IBM DB2 Analytics Accelerator for z/OS Version 4.1.0

Installation Guide



Note

Before using this information and the product it supports, read the information in "Notices" on page 129.

Second Edition, March 2014

This edition applies to Version 4.1.0 of IBM DB2 Analytics Accelerator for z/OS, program number 5697-DAB, and to all subsequent releases and modifications until otherwise indicated in new editions. This edition replaces SH12-7038-00. Changes to this edition are marked with a vertical bar.

This product includes software developed by Macromedia, Inc. (http://www.macromedia.com).

© Copyright IBM Corporation 2009, 2014.

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Figures v
About this book
Who should read this book ix
What's new in version 4? xi
How to send your comments xiii
Chapter 1. Migrating from earlierversions1Migrating HPSS partitions4Fixing HPSS issues after the migration4
Chapter 2. Configuring TCP/IP for connections between IBM System z and the IBM PureData System for Analytics . 7
Chapter 3. Security 9
Chapter 4. Installing IBM DB2 Analytics Accelerator Studio
Chapter 5. Enabling an existing DB2 subsystem for IBM DB2 Analytics Accelerator for z/OS
Chapter 6. Setting up IBM DB2 Analytics Accelerator for z/OS 27 Setting up a WLM application environment for IBM DB2 Analytics Accelerator for z/OS stored procedures

Defining WLM performance goals for IBM DB2 Analytics Accelerator for z/OS stored procedures Adjusting WLM performance goals for	. 32
SYSPROC.ACCEL_LOAD_TABLES	. 32
AOTTIISP	34
Customizing and munning AOTTUCD	. 01 25
	. 33
Verifying the installation of IBM DB2 Analytics	
Accelerator for z/OS stored procedures	. 36
Customizing and running AQTSJI00	. 37
Customizing and running AQTSJI01	. 38
Chapter 7 Connecting IBM DB2	
Analytics Accelerator for z/OS and DB2	39
Required authorizations	30
Required authorizations	. 39
Binding DB2 packages and granting user privileges	40
Creating EXPLAIN tables.	. 41
Creating a database connection profile	. 42
Testing the connection from IBM DB2 Analytics	
Accelerator Studio	. 44
Obtaining the pairing code for authentication	45
Completing the authentication using the Add	. 10
A sealerster scient	47
	. 4/
Testing the stored procedures	. 48
Customizing AQTSJI02	. 50
Customizing AQTSXTCO.	. 50
Customizing AOTSCI02	. 50
Running AOTSII02	E1
	<u> </u>
	. 51
Chapter 8. Testing query acceleration	53
Chapter 8. Testing query acceleration	51 53
Chapter 8. Testing query acceleration Selecting a test query	53 . 53
Chapter 8. Testing query acceleration Selecting a test query Defining tables Loading and apphling tables	53 . 53 . 53
Chapter 8. Testing query acceleration Selecting a test query Defining tables Loading and enabling tables	53 53 53 53 53
Chapter 8. Testing query acceleration Selecting a test query Defining tables Loading and enabling tables Running a test query	53 53 53 53 53 53 53
Chapter 8. Testing query acceleration Selecting a test query Defining tables Loading and enabling tables Running a test query Chapter 9. Incremental updates	53 53 53 53 53 53 54 57
Chapter 8. Testing query acceleration Selecting a test query Defining tables Loading and enabling tables Running a test query Chapter 9. Incremental updates Concepts and architecture	53 . 53 . 53 . 53 . 53 . 54 57
Chapter 8. Testing query acceleration Selecting a test query Defining tables Loading and enabling tables Running a test query Chapter 9. Incremental updates Integration	53 53 53 53 53 53 54 57 59 60
Chapter 8. Testing query acceleration Selecting a test query Defining tables Loading and enabling tables Running a test query Chapter 9. Incremental updates Integration Concepts and architecture Selecting Concepts and architecture Concepts	53 53 53 53 53 53 54 57 59 60
Chapter 8. Testing query acceleration Selecting a test query Defining tables Loading and enabling tables Running a test query Chapter 9. Incremental updates Integration Completing the installation and configuring CDC Configuration settings for IBM InfoSphere	53 53 53 53 53 53 54 57 59 60 61
Chapter 8. Testing query acceleration Selecting a test query Defining tables Loading and enabling tables Running a test query Chapter 9. Incremental updates Integration Concepts and architecture Integration Completing the installation and configuring CDC Configuration settings for IBM InfoSphere Change Data Capture for z/OS.	53 53 53 53 53 53 53 54 57 57 60 . 61
Chapter 8. Testing query acceleration Selecting a test query Defining tables Loading and enabling tables Running a test query Chapter 9. Incremental updates Concepts and architecture Integration Completing the installation and configuring CDC Configuration settings for IBM InfoSphere Change Data Capture for z/OS Chapting DATA CAPTURE CHANGES	53 53 53 53 53 53 53 53 54 57 59 60 61 61
Chapter 8. Testing query acceleration Selecting a test query Defining tables Loading and enabling tables Running a test query Chapter 9. Incremental updates Concepts and architecture Integration Completing the installation and configuring CDC Configuration settings for IBM InfoSphere Chapter Data Capture for z/OS Enabling DATA CAPTURE CHANGES	53 53 53 53 53 53 53 53 54 57 60 61 61 62 64
Chapter 8. Testing query acceleration Selecting a test query Defining tables Loading and enabling tables Running a test query Chapter 9. Incremental updates Concepts and architecture Integration Completing the installation and configuring CDC Configuration settings for IBM InfoSphere Chapter Data Capture for z/OS Enabling log caching Enabling log caching	53 53 53 53 53 53 53 53 53 54 57 59 60 61 61 62 64 64
Chapter 8. Testing query acceleration Selecting a test query Defining tables Loading and enabling tables. Running a test query Chapter 9. Incremental updates Integration Concepts and architecture Integration Completing the installation and configuring CDC Configuration settings for IBM InfoSphere Change Data Capture for z/OS. Enabling DATA CAPTURE CHANGES Performance tuning.	53 53 53 53 53 53 53 53 53 53 53 53 53 54 57 60 61 . 62 . 64 . 64 . 64
Chapter 8. Testing query acceleration Selecting a test query Defining tables Loading and enabling tables Running a test query Chapter 9. Incremental updates Integration Concepts and architecture Integration Completing the installation and configuring CDC Configuration settings for IBM InfoSphere Change Data Capture for z/OS Enabling log caching Performance tuning. Starting CDC	53 53 53 53 53 53 53 53 54 57 59 60 61 61 62 64 64 64 64 64 5
Chapter 8. Testing query acceleration Selecting a test query Defining tables Loading and enabling tables Running a test query Chapter 9. Incremental updates Integration Concepts and architecture Integration Completing the installation and configuring CDC Configuration settings for IBM InfoSphere Change Data Capture for z/OS Enabling log caching Performance tuning. Starting CDC Installing and activating the Access Server and the	53 53 53 53 53 53 53 54 57 59 60 61 61 62 64 64 64 64 65
Chapter 8. Testing query acceleration Selecting a test query Defining tables Loading and enabling tables Running a test query Chapter 9. Incremental updates Integration Completing the installation and configuring CDC Configuration settings for IBM InfoSphere Chapting DATA CAPTURE CHANGES Enabling log caching Performance tuning. Starting CDC Installing and activating the Access Server and the	53 53 53 53 53 53 53 54 57 60 61 61 62 64 64 64 64 65 66
Chapter 8. Testing query acceleration Selecting a test query Defining tables Loading and enabling tables Running a test query Chapter 9. Incremental updates Integration Completing the installation and configuring CDC Completing the installation and configuring CDC Configuration settings for IBM InfoSphere Chapting DATA CAPTURE CHANGES Enabling log caching Performance tuning. Starting CDC Installing and activating the Access Server and the replication engine	53 53 53 53 53 53 53 54 57 60 61 61 62 64 64 64 64 65 66
Chapter 8. Testing query acceleration Selecting a test query Defining tables Loading and enabling tables. Running a test query Chapter 9. Incremental updates Concepts and architecture Integration Completing the installation and configuring CDC Configuration settings for IBM InfoSphere Chapting DATA CAPTURE CHANGES Enabling log caching Performance tuning. Starting CDC Installing and activating the Access Server and the replication engine	53 53 53 53 53 53 53 53 53 54 57 60 61 61 64 64 64 64 65 66 66
Chapter 8. Testing query acceleration Selecting a test query Defining tables Loading and enabling tables Running a test query Chapter 9. Incremental updates Concepts and architecture Integration Completing the installation and configuring CDC Configuration settings for IBM InfoSphere Chapting DATA CAPTURE CHANGES Enabling log caching Performance tuning. Starting CDC Installing and activating the Access Server and the replication engine Transferring packages for the Access Server or the replication engine	53 53 53 53 53 53 53 53 53 53 57 57 60 61 61 62 64 64 64 65 66 66
Chapter 8. Testing query acceleration Selecting a test query Defining tables Loading and enabling tables Running a test query Chapter 9. Incremental updates Concepts and architecture Integration Completing the installation and configuring CDC Configuration settings for IBM InfoSphere Chapting DATA CAPTURE CHANGES Enabling log caching Performance tuning. Starting CDC Installing and activating the Access Server and the replication engine Activating the Access Server and the replication	53 53 53 53 53 53 53 53 53 53 53 53 53 53
Chapter 8. Testing query acceleration Selecting a test query Defining tables Loading and enabling tables Running a test query Chapter 9. Incremental updates Concepts and architecture Integration Completing the installation and configuring CDC Configuration settings for IBM InfoSphere Chapter Data Capture for z/OS Enabling DATA CAPTURE CHANGES Enabling log caching Performance tuning. Starting CDC Installing and activating the Access Server and the replication engine Transferring packages for the Access Server or the replication engine	53 53 53 53 53 53 53 53 53 53 53 53 53 53
Chapter 8. Testing query acceleration Selecting a test query Defining tables Loading and enabling tables Running a test query Chapter 9. Incremental updates Concepts and architecture Integration Completing the installation and configuring CDC Configuration settings for IBM InfoSphere Chapter Data Capture for z/OS Enabling log caching Performance tuning. Starting CDC Installing and activating the Access Server and the replication engine Activating the Access Server and the replication engine Consignering packages for the Access Server or the replication engine Starting the Access Server and the replication Installing or disabling DB2 subsystems	53 53 53 53 53 53 53 53 53 53 53 53 53 53
Chapter 8. Testing query acceleration Selecting a test query Defining tables Loading and enabling tables Running a test query Chapter 9. Incremental updates Running a test query Concepts and architecture Integration Completing the installation and configuring CDC Configuration settings for IBM InfoSphere Chapting DATA CAPTURE CHANGES Enabling log caching Performance tuning. Starting CDC Installing and activating the Access Server and the replication engine Activating the Access Server and the replication engine Activating the Access Server and the replication engine Activating the Access Server and the replication engine Activating the Access Server and the replication engine Enabling or disabling DB2 subsystems Enabling incremental updates for a DB2	53 53 53 53 53 53 53 53 53 53 53 53 53 53

Disabling incremental updates for a DB2	
subsystem	70
Chapter 10 Beyond the basics	72
Alle estima recommende a la constante	72
Allocating resources of shared accelerators	73
mapping of wLM importance levels to Netezza	75
priorities	75
Dischling the conversion of 24.00.00 time column	70
Disabling the conversion of 24:00:00 time values	78
Chapter 11. Installing updates.	81
Installing prerequisite PTFs for DB2 10 for z/OS or	
DB2 11 for z/OS	82
Updating IBM DB2 Analytics Accelerator Studio	82
Using IBM Installation Manager to add or update	
IBM DB2 Analytics Accelerator Studio from a	
workstation without an internet connection	82
Using IBM Installation Manager to add or update	
IBM DB2 Analytics Accelerator Studio from a	
workstation with an internet connection	83
Differences to earlier versions and customization	83
Updating the IBM DB2 Analytics Accelerator	
software, the Netezza Performance Server (NPS), the	
Access Server, or the replication engine	84
Installing update packages for the accelerator	84
Transferring update packages for the accelerator	85
Switching between available versions of	
accelerator software components	87
Updating other Netezza software	89
Transferring Netezza Firmware (FDT) updates or	
Netezza Host Platform (HPF) updates	89
Removing obsolete software packages from an	
accelerator	90
Updating the z/OS components of IBM DB2	
Analytics Accelerator	91
Running different versions of IBM DB2 Analytics	
Accelerator for z/OS at the same time \ldots \ldots	93
Chapter 10 Troublachasting	05
Chapter 12. Troubleshooting	95
Message AQT102001 and SQL code -206 when	~
running stored procedures after migration	95
Message DSNUTILU NOT INVOKED APF AUTHORIZED	95
Error during bind step for stored procedure or	
	96
No CEEDUMPs	96
PKIQ value too high when creating table spaces	97
The DB2 command -DIS ACCEL does not work	97
Connection authorization failure (error -4214)	97
Errors during ZPAKM compilation	98
DKDA connection does not work	98
kemoving orphaned system-table entries and	00
catalog-table entries	98

Т

Package not found when running a stored procedure from IBM DB2 Analytics Accelerator	
Studio . Packet delivery failure	99 100
for z/OS stored procedures.	101
DSNAQT schema	102
error -471 and reason code E790002 for DSNUTILU Load of partitioned tables freezes during unload	102
phase	103
COLUMN DEFINITIONS	103
Repeated IBM InfoSphere Change Data Capture for z/OS errors after abnormal end of DB2 for z/OS .	104
Appendix A. Sample TCP/IP	
configuration 1	05
Appendix B. Members of SAQTSAMP 1	07
Appendix C. Environment variables	111
Appendix D. Using the DB2 command	
line processor from UNIX System	115
	15
Appendix E. Disabling accelerators in a DB2 subsystem	117
Appendix F. Enabling accelerators in a DB2 subsystem 1	19
Glossary 1	21
A	121 121
L	121
M	121
P	122
Τ	122
V	122
· ·	125
Index 1	25
Notices 1	29

Figures

- 1. Possible upgrades and downgrades 2
- 2. Assigning a medium-to-high priority service class to the SYSPROC.ACCEL_LOAD_TABLES stored procedure in the WLM ISPF application. 34
- 3. Interaction diagram of the components that are involved in the incremental update process. . . 60
- SQL code -430 message window 101
 Recommended connections between IBM
 - zEnterprise server and IBM PureData System for Analytics (including switches). 105

About this book

This book provides information about the following subjects:

- Installation of the software package for this product on a z/OS[®] data server and subsequent transfer of operating system and other software to an IBM[®] PureData[™] System for Analytics.
- Installation of a DB2[®] for z/OS Program Temporary Fix (PTF), which includes the DB2 software that is required to integrate this product into your DB2 environment.
- Configuration of the product from DB2 (invocation of stored procedures).
- Installation of the IBM DB2 Analytics Accelerator for z/OS client software.

Who should read this book

This book is intended for the following audiences:

- IBM Customer Service Centers (support personnel) that need to install this product at the client site.
- z/OS administrators who need to install the software for this product using SMP/E.
- DB2 for z/OS administrators who need to set up and configure this product for use with a DB2 for z/OS data server.
- Other system administrators who need to install and maintain the client software for this product on Linux or Microsoft Windows computers.

What's new in version 4?

This version of IBM DB2 Analytics Accelerator for z/OS, when used with the latest DB2 for z/OS support, includes the following new features:

- Support for acceleration of static SQL queries in DB2 for z/OS applications.
- DB2 for z/OS multi-row fetching of results from accelerated queries (a configurable number of rows can be retrieved in a single time unit).
- Workload balancing of accelerated queries. If a query can be executed on more than one accelerator, DB2 for z/OS automatically distributes the workload among qualifying accelerators. Accelerators with a low utilization are favored.
- High-performance storage saver (HPSS) improvements. A full-scale restore function has been added to the HPSS scope. This function takes advantage of the newest DB2 for z/OS Utilities support. In addition, an enhanced DB2 for z/OS Utility function now sets the affected table space to a persistent read-only state when the HPSS is invoked to move partition data. Furthermore, the same partition data can now be moved to multiple accelerators.
- Table-load performance improvements, resulting in a two-digit percentage of CPU-cost reduction.
- Incremental updates configurable for more than two DB2 subsystems.
- Automated transfer and installation (apply action) of updates for the Netezza Performance Server[®] (NPS[®]).
- Incremental updates continue while tables are being reloaded.
- Workload Manager (WLM) support for local DB2 for z/OS applications. That is, WLM priorities can be assigned to local applications to support your processing preferences.
- Support for DB2 11 for z/OS.

Enhancements in version 4 PTF-2:

PTF-2 of the product introduces the following enhancements:

- Migration stored procedure that adjusts table names in accordance with new naming scheme (required for workload balancing and the acceleration of static SQL queries)
- Support for TIME and TIMESTAMP columns containing 24:00:00 values (conversion no longer required)
- Support for FOR BIT DATA values in single-byte character EBCDIC tables
- Elimination of cast problems related to character data in EBCDIC and Unicode tables located on the same accelerator
- Support for lock-free table loading of incrementally updated tables, which allows INSERTs, UPDATEs, and DELETEs to continue
- Template support for temporary data sets created by the High Performance Storage Server (HPSS)

I

L

How to send your comments

Your feedback is important in helping to provide the most accurate and high-quality information.

If you have any comments about this book or any other IBM DB2 Analytics Accelerator for z/OS documentation:

• Go to the IBM DB2 Analytics Accelerator for z/OS support home page at: http://www.ibm.com/support/entry/portal/Overview/Software/ Information_Management/DB2_Analytics_Accelerator_for_z~OS

The link **Support feedback** > **Help us improve online support**, which is usually on the lower right, opens the feedback form, in which you can enter your comments.

- Send your comments by email to swsdid@de.ibm.com. Be sure to include the name of the book, the publication number of the book, the version of IBM DB2 Analytics Accelerator for z/OS, and, if applicable, the specific location of the text you are commenting on (for example, a page number or table number).
- Complete one of the forms at the back of this book and send it by fax, or give it to an IBM representative.

Chapter 1. Migrating from earlier versions

For a migration from version 3.1.0 to 4.1.0, update the individual components of IBM DB2 Analytics Accelerator for z/OS in the order that is suggested here.

You can only migrate from version 3.1.0 to 4.1.0. To upgrade version 2.1.x to 4.1.0, you must first migrate to version 3.1.0. The following diagram shows all upgrading and downgrading possibilites:



Figure 1. Possible upgrades and downgrades

Backward migration or downgrade

I

I I I I I Т

L I

Т As a general rule, a backward migration between product releases is supported unless you have used persistent features of the newer release. After migrating or updating IBM DB2 Analytics Accelerator for z/OS do not attempt a backward migration or downgrade if any of the conditions in Table 1 applies.

Reasons against a backward migration from version 4.1.0 PTF- <i>x</i> to 3.1.0 PTF- <i>x</i>	Reasons against a downgrade from version 4.1.0 PTF- <i>y</i> to 4.1.0 PTF- <i>x</i>
You have used the High Performance Storage Saver (HPSS) to move partitions of the same table to different accelerators.	The Netezza Performance Server (NPS) has already been updated to version 7.0.4 or higher (requirement for a successful update to product version 4.1.0 PTF-2 or higher).
You have defined tables with columns encoded in UNICODE and EBCDIC on the same accelerator with the help of the AQT_ENABLE_MULTIPLE_ENCODING environment variable.	
The Netezza Performance Server (NPS) has already been updated to version 7.0.4 or higher (requirement for a successful migration to product version 4.1.0 PTF-2 or higher).	

Table 1. Reasons against a backward migration or downgrade

In addition, if partitions have been unsuccessfully moved or restored in version 4.1.0, a backward migration is blocked. You must first solve the associated problems and then complete the move or restore operations to bring your installation into a state that allows a backward migration.

