acrobat, PostScript, and Reader are either trademarks or registered trademarks of Adobe Systems Incorporated in the United States 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, ICA, Program Neighborhood, MetaFrame, WinFrame, VideoFrame, and MultiWin are trademarks or registered trademarks of Citrix Systems, Inc.

HTML, XML, XHTML and W3C 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, R/3, SAP NetWeaver, Duet, PartnerEdge, ByDesign, SAP Business ByDesign, 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 other countries.

Business Objects and the Business Objects logo, BusinessObjects, Crystal Reports, Crystal Decisions, Web Intelligence, Xcelsius, and other Business Objects products and services mentioned herein as well as their respective logos are trademarks or registered trademarks of Business Objects S.A. in the United States and in other countries. Business Objects is an SAP company.

All other product and service names mentioned are the trademarks of their respective companies. Data contained in this document serves informational purposes only. National product specifications may vary.

These materials are subject to change without notice. These materials are provided by SAP AG and its affiliated companies ("SAP Group") for informational purposes only, without representation or warranty of any kind, and SAP Group shall not be liable for errors or omissions with respect to the materials. The only warranties for SAP Group 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.