To manually restore partitions that were moved with version 4.1.0, see the *Restore* without active accelerator section in the description of the SYSPROC.ACCEL_RESTORE_TABLES stored procedure. You find it in the IBM DB2 Analytics Accelerator for z/OS: Stored Procedures Reference.

High Performance Storage Saver

If you have used the High Performance Storage Saver (HPSS) of version 3.1.0, make sure to read the following topics before the migration:

- 1. "Migrating HPSS partitions" on page 4
- 2. "Fixing HPSS issues after the migration" on page 4

Migrating

For instructions on how to migrate, see Chapter 11, "Installing updates," on page 81. Information that is specific to the migration from version 2.1.x to 3.1.0 or from 3.1.0 to 4.1.0 has been added to the *Before you begin* and *About this task* sections of the relevant topics.

Related information:



IBM DB2 Analytics Accelerator for z/OS: Stored Procedures Reference

Migrating HPSS partitions

If you have moved the data of partitions or entire tables with the High Performance Storage Saver (HPSS) that was introduced with version 3.1.0 of IBM DB2 Analytics Accelerator for z/OS, you must take a few preparatory steps before a migration to version 4.1.0

Before you begin

If you have moved tables or partitions to an accelerator with the High Performance Storage Saver (HPSS) of IBM DB2 Analytics Accelerator for z/OS Version 3.1.0, make sure that the move operations have been completed successfully and that none of the affected table spaces are in a state that prevents a successful migration to version 4.1.0, such as RECOVER PENDING.

Procedure

- 1. Identify tables that have been moved to an accelerator, either fully or partially. Use one of the following methods:
 - Open the appropriate Accelerator views in IBM DB2 Analytics Accelerator Studio. These views list all tables on an accelerator. The **Storage Saver Partitions** column in the **Tables** section indicates whether partitions of a table have been moved.
 - Run the SYSPROC.ACCEL_GET_TABLES_INFO stored procedure with the value NULL for the *table_set* parameter (this selects all tables) and check the <tableInformation> elements in the output for tables with an attribute setting of archiveStatus=true; this indicates that tables or partitions of the tables have been moved.
- 2. Find out which partitions have been moved. Use one of the following methods:
 - In a relevant Accelerator view, in the **Tables** section, click **Storage Saver**on the toolbar and select **Move Partitions to Accelerator** or **Restore Partitions to DB2**. The window that opens in either case gives you the required information.
 - Run the SYSPROC.ACCEL_GET_TABLES_DETAILS stored procedure against all tables with moved partitions to find out which partitions have been moved.
- 3. Run the DB2 command DISPLAY DB(dbname) SP(tsname) RESTRICT on the identified tables, where dbname is the database name and tsname is the table space name. This gives you the tables that are in an adverse state.
- 4. Take appropriate action to remove the adverse states.
- 5. Set the affected table spaces to the persistent read-only (PRO) state that was introduced with DB2 10 for z/OS. This state protects the partitions in DB2 against accidental changes after a move. The state was not automatically applied by the HPSS of IBM DB2 Analytics Accelerator for z/OS 3.1.0, so you must now set it manually before the migration.
- 6. Migrate to version 4.1.0.

Fixing HPSS issues after the migration

Follow the steps here to fix High Performance Storage Saver (HPSS) issues that resulted from or existed before a migration to IBM DB2 Analytics Accelerator for z/OS Version 4.1.0.

Procedure

- 1. Identify the HPSS partitions that are in an error state. Use one of the following methods:
 - In a relevant Accelerator view, in the **Tables** section, click **Storage Saver**on the toolbar and select **Move Partitions to Accelerator** or **Restore Partitions to DB2**. In the upcoming window, the partitions in an error state are marked with a warning icon.
 - Run the SYSPROC.ACCEL_GET_TABLES_INFO stored procedure with the value NULL for the *table_set* parameter (this selects all tables) and check the <tableInformation> elements in the output for tables with an attribute setting of archiveErrorStatus=DataInArchivedPartition or archiveErrorStatus=ArchivedPartitionNotReadonly; these are DB2 partitions that were accidentally updated after a move or that are not protected by the partition read-only (PRO) state.
- 2. Take appropriate action to resolve the conflicts:
 - In DB2 for z/OS, remove the data from each moved partition by calling the LOAD REPLACE utility with an empty data set as input.
 - For each moved partition not in the PRO state (not read-only), set the PRO state by running the REPAIR utility.

Chapter 2. Configuring TCP/IP for connections between IBM System z and the IBM PureData System for Analytics

To transfer data between your database management system and IBM DB2 Analytics Accelerator for z/OS, TCP/IP connections must exist between your IBM System z[®] and your IBM PureData System for Analytics.

Before you begin

- Appropriate cable connections must exist between the OSA-Express ports of your IBM System z and your IBM PureData System for Analytics.
- IBM System z and z/OS must be operational.
- The IBM PureData System for Analytics must be operational.

About this task

A private network must exist between System z and IBM DB2 Analytics Accelerator for z/OS. This means that you cannot reuse existing subnets that are already defined on other interfaces. The number of required TCP/IP addresses depends on the configuration. For the minimum configuration, provide at least four TCP/IP addresses in the same subnet. For the recommended configuration, provide six IP addresses.

- One IP address for the 10 GbE interface of Netezza® host 1
- One IP address for the 10 GbE interface of Netezza host 2
- One IP address (also called wall IP address) for DRDA[®] connectivity. This is a floating IP address, which is provided by the active Netezza host.
- One IP address for the OSA-Express card facing the IBM PureData System for Analytics.

The recommended configuration uses two OSA-Express cards, each of which using its own IP address in combination with a virtual IP address (VIPA).

To view a sample configuration, click the **Related reference** link at the end.

Related reference:

Appendix A, "Sample TCP/IP configuration," on page 105 In this section, you find a sample VTAM[®] definition and a sample profile for the TCP/IP setup on your System z.

Chapter 3. Security

The accelerator contains a copy of DB2 data for selected tables. It is important to protect that data against unauthorized access. A properly configured system environment makes it very hard for attackers from the outside to get into the system and manipulate, damage, or steal data. See which security measures are in place and which measures you can take to increase the system security.

Product security

The system environment of IBM DB2 Analytics Accelerator for z/OS can be considered safe for the following reasons:

- If properly configured, the System z data server and the IBM PureData System for Analytics (accelerator) are connected over a private data network (PDN), which does not allow access from the internet or your organization's intranet.
- Only TSO users with sufficient access rights for the relevant logical partition (LPAR) in z/OS can establish a TCP/IP connection from the LPAR to the accelerator over the PDN. A system authorization facility (SAF) like RACF[®] allows you to restrict this right to just a few selected users.
- Communication with the accelerator is always started from z/OS or DB2 for z/OS rather than vice versa.
- The transfer, manipulation or extraction of data from the accelerator is carried out by the stored procedures that come with the product. The authorizations required to execute these tasks lie, to a large extent, with the stored procedures, so that the rights of the user executing the stored procedures can be restricted.
- Queries just read the data on an accelerator. They do no change it.
- The pairing mechanism exclusively links a DB2 subsystem to an accelerator. A user with access to one subsystem cannot view or otherwise access data from another subsystem just because the other subsystem is linked to the same accelerator.
- The Linux system on the accelerator machine does not allow a direct remote connection to the machine. The built-in private authentication module (PAM) is configured to prevent this. On-site access is possible, but limited through the use of service passwords with temporary validity. In addition, the service passwords can only be used for a single machine, as they are bound to the serial number. You must request a service password from IBM support. The facility to request a password is not publicly available on the internet. IBM support checks if the requesting party is a registered customer or licensee of the product.

What you can do to guarantee the maximum system security

Many security features are provided by the product, by the hardware, by DB2 for z/OS, by z/OS, or by your system authorization facility (SAF). Often, these features cannot be enabled automatically, but require an intervention on your part. Check the following list carefully, especially for production systems, and take appropriate action if one of the items reveals a security gap:

• Because TSO users, DB2 users, and users of IBM DB2 Analytics Accelerator Studio can be seen as the highest security risks, select these users carefully. Grant users access rights as required by the roles that they play. Do not give them more rights than needed. Where advantageous, combine access rights with user groups and grant access rights to users via group membership. For example, restrict the rights of the installer. This is the user who runs the AQTTIJSP job to create and bind stored procedures. The access rights that this user needs are listed in a separate section.

Grant users of IBM DB2 Analytics Accelerator Studio just the minimum set of access rights.

Rights for the SYSPROC.ACCEL_CONTROL_ACCELERATOR stored procedure can be granted for each XML input element of the **command** parameter because each element invokes a user-defined function. Make use of this feature.

A power user is helpful for first-time installations and system verification tests, but not necessarily required after that. Consider removing power users or revoking some of their access rights.

- Change the password for the IBM DB2 Analytics Accelerator Console at the first logon.
- Make sure that the hardware is located in an access-controlled area. Although direct remote connections to the accelerator are not possible, you can open a terminal session on-site with the proper user ID and a service password.
- If available, use an access auditing solution that protocols access to the z/OS system. Access auditing can be enabled in RACF.
- Protect the PDN against unauthorized access. Follow the **Related information** link at the end of this topic for instructions.
- Make sure that regular DB2 users do not have SELECT authorization for the SYSIBM.USERNAMES catalog table in DB2 for z/OS because this would allow them to read the authentication tokens, which are created as a result of the pairing process.
- The data on damaged hard drives of the IBM PureData System for Analytics is unencrypted. If you must replace a disk, degauss the damaged disk or destroy it physically.
- IBM recommends two redundant PDN cable connections to support failover scenarios. For these connections, two network switches are required. Usually, such switches are configured from a web-browser interface. Make sure that access to this interface is limited, that the switches restrict routing to the confines of the PDN, and that the initial password has been changed. Someone with access to the interface can change the configuration of the network and might, for example, open the PDN to the intranet or internet.
- The trace function of the product allows you to select components for tracing that might disclose sensitive information in the trace file or result set. Carefully decide whether this type of information is really required. If so, avoid to send this data to parties that have no need to know.
- To be able to trace data manipulation of tables on the accelerator, enable auditing for the SYSIBM.SYSACCELERATEDTABLES table by adding a corresponding row to the SYSIBM.SYSAUDITPOLICIES table. This way, the DB2 transaction log will tell you which accelerator tables have been modified. Consider adding SYSIBM.USERNAMES as well.

Related reference:

"Setting access rights for the user who runs AQTTIJSP" on page 34 The AQTTIJSP installation job, which you need to run in a later step, requires a user ID with sufficient DB2 access rights. These access rights are listed here.

"Required authorizations" on page 39

The various IBM DB2 Analytics Accelerator for z/OS components require different authorizations. It is useful to create at least one power user with extensive authorizations, that is, a user who can run all IBM DB2 Analytics Accelerator for z/OS functions and thus control all components. This section lists the required DB2, RACF, and file-system authorizations for such a power user. In subsequent chapters of this manual, it is expected that the required authorizations have already been granted.

Related information:

Protecting the private data network against unauthorized access

Chapter 4. Installing IBM DB2 Analytics Accelerator Studio

You must install the administration client, IBM DB2 Analytics Accelerator Studio, because the installation wizard contains all product license texts for IBM DB2 Analytics Accelerator for z/OS. You must accept these licenses before you can continue with the installation. Starting with version 3.2 of IBM Data Studio, which is the required framework for the IBM DB2 Analytics Accelerator Studio plugin, you must use IBM Installation Manager for the installation.

About this task

You can install IBM DB2 Analytics Accelerator Studio using one of the following methods:

- · From a workstation with an internet connection
- From a workstation without direct connection to the internet

The first method is recommended for individual users who want to install the product on their workstations.

The second method is recommended for users who do not have a direct internet connection from their workstations or for administrations who want to install the product on many workstations.

Installing IBM DB2 Analytics Accelerator Studio from a workstation without an internet connection

Following the steps in this topic, you download an installation package from the internet to a workstation. Users with a need to install can then copy the package or start the installation program remotely, provided that an intranet connection exists.

About this task

The procedure described here will download a package that contains all required software components. It does not matter whether one or more of these components are already installed on a target system. If one or more components are already present, only the remaining components will be installed. Sometimes, an update is offered for already installed components. Currently, the installation requires 1.5 GB of free disk space.

Attention: If possible, let users install the components on a local disk. Installing these on a shared network drive bears the potential risk of damaging or even destroying the Eclipse workspace.

A damage is likely to occur if you update or uninstall components while instances of the program are still running on connected computers. So if an installation on a shared network drive cannot be avoided, make sure that users close all running instances of IBM Installation Manager, IBM Data Studio, or IBM DB2 Analytics Accelerator Studio before they update or uninstall components.

Your operating system might issue several security warnings while you are following the steps in this procedure. Always grant permission, that is, click **Accept**, **OK**, or **Run** when asked whether you want to allow programs or

processes to make changes to your system or to access the internet.

Procedure

- From a workstation with a direct internet connection, download the installation package, which is found at: https://www14.software.ibm.com/webapp/iwm/web/ preLogin.do?source=swg-db2aas
- 2. Type your IBM ID and password in the appropriate fields and click **Sign in**. If you do not have an IBM ID, click **Get an IBM ID** and follow the registration instructions on the form. After finishing this task, you return to this page where you can sign in using your new IBM ID and password.
- 3. Download the installation package.
- 4. Depending on how you want to proceed, follow one of these methods.
 - If you want users to start the installation from a remote machine, extract the downloaded, compressed archive.
 - If you want users to copy the package so that they can start a local installation, provide them with information on the package location, how to access it, how to extract the package, and how to start the installation.
- 5. Depending on your choice in step 4, extract the package on the download machine, or let your users extract it on their local machines.
- 6. Let your users start one of the installation programs, that is, setup.exe or setup_win7_noadmin64.exe The setup.exe program launches an administrative installation. The setup_win7_noadmin64.exe program launches a restricted installation and only works on 64-bit Windows 7 operating systems. Description of these installation types:

Administrative installation

Selecting this installation, users can install the product for all users of their local workstation, and in any folder they prefer. This installation requires administrative rights on the workstation.

Restricted installation

Selecting this installation, users can install the product only for a single user of their local workstation (the user who is logged in and starts the installation program). The choice of an installation folder is restricted to the folders that this user has write access to. This installation does not require administrative rights on the workstation.

7. Let your users follow the instructions in the installation wizard.

Installing IBM DB2 Analytics Accelerator Studio from a workstation with an internet connection

If the workstation on which IBM DB2 Analytics Accelerator Studio is to be installed has a direct connection to the internet, you (as the administrator) or other users can install the product by themselves.

Before you begin

To follow the procedure in this topic successfully, you need an IBM ID (email address where you can be reached) and a password. If you do not have such credentials, you can register for an IBM ID from the window that opens as you complete step 6 on page 15.

About this task

The procedure described here will download a package that contains all required software components. It does not matter whether one or more of these components are already installed on your system. If one or more components are already present, only the remaining components will be installed. Sometimes, an update is offered for already installed components. Currently, the installation requires 1.5 GB of free disk space.

Attention: If possible, install the components on a local disk. Installing these on a shared network drive bears the potential risk of damaging or even destroying the Eclipse workspace.

A damage is likely to occur if you update or uninstall components while instances of the program are still running on connected computers. So if you cannot avoid an installation on a shared network drive, make sure that all running instances of IBM Installation Manager, IBM Data Studio, or IBM DB2 Analytics Accelerator Studio are closed before you update or uninstall components.

Your operating system might issue several security warnings while you are following the steps in this procedure. Always grant permission, that is, click **Accept**, **OK**, or **Run** when asked whether you want to allow programs or processes to make changes to your system or to access the internet.

Procedure

- 1. Uninstall any older version of IBM Data Studio or IBM DB2 Analytics Accelerator Studio from the workstation that you want to use.
- 2. Open your web browser and go to http://www.ibm.com/software/products/ us/en/data-studio.
- 3. Click Download at no charge.
- 4. On the Try IBM Data Studio: Available at no charge page, click Start a trial.
- In the Choose your preferred platform window, find the entry (Optional) Data Studio V4.1 client with IBM DB2 Analytics Accelerator Studio V4.1: Windows (HTTP or Download Director) and click the Download button next to it.
- 6. Type your IBM ID and password in the appropriate fields and click **Sign in**. If you do not have an IBM ID, click **Get an IBM ID** and follow the registration instructions on the form. After finishing this task, you return to this page where you can sign in using your new IBM ID and password.
- 7. On the IBM developerWorks site, scroll down until you find **Data Studio client and 3rd-party product extensions**.
- 8. Click the Download button next to this entry.
- **9**. On the IBM Installation Manager for Trial and Beta Products page, select your preferences. Make sure that you select **I agree** to agree to the terms and conditions.
- 10. Click I confirm.
- 11. You see another page with the heading IBM Installation Manager for Trial and Beta Products. A tabbed notebook with a page headed **Download using Download Director** page is in front. Select the appropriate check box to choose the correct version of the product (for Linux or for Windows).
- 12. Click the **Download now** button. The IBM Download Director window opens and the package is downloaded.

- **13**. When finished, a Question window opens asking you Do you want to launch this file?. Click **OK**. IBM IBM Installation Manager starts.
- 14. Type your IBM ID and password in the appropriate fields when prompted and click **OK**.
- **15**. IBM Installation Manager lists the installable packages. Make sure that the following packages are selected:
 - IBM Installation Manager
 - IBM Data Studio client
 - IBM DB2 Analytics Accelerator Studio

If any of these are already installed, they will be greyed out. Sometimes, IBM Installation Manager will offer to install a newer version.

- 16. Click Next to continue to the next page.
- 17. Select I accept the terms in the license agreements on the left and click Next.
- 18. The next page asks for a Shared Resources Directory and an Installation Manager Directory. You can change the default settings, but have to be aware that both directories must be different. Click Next when finished.
- 19. Accept the default settings on all pages that follow by clicking Next.
- 20. On the final page of the wizard, click Install.
- 21. Click **Finish** when the installation has been completed.

Chapter 5. Enabling an existing DB2 subsystem for IBM DB2 Analytics Accelerator for z/OS

To add IBM DB2 Analytics Accelerator for z/OS to your DB2 for z/OS environment, you must update the DB2 libraries supporting IBM DB2 Analytics Accelerator for z/OS as well as create and bind several stored procedures.

Software prerequisites for the DB2 data server

Before you begin with the installation, make sure that you meet the software requirements for the DB2 data server.

The prerequisites are listed at:

http://www.ibm.com/support/docview.wss?uid=swg27035960

Installing libraries with IBM DB2 Analytics Accelerator for z/OS support

Use SMP/E to install the product packages. Follow the installation steps in the package description.

Before you begin

- 1. Identify a set of queries to be used for installation verification.
- 2. Run these queries. The queries must return the same results before and after the installation of the DB2 libraries.
- **3.** Save or print the results so that you can compare them with the results obtained after the installation.

About this task

In the following text, <HLQBASE> is used as a placeholder for the high-level qualifier (HLQ) of your DB2 libraries. Replace it with the actual HLQ that is used in your system.

Procedure

1. Use the SMP/E *Apply* function to install the following packages:

FMID HAQT410

Includes the base product (DB2 for z/OS libraries)

FMID HCHCA21

Includes IBM InfoSphere[®] Change Data Capture for z/OS components

Important: It is required to install all packages, even if you do not want to use the incremental update function.

- 2. Use the SMP/E *Apply* function to install the IBM DB2 Analytics Accelerator for z/OS program temporary fixes (PTFs) that came with the product package.
- 3. If you want to use the incremental update function, follow these steps:
 - a. Complete the installation of IBM InfoSphere Change Data Capture for z/OS.
 - b. Configure IBM InfoSphere Change Data Capture for z/OS for use with IBM DB2 Analytics Accelerator for z/OS.

For instructions, follow the links at the end of this topic.

Related information:

Installing or upgrading InfoSphere CDC

"Configuring the incremental update function" on page 25 The incremental update function of IBM DB2 Analytics Accelerator for z/OS requires components of IBM InfoSphere Change Data Capture for z/OS (CDC) and additional components on the accelerator, called the Access Server and the replication engine. If you want to use this feature, you are advised to configure at least the CDC components at this point. To read an introduction to the incremental update function or go straight away to the CDC configuration section, follow the appropriate link at the end of this topic.

Creating the IBM DB2 Analytics Accelerator database

Customize and submit the DSNTIJAS job to create the database and tables for IBM DB2 Analytics Accelerator for z/OS in DB2 for z/OS.

About this task

Complete this tasks before you create the IBM DB2 Analytics Accelerator stored procedures (job AQTTIJSP).

Procedure

- 1. Copy and customize the DSNTIJAS sample job member in the <HLQBASE>.<SDSNSAMP> library according to your needs.
- 2. Submit DSNTIJAS.

Setting ZPARMs for IBM DB2 Analytics Accelerator for z/OS

Set the ZPARMs for IBM DB2 Analytics Accelerator for z/OS according to your needs.

Setting ZPARMs for IBM DB2 Analytics Accelerator in DB2 11 for z/OS

In DB2 11 for z/OS, you can set the ZPARMs for query acceleration on the DSNTIP82 and DSNTIP8A installation panels.

About this task

The following ZPARMs are available:

ACCEL=COMMAND | AUTO | NO

COMMAND

To start the accelerator by manually invoking the **-start ACCEL <name>** command, where **<name>** is the name of the accelerator.

- **AUT0** To automatically start the accelerator when the DB2 for z/OS subsystem starts.
- **NO** To specify that the accelerator cannot be used with this DB2 subsystem.

GET_ACCEL_ARCHIVE NO | YES

Specifies whether data that has been moved by the High-Performance Storage Saver (HPSS) is searched when a query is executed.

- **YES** Moved data is searched.
- **NO** Moved data is not searched.

QUERY_ACCELERATION

The treatment of incoming queries depends, among other factors, on the setting of the CURRENT QUERY ACCELERATION special register, which is a DB2 for z/OS special register that was introduced for IBM DB2 Analytics Accelerator for z/OS.

The value of the QUERY_ACCELERATION ZPARM provides the default setting for the CURRENT QUERY ACCELERATION special register. Both, the ZPARM and the special register accept the following values:

1 (NONE)

No routing of dynamic SQL queries to an accelerator. Queries will be processed by DB2 for z/OS only (inhouse query processing).

2 (ENABLE)

A dynamic SQL query will be routed to an accelerator if it fulfills all required conditions. An incoming query is tested against a set of heuristics, which include the table size and a response time estimate based on cost information from the

SYSIBM.DSN_PROFILE_ATTRIBUTES table. Both tests ensure that a query will only be routed to an accelerator if the query can be expected to run faster than in DB2 for z/OS. However, if an error occurs while the query is being processed by the accelerator, DB2 for z/OS will return a negative SQLCODE to the application and query processing will stop.

3 (ENABLE WITH FAILBACK)

Dynamic queries are accelerated only if DB2 for z/OS determines that it is advantageous to do so. If an accelerator returns an error during the PREPARE phase or when first opening (OPEN) the query, the query is processed by DB2 for z/OS rather than sent to the accelerator. If the accelerator returns an error during a FETCH operation or a subsequent OPEN operation, DB2 for z/OS returns an error to the user and the query ends abnormally.

4 (ELIGIBLE)

Dynamic queries are accelerated if they are eligible for acceleration. DB2 for z/OS does not use cost information to determine whether to accelerate the queries. Queries that are not eligible for acceleration are executed by DB2 for z/OS. If an accelerator fails while a query is running, or if the accelerator returns an error, DB2 for z/OS returns a negative SQL code to the application.

5 (ALL)

A dynamic query will always be routed to an accelerator, no matter if it fulfills the conditions or not. If processing cannot start or continue because an incoming query fails to fulfill all the conditions for accelerated query processing, DB2 for z/OS returns a negative SQLCODE to the application and query processing ends abruptly. That is, the query will not be processed at all.

Exceptions: The following dynamic queries are always processed by DB2 for z/OS, regardless of any relevant ZPARM settings (QUERY_ACCLERATION or GET_ACCEL_ARCHIVE):

• If all the tables referenced in the query have the qualifier: - SYSIBM **Exception:** Test queries against SYSIBM.SYSDUMMY1, SYSIBM.SYSDUMMYU and SYSDUMMYA are supported.

- SYSACCEL
- DB2GSE
- SYSXSR
- DGTT
- Dynamic queries whose top query block is pruned, and which therefore return an empty result set. To check whether a query falls into this category, follow these steps:
 - Explain the query using the DB2 EXPLAIN function. Use the following special register setting in the SQL statement: SET CURRENT QUERY ACCELERATION = NONE
 - 2. Check the PLAN_TABLE. If the top query block has been pruned, the entry for the query block shows PRUNED in the QBLOCK_TYPE column.

QUERY_ACCEL_OPTIONS=NONE | 1 | 2 | 3 | 4

You can also use any combination of 1, 2, 3, and 4. To do so, enclose the values in parentheses and separate them by a comma, for example (1,3). Meaning of the values:

NONE (default)

Т

Queries that can be enabled for query routing to an accelerator by the other options $(1 \mid 2 \mid 3 \mid 4)$ are disabled.

- Queries that include data encoded in multibyte EBCDIC are not blocked from processing by IBM DB2 Analytics Accelerator for z/OS, although IBM DB2 Analytics Accelerator for z/OS uses a different encoding scheme (Unicode, UTF-8) for the same data. Consequently, a query processed by DB2 might return a result set that is different from the result set that IBM DB2 Analytics Accelerator for z/OS returns for the same query.
- 2 For an INSERT operation that includes a SELECT statement, this option causes the SELECT portion of the statement to be routed to IBM DB2 Analytics Accelerator for z/OS and the INSERT operation to be performed by DB2 for z/OS.

Notes:

- The referenced tables on the accelerator might not be in sync with the tables in DB2 for z/OS.
- DB2 for z/OS does not route the SELECT portion of the statement if the target table in the INSERT statement uses an encoding scheme that is different from the scheme of the tables in the SELECT statement.
- 3 Allows query routing to IBM DB2 Analytics Accelerator for z/OS for queries that include byte-based string functions on data encoded by multibyte character set (MBCS) encoding schemes (like Unicode), although IBM DB2 Analytics Accelerator for z/OS supports only character-based string functions. If the string function operates on data that contains only single-byte characters, the query results returned by DB2 andIBM DB2 Analytics Accelerator for z/OS will be the same, irrespective of the encoding scheme that is used for the data. However, if the data contains multibyte characters, the results might be different.

4 Allows the execution of queries on IBM DB2 Analytics Accelerator for z/OS for queries that use an expression with a DATE data type in a LOCAL format.

IBM DB2 Analytics Accelerator for z/OS uses the dd/mm/yyyy format to interpret the input and the output date values.

Use this option only if the DATE FORMAT field of install panel DSNTIP4 specifies LOCAL or if application programs that process SQL on DB2 have been precompiled with the DATE(LOCAL) option. In either case, the LOCAL date exit routine must also define the specific dd/mm/yyyy date format. If the LOCAL format is not defined as dd/mm/yyyy, queries might return unpredictable results.

Recommendation: If you want to use this option in connection with a DB2 data sharing group, specify it on all members of the group.

Important:

- You can change the value of QUERY_ACCEL_OPTIONS online.
- If you use IBM DB2 Analytics Accelerator for z/OS with a DB2 for z/OS data sharing group, make sure that all members of the data sharing group use the same setting for QUERY_ACCEL_OPTIONS.
- For restrictions and other information regarding option 1 of QUERY_ACCEL_OPTIONS (support for mixed and double-byte EBCDIC), see *Conditions that prevent query routing to an accelerator* in the *IBM DB2 Analytics Accelerator Studio: User's Guide.*

Procedure

1. Set the ZPARMs for query acceleration on the DSNTIP82 and DSNTIP8A panels. The DSNTIP82 panel looks like this:

```
DSNTIP82 INSTALL DB2 - QUERY ACCELERATOR PREFERENCES
===>
Enter query accelerator options below:
1 ACCELERATOR STARTUP ===> N0 N0, COMMAND, or AUTO
2 GET ACCEL ARCHIVE ===> N0 N0 or YES
3 ACCELERATION OPTIONS ===> NONE NONE or YES
Enter CURRENT QUERY ACCELERATION special register option:
4 CURRENT QUERY ACCEL ===> 1 1 = NONE
2 = ENABLE
3 = ENABLE
3 = ENABLE WITH_FAILBACK
4 = ELIGIBLE
5 = ALL
PRESS: ENTER to continue RETURN to exit HELP for more information
```

DSNTIP8A is a popup panel that opens on top of DSNTIP82 and allows you to select values for the QUERY_ACCEL_OPTIONS parameter. It is not available if NONE is specified for 3 ACCELERATION OPTIONS.

- 2. Stop DB2.
- 3. Restart DB2 so that the changes can take effect.
- 4. Run sample queries or a test load to verify the proper functioning of DB2. **Related information**:

CURRENT QUERY ACCELERATION in DB2 10 for z/OS information center

🔄 Subsystem parameters that are not on installation panels

Using a sample job to set ZPARMs for IBM DB2 Analytics Accelerator for z/OS

If you do not use DB2 11 for z/OS or prefer a different method, you can customize and run a sample job to modify the ZPARM settings for IBM DB2 Analytics Accelerator for z/OS.

About this task

The following ZPARMs are available:

ACCEL=COMMAND | AUTO | NO

COMMAND

To start the accelerator by manually invoking the **-start ACCEL <name>** command, where **<name>** is the name of the accelerator.

- **AUT0** To automatically start the accelerator when the DB2 for z/OS subsystem starts.
- **NO** To specify that the accelerator cannot be used with this DB2 subsystem.

GET_ACCEL_ARCHIVE NO | YES

Specifies whether data that has been moved by the High-Performance Storage Saver (HPSS) is searched when a query is executed.

YES Moved data is searched.

NO Moved data is not searched.

QUERY_ACCELERATION

The treatment of incoming queries depends, among other factors, on the setting of the CURRENT QUERY ACCELERATION special register, which is a DB2 for z/OS special register that was introduced for IBM DB2 Analytics Accelerator for z/OS.

The value of the QUERY_ACCELERATION ZPARM provides the default setting for the CURRENT QUERY ACCELERATION special register. Both, the ZPARM and the special register accept the following values:

1 (NONE)

No routing of dynamic SQL queries to an accelerator. Queries will be processed by DB2 for z/OS only (inhouse query processing).

2 (ENABLE)

A dynamic SQL query will be routed to an accelerator if it fulfills all required conditions. An incoming query is tested against a set of heuristics, which include the table size and a response time estimate based on cost information from the

SYSIBM.DSN_PROFILE_ATTRIBUTES table. Both tests ensure that a query will only be routed to an accelerator if the query can be expected to run faster than in DB2 for z/OS. However, if an error occurs while the query is being processed by the accelerator, DB2 for z/OS will return a negative SQLCODE to the application and query processing will stop.

3 (ENABLE WITH FAILBACK)

Dynamic queries are accelerated only if DB2 for z/OS determines that
it is advantageous to do so. If an accelerator returns an error during the PREPARE phase or when first opening (OPEN) the query, the query is processed by DB2 for z/OS rather than sent to the accelerator. If the accelerator returns an error during a FETCH operation or a subsequent OPEN operation, DB2 for z/OS returns an error to the user and the query ends abnormally.

4 (ELIGIBLE)

Dynamic queries are accelerated if they are eligible for acceleration. DB2 for z/OS does not use cost information to determine whether to accelerate the queries. Queries that are not eligible for acceleration are executed by DB2 for z/OS. If an accelerator fails while a query is running, or if the accelerator returns an error, DB2 for z/OS returns a negative SQL code to the application.

5 (ALL)

A dynamic query will always be routed to an accelerator, no matter if it fulfills the conditions or not. If processing cannot start or continue because an incoming query fails to fulfill all the conditions for accelerated query processing, DB2 for z/OS returns a negative SQLCODE to the application and query processing ends abruptly. That is, the query will not be processed at all.

Exceptions: The following dynamic queries are always processed by DB2 for z/OS, regardless of any relevant ZPARM settings (QUERY_ACCLERATION or GET_ACCEL_ARCHIVE):

- If all the tables referenced in the query have the qualifier:
 - SYSIBM

Exception: Test queries against SYSIBM.SYSDUMMY1, SYSIBM.SYSDUMMYU and SYSDUMMYA are supported.

- SYSACCEL
- DB2GSE
- SYSXSR
- DGTT
- Dynamic queries whose top query block is pruned, and which therefore return an empty result set. To check whether a query falls into this category, follow these steps:
 - Explain the query using the DB2 EXPLAIN function. Use the following special register setting in the SQL statement: SET CURRENT QUERY ACCELERATION = NONE
 - 2. Check the PLAN_TABLE. If the top query block has been pruned, the entry for the query block shows PRUNED in the QBLOCK_TYPE column.

QUERY_ACCEL_OPTIONS=NONE | 1 | 2 | 3 | 4

You can also use any combination of 1, 2, 3, and 4. To do so, enclose the values in parentheses and separate them by a comma, for example (1,3). Meaning of the values:

NONE (default)

I

Queries that can be enabled for query routing to an accelerator by the other options $(1 \mid 2 \mid 3 \mid 4)$ are disabled.

1 Queries that include data encoded in multibyte EBCDIC are not blocked from processing by IBM DB2 Analytics Accelerator for z/OS, although IBM DB2 Analytics Accelerator for z/OS uses a different encoding scheme (Unicode, UTF-8) for the same data. Consequently, a query processed by DB2 might return a result set that is different from the result set that IBM DB2 Analytics Accelerator for z/OS returns for the same query.

2 For an INSERT operation that includes a SELECT statement, this option causes the SELECT portion of the statement to be routed to IBM DB2 Analytics Accelerator for z/OS and the INSERT operation to be performed by DB2 for z/OS.

Notes:

- The referenced tables on the accelerator might not be in sync with the tables in DB2 for z/OS.
- DB2 for z/OS does not route the SELECT portion of the statement if the target table in the INSERT statement uses an encoding scheme that is different from the scheme of the tables in the SELECT statement.
- 3 Allows query routing to IBM DB2 Analytics Accelerator for z/OS for queries that include byte-based string functions on data encoded by multibyte character set (MBCS) encoding schemes (like Unicode), although IBM DB2 Analytics Accelerator for z/OS supports only character-based string functions. If the string function operates on data that contains only single-byte characters, the query results returned by DB2 andIBM DB2 Analytics Accelerator for z/OS will be the same, irrespective of the encoding scheme that is used for the data. However, if the data contains multibyte characters, the results might be different.
- 4 Allows the execution of queries on IBM DB2 Analytics Accelerator for z/OS for queries that use an expression with a DATE data type in a LOCAL format.

IBM DB2 Analytics Accelerator for z/OS uses the dd/mm/yyyy format to interpret the input and the output date values.

Use this option only if the DATE FORMAT field of install panel DSNTIP4 specifies LOCAL or if application programs that process SQL on DB2 have been precompiled with the DATE(LOCAL) option. In either case, the LOCAL date exit routine must also define the specific dd/mm/yyyy date format. If the LOCAL format is not defined as dd/mm/yyyy, queries might return unpredictable results.

Recommendation: If you want to use this option in connection with a DB2 data sharing group, specify it on all members of the group.

Important:

- You can change the value of QUERY_ACCEL_OPTIONS online.
- If you use IBM DB2 Analytics Accelerator for z/OS with a DB2 for z/OS data sharing group, make sure that all members of the data sharing group use the same setting for QUERY_ACCEL_OPTIONS.
- For restrictions and other information regarding option 1 of QUERY_ACCEL_OPTIONS (support for mixed and double-byte EBCDIC), see *Conditions that prevent query routing to an accelerator* in the *IBM DB2 Analytics Accelerator Studio: User's Guide.*

Procedure

- 1. Add or change the listed parameters on the DSN6SPRM panel of the DSNTIJUZ sample job in your working libraries.
- 2. Stop DB2.
- 3. Submit DSNTIJUZ to assemble the new ZPARM load member.
- 4. Restart DB2 so that the changes can take effect.
- **5**. Run sample queries or a test load to verify the proper functioning of DB2. **Related information**:

CURRENT QUERY ACCELERATION in DB2 10 for z/OS information center

Subsystem parameters that are not on installation panels

Configuring the incremental update function

The incremental update function of IBM DB2 Analytics Accelerator for z/OS requires components of IBM InfoSphere Change Data Capture for z/OS (CDC) and additional components on the accelerator, called the Access Server and the replication engine. If you want to use this feature, you are advised to configure at least the CDC components at this point. To read an introduction to the incremental update function or go straight away to the CDC configuration section, follow the appropriate link at the end of this topic.

Related tasks:

"Completing the installation and configuring CDC" on page 61 Read how to complete the installation and configure IBM InfoSphere Change Data Capture for z/OS (CDC) for an optimal performance of the incremental update function.

Related information:

Chapter 9, "Incremental updates," on page 57

The incremental update function of IBM DB2 Analytics Accelerator for z/OS allows you to update accelerator tables continuously. Changes to the data in original DB2 for z/OS tables are thus propagated to the corresponding accelerator tables with a high frequency and just a brief delay. This way, query results from an accelerator are always extracted from recent, close-to-realtime data.

Chapter 6. Setting up IBM DB2 Analytics Accelerator for z/OS

Set up the IBM DB2 Analytics Accelerator for z/OS by completing the following tasks.

Setting up a WLM application environment for IBM DB2 Analytics Accelerator for z/OS stored procedures

Follow the steps in this section to set up a suitable Workload Manager (WLM) application environment.

Before you begin

Make sure that the following components have been installed:

- DB2 with the required DB2 program temporary fix (PTF) for IBM DB2 Analytics Accelerator for z/OS support
- DB2 command line processor for calling and verifying the IBM DB2 Analytics Accelerator for z/OS from the local environment

Note: The DB2 command line processor is a JavaTM application that requires IBM Data Server drivers for JDBC.

If you have not installed all of these components yet, follow the **Related tasks** link at the end of this section for instructions.

About this task

The following placeholders are used for specific high-level qualifiers in the steps and examples that follow. Replace these with the actual high-level qualifiers used in your system.

<hr/> <hr/>

HLQ for your DB2 libraries

<HLQSP>

HLQ for the IBM DB2 Analytics Accelerator for z/OS stored-procedure libraries

<HLQDB2SSN>

HLQ for DB2 subsystem-specific libraries

<HLQXML4C1>

HLQ for the XML toolkit

<HLQACTIVE>

A suggested HLQ for a copies that are independent of the original libraries under SMP/E control

Having completed the SMP/E *Apply* steps, the parts for the stored procedures can be found in the following libraries:

<HLQSP>.SAQTSAMP

Contains a job for the installation of the stored procedures, installation verification jobs, sample jobs for calling stored procedures, and XML samples as input for the stored procedures.

Contains database request modules (DBRMs) that must be bound to DB2.

Contains shared libraries and load modules for the stored procedures.

<hr/>

Contains the preprocessed Optimized Schema Representations (OSRs) for the XML System Services parser.

Procedure

- Copy the <HLQSP>.SAQTMOD load-module data-set as <HLQACTIVE>.SAQTMOD for reference in your WLM procedure. Also copy the <HLQSP>.SAQTSAMP(AQTENV) member as <HLQACTIVE>.SAQTSAMP(AQTENV). This way, you can install updates on the data sets that are controlled by SMP/E under <HLQSP> without affecting your running database environment.
- 2. Create a separate Workload Manager (WLM) environment for the IBM DB2 Analytics Accelerator for z/OS stored procedures. Use the following properties:

Appl Environment Name	•	•	DSNWLMV9
Description	•	•	DB2 V10 default Stored Procedures for IDAA
Subsystem type			DB2
Procedure name		•	DSNWLM
Start parameters		•	DB2SSN=&IWMSSNM,APPLENV=DSN
	•	•	WLMV9

Important:

- DSNWLMV9 is an example. The value that you must enter here is the one used for the !WLMENV! placeholder in the AQTTIJSP job.
- The procedure name (in this example: DSNWLM) must match the name of the defined procedure that you use to start the WLM-managed address space.
- The task that is started by the WLM-managed address space is run under a certain user ID. An OMVS segment must be defined for this user ID.
- Do not specify a value for NUMTCB in the Start parameters section of the definition because this value takes precedence, and thus makes it impossible to set the value by running the JCL that is discussed in the next paragraph.

You might want to modify the following template, which contains a procedure for the WLM-managed address space started task. The template includes the required STEPLIB and DD names.

```
//* PROCEDURE NAME = DSNWLM
//*
//*
      JCL FOR RUNNING THE WLM-ESTABLISHED STORED PROCEDURES
//*
      ADDRESS SPACE
//*
         RGN -- THE MVS REGION SIZE FOR THE ADDRESS SPACE.
//*
         DB2SSN -- THE DB2 SUBSYSTEM NAME.
//*
         NUMTCB -- THE NUMBER OF TCBS USED TO PROCESS
//*
                  END USER REQUESTS.
//*
         APPLENV -- THE MVS WLM APPLICATION ENVIRONMENT
//*
                   SUPPORTED BY THIS JCL PROCEDURE.
//*
//*
         <HLQXML4C1.10> is the HLQ where you have installed the
//*
                  XML Toolkit for z/OS//*
//*
         DB2VERS -- DB2-VERSION (I.E. V910)
                   SET BY APPLICATION ENVIRONMENT
//*
//*
                       DSNWLMV9 ==> V910
//*
//*
      The user ID that is used to start the task must have
//*
      read access to the <HLQs> in the STEPLIB statement
```

//*		
//*******	****	******
//DSNWLMV9	PR	<pre>DC RGN=0K,APPLENV=DSNWLMV9,NUMTCB=15,</pre>
//IEFPROC	EXE	C PGM=DSNX9WLM,REGION=&RGN,TIME=NOLIMIT,
//	PARI	M='&DB2SSN,&NUMTCB,&APPLENV'
//STEPLIB	DD	DISP=SHR,DSN= <hlqdb2ssn>.SDSNEXIT</hlqdb2ssn>
//	DD	DISP=SHR,DSN= <hlqbase>.SDSNLOAD</hlqbase>
//	DD	DISP=SHR,DSN= <hlqbase>.SDSNLOD2</hlqbase>
//	DD	DISP=SHR,DSN= <hlqactive>.SAQTMOD</hlqactive>
//	DD	DISP=SHR,DSN= <hlqxml4c1.10>.SIXMLOD1</hlqxml4c1.10>
//SYSTSPRT	DD	SYSOUT=A
//CEEDUMP	DD	SYSOUT=H
//OUT1	DD	SYSOUT=A
//UTPRINT	DD	SYSOUT=A
//DSSPRINT	DD	SYSOUT=A
//SYSPRINT	DD	SYSOUT=A
//AQTENV	DD	DSN= <hlqactive>.SAQTSAMP(AQTENV),DISP=SHR</hlqactive>
//AQTOSR	DD	DSN= <hlqactive>.SAQTOSR(AQTOSR),DISP=SHR</hlqactive>

Important:

• To avoid conflicts with environment variables that are set for stored procedures of other applications, create a WLM application environment that is exclusively used by the IBM DB2 Analytics Accelerator for z/OS stored procedures.

The stored procedures occasionally use the SYSPRINT output for diagnostic messages. For example, restart attempts for DSNUTILU, which might indicate inappropriately configured workload classes, are recorded here. Make sure to define a valid destination for SYSPRINT that grants write access to all users of the stored procedures. Otherwise, authorization failures will be recorded even if SYSPRINT was not used at all.

- The IBM DB2 Analytics Accelerator for z/OS stored procedures call the following DB2 for z/OS stored procedures:
 - SYSPROC.ADMIN_INFO_SYSPARM
 - SYSPROC.DSNUTILU
 - SYSPROC.ADMIN_COMMAND_DB2,

Run each of these DB2-supplied stored procedures in separate WLM application environments, and make sure that none of these runs in the same WLM environment as the IBM DB2 Analytics Accelerator for z/OS stored procedures. For more information, follow the appropriate links under **Related information**.

• If your system has more than one IP stack, you must unequivocally identify the stack that IBM DB2 Analytics Accelerator for z/OS is supposed to use. To do so, add the following statement to the procedure that starts the address space:

//SYSTCPD DD DISP=SHR,DSN=<TCPIP.DATA file>

For more information about the TCPIP.DATA data set, follow the **Related information** link at the end of this section.

Make sure that the z/OS UNIX System Services are configured to use the same IP stack. Connectivity from UNIX System Services to the accelerator is required for diagnostic and service purposes.

• You might want to change the default settings for IBM DB2 Analytics Accelerator for z/OS stored procedures, especially if you want to use parallel processing for loading tables. To do so, you must set the environment variables in the AQTENV data set accordingly.

Notes:

- If you must set environment variables, use the samle AQTENV data set that came with IBM DB2 Analytics Accelerator for z/OS Version 4.1. Do not re-customize an older version of AQTENV because fundamental settings have changed and using an old configuration might lead to problems.
- Be sure to adjust the NUMTCB value if you want to increase the value of the AQT_MAX_UNLOAD_IN_PARALLEL environment variable. The value of NUMTCB must be at least three times the value of AQT_MAX_UNLOAD_IN_PARALLEL plus 1:

<value of NUMTCB> > = 3 \times <value of AQT MAX UNLOAD IN PARALLEL> + 1

Thus, if you increase the value of AQT_MAX_UNLOAD_IN_PARALLEL to 6, you must set NUMTCB to 19 at least.

- To use the High Performance Storage Saver, you must set at least one of the following environment variables:
 - AQT_ARCHIVE_COPY1
 - AQT_ARCHIVE_COPY2
 - AQT_ARCHIVE_RECOVERYCOPY1
 - AQT_ARCHIVE_RECOVERYCOPY2
- The update transfer function of IBM DB2 Analytics Accelerator Studio only works if the values of AQT_HOST_PACKAGE_DIRECTORY and AQT_INSTALL_PREFIX are set correctly.

For more information, in particular about the relationship between NUMTCB and AQT_MAX_UNLOAD_IN_PARALLEL, follow the **Related reference** link at the end.

Note: Do not use the NUM ON option in the ISPF editor when modifying the AQTENV data set because this makes the line numbers in the columns from 72 to 80 part of the variable value. If this has happened, an error message similar to the following one is displayed when you run the **Transfer new** function in IBM DB2 Analytics Accelerator Studio:

The ACCEL_UPDATE_SOFTWARE procedure invoked by the "Transfer new software" function in the GUI returns a file open error, because a line number (here: 00360003) was considered part of the file path name:

AQT10206I - The OPEN operation on the "/SYSTEM/local/dwatest/swupdate_smpe 00360003/usr/lpp/aqt/packages" file, data set or pipe failed . Diagnostic information: Errno is 129

Related tasks:

|

1

"Creating the IBM DB2 Analytics Accelerator database" on page 18 Customize and submit the DSNTIJAS job to create the database and tables for IBM DB2 Analytics Accelerator for z/OS in DB2 for z/OS.

"Installing libraries with IBM DB2 Analytics Accelerator for z/OS support" on page 17

Use SMP/E to install the product packages. Follow the installation steps in the package description.

Chapter 4, "Installing IBM DB2 Analytics Accelerator Studio," on page 13 You must install the administration client, IBM DB2 Analytics Accelerator Studio, because the installation wizard contains all product license texts for IBM DB2 Analytics Accelerator for z/OS. You must accept these licenses before you can continue with the installation. Starting with version 3.2 of IBM Data Studio, which is the required framework for the IBM DB2 Analytics Accelerator Studio plugin, you must use IBM Installation Manager for the installation.

Related reference:

Appendix C, "Environment variables," on page 111 The job control language (JCL) for the configuration of the Workload Manager (WLM) environment for IBM DB2 Analytics Accelerator for z/OS stored procedures contains a data definition (DD) "AQTENV". This data definition includes a data set in which environment variables are defined. These variables control the behavior of some stored procedures.

Related information:

- Creating TCPIP.DATA
- Using the DB2 command line processor
- Java packages for JDBC support
- DB2 10 for z/OS: DB2-supplied stored procedures and user-defined functions
- DB2 Version 9.1 for z/OS: Enabling DB2-supplied routines

Verifying the correct setup of DB2-supplied stored procedures

The DB2 for z/OS stored procedures SYSPROC.ADMIN_INFO_SYSPARM, SYSPROC.DSNUTILU, and SYSPROC.ADMIN_COMMAND_DB2 must run in different Workload Manager (WLM) environments that are separate from the one used by the IBM DB2 Analytics Accelerator for z/OS stored procedures. Verify that this and a few other requirements are met by following the steps here.

Procedure

- Verify that each DB2-supplied stored procedure, SYSPROC.ADMIN_INFO_SYSPARM, SYSPROC.DSNUTILU and SYSPROC.ADMIN_COMMAND_DB2, runs in a different WLM environment.
- 2. Make sure that NUMTCB is set to 1 (NUMTCB=1) for the SYSPROC.ADMIN_INFO_SYSPARM and SYSPROC.DSNUTILU WLM environments.
- **3**. Verify that all WLM environments include the following libraries in their STEPLIB statements:

//STEPLIB	DD	DISP=SHR,DSN= <hlqdb2ssn>.SDSNEXIT</hlqdb2ssn>
//	DD	DISP=SHR,DSN= <hlqbase>.SDSNLOAD</hlqbase>

// DD DISP=SHR,DSN=<HLQBASE>.SDSNLOD2

Related tasks:

"Creating the IBM DB2 Analytics Accelerator database" on page 18 Customize and submit the DSNTIJAS job to create the database and tables for IBM DB2 Analytics Accelerator for z/OS in DB2 for z/OS.

Related information:

DB2 10 for z/OS: DB2-supplied stored procedures and user-defined functions

DB2 Version 9.1 for z/OS: Enabling DB2-supplied routines

Defining WLM performance goals for IBM DB2 Analytics Accelerator for z/OS stored procedures

It is important to define Workload Manager (WLM) performance goals in such a way that the WLM service class for the IBM DB2 Analytics Accelerator for z/OS stored procedures can provide a sufficient number of additional WLM address spaces in a timely manner when needed.

About this task

IBM DB2 Analytics Accelerator for z/OS stored procedures are called from a remote graphical user interface. This requires that a sufficient number of address spaces is available or can be started with minimum delay. To ensure such conditions, the goals of the service class for DDF transactions must be defined accordingly. Under favorable conditions, the starting of an address space takes two seconds. Under good conditions, this action takes about 10 seconds. However, if the workload is very high, the time needed to start an address space can be considerably longer.

Procedure

- 1. Classify your DDF transactions explicitly.
- 2. Assign the DDF transactions to a WLM service class.
- **3**. Make sure that the performance objectives of this service class are in accordance with the objectives for the rest of the workload on your system. The service class for IBM DB2 Analytics Accelerator for z/OS stored procedures must have at least medium priority.

Important: If classification rules do not exist to classify some or all of your DDF transactions into service classes, the unclassified transactions are assigned the SYSOTHER service class. This service class has no performance goal and is even lower in priority than a service class with a discretionary goal.

- 4. Assign the address spaces for the stored procedures to a separate service class for started tasks (STC). This ensures that the address spaces can be started before DDF transactions (stored procedures) start running. For more information, read the following articles:
 - Using z/OS Workload Manager to set performance objectives in the DB2 Version 9.1 for z/OS information center
 - Page 414 in DB2 9 for z/OS Stored Procedures: Through the CALL and Beyond
 - Chapter 20, server address space management (pages 423-434) in DB2 9 for z/OS Stored Procedures: Through the CALL and Beyond

Adjusting WLM performance goals for SYSPROC.ACCEL_LOAD_TABLES

The SYSPROC.ACCEL_LOAD_TABLES stored procedure is a special case because it starts one or more instances of the SYSPROC.DSNUTILU stored procedure (the DB2 Unload Utility) in turn. To start these procedures without delay, you must classify their workload accordingly.

About this task

By default, all DDF transactions are assigned to the SYSOTHER service class. The priority of this service class is too low. The WLM would delay or even prevent the parallel start of nested calls of the DB2 Unload Utility.

In addition, special attention is required if one or all of the following conditions apply:

- Logical Partition (LPAR) CPU capping is active
- Million-of-service-units (MSU) capping is active

In all of these cases, you will see a performance degradation with regard to the SYSPROC.ACCEL_LOAD_TABLES stored procedure and possibly other IBM DB2 Analytics Accelerator for z/OS stored procedures. With SYSPROC.ACCEL_LOAD_TABLES, you might also run into SQLCODE = -471 E790002 errors. This indicates that the time limit defined at installation time expired before the WLM could assign the request to a TCB in the address space for SYSPROC.DSNUTILU.

To avoid load processes that do not run to completion, adjust the WLM service class definitions or adjust the values of the following IBM DB2 Analytics Accelerator for z/OS environment variables:

AQT_MAX_UNLOAD_IN_PARALLEL

Decrease the value of this variable to reduce performance requirements

AQT_MAX_RETRIES_DSNUTILU

Increase the value of this variable to reduce performance requirements

AQT_SECONDS_BEFORE_RETRY_DSNUTILU

Increase the value of this variable to reduce performance requirements

Classifying the DDF workload for remote invocations of SYSPROC.ACCEL_LOAD_TABLES

To call the SYSPROC.ACCEL_LOAD_TABLES stored procedure from a remote environment, such as IBM DB2 Analytics Accelerator Studio, you must explicitly classify your DDF workload.

Procedure

- 1. Create an additional classification rule with the PR attribute, for example by creating a sub-rule to an already existing default rule.
- 2. To this classification rule, assign a service class with *medium-to-high* priority.

Example

In the following example, a sub-rule is assigned to an existing rule with the name DDFLOAD. The DDFLOAD service class already has medium-to-high priority. The name of the sub-rule is ACCEL_L*, which ensures that the sub-rule is applied when a stored procedure is invoked whose name starts with ACCEL_L, just like SYSPROC.ACCEL_LOAD_TABLES.



Figure 2. Assigning a medium-to-high priority service class to the SYSPROC.ACCEL_LOAD_TABLES stored procedure in the WLM ISPF application

Classifying the workload for local invocations of SYSPROC.ACCEL_LOAD_TABLES

If the SYSPROC.ACCEL_LOAD_TABLES stored procedure is called from a local z/OS environment, you must also ensure that a service class with medium-to-high priority is assigned to this stored procedure.

Example

Т

T

If you call the SYSPROC.ACCEL_LOAD_TABLES stored procedure from an application that is part of a batch process, and your TSO batch (JES) workload is assigned a service class with medium priority, you must create an additional classification rule that assigns a service class with medium-to-high priority to the batch jobs. (The batch jobs invoke the application that calls the SYSPROC.ACCEL_LOAD_TABLES stored procedure.)

Setting access rights for the user who runs AQTTIJSP

The AQTTIJSP installation job, which you need to run in a later step, requires a user ID with sufficient DB2 access rights. These access rights are listed here.

Note: A user with SYSADM authority has all the required privileges.

A user with an authority that is lower than SYSADM requires the following access rights:

- Privilege to create and drop global temporary tables with the qualifier DSNAQT, for example, CREATETMTAB
- Privilege to create views on the SYSACCEL.SYSACCELERATORS table, for example, DBADM on the DSNACCEL database
- CREATEIN and DROPIN on the SYSPROC schema. The CREATEIN privilege is required for the creation of IBM DB2 Analytics Accelerator for z/OS stored procedures. DROPIN is required for dropping previously installed stored procedures.

- CREATEIN and DROPIN on the DSNAQT schema or an equivalent authorization to create or drop a sequence in the DSNAQT schema.
- SELECT, INSERT, UPDATE, and DELETE on the following tables:
 - SYSACCEL.SYSACCELERATEDPACKAGES
 - SYSACCEL.SYSACCELERATEDTABLES
 - SYSACCEL.SYSACCELERATORS
 - SYSIBM.IPNAMES
 - SYSIBM.LOCATIONS
 - SYSIBM.USERNAMES
- SELECT on the following tables:
 - SYSIBM.SYSCOLUMNS
 - SYSIBM.SYSDATATYPES
 - SYSIBM.SYSDUMMY1
 - SYSIBM.SYSFOREIGNKEYS
 - SYSIBM.SYSINDEXES
 - SYSIBM.SYSINDEXPART
 - SYSIBM.SYSKEYCOLUSE
 - SYSIBM.SYSKEYS
 - SYSIBM.SYSRELS
 - SYSIBM.SYSTABCONST
 - SYSIBM.SYSTABLEPART
 - SYSIBM.SYSTABLES
 - SYSIBM.SYSTABLESPACE
 - SYSIBM.SYSTABLESPACESTATS
- Privileges to call the following DB2 stored procedures (for example, EXECUTE on the DSNUTILU package):
 - SYSPROC.ADMIN_COMMAND_DB2
 - SYSPROC.DSNUTILU
 - SYSPROC.ADMIN_INFO_SYSPARM

Related tasks:

"Creating the IBM DB2 Analytics Accelerator database" on page 18 Customize and submit the DSNTIJAS job to create the database and tables for IBM DB2 Analytics Accelerator for z/OS in DB2 for z/OS.

Customizing and running AQTTIJSP

Customize the AQTTIJSP job member (JCL) for the installation of IBM DB2 Analytics Accelerator stored procedures before you run the job.

Before you begin

Make sure that you have already created the IBM DB2 Analytics Accelerator database and tables by running the DSNTIJAS job.

Procedure

- 1. Replace the DB2 subsystem name and the other placeholders with the actual names as described in the customization notes within the AQTTIJSP job member.
- **2.** Review and optionally update the GRANT statements to conform to the authorization policy at your site.
- **3**. Recommendation: Copy the AQTTIJSP job to complete the following steps on the copy.

- 4. Modify and run just the job steps that are required in your situation:
 - For an initial installation:
 - a. AQTTIC03
 - b. AQTTICO4
 - c. AQTTIND
 - d. AQTTIGR (+)
 - e. For DB2 11 for z/OS also: AQTTINDA
 - f. For DB2 11 for z/OS also: AQTTIGRA
 - For an upgrade from IBM DB2 Analytics Accelerator for z/OS Version 3.1.0:
 - a. AQTTICO4
 - b. AQTTIND
 - c. AQTTIGR (+)
 - d. For DB2 11 for z/OS also: AQTTINDA
 - e. For DB2 11 for z/OS also: AQTTIGRA
 - For a clean re-installation:
 - a. AQTTIDO
 - b. AQTTIC03
 - c. AQTTICO4
 - d. AQTTIND
 - e. AQTTIGR (+)
 - f. For DB2 11 for z/OS also: AQTTINDA
 - g. For DB2 11 for z/OS also: AQTTIGRA
 - After the installation of a DB2 10 for z/OS fix pack: a. AQTTIND

unless you are directed otherwise by the ++HOLD information

- After the installation of a DB2 11 for z/OS fix pack:
 - a. AQTTIND
 - b. AQTTINDA

unless you are directed otherwise by the ++HOLD information.

- 5. Verify that you have created all required IBM DB2 Analytics Accelerator databases and tables before you submit the job. That is, check whether the DSNTIJAS job has been run successfully.
- 6. Submit the customized AQTTIJSP job.

Related tasks:

"Creating the IBM DB2 Analytics Accelerator database" on page 18 Customize and submit the DSNTIJAS job to create the database and tables for IBM DB2 Analytics Accelerator for z/OS in DB2 for z/OS.

Verifying the installation of IBM DB2 Analytics Accelerator for z/OS stored procedures

Verifying the installation of IBM DB2 Analytics Accelerator for z/OS stored procedures does not require a connection to IBM DB2 Analytics Accelerator for z/OS on an IBM PureData System for Analytics.

About this task

The verification consists of the following steps:

- Collecting information about the environment and the setup
- Verifying that required DB2 stored procedures can be called
- Verifying that IBM DB2 Analytics Accelerator for z/OS stored procedures can be called in *versionOnly* mode.

To this end, you must customize the AQTSJI00 and AQTSJI01 job members (JCLs) and run these job members.

Related tasks:

"Creating the IBM DB2 Analytics Accelerator database" on page 18 Customize and submit the DSNTIJAS job to create the database and tables for IBM DB2 Analytics Accelerator for z/OS in DB2 for z/OS.

Related reference:

Appendix B, "Members of SAQTSAMP," on page 107 The SAQTSAMP data set contains various samples for installing, configuring or running the IBM DB2 Analytics Accelerator for z/OS stored procedures. The following table lists all these members and provides brief descriptions of their functions.

Customizing and running AQTSJI00

Customize the AQTSJI00 job member (JCL) as described in the job notes before running SAQTSAMP(AQTSJI00).

Before you begin

Make sure that the user running AQTSJI00 has the privileges to read the following tables:

- DB2 catalog tables
- SYSACCEL.SYSACCELERATORS
- SYSACCEL.SYSACCELERATEDTABLES

Procedure

- 1. In the AQTSJI00 JCL, replace all instances of DSN!!0 with the name of the library that contains SDSNLOAD.
- 2. Replace all instances of DSNTEP!! with the name of the plan of the DSNTEPx sample program.
- **3**. Replace all instances of !DSN! with the name of the DB2 subsystem in which to run the IBM DB2 Analytics Accelerator for z/OS stored procedures.
- 4. Submit the AQTSJI00 JCL by running SAQTSAMP(AQTSJI00).

Results

The job returns the following information:

• The contents of the SYSACCEL.SYSACCELERATORS table.

If the job was run immediately after the installation, the SYSACCEL.SYSACCELERATORS table is shown, but it does not contain entries for accelerators.

If the job was run later to collect diagnostic information, the SYSACCEL.SYSACCELERATORS table lists all defined accelerators.

- Does the DB2 Communication Database (CDB) exist?
- Do the following tables exist?
 - ACCEL_NAMES
 - ACCEL_QUERY_INFO
 - ACCEL_TRACE_ACCELERATOR
- Have all stored procedures been defined that are used or provided by IBM DB2 Analytics Accelerator for z/OS, and if so, what are their run options (RUNOPTS) and Workload Manager (WLM) settings?

What to do next

Save the job output. It contains important information that might be required to solve installation problems.

Customizing and running AQTSJI01

Customize the AQTSJI01 job member (JCL) as described in the job notes before running SAQTSAMP(AQTSJI01).

Before you begin

- Make sure that the DB2 command line processor is installed.
- Make sure that the user who runs AQTSJI01 has the privileges to call the following stored procedures:
 - SYSPROC.DSNUTILU
 - SYSPROC.ADMIN_INFO_SYSPARM
 - SYSPROC.ADMIN_COMMAND_DB2
 - SYSPROC.ACCEL*

Procedure

- 1. In the AQTSJI01 JCL, replace all instances of !SAQTSAMP! with the location of the SAQTSAMP data set.
- 2. In the AQTSCI01 member (DB2 command script), replace all instances of !DB2ALIAS! with the connection alias as defined in the properties file of the DB2 command line processor.
- 3. Submit the AQTSJI01 JCL by running SAQTSAMP(AQTSJI01).

Chapter 7. Connecting IBM DB2 Analytics Accelerator for z/OS and DB2

For security reasons, communication between a DB2 subsystem and an accelerator requires an authentication of the DB2 subsystem. Follow the steps here to enable communication between these components.

Required authorizations

The various IBM DB2 Analytics Accelerator for z/OS components require different authorizations. It is useful to create at least one power user with extensive authorizations, that is, a user who can run all IBM DB2 Analytics Accelerator for z/OS functions and thus control all components. This section lists the required DB2, RACF, and file-system authorizations for such a power user. In subsequent chapters of this manual, it is expected that the required authorizations have already been granted.

If you want to create users who are permitted to run particular stored procedures only, look up the stored procedures in question in the *IBM DB2 Analytics Accelerator for z/OS: Stored Procedures and Reference*.

Attention: Do not give ordinary users SELECT authorization on the SYSIBM.USERNAMES table because this allows the users to see the authentication information in clear text in the SYSIBM.USERNAMES.NEWAUTHID column.

Required power-user authorizations in DB2 for z/OS

A power user requires the following authorizations in DB2 for z/OS:

- EXECUTE on the SYSPROC.* stored procedures
- EXECUTE on the DSNAQT.* functions
- EXECUTE on the SYSACCEL.* packages
- MONITOR1 privilege (needed to call the ADMIN_INFO_SYSPARM stored procedure internally)
- TRACE privilege
- DISPLAY privilege
- SYSOPR authorization to start and stop accelerators

Required z/OS power-user access rights

A power user requires the following access rights in RACF and in the hierarchical file system (HFS):

- An OMVS segment is required for the user ID.
- Write access to the /tmp directory (UNIX System Services pipes are created in this directory).
- Write access to the <prefix>/usr/lpp/aqt/packages (only needed if the z/OS power user installs IBM DB2 Analytics Accelerator for z/OS update packages using SMP/E, such as ++APARs and PTFs. Otherwise, read access is sufficient). This is the directory in which downloaded software updates are stored. The <prefix> part of the file path is the value of the AQT_INSTALL_PREFIX environment variable. You set this variable in the data set that the AQTENV DD

statement for the WLM environment refers to. See the template in step 2 on page 28. By default, no <prefix> is defined, meaning that the path for software updates is /usr/lpp/aqt/packages. You can adjust the prefix at installation time.

Note: Do not use the NUM ON option in the ISPF editor when modifying the AQTENV data set because this makes the line numbers in the columns from 72 to 80 part of the variable value. If this has happened, an error message similar to the following one is displayed when you run the **Transfer new** function in IBM DB2 Analytics Accelerator Studio:

The ACCEL_UPDATE_SOFTWARE procedure invoked by the "Transfer new software" function in the GUI returns a file open error, because a line number (here: 00360003) was considered part of the file path name:

AQT10206I - The OPEN operation on the "/SYSTEM/local/dwatest/swupdate_smpe 00360003/usr/lpp/aqt/packages" file, data set or pipe failed . Diagnostic information: Errno is 129

• Read access to all subdirectories of <prefix>/usr/lpp/aqt/packages directory.

DB2 for z/OS power-user authorizations for IBM DB2 Analytics Accelerator Studio

You might want to enable your power user to run IBM DB2 Analytics Accelerator Studio. If so, give the power user the following authorizations in DB2 for z/OS:

- SELECT on the DSNAQT.ACCEL_NAMES view. This privilege is required for the enumeration of accelerators.
- SELECT on the SYSACCEL* tables. This privilege is required to associate the tables with an accelerator pairing code.
- SELECT on the following catalog tables of the database management system:
- SYSIBM.SYSCOLUMNS
- SYSIBM.SYSCONTROLS
- SYSIBM.SYSDATABASE
- SYSIBM.SYSDUMMY1
- SYSIBM.SYSINDEXES
- SYSIBM.SYSRELS
- SYSIBM.SYSTABLEPART
- SYSIBM.SYSTABLES
- SYSIBM.SYSVIEWS

Read access to these tables is required for the creation of accelerator tables and the calculation of the overall table size.

Binding DB2 packages and granting user privileges

To enable access to DB2 for z/OS from IBM DB2 Analytics Accelerator for z/OS, IBM Data Studio, or IBM Optim[™] Query Tuner, you must create and bind certain DB2 packages and grant the EXECUTE privilege to the users of these applications.

About this task

The user ID under which the bind task is carried out automatically gains the EXECUTE privilege on the packages. However, at this stage, this is the only user having this privilege. Other users who run IBM DB2 Analytics Accelerator Studio or IBM Optim Query Tuner require the EXECUTE privilege as well. The creator or binder must therefore grant the EXECUTE privilege to the others users.

Note: The IBM Optim Query Tuner functions for single query tuning and Visual Explain (access plan graph) have been integrated into the IBM Data Studio, which serves as the basis for IBM DB2 Analytics Accelerator Studio. These modules can be used to compare the DB2 access plans with and without an accelerator, a functionality which allows you to see whether a query can be accelerated.

Procedure

- To create and bind the DB2 packages, follow the instructions on this website: http://publib.boulder.ibm.com/infocenter/dstudio/v4r1/topic/ com.ibm.datatools.qrytune.configothers.doc/topics/ enabledb2zfromclient_ds.html
- 2. To grant the EXECUTE privilege to other users, proceed as follows:
 - a. Select Analyze and Tune > Configure for Tuning > Advanced Configuration and Privilege Management as described on the following website:

http://publib.boulder.ibm.com/infocenter/dstudio/v4r1/topic/ com.ibm.datatools.qrytune.workloadtune.doc/topics/verify.html

b. Click **Manage Package Privileges** to display and modify the authorization IDs that can execute the tuning packages.

Creating EXPLAIN tables

To be able to display an access plan graph of your queries (accelerated or not) in IBM DB2 Analytics Accelerator Studio, you must create certain EXPLAIN tables in the DB2 subsystems that you use.

About this task

The instructions on the following web pages were originally written for IBM Query Tuner. However, the steps can be applied to IBM DB2 Analytics Accelerator Studio likewise.

Procedure

1. To create the basic EXPLAIN tables, follow these the instructions on this web page:

http://pic.dhe.ibm.com/infocenter/idm/v2r2/index.jsp?topic=/ com.ibm.datatools.qrytune.configdb2z.doc/topics/enabledb2zfromclient.html In addition to these, you need an EXPLAIN table called DSN_QUERYINFO_TABLE. This table is described here:

http://publib.boulder.ibm.com/infocenter/dzichelp/v2r2/index.jsp?topic= %2Fcom.ibm.db2z10.doc.perf%2Fsrc%2Ftpc%2Fdb2z_dsnqueryinfotable.htm

2. To create it, you find a suitable CREATE TABLE statement in the DSNTESC member of the SDSNSAMP library. You can also copy the sample DDL statements in the appendix of the *IBM DB2 Analytics Accelerator Studio: User's Guide*.

Creating a database connection profile

Create a database connection profile to gain access to a DB2 subsystem on a database server. A DB2 subsystem houses one or more databases, in which the source data for query acceleration (schemas and tables) is kept. To authenticate the DB2 subsystem to IBM DB2 Analytics Accelerator for z/OS, you must start the Add New Accelerator wizard. However, you can only start this wizard after connecting to a DB2 subsystem.

About this task

In IBM DB2 Analytics Accelerator Studio, the connection information is stored in profiles for reuse. Having created a profile, you can reconnect to a database by double-clicking the icon representing the profile in the Administration Explorer.

Procedure

- 1. Start IBM DB2 Analytics Accelerator Studio.
- 2. If the Welcome screen is displayed, close it.
- **3**. On the header of the Administration Explorer on the left, click the downward-pointing arrow next to **New** and select **New Connection to a Database**.
- 4. In the New Connection window, decide how to name the database connection profile:
 - To use the name of the database server that you want to connect to, leave **Use default naming convention** selected.
 - To choose a different name, clear **Use default naming convention**, and type the name in the **Connection Name** field.
- 5. From the **Select a database manager** list, select **DB2 for z/OS**.
- 6. Make sure that in the JDBC driver drop-down list, IBM Data Server Driver for JDBC and SQLJ (JDBC 4.0) Default is selected.
- 7. In the **Location** field, type the name of the database server that you want to connect to.

Tip: To determine the **Location**, **Host** name, and **Port number**, a DB2 for z/OS systems programmer or database administrator can issue a DIS DDF command.

- 8. In the **Host** field, type the host name or IP address of the data server on which the database server is located.
- **9**. In the **Port number** field, you see that port number 446 is selected by default. Leave this setting unless the database server uses another port.
- 10. Select **Retrieve objects created by this user only** if you want to restrict database access to the databases, schemas, tables, and other objects that were created by the logon user. If you do not select this option (default), IBM DB2 Analytics Accelerator Studio will show and make selectable all databases, schemas, and tables that the logon user has access to, including those to which this user might have only read access.
- 11. In the **User name** field, type the user ID that you want to use to log on to the database server. Note that you can only use IBM DB2 Analytics Accelerator Studio successfully if this user has sufficient rights to run the stored procedures behind the IBM DB2 Analytics Accelerator Studio functions. The section *Appendix C. Required access rights* in the *IBM DB2 Analytics Accelerator*

for z/OS: Stored Procedures Reference lists the privileges that are required to run a particular stored procedure. If you are uncertain, use an ID with SYSADMIN authority.

In many organizations, it is a common practice to have personal user IDs with restricted authority and special-purpose user IDs (groups in most cases) with extensive privileges in a certain field. IBM DB2 Analytics Accelerator Studio supports this practice in that you can specify a secondary user ID, which might have the privileges that your logon user ID lacks, such as the privilege to run stored procedures. If the secondary ID is a group user ID, the logon user must of course be a member of that group. To specify a secondary user ID, follow these steps:

- a. On the Connection Parameters page, click the **Optional** tab.
- b. In the **Property** field, type the following statement: currentSQLID
- c. In the Value field, type the secondary user ID.
- d. Click Add.
- e. Click the General tab to return to that page and complete the logon.
- 12. In the **Password** field, type the password belonging to the logon user ID.
- **13**. Leave the **Save password** check box deselected.

Attention: You can select **Save password** to avoid having to enter the password each time that you want to work with the database server. This, however, is not recommended because only a lightweight encryption is applied when the password is stored on your local hard disk.

- 14. Leave the **Default schema** field blank.
- 15. Click Test Connection to check if you can log on to the database server.
- 16. Click Finish.

Results

After creating the profile, IBM DB2 Analytics Accelerator Studio automatically connects to the DB2 subsystem.

What to do next

If it takes too long to load all objects of the DB2 subsystem into the Administration Explorer (more than one minute), you can set a filter to limit the number of schemas to be loaded:

- 1. In the Administration Explorer, right-click the icon representing the DB2 subsystem (database symbol).
- 2. Select **Properties** from the menu.
- 3. In the Properties for ... window, select **Default Schema Filter**.
- 4. Clear the **Disable filter** check box. This activates the filter controls.
- 5. From the **Name** drop-down list, select a suitable filter mask. In the adjacent text field, type the filter string. For example, to exclude all schemas whose names starts with the characters BLU:
 - a. From the Name drop-down list, select Does not start with the characters.
 - b. In the text field, type BLU.
- 6. Click Apply.

Related reference:

"Required authorizations" on page 39

The various IBM DB2 Analytics Accelerator for z/OS components require different authorizations. It is useful to create at least one power user with extensive authorizations, that is, a user who can run all IBM DB2 Analytics Accelerator for z/OS functions and thus control all components. This section lists the required DB2, RACF, and file-system authorizations for such a power user. In subsequent chapters of this manual, it is expected that the required authorizations have already been granted.

Testing the connection from IBM DB2 Analytics Accelerator Studio

Follow the steps in this section to see if you can connect to a DB2 subsystem from IBM DB2 Analytics Accelerator Studio.

Procedure

Double-click the icon representing the database connection profile in the Administration Explorer of IBM DB2 Analytics Accelerator Studio. You find the Administration Explorer on the left.

Results

When the connection was successful, the icon representing the subsystem changes and the database object types, such as table spaces or tables, are displayed in a folder hierarchy in the Administration Explorer.

Example



What to do next

If you can connect to the database, check if you can navigate between the objects in the database. To do so, explore the folder structure in the Administration Explorer. For example, select the **Schemas** folder and check whether the schemas are displayed in the Object List Editor on the right.

Obtaining the pairing code for authentication

Communication between an accelerator and a DB2 subsystem requires both components to share credentials. These credentials are generated after you submit a temporarily valid pairing code. This step is required each time you add a new accelerator. The following steps describe how to obtain the pairing code.

About this task

Note: You can renew the authentication for an existing accelerator without having to use a new pairing code. To do so, click the **Update** link in the Accelerator view.

The steps *Obtaining the pairing code for accelerator authentication* and *Completing the authentication using the Add New Accelerator wizard* (next topic) belong together, but are seldom carried out by the same person. Since the pairing code obtained from the IBM DB2 Analytics Accelerator Console is only valid for a limited time (30 minutes by default), the persons operating the console and IBM DB2 Analytics Accelerator Studio must coordinate the steps.

Procedure

- 1. Ask the network administrator or the person who did the TCP/IP setup for the IP address (virtual IP or wall IP address) of the accelerator. Make a note of this information. You need to enter it as you complete the steps that follow.
- 2. Start a client or emulator session (using, for example, IBM Personal Communications) to communicate with the z/OS system on which your DB2 subsystem is located.
- 3. Log on to TSO/ISPF.
- Enter the following command: tso telnet <hostname> 1600

where

<hostname>

Is the IP address of the accelerator that is connected to the DB2 for z/OS data server.

1600

Is the number of the port configured for accessing the IBM DB2 Analytics Accelerator Console using a telnet connection between the DB2 for z/OS data server and the accelerator.

For example:

tso telnet 10.101.8.8 1600

- 5. When prompted, enter the console password. The initial password is dwa-1234. You must change this password at the first logon.
- 6. Press the Pause key, then Enter to display the following screen:

7. Type 3 and press Enter to display the submenu:

main -> 3
You have the following options:
(0) - Go back one level
Obtain pairing code, IP address, and port
(2) - List paired DB2 subsystems
(3) - Set resource limits for DB2 subsystems
(4) - Clear query history
(5) - Specify the priority of maintenance tasks
(6) - Set the DB2 subsystem for time synchronization
(7) - Restart accelerator process
(8) - Disable the conversion mode for 24:00:00

- 8. Type 1 and press Enter:
- 9. When the message Specify for how long you want the pairing code to be valid. is displayed, enter an appropriate integer to specify the validity period in minutes. The time that you choose must be sufficient for you or a coworker to go to the workstation that runs IBM DB2 Analytics Accelerator Studio, start the Add New Accelerator wizard, and enter the information that is returned by the console. Values from 5 to 1440 are allowed. If you just press Enter, you accept the default of 30 minutes.

```
Press <return> to accept the default of 30 minutes.
Cancel the process by entering 0.
Accelerator pairing information:
Pairing code : 6048
IP address : 9.152.85.192
Port : 1400
Valid for : 30 minutes
Press <return> to continue
```

Important: A pairing code is valid for a single try only. Furthermore, the code is bound to the IP address that is displayed on the console.

- 10. Make a note of the following information on the console:
 - Pairing code
 - IP address
 - Port
- 11. Press Enter to return to the main menu of the console.
- 12. Type x and press Enter to exit the console and close the telnet session.

Completing the authentication using the Add Accelerator wizard

To complete the authentication, you enter the IP address, port number, and the pairing code in the Add Accelerator wizard.

Before you begin

Make sure that the following conditions apply:

- You need privileges to run DB2 administration commands and stored procedures on z/OS. If you created a power user as suggested, the power user will have the required privileges. For more information, follow the **Related information** link at the end of this topic.
- You have a valid pairing code. The pairing code, which is of temporary validity, can be obtained by using the IBM DB2 Analytics Accelerator Console. For more information see the **Related tasks** section at the end of this topic.

Attention: Do not give ordinary users SELECT authorization on the SYSIBM.USERNAMES table because this allows the users to see the authentication information in clear text in the SYSIBM.USERNAMES.NEWAUTHID column.

About this task

You can renew the authentication for an existing accelerator without having to use a new pairing code. To do so, click the **Update** link in the Accelerator view.

Attention: Making a new backup of your DB2 catalog tables is strongly recommended after each authentication update because restoration processes in your DB2 subsystem can make an accelerator unusable. This happens if you must restore your DB2 catalog and the backup of the catalog was made before the last update of the accelerator credentials. In this case, the latest authentication information will not be in the catalog tables of the backup, and so the accelerator can no longer be used.

For the completion of this task, the following stored procedures are run on your data server:

- SYSPROC.ACCEL_TEST_CONNECTION
- SYSPROC.ACCEL_ADD_ACCELERATOR

For information about the privileges that are required to run these procedures and further details, see the appropriate section in the *IBM DB2 Analytics Accelerator for* z/OS: *Stored Procedures Reference*. A link to this document is provided under **Related information** at the end of this section.

Procedure

- 1. Start IBM DB2 Analytics Accelerator Studio.
- 2. Select the Accelerators folder in the Administration Explorer.
- **3**. On the menu bar of the Object List Editor, click the downward-pointing arrow next to the green plus sign.
- 4. From the drop-down menu, select Add Accelerator.
- 5. In the **Name** field, type a name for the accelerator. This name is automatically copied to the **Location** field.

The location name is the unique name of the accelerator in the SYSIBM.LOCATIONS table. Mostly, this is the same name as the accelerator name.

Restriction: An accelerator cannot be shared between two or more DB2 subsystems if the subsystems use the same location name. If you copy an entire subsystem, make sure to change the location name of the copy afterwards.

- 6. In the **Pairing code** field, type the pairing code.
- 7. In the IP address field, type the IP address of the accelerator.
- 8. In the **Port** field, type 1400. This is the fixed port for network communication between the z/OS data server and the accelerator.
- **9**. Click **Test Connection** to check whether the accelerator with the given address can be connected to.
- 10. Click OK.

Related tasks:

"Obtaining the pairing code for authentication" on page 45 Communication between an accelerator and a DB2 subsystem requires both components to share credentials. These credentials are generated after you submit a temporarily valid pairing code. This step is required each time you add a new accelerator. The following steps describe how to obtain the pairing code.

Related reference:

"Required authorizations" on page 39

The various IBM DB2 Analytics Accelerator for z/OS components require different authorizations. It is useful to create at least one power user with extensive authorizations, that is, a user who can run all IBM DB2 Analytics Accelerator for z/OS functions and thus control all components. This section lists the required DB2, RACF, and file-system authorizations for such a power user. In subsequent chapters of this manual, it is expected that the required authorizations have already been granted.

Related information:

"DRDA connection does not work" on page 98

You can ping the accelerators, but you cannot establish a distributed relational database access (DRDA) connection between your database management system and the IBM PureData System for Analytics.

Testing the stored procedures

Run a few more tests to check whether the stored procedures behave as expected.

Before you begin

Make sure that the following conditions apply:

- The DB2 libraries for IBM DB2 Analytics Accelerator for z/OS have been installed.
- The IBM DB2 Analytics Accelerator for z/OS stored procedures have been installed.
- The DB2 command line processor has been installed.

Note: The DB2 command line processor is a Java application that requires IBM Data Server drivers for JDBC.

• The accelerator has been successfully added to the DB2 subsystem from IBM DB2 Analytics Accelerator Studio.

Tips:

• Create a separate DB2 command line processor properties-file for the user who will run the tests. Remove read access to this file for everybody else. This way, you can safely add a connection alias for the accelerator network connection to the properties file, including the password. Use the following syntax to add this information to the properties file:

<connection_alias>=<db2host>:<db2port>/<db2location>,<uid>,<password>

where

<db2host>

Is the host name or IP address of the System z server on which DB2 runs

<db2port>

Is the port for network connections between the DB2 command-line client and the DB2 host. To identify the DB2 port, run the following DB2 command:

- <ssid> DIS DDF

where <ssid> is the DB2 subsystem ID.

<db2location>

Is the location of the DB2 subsystem that is supposed to interact with an accelerator.

<uid>

Is the DB2 ID of the user running the tests

<password>

Is the password of the user running the tests

• It is easiest to give the testing user SYSADM authority in DB2. If this should violate a security policy of your organization, grant the testing user the required minimum access rights for each stored procedure. For more information, see *Required access rights* in the *IBM DB2 Analytics Accelerator for z/OS: Stored Procedures Reference*.

About this task

The AQTSJI02 JCL executes the test and calls a DB2 command script named AQTSCI02, which cleans up the SYSACCEL* DB2 tables after running the test. You must customize both, the JCL and the DB2 command script before you can submit the AQTSJI02 JCL.

Related reference:

Appendix B, "Members of SAQTSAMP," on page 107 The SAQTSAMP data set contains various samples for installing, configuring or running the IBM DB2 Analytics Accelerator for z/OS stored procedures. The following table lists all these members and provides brief descriptions of their functions.

Related information:

Using the DB2 command line processor

Customizing AQTSJI02

Customize the AQTSJI02 job member (JCL) as described in the job notes before submitting this JCL.

Procedure

- 1. In the AQTSJI02 JCL, replace all instances of !SAQTSAMP! with the location of the SAQTSAMP data set on your system.
- 2. If you change the number of IBM DB2 Analytics Accelerator for z/OS stored procedure calls in the SAQTSAMP(AQTSCI02) input data set for the DB2 command line processor, you must set the expected number of successful IBM DB2 Analytics Accelerator for z/OS stored procedure calls accordingly in the CHKCALL job step. See the following example, in which 15 stored procedure calls are expected:

//CHKCALL EXEC PGM=BPXBATCH,REGION=0M, // PARM='SH /tmp/ivp/AQTSSCHK /tmp/ivp/out 15'

Customizing AQTSXTCO

Customize the AQTSXTCO job member (XML input parameter for SYSPROC.ACCEL_TEST_CONNECTION) as described in the notes in XML comments before submitting the JCL SAQTSAMP(AQTSJI02).

About this task

The AQTSXTCO member contains XML code in UTF-8 encoding that is used to test the connection to an accelerator.

Procedure

In the AQTSXTCO member, replace all instances of !IDAAA! with a valid accelerator name.

Customizing AQTSCI02

Customize the DB2 command script AQTSCI02 as described before running SAQTSAMP(AQTSJI02).

Procedure

- 1. In the AQTSCI02 script, replace all instances of !DB2ALIAS! with the connection alias that is defined in the DB2 command line processor properties-file.
- 2. To test adding and removing an accelerator, uncomment (remove leading dashes) before the SYSPROC.ACCEL_ADD_ACCELERATOR and SYSPROC.ACCEL_REMOVE_ACCELERATOR calls and adjust the CHKCALL job step in AQTSJI02 accordingly. To test an accelerator that was recently added, additionally complete these steps:
 - a. Replace all instances of !IDAAIP! with the IP address of the accelerator.
 - b. Replace all instances of !IDAAPORT! with the listening port number of the accelerator (by default, this is 1400).
 - c. Replace all instances of !IDAAPIN! with a valid pairing code that was created by the IBM DB2 Analytics Accelerator Console. Bear in mind that this code is valid for a limited period of time only. For more information on how to obtain a pairing code, follow the **Related tasks** link at the end of this topic.
 - d. Replace all instances of !IDAAA! with a valid accelerator name.

Related tasks:

"Transferring update packages for the accelerator" on page 85 Transfer update packages for IBM DB2 Analytics Accelerator for z/OS, the Access Server, the replication engine, or the Netezza Performance Server (NPS) by completing the steps in this section.

"Obtaining the pairing code for authentication" on page 45 Communication between an accelerator and a DB2 subsystem requires both components to share credentials. These credentials are generated after you submit a temporarily valid pairing code. This step is required each time you add a new accelerator. The following steps describe how to obtain the pairing code.

Running AQTSJI02

To complete the test, submit the AQTSJI02 JCL after customizing the JCL itself and the AQTSCI02 DB2 command script.

Procedure

Submit the AQTSJI02 JCL by running SAQTSAMP(AQTSJI02).

Results

If all steps in AQTSJI02 end with return code 0, the IBM DB2 Analytics Accelerator for z/OS stored procedures are set up correctly and network communication with the accelerator works.

Chapter 8. Testing query acceleration

Run a test query to verify that all required installation and configuration steps have been completed successfully and that the entire system is in an operable state.

Related information:

Running an SQL script from IBM DB2 Analytics Accelerator Studio Loading tables

Selecting a test query

For a simple test, select a query that runs against a single table and uses an aggregation function, such as SELECT COUNT(*) FROM . To qualify for query acceleration, the table must occupy at least 500 pages in DB2 for z/OS.

Defining tables

Define the tables that are referenced in your test query on the accelerator.

Procedure

- 1. In the Administration Explorer (upper left) of IBM DB2 Analytics Accelerator Studio, select the Accelerators folder.
- 2. In the list of the Object List Editor on the right, double-click the name of the accelerator. This opens the Accelerator view. It is empty by default, that is, it does not contain any tables.
- 3. Click the Add button.
- 4. In the Add Tables wizard, select the tables that are referenced in your test query.
- 5. Click OK.

Results

The selected tables are shown in the Accelerator view.

Loading and enabling tables

Before you can run your test query, you must load the currently empty tables with data and enable these for query acceleration.

About this task

Loading or updating the data in tables for query acceleration requires IBM DB2 Analytics Accelerator Studio to hold the network connection until the process has finished.

An option that circumvents this problem is to invoke IBM DB2 Analytics Accelerator for z/OS stored procedures directly from JCL batch jobs. Such jobs can even be scheduled by using a third-party tool.

Procedure

- 1. In the Accelerator view, select the tables that are referenced in your test query and click **Load**.
- 2. Do not change value of the Lock original database tables while loading drop-down list. Leave the setting None.
- 3. Leave After the load, enable acceleration for disabled tables selected.
- 4. Click OK.

Note: The processes can take a couple of minutes to complete. If the load fails, first check whether DSNUTIL was started in DB2 for z/OS.

Running a test query

Run a simple test query against the selected tables to see whether the setup works.

Before you begin

- 1. Make sure that the tables referenced by the query exist on the accelerator, are loaded, and enabled for query acceleration.
- 2. Make sure that the accelerator has been started and is online. The status of the accelerator is shown on top of the Accelerator view.

About this task

If you followed the instructions in the previous sections, the query should run successfully, that is, it should not fail and return the intended result.

Procedure

- 1. On the toolbar of the Administration Explorer, click the downward-pointing arrow next to the **New** button.
- 2. Select New SQL Script from the menu.
- 3. In the blank space of the Script<x>.sql workspace that opens on the upper right, type the following statement on the first line:
 - SET CURRENT QUERY ACCELERATION = ENABLE;

Note: <x> stands for a counting number (integer). This means that your first SQL script is named Script1.sql, the second Script2.sql, and so on.

- Type your query. Use a simple query, such as SELECT SUM(QUANTITY) FROM SALES.FIGURES_FACT.
- 5. Select **Script** > **Run SQL** from the main menu. The SQL Results pane in the lower left shows you whether the query ran successfully.
- 6. To check whether the accelerator was used, follow these steps:
 - a. Return to the accelerator view by clicking the tab with the name of the accelerator.
 - b. Click the twistie next to the heading **Query Monitoring**.

The section that unfolds shows a table that lists the most recent queries. These are only the queries that were run on the selected accelerator. Inhouse DB2 queries are not listed. Your test query should appear at the top of the table. Your query might not be listed there for various reasons, for example:

- The installation of IBM DB2 Analytics Accelerator for z/OS is incomplete.
- The DB2 optimizer did not route the query to the accelerator because the query does not qualify.

• One of the tables in the query is not defined on the accelerator.

To determine the cause for the failure, use the DB2 EXPLAIN function and look up the reason code in DSN_QUERYINFO_TABLE. For more information, read the section EXPLAIN information in the IBM DB2 Analytics Accelerator Studio: User's Guide. You find a link under Related information at the end of this topic.

Related information:



IBM DB2 Analytics Accelerator for z/OS: User's Guide

Chapter 9. Incremental updates

The incremental update function of IBM DB2 Analytics Accelerator for z/OS allows you to update accelerator tables continuously. Changes to the data in original DB2 for z/OS tables are thus propagated to the corresponding accelerator tables with a high frequency and just a brief delay. This way, query results from an accelerator are always extracted from recent, close-to-realtime data.

If you followed the instructions in this guide, you already installed the necessary IBM InfoSphere Change Data Capture for z/OS (CDC) components by completing the steps in "Installing libraries with IBM DB2 Analytics Accelerator for z/OS support" on page 17. However, for updates or more detailed information, follow the link to the installation instructions in the IBM InfoSphere Change Data Capture for z/OS information center at the end of this topic.

Attention:

- It is not possible to integrate and use an existing CDC installation. You must use a dedicated CDC that was installed from the SMP/E installation package delivered with IBM DB2 Analytics Accelerator for z/OS.
- A risk of data loss exists if you include tables without a unique constraint, such as a primary key or a unique index. To identify rows in such a table, the values of all table columns are used. It is possible that rows with exactly the same values exist, in which case it is impossible to identify a row unambiguously. Hence a deletion of a single row from a table in DB2 for z/OS results in the deletion of this row and all its duplicates from the corresponding accelerator table. This, in turn, leads to incorrect results for accelerated queries.
- If you update DB2 for z/OS tables by running the LOAD utility rather than an INSERT, UPDATE, or DELETE operation, you must set the following parameters (for the LOAD utility):
 - SHRLEVEL CHANGE
 - LOG YES

Otherwise, the changes that were made by the LOAD utility are not detected by the incremental update function, and will thus not be reflected in your accelerator tables.

- With the introduction of IBM DB2 Analytics Accelerator for z/OS Version 3.1, the DB2 attribute DATA CAPTURE is added to the tables on an accelerator. The attribute can carry the value Y or N (default), for *yes* or *no*. When incremental updates are enabled for a table, the DATA CAPTURE attribute of the table is set to the value Y. Once set, this attribute value persists, even if the table is disabled at a later time. Bear this in mind, especially if you run applications that use the DATA CAPTURE attribute.
- In general, an ALTER TABLE operation on a DB2 table prevents query routing to an accelerator if the altered table is referenced in the query (reason code 14). This is independent of any incremental update configuration that might exist. Therefore, you must complete the following steps after an ALTER TABLE operation:
 - 1. Remove the table from all accelerators on which it is defined.
 - 2. Redefine the table on the accelerators from which you previously removed it.
 - 3. Enable incremental updates for the redefined table on all accelerators.
 - 4. Load the redefined table on all accelerators.
 - 5. Enable query acceleration for the redefined table on all accelerators.

Special attention must be paid to tables that have been altered at any time before enabling incremental updates for these tables. You must ensure that such tables have been reorganized after the last alteration, and before the incremental update on them is enabled. If the reorganization did not happen, you must reorganize such tables before enabling incremental updates for them.

Exception: You need not follow the steps previously described after an ALTER TABLE ROTATE PARTITION operation. The CDC capture agent automatically detects a rotation in the DB2 transaction log. If a partition has been rotated, the agent generates a DELETE DML statement for the records in the to-be-deleted partition (rolled-off partition). It will also propagate new records that were inserted into the rolled-on partition. So there is no need for a table removal, redefinition, re-enablement, or reload.

Restriction: It is not possible to enable tables if the table name or the schema name contains GB18030 characters (Simplified Chinese) of Unicode plane 2

1

Т

T

Т

1
(U+20000-U+2FFFF: Supplementary Ideographic Plane). Trying to do so results in an error.

Related information:

IBM DB2 Analytics Accelerator for z/OS: User's Guide

Concepts and architecture

The incremental update function is deeply integrated into the IBM DB2 Analytics Accelerator for z/OS solution.

You administer incremental updates from IBM DB2 Analytics Accelerator Studio or call the corresponding IBM DB2 Analytics Accelerator for z/OS stored procedures directly.

Components

IBM InfoSphere Change Data Capture for DB2 for z/OS is required for an implementation of the incremental update function. For further details, refer to the list of IBM DB2 Analytics Accelerator for z/OS prerequisites on the web.

System architecture

IBM DB2 Analytics Accelerator Studio communicates with the IBM DB2 Analytics Accelerator for z/OS server process on the Netezza host by means of the SYSPROC.ACCEL_SET_TABLES_REPLICATION and SYSPROC.ACCEL_CONTROL_ACCELERATOR stored procedures. Requests for incremental updates are forwarded to an automation program that interfaces with the application programming interface (API) of the IBM DB2 Analytics Accelerator Access Server. The Access Server executes the requested action, such as adding or removing a table from the selection list for incremental updates (called subscription).

Note: IBM DB2 Analytics Accelerator for z/OS automates the setup, but is not involved when data is transferred from the source to the target tables. Data flows directly from the CDC for DB2 for z/OS Agent to the IBM DB2 Analytics Accelerator Access Server, without IBM DB2 Analytics Accelerator for z/OS components in between (see the blue arrows in Figure 3 on page 60). In addition, the table setup is a simple 1:1 column mapping. Tables or data are neither converted nor transformed in the process, and user-exit programs are not employed.

The process flow is illustrated in the following diagram:



(private network)

Figure 3. Interaction diagram of the components that are involved in the incremental update process.

Integration

The deep integration of the incremental update function into the IBM DB2 Analytics Accelerator for z/OS framework ensures data integrity, system stability, and uniform monitoring and recovery processes.

IBM DB2 Analytics Accelerator for z/OS is always aware of ongoing incremental update processes and can thus prevent actions that would otherwise harm the consistency of table data. For example, it is not possible to change a distribution key or organizing keys or load or update a table while an incremental update process is running.

The failover mechanism for the Netezza hosts remains in place. If a host fails during an incremental update, the inactive host will take over, and incremental updates will continue.

Furthermore, you can use the existing IBM DB2 Analytics Accelerator for z/OS functions for logging, tracing, monitoring, and software updates.

Tracing

Incremental update information is included in the collection of trace data.

Logging

Warning or errors related to the incremental update function are reported on the z/OS operator console using the uniform message ID DSNX881I.

Event monitoring

Incremental update events can be viewed in the event history view provided by IBM DB2 Analytics Accelerator Studio.

Software updates

Software updates for the incremental update function are delivered as part of regular product updates, and are uploaded, transferred, and applied through the same channels.

Completing the installation and configuring CDC

Read how to complete the installation and configure IBM InfoSphere Change Data Capture for z/OS (CDC) for an optimal performance of the incremental update function.

Procedure

- 1. Check and make sure that you meet the system requirements for CDC at: System requirements for InfoSphere CDC for z/OS. This includes the following steps (among others):
 - a. Creating a z/OS user ID for running the CDC started task.

Tip: Set the password of this user ID so that it never expires. This prevents unwanted interruptions of the incremental update process in the future.

b. Defining an OMVS segment for this user ID because the ID needs to access TCP/IP services.

This user ID becomes the owner of CDC metadata in DB2 for z/OS and requires SYSCTRL authorization in addition to other authorizations, including SELECT authorization for all tables to be processed by the incremental update function.

- 2. Complete the necessary preinstallation steps as described in Before you install InfoSphere CDC for z/OS. In one of the steps, you define a port number for the CDC address space used for communication between the CDC agents on z/OS and the accelerator. The default port is 5999.
- **3.** Using the SMP/E Apply function, add required program temporary fixes (PTFs) to your base installation. (You installed the CDC base as you completed step 1 on page 17). The required PTFs are listed in:

http://www.ibm.com/support/docview.wss?uid=swg27039487Prerequisites and Maintenance for IBM DB2 Analytics Accelerator for z/OS Version 4.1.

- 4. Complete the steps in section Completing the installation using the distributed sample jobs of the CDC information center. Basically, this task comprises the following steps:
 - a. Creating the configuration control data set and copying members into this data set. The settings in this control data set are read and processed during the initialization of the CDC address space.
 - b. Defining a security identifier for use of the CDC address space to ensure a controlled access to the relevant DB2 subsystems and the CDC metadata tables.
 - c. Customizing and running a number of jobs to prepare your DB2 subsystems for CDC. For example, the CHCMDMUT job creates the metadata tables for CDC. You must customize this job and specify the security identifier as the owner of these tables.

What to do next

Follow the instructions in the following topics, in the order indicated:

Configuration settings for IBM InfoSphere Change Data Capture for z/OS

As you enable incremental updates, IBM DB2 Analytics Accelerator for z/OS uses the default values in the various configuration control data set members of IBM InfoSphere Change Data Capture for z/OS (CDC). Not all of these default values are suitable for the configuration of incremental updates with IBM DB2 Analytics Accelerator for z/OS. Therefore, you must change a few of these values.

1. Change values of the following keywords in the TCP/IP statement of the CHCCMMxx member as follows:

SERVICENAME=<port_no>

where <port_no> matches the port number or assigned service name that you have defined in step 2 on page 61. By default, this statement is set to the value 5999.

2. Change the values of the following keywords in the CONFIG statement of the CHCCFGxx member as follows:

```
PALCLEANUPTIME=23:59,
TIMEZONE=<value>,
ADMININACTTIMEOUT=15,
AUTORESTARTINTERVAL=2
```

3. Change the values of the following keywords in the DB2 statement of the CHCDBMxx member as shown:

```
DB2 SSID=<value>,
PLANSUFFIX=<value>,
ONUTILITYACTION=IDLE,
ONDECOMPRESSIONERROR=(300,IDLE),
ONSCHEMACHANGE=IDLE,
ADDCOLUMNISSCHEMACHANGE=NO,
LOGPOLLINTERVAL=(3,ALWAYS),
LOGCACHEDELAY=5,
CACHEBLOCKSIZE=2,
CACHELEVEL1RESERVED=400,
CACHELEVEL1SIZE=500
```

Important: If you use the High Performance Storage Saver (HPSS), you must set ONUTILITYACTION=IGNORE. See the parameter description in the following section.

Description

T

L

Т

Here is a brief description of the keywords (in the order of appearance):

SERVICENAME

Specifies the name of the TCP/IP service that the InfoSphere CDC address space provides to the Management Console and to other InfoSphere CDC servers that require replication services.

PALCLEANUPTIME

Specifies the time at which IBM InfoSphere Change Data Capture for z/OS will initiate a daily cleanup of the event log.

TIMEZONE

Specifies the local time zone of the system on which IBM InfoSphere Change Data Capture for z/OS is executing, for example Europe/Berlin. You can find the correct time zone for your system in the CHCTMZON member.

ADMININACTTIMEOUT

Specifies for how long (in minutes) a Management Console user connection can be inactive before the connection is severed.

AUTORESTARTINTERVAL

Sets a period in minutes after which an incremental update process is restarted in case of a normal or abnormal termination.

SSID

Specifies the name of the DB2 subsystem that CDC is supposed to work on.

PLANSUFFIX

Specifies two characters to associate a set of DB2 Plans unequivocally to CDC metadata tables. This is required because more than one instance of CDC might access the same DB2 subsystem.

ONUTILITYACTION

Specifies the behavior of IBM InfoSphere Change Data Capture for z/OS after detecting a DB2 utility (such as LOAD or RECOVER) has run on a table space containing source tables that are being mirrored.

Important: The setting ONUTILIYACTION=IDLE leads to a conflict if you intend to use the High Performance Storage Saver (HPSS) because IDLE suspends the propagation of data changes for the entire table, and not just the partitions that have been moved by the HPSS. Hence the parts that have not been moved will not be synchronized anymore. To avoid this conflict, set the keyword to the value IGNORE.

ONSCHEMACHANGE

Specifies which action to take if the schema of a table to be updated incrementally has changed such that changes to the table can no longer be propagated correctly. For IBM DB2 Analytics Accelerator for z/OS, this keyword must be set to the value IDLE, which means that the change propagation continues for all but the altered tables.

ADDCOLUMNISSCHEMACHANGE

Specifies whether the adding of a column is treated like any other schema change or as a special case. IBM DB2 Analytics Accelerator for z/OS requires this value to be set to NO, which means that adding a column is treated as a special case.

LOGPOLLINTERVAL

|

1

Specifies how frequently the log profiling task performs unprompted reads in a DB2 data sharing group when the local DB2 log is idle (that is, it is not actively being extended by DB2 activity). The value for this keyword is the number of seconds the local DB2 log must have been idle before the log profiling task autonomously requests additional DB2 log data.

LOGCACHEDELAY

Specifies the delay in seconds that is imposed before newly scraped DB2[®] log data is added to the InfoSphere CDC log cache, unless the newly scraped DB2 log data fills a complete InfoSphere CDC log cache buffer (whose size is determined by the CACHEBLOCKSIZE keyword).

CACHEBLOCKSIZE

Specifies the maximum size, in MB, of each buffer in the level 1 and level 2 DB2 log caches.

CACHELEVEL1RESERVED

Specifies the size, in MB, of the reserved portion of the level 1 DB2 log cache. The reserved portion maintains records that were recently placed in the DB2 log.

CACHELEVEL1SIZE

Specifies the size, in MB, of the level 1 DB2 log cache.

Important: These settings are different from those of IBM InfoSphere Change Data Capture for z/OS. For example, IBM InfoSphere Change Data Capture for z/OS sets the ADMININACTTIMEOUT and the AUTORESTARTINTERVAL values to 0, which means that Management Console connections are allowed to run endlessly, and mirroring processes are never restarted automatically. For more information, especially in case that you want to change these values, follow the links under **Related information**.

Enabling DATA CAPTURE CHANGES

Enabling DATA CAPTURE CHANGES for the SYSIBM.SYSTABLES table has the effect that data definition (DDL) changes to source tables are recorded in the DB2 for z/OS and hence do not go unnoticed in cases where they affect incremental updates. In some cases, the enablement of DATA CAPTURE CHANGES even allows incremental updates to continue. Therefore, enabling DATA CAPTURE CHANGES is highly recommended. Follow the link to the CDC documentation for further information and instructions.

Enabling log caching

Т

1

Т

Т

Т

|

1

T

If you want to use the incremental update function with multiple accelerators that are attached to the same DB2 subsystem, enable log caching for the capturing agent.

About this task

When the incremental update process is restarted after a stop, the DB2 log reader must resume reading from the point of interruption. The size of the DB2 active log and the incremental update log cache determine for how long incremental updates can be stopped before archived log records must be requested to locate the correct position for a resumption. To request and search archived logs takes longer, and the log reader might be unable to find the correct position within the allotted time. If the correct position can not be determined, a resumption is not possible. In this case, the only way to restart an incremental update process is to set it up anew, which involves a reloading of all the tables that take part in the process.

Recommendation:

Make the DB2 active log so large that it can hold about one day of DB2 log activity. This means that the log reader can be stopped for one day if necessary. Additionally, a log cache has two advantages, and is generally recommended for all installations of IBM DB2 Analytics Accelerator for z/OS. If a log cache is in place, the log must not be read multiple times if data must be updated on more than one accelerator. A log cache also significantly increases the ability of the log reader to catch up after an interruption of the incremental update process.

Performance tuning

The speed at which IBM InfoSphere Change Data Capture for DB2 for z/OS reads from the DB2 for z/OS log affects the overall performance of the incremental update function. The DB2 log buffer and the buffer threshold are parameters that you can tweak to optimize the log reading speed.

IBM InfoSphere Change Data Capture for DB2 for z/OS (the capturing agent) uses the instrumentation facility interface (IFI) of DB2 for z/OS to access the DB2 log. Through the IFI, the agent reads the log entries for all tables for which DATA CAPTURE CHANGES have been enabled. This includes log entries for tables that have been enabled for use with other replication products. In a second step, the agent filters all the entries it has read, and keeps only those that are relevant. The filtering can be quite resource-intensive. It is therefore recommended that you use the following options for optimizing the log reading speed.

DB2 log buffer

The size of the DB2 log buffer determines how often the capturing agent can read records directly from the system memory within a fixed timeframe. Increasing the size of the log buffer increases the log reading speed because reading from memory is much faster than reading from a data set on a hard disk.

For information on how to increase the size of the DB2 log buffer, follow this link to the CDC documentation:

Related information:

http://publib.boulder.ibm.com/infocenter/dzichelp/v2r2/index.jsp?topic= %2Fcom.ibm.db2z10.doc.inst%2Fsrc%2Ftpc%2Fdb2z_ipf_outbuff.htm

DB2 log buffer threshold

The BUFTHRESHOLD keyword specifies a minimum amount of data that must be stored in the DB2 log buffer before the capturing agent reads from the buffer. Setting the threshold higher thus prolongs the interval between the read operations and causes less processing resources to be spent on log reading.

The following links lead to sections in the CDC documentation that describe the BUFTHRESHOLD keyword and explain in more detail how the setting of this keyword can reduce the CPU workload with regard to the capturing process.

Related information:

http://publib.boulder.ibm.com/infocenter/cdc/v6r5m2/topic/ com.ibm.cdcdoc.cdcforzos.doc/concepts/bufthresholdvalueconsiderations.html

http://publib.boulder.ibm.com/infocenter/cdc/v6r5m2/topic/ com.ibm.cdcdoc.cdcforzos.doc/concepts/cpuresourcesconsiderations.html

Starting CDC

Having finished the configuration, start IBM InfoSphere Change Data Capture for z/OS as a z/OS started task.

Procedure

Start CDC as a z/OS started task. For more information, see Executing InfoSphere CDC

If you run into RACF errors that report insufficient access rights for the SO_BROADCAST socket option, proceed as follows:

- 1. Stop the CDC started task.
- Add the following keyword to the TCP/IP statement in the CHCCMMxx member of the configuration control data set: AUTODISCEXCLUDE=<value>

where <value> is a configured IP interface name or an IP address. Using this keyword, you can exclude certain IP interfaces from auto-discovery broadcasting, which is the reason for the error messages.

3. Restart the CDC started task.

Installing and activating the Access Server and the replication engine

Transfer (install) and apply the Access Server and the replication engine for incremental updates like any other accelerator software package.

Transferring packages for the Access Server or the replication engine

Software packages for the Access Server and the replication engine are extra components, which are not automatically installed on the accelerator. Hence you must first transfer these packages from the hierarchical file system (HFS) of your System z data server to the accelerator and apply these in a later step.

About this task

The update packages are copied from the hierarchical file system (HFS) in z/OS to the accelerator, but are not yet activated. However, to actually use a new version, you must activate it. How to do this is described in a later topic. See the hint at the end of this one.

Installing a software update invokes the **SYSPROC.ACCEL_UPDATE_SOFTWARE** stored procedure on your data server. For information about the privileges that are required to run this procedure and further details, see the appropriate section in the *IBM DB2 Analytics Accelerator for z/OS: Stored Procedures Reference*. A link to this document is provided under **Related reference** at the end of this section.

Procedure

- 1. Start IBM DB2 Analytics Accelerator Studio.
- 2. In the Administration Explorer, select the Accelerators folder.
- 3. In the Object List Editor on the right, double-click the accelerator.
- 4. If necessary, expand the **About** section.
- 5. In the About section, click the Transfer updates link.
- 6. In the Transfer Software Versions window, you can see all software packages that are available in the HFS of your z/OS data server. Select the appropriate check boxes in the first column of the table to mark the packages that you want to transfer.

Attention: Make sure that you select the proper packages, that is, packages belonging to the release level that you want to upgrade to. The list in the Transfer Software Versions window might be confusing, especially if it also contains older packages. To find the correct package numbers, see the closing information for the latest program temporary fix (PTF). You find the closing information for a PTF on the support home page, in the category *Plan and install documentation* (see link under **Related information** at the end of this topic). If a PTF was shipped with a major product release, then, in general, you must transfer the packages included in the PTF rather than the packages in the base version.

7. Click **Transfer** to complete the installation.

Related information:

IBM DB2 Analytics Accelerator for z/OS: Stored Procedures Reference

Activating the Access Server and the replication engine

To use the recently transferred packages for the Access Server and the replication engine, you must first activate these packages.

About this task

- The application of an update affects all DB2 subsystems that are connected to an accelerator.
- You must activate the packages for Access Server and the replication engine separately.
- You cannot activate earlier versions of the Access Server or the replication engine. Just upgrades are possible here.

Activating a software update invokes the

SYSPROC.ACCEL_UPDATE_SOFTWARE stored procedure on your data server. For information about the privileges that are required to run this procedure and further details, see the appropriate section in the *IBM DB2 Analytics Accelerator for z/OS: Stored Procedures Reference*. A link to this document is provided under **Related reference** at the end of this section.

Attention: The sequence in which you activate the different software components is important. The sequence depends on the source and on the target version. To find the proper sequence for your particular update, see Table 5 on page 81.

Procedure

- 1. In the Administration Explorer, select the Accelerators folder.
- 2. In the Object List Editor on the right, double-click the accelerator.
- **3**. In the Accelerator view, from the **Refresh** drop-down list in the upper right, select **Automatic off**. Otherwise, you might see warnings during the activation of the new software saying that the accelerator cannot be contacted.
- 4. If necessary, expand the About section.
- 5. In the About section, click the Apply other software version link.
- 6. On the first page of the Apply Software Version wizard, select one of the component that you want to activate:
 - Access Server
 - Replication engine
- 7. Click Next.
- 8. On the second page of the Apply Software Version wizard, you can see all software packages on the accelerator that are currently available for the selected component. To read information about a particular version before you activate it, select the appropriate entry in the list. The information is provided in the **Details of selected version** text box at the bottom.
- **9**. Activate a version by selecting the appropriate radio button in the **Switch To** column.
- 10. Click Finish.
- 11. Repeat steps 4 through 10 for the other component.

Results

When you have successfully activated a package, this or a similar message is displayed:

New software version was successfully activated.

Related information:

IBM DB2 Analytics Accelerator for z/OS: Stored Procedures Reference

Enabling or disabling DB2 subsystems

After installing the capturing agent on your DB2 for z/OS data server, you must still enable incremental updates for all the DB2 subsystems that you want to use the function for. This gives you additional control and security because the enablement process requires a distinct user ID to be specified for running incremental update processes. Furthermore, if the function must be enabled, this also means that it can be disabled, which has the advantage that the capturing agent does not have to be uninstalled for maintenance, error analysis, or error recovery.

Enabling incremental updates for a DB2 subsystem

Follow the steps in this topic to enable incremental updates for a single DB2 subsystem from the IBM DB2 Analytics Accelerator Console. Repeat the process for other subsystems if required.

Before you begin

Make sure that the authentication process (pairing) between the accelerator and the DB2 subsystem that you want to enable has been completed successfully.

About this task

Restriction: You can enable incremental updates for a maximum of ten DB2 subsystems that are connected to one IBM PureData System for Analytics.

Procedure

- 1. Ask the network administrator or the person who did the TCP/IP setup for the IP address (virtual IP or wall IP address) of the accelerator. Make a note of this information. You need to enter it as you complete the steps that follow.
- Start a client or emulator session (using, for example, IBM Personal Communications) to communicate with the z/OS system on which your DB2 subsystem is located.
- 3. Log on to TSO/ISPF.
- Enter the following command: tso telnet <hostname> 1600

where

<hostname>

Is the IP address of the accelerator that is connected to the DB2 for z/OS data server.

1600

Is the number of the port configured for accessing the IBM DB2 Analytics Accelerator Console using a telnet connection between the DB2 for z/OS data server and the accelerator.

For example:

tso telnet 10.101.8.8 1600

- 5. When prompted, enter the console password. The initial password is dwa-1234. You must change this password at the first logon.
- 6. Press the Pause key, then Enter to display the following screen:

7. Type 5 and press Enter to display the submenu:

```
main -> 5

You have the following options:

(0) - Go back one level

(1) - Enable incremental updates

(2) - Disable incremental updates

(3) - Update DB2 subsystem credentials

(4) - Restart replication processes

(Default 0) >
```

- 8. Type 1 and press Enter.
- **9**. The enablement setup process asks you for information. Type the requested information at each prompt and press the Enter.
 - a. Select the DB2 subsystem to be enabled for replication: Type the number of the DB2 subsystem that you want to enable and press Enter.
 - b. Enter the Capture Agent IP address on z/OS ('' or '0' to exit): Enter the virtual IP address of the logical partition (LPAR) on which the CDC Capture Agent is running. Just pressing the Enter key or entering 0 ends the process and you return to the IBM DB2 Analytics Accelerator Console.
 - c. Enter the Capture Agent TCP port ('' or '0' to exit): Enter the port number of the CDC Capture Agent that you specified in step 2 on page 61. By default, this is 5999. The port must be open and must not be blocked by a firewall. Just pressing the Enter key selects the default. Entering 0 ends the process and you return to the IBM DB2 Analytics Accelerator Console.
 - d. Enter the DB2 UserID for replication: Incremental update processes are run under a certain DB2 user ID. This user ID must have SELECT authorization for all tables that you want to enable.

Tip: Select the same user ID for the execution of incremental update jobs as for running the IBM InfoSphere Change Data Capture for z/OS started task. Set the password for this user ID so that it never expires. This has the advantage that you need not worry about expired CDC passwords when running incremental update jobs; you will not have to use options (5) - Manage Incremental Updates and (3) - Update DB2 subsystem credentials on the IBM DB2 Analytics Accelerator Console to correct an error that is due to an expired CDC password.

For details, follow the **Related information** link at the end of this topic.

e. Enter the password: Enter the password of the user entered in the previous step.

Important: If you enter a wrong user ID and password combination, you receive a CHC0025W error. In such a case, proceed as follows:

- 1) Quit the enablement process.
- Disable incremental updates for the selected subsystem. This is option
 9 on the IBM DB2 Analytics Accelerator Console. The entire process is described in a later topic.
- 3) Reenter the enablement setup (option 8 on the console).
- f. Enter the new password again to confirm: Repeat the password and press Enter to confirm it.
- g. Press 'y' to register DB2 subsystem 'XYZ' for replication: Type y, then Enter. When this step has been completed successfully, the word done. is displayed on the screen.

If the enablement fails, proceed as described in the **Important** note to step 9e.

- h. Press the Enter key repeatedly until you reach the main menu of the IBM DB2 Analytics Accelerator Console.
- 10. Type x and press Enter to exit the console.

What to do next

Having configured a DB2 subsystem for incremental updates, you can start or stop incremental updates for the accelerators that are attached to the DB2 subsystem. You can also include or exclude certain tables from the process. You can complete these tasks from IBM DB2 Analytics Accelerator Studio.

Disabling incremental updates for a DB2 subsystem

The disablement of incremental updates might be required for maintenance or error recovery tasks. It also a best practice in case of an enablement failure.

Procedure

- Start a client or emulator session (using, for example, IBM Personal Communications) to communicate with the z/OS system on which your DB2 subsystem is located.
- 2. Log on to TSO/ISPF.
- **3**. Enter the following command:

tso telnet <hostname> 1600

where

<hostname>

Is the IP address of the accelerator that is connected to the DB2 for z/OS data server.

1600

Is the number of the port configured for accessing the IBM DB2 Analytics Accelerator Console using a telnet connection between the DB2 for z/OS data server and the accelerator.

For example:

tso telnet 10.101.8.8 1600

- 4. When prompted, enter the console password. The initial password is dwa-1234. You must change this password at the first logon.
- 5. Press the Pause key, then Enter to display the following screen:

6. Type 5 and press Enter to display the submenu:



- 7. Type 2 and press Enter.
- 8. Provide the requested information:
 - a. Select the DB2 subsystem where replication should be disabled: Type the name of the DB2 subsystem that you want to disable and press Enter.
 - b. Press 'y' to disable replication on DB2 system 'XYZ': Type y, then Enter. When this step has been completed successfully, the phrase done. is displayed on the screen.
 - c. Press the Enter key repeatedly until you reach the main menu of the IBM DB2 Analytics Accelerator Console.
- 9. Type x and press Enter to exit the console.

Chapter 10. Beyond the basics

This chapter contains information about resource allocation and job prioritization, including tips on how to fine-tune IBM DB2 Analytics Accelerator for z/OS with regard to these issues.

Allocating resources of shared accelerators

If an accelerator is shared by multiple DB2 subsystems, you can distribute the available processing resources of the accelerator among the subsystems. This allows you to adapt the processing resources to the size of the query workloads generated by each subsystem.

Before you begin

The authentication process (pairing) must have been completed for the accelerators whose resources you want to allocate.

About this task

By default, the processing resources of an accelerator are distributed equally. Each connected subsystem gets the same share. This prevents a resource drain, that is, a situation in which a single DB2 subsystem occupies all available resources and blocks resource access for other subsystems.

Resources are allocated as a percentage of all available resources, which always add up to 100 percent. If you allocate more than 100 percent in sum, the quotas are reduced proportionally. For example, if two DB2 subsystems are attached to one accelerator, and you assign 80 percent to subsystem A and 30 percent to subsystem B, the result will be that 73 percent of the resources are reserved for subsystem A, and the remaining 27 percent for subsystem B.

You assign accelerator processing resources on the IBM DB2 Analytics Accelerator Configuration Console. The values that you enter are then mapped to guaranteed resource assignment (GRA) definitions for the IBM PureData System for Analytics.

A resource allocation is valid only for a single accelerator and the DB2 subsystems that it is connected to. That is, if you have more than one accelerator attached to the same DB2 subsystem, you must allocate resources for each accelerator separately.

Procedure

- 1. Ask the network administrator or the person who did the TCP/IP setup for the IP address (virtual IP or wall IP address) of the accelerator. Make a note of this information. You need to enter it as you complete the steps that follow.
- 2. Start a client or emulator session (using, for example, IBM Personal Communications) to communicate with the z/OS system on which your DB2 subsystem is located.
- 3. Log on to TSO/ISPF.
- Enter the following command: tso telnet <hostname> 1600

where

<hostname>

Is the IP address of the accelerator that is connected to the DB2 for z/OS data server.

1600

Is the number of the port configured for accessing the IBM DB2 Analytics Accelerator Console using a telnet connection between the DB2 for z/OS data server and the accelerator.

For example:

tso telnet 10.101.8.8 1600

- 5. When prompted, enter the console password. The initial password is dwa-1234. You must change this password at the first logon.
- 6. Press the Pause key, then Enter to display the following screen:

7. Type 3 and press Enter to display the submenu:

main -> 3
 (0) - Go back one level (1) - Obtain pairing code, IP address, and port (2) - List paired DB2 subsystems (3) - Set resource limits for DB2 subsystems (4) - Clear query history (5) - Specify the priority of maintenance tasks (6) - Set the DB2 subsystem for time synchronization (7) - Restart accelerator process (8) - Disable the conversion mode for 24:00:00

8. Type 3 and press Enter. You see a screen similar to the following:

Index	Location	Minimum	Minimum	Maximum	Maximum
	Name	Resources	Relative	Resources	Relative
1	TEST	100	50%	100	
2	PROD	100	50%	100	

9. Type the number of the subsystem that you want to allocate resources to (one of the numbers in the Index column) and press Enter.

```
Maximum resource allocation for system with location name 'TEST' (1-100, 0 or empty to abort):
```

Important: Do not set a maximum unless you are advised to do so by IBM support. Maximum settings cause problems under certain conditions.

10. Press Enter to bypass the maximum setting. You come to the next question:

Minimum resource allocation for system with location name 'TEST' (1-100, 0 or empty to abort):

11. Type an appropriate percentage value and press Enter. For example: 30. The following message is displayed:

Successfully updated minimum resource allocation for system with location name 'TEST' to 30.

Press <return>

12. Press Enter. You see a screen similar to the following:

Index	Location Name	Minimum Resources	Minimum Relative	Maximum Resources	Maximum Relative	
1	TEST	30	23.08%	100	50%	
2	PROD	100	76.92%	100	50%	

- **13.** Repeat steps 8 on page 74 through 12 for other DB2 subsystems as often as required.
- 14. When finished, press just the Enter key repeatedly until you reach the main menu of the console.
- 15. Type x and press Enter to exit the console.

Mapping of WLM importance levels to Netezza priorities

During query execution, z/OS importance levels are automatically converted. Queries thus run with the corresponding Netezza priority on an accelerator.

The priority of DB2 for z/OS queries is handled by WLM service classes of the DDF Classification Rule. The importance levels that are used map to Netezza priorities as shown in Table 2.

WLM importance level	Netezza priority
System	Critical
Importance 1	Critical
Importance 2	Critical
Importance 3	High
Importance 4	Normal
Importance 5	Low
Discretionary	Low

Table 2. Mapping of WLM importance levels to Netezza priorities

Important:

• The allocated accelerator resources set the boundaries for the validity of job priorities. If, for example, two subsystems share an accelerator, and subsystem A is allowed to use 90 percent of the resources, while subsystem B gets only ten percent, priority settings for queries from subsystem B will be considered only within the ten-percent-share. This means that a query with a low priority from

subsystem A might still be processed faster than a query from subsystem B with a high priority. The actual outcome, however, depends on the workload that is assigned to the shares.

• The *period duration* feature of the WLM service classes, that is, the change of importance levels over time, is not supported by Netezza.

Related tasks:

"Allocating resources of shared accelerators" on page 73 If an accelerator is shared by multiple DB2 subsystems, you can distribute the available processing resources of the accelerator among the subsystems. This allows you to adapt the processing resources to the size of the query workloads generated by each subsystem.

Setting the priority of data maintenance tasks

You can set the priority of maintenance tasks, such as table loading, groom or genstats on the IBM DB2 Analytics Accelerator Configuration Console. (The groom and genstats commands are comparable to REORG and RUNSTATS in DB2 for z/OS.)

About this task

The task priorities of IBM DB2 Analytics Accelerator for z/OS are similar to the importance levels of Workload Manager (WLM) service classes. They map to Netezza priorities in the same way. When data maintenance jobs are run on the accelerator, the corresponding Netezza priorities are applied.

IBM DB2 Analytics Accelerator	Netezza priority
SYSTEM	Critical
HIGHEST	Critical
HIGH	High
NORMAL	Normal
LOW	Normal
LOWEST	Normal
DISCRETIONARY	Low

Table 3. Mapping of IBM DB2 Analytics Accelerator priorities for maintenance tasks to Netezza priorities

A default setting exists as well (DEFAULT). It translates as follows:

Table 4. Default priorities of maintenance tasks

Maintenance job type	Priority
Table loading	NORMAL
genstats (equivalent to RUNSTATS in DB2 for z/OS)	LOW
groom (equivalent to REORG in DB2 for z/OS)	LOW

To change the default priority of maintenance tasks, follow these steps:

Procedure

- 1. Ask the network administrator or the person who did the TCP/IP setup for the IP address (virtual IP or wall IP address) of the accelerator. Make a note of this information. You need to enter it as you complete the steps that follow.
- 2. Start a client or emulator session (using, for example, IBM Personal Communications) to communicate with the z/OS system on which your DB2 subsystem is located.
- 3. Log on to TSO/ISPF.
- 4. Enter the following command:

tso telnet <hostname> 1600

where

<hostname>

Is the IP address of the accelerator that is connected to the DB2 for z/OS data server.

1600

Is the number of the port configured for accessing the IBM DB2 Analytics Accelerator Console using a telnet connection between the DB2 for z/OS data server and the accelerator.

For example:

tso telnet 10.101.8.8 1600

- 5. When prompted, enter the console password. The initial password is dwa-1234. You must change this password at the first logon.
- 6. Press the Pause key, then Enter to display the following screen:

7. Type 3 and press Enter to display the submenu:

main -> 3
You have the following options:
 (0) - Go back one level (1) - Obtain pairing code, IP address, and port (2) - List paired DB2 subsystems (3) - Set resource limits for DB2 subsystems (4) - Clear query history (5) - Specify the priority of maintenance tasks (6) - Set the DB2 subsystem for time synchronization (7) - Restart accelerator process (8) - Disable the conversion mode for 24:00:00

8. Type 5 and press Enter.

 9. On the following screen, type one of the numbers in front of the listed DB2 subsystems and press Enter. Example:

```
Select a database system:

1 : DZBD912

2 : DZBD913

3 : DZBD924

Select database system by name or id (0 to go back):
```

In this example, you would type 3 and press Enter to change the priority of maintenance tasks for the accelerator in relation to the query workload that is received from the DZBD924 subsystem.

10. Set the priority by typing the appropriate number. Then press Enter.

```
Queries are run with the priority that is set in the corresponding
System z Workload Manager (WLM) environment. You can set the priority
of maintenance operations, such as loading data, analogously for
the 'DWBDA21' subsystem:
* 1: DEFAULT
2: SYSTEM
3: HIGHEST
4: HIGH
5: NORMAL
6: LOW
7: LOWEST
8: DISCRETIONARY
Enter the appropriate number to set the priority level.
```

For example, type 4 and press Enter to run maintenance tasks with a priority of HIGH. Note that this setting assigns a priority to data maintenance tasks that might be as high or even exceed the priority of incoming queries. If this is the case, query processing can be slower than before.

- 11. Repeat steps 8 on page 77 through 10 for another DB2 subsystem if necessary.
- **12.** When finished, press just the Enter key repeatedly until you reach the main menu of the console.
- **13**. Type x and press Enter to exit the console.

Related tasks:

"Allocating resources of shared accelerators" on page 73 If an accelerator is shared by multiple DB2 subsystems, you can distribute the available processing resources of the accelerator among the subsystems. This allows you to adapt the processing resources to the size of the query workloads generated by each subsystem.

Disabling the conversion of 24:00:00 time values

In the past, it was sometimes necessary to convert 24:00:00 values in TIME or TIMESTAMP columns to 23:59:59 by setting the environment variable LOAD_ENABLE_HOUR24_CONVERSION. With Netezza Performance Server (NPS) 7.0.2 P11, 7.0.4 P3, 7.1, and later versions of the 7.0.2, 7.0.4, and 7.1 series, the Netezza database supports the 24:00:00 representation. Because this makes the conversion superfluous, a function has been added to the IBM DB2 Analytics Accelerator Console so that you can switch the conversion off.

T

1

Т

1

Т

T

1

Before you begin

Make sure that Netezza Performance Server (NPS) 7.0.2 P11, 7.0.4 P3, or 7.1 is installed on your IBM PureData System for Analytics.

Procedure

I

I

I

I

T

T

Т

- 1. Ask the network administrator or the person who did the TCP/IP setup for the IP address (virtual IP or wall IP address) of the accelerator. Make a note of this information. You need to enter it as you complete the steps that follow.
- 2. Start a client or emulator session (using, for example, IBM Personal Communications) to communicate with the z/OS system on which your DB2 subsystem is located.
- 3. Log on to TSO/ISPF.
- 4. Enter the following command:

tso telnet <hostname> 1600

where

<hostname>

Is the IP address of the accelerator that is connected to the DB2 for z/OS data server.

1600

Is the number of the port configured for accessing the IBM DB2 Analytics Accelerator Console using a telnet connection between the DB2 for z/OS data server and the accelerator.

For example:

tso telnet 10.101.8.8 1600

- 5. When prompted, enter the console password. The initial password is dwa-1234. You must change this password at the first logon.
- 6. Press the Pause key, then Enter to display the following screen:

7. Type 3 and press Enter to display the submenu:

main -> 3

```
You have the following options:
(0) - Go back one level
(1) - Obtain pairing code, IP address, and port
(2) - List paired DB2 subsystems
(3) - Set resource limits for DB2 subsystems
(4) - Clear query history
(5) - Specify the priority of maintenance tasks
(6) - Set the DB2 subsystem for time synchronization
(7) - Restart accelerator process
(8) - Disable the conversion mode for 24:00:00
```

8. Type 8 and press Enter.

Т

Т

|

- **9**. When finished, press just the Enter key repeatedly until you reach the main menu of the console.
- 10. Type x and press Enter to exit the console.

Chapter 11. Installing updates

Refer to the appropriate sections to update individual components or migrate from one version to another.

It is rarely necessary to update all listed components. However, if you must update more than one component, follow the suggested order. The order for the individual components depends on the product version or the PTF level.

Important: You cannot upgrade components of IBM DB2 Analytics Accelerator for z/OS directly from version 2.1x to version 4.1.0 or to 3.1.0 PTF levels higher than 3.1.0 PTF-1. In all these cases, you must first migrate to 3.1.0 PTF-1.

Migration from 2.1. <i>x</i> to 3.1.0 PTF-1	Update from 3.1.0 PTF- <i>x</i> to 3.1.0 PTF- <i>y</i>	Migration from 3.1.0 PTF- <i>x</i> to 4.1.0 PTF- <i>x</i>	Update from 4.1.0 PTF- <i>x</i> to 4.1.0 PTF- <i>y</i>	
 IBM DB2 Analytics Accelerator for z/OS Version 3.1.0 FMIDs and program temporary fixes (PTFs) via SMP/E Prerequisite PTFs for DB2 10 for z/OS IBM DB2 Analytics Accelerator Studio Netezza firmware (FDT) Netezza Host Platform (HPF) Netezza Performance Server (NPS) IBM DB2 Analytics Accelerator software IBM DB2 Analytics Accelerator software IBM DB2 Analytics Accelerator stored procedures IBM DB2 Analytics Accelerator Access Server IBM DB2 Analytics Accelerator replication engine IBM InfoSphere Change Data Capture for z/OS Capture Agent 	 Prerequisite program temporary fixes (PTFs) for DB2 10 for z/OS IBM DB2 Analytics Accelerator Studio Netezza Firmware (FDT) Netezza Host Platform (HPF) Netezza Performance Server (NPS) IBM DB2 Analytics Accelerator Access Server IBM DB2 Analytics Accelerator replication engine IBM DB2 Analytics Accelerator software IBM DB2 Analytics Accelerator software IBM DB2 Analytics Accelerator software IBM DB2 Analytics Accelerator stored procedures IBM InfoSphere Change Data Capture for z/OS Capture Agent 	 IBM DB2 Analytics Accelerator for z/OS Version 4.1.0 FMIDs and program temporary fixes (PTFs) via SMP/E Prerequisite PTFs for DB2 10 for z/OS or DB2 11 for z/OS IBM DB2 Analytics Accelerator Studio IBM DB2 Analytics Accelerator Access Server IBM DB2 Analytics Accelerator Access Server IBM DB2 Analytics Accelerator replication engine IBM DB2 Analytics Accelerator software IBM DB2 Analytics Accelerator software IBM DB2 Analytics Accelerator software IBM DB2 Analytics Accelerator stored procedures Netezza Firmware (FDT) (requires assistance from IBM support) Netezza Host Platform (HPF) (requires assistance from IBM support) Netezza Performance Server (NPS) IBM InfoSphere Change Data Capture for z/OS Capture Agent 	 IBM DB2 Analytics Accelerator for z/OS Version 4.1.0 FMIDs and program temporary fixes (PTFs) via SMP/E Prerequisite PTFs for DB2 10 for z/OS or DB2 11 for z/OS IBM DB2 Analytics Accelerator Studio IBM DB2 Analytics Accelerator Access Server IBM DB2 Analytics Accelerator replication engine IBM DB2 Analytics Accelerator replication engine IBM DB2 Analytics Accelerator software IBM DB2 Analytics Accelerator software IBM DB2 Analytics Accelerator stored procedures Netezza Performance Server (NPS) IBM InfoSphere Change Data Capture for z/OS Capture Agent 	

Table 5. Order in which to update components

Important: The IBM DB2 Analytics Accelerator for z/OS software, the Access Server and the replication engine are updated in the same manner, which is why the update of these components is described in the same topics. The order of topic presentation in this manual therefore deviates from the recommended update installation order. For the update of individual components, strictly follow the order in this topic. Jump between the instruction topics where necessary.

Installing prerequisite PTFs for DB2 10 for z/OS or DB2 11 for z/OS

When you update your existing installation, first apply the prerequisite program temporary fixes (PTFs) for DB2 10 for z/OS or DB2 11 for z/OS.

Before you begin

Before you migrate from IBM DB2 Analytics Accelerator for z/OS Version 3.1.0 to Version 4.1.0:

- Make sure that the database management system is DB2 10 for z/OS or DB2 11 for z/OS. DB2 for z/OS Version 9 is not supported.
- Make sure that you meet the software requirements of DB2 for IBM DB2 Analytics Accelerator for z/OS Version 4.1.0. For more information, click the link at the end of this topic.

Note: IBM DB2 Analytics Accelerator for z/OS Version 3.1.0 continues to work after installing the required DB2 for z/OS PTFs, meaning that you can migrate or update and work with the new accelerator version without also having to update the IBM DB2 Analytics Accelerator for z/OS stored procedures.

Procedure

Install the latest prerequisite DB2 for z/OS PTFs for IBM DB2 Analytics Accelerator for z/OS Version 4.1.0.

Related information:

http://www.ibm.com/support/docview.wss?uid=swg27039487

Updating IBM DB2 Analytics Accelerator Studio

The workstation on which IBM DB2 Analytics Accelerator Studio is installed might have an Internet connection or no Internet connection. Refer to the appropriate section for adding or updating IBM DB2 Analytics Accelerator Studio.

About this task

You can use the newest version (4.1.0) of IBM DB2 Analytics Accelerator Studio with older versions of accelerators and stored procedures. Version 2.1.x and 3.1.x of IBM DB2 Analytics Accelerator for z/OS continue to work even if you have installed IBM DB2 Analytics Accelerator Studio Version 4.1.0. However, the user interface does not show the new version 3 or version 4 functions until you update the stored procedures to version 3.1.0 (for version 3 functions) or version 4.1.0 (for version 4 functions).

Using IBM Installation Manager to add or update IBM DB2 Analytics Accelerator Studio from a workstation without an internet connection

Follow the steps in this section if you want to add or update IBM DB2 Analytics Accelerator Studio from a workstation without an internet connection.

Before you begin

Because you cannot update installed products directly by connecting to a repository location on the internet, an administrator must provide relevant update repositories on a computer that you can access

(shared network drive or intranet server for example). Administrators can download update repositories for IBM DB2 Analytics Accelerator Studio from the following website:

IBM Support: Fix Central - Select fixes

The administrator requires an IBM ID and password to log on to the download site. It is possible to register for such an ID during the logon. After the necessary credentials have been provided, IBM Installation Manager will connect to the download site.

After the download, the administrator must extract the compressed archive to a location that you can access.

About this task

If IBM Installation Manager is installed on a workstation without an internet connection, it cannot locate the repositories for updates or additional plugins of already installed products automatically. To update products successfully, you must therefore point IBM Installation Manager to the proper repository location.

Procedure

- 1. Start IBM Installation Manager.
- 2. From the main menu, select File > Preferences > Repositories.
- 3. Click Add Repository.
- 4. Click **Browse** and navigate to the directory that contains the extracted update package.
- 5. Select the repository.config file and click **Open**.
- 6. Click **OK** twice to return the main window of IBM Installation Manager.
- 7. Click Install.
- 8. Follow the instructions in the installation wizard.

Using IBM Installation Manager to add or update IBM DB2 Analytics Accelerator Studio from a workstation with an internet connection

Follow the steps in this section to add or update an existing version of the product from a workstation with an internet connection.

Procedure

- 1. Start IBM Installation Manager.
- 2. Click Update.
- 3. Follow the instructions in the wizard. Deselect components that you do not want to update.

Differences to earlier versions and customization

IBM DB2 Analytics Accelerator Studio 4.1 differs in some respects from its predecessors. You might want to customize your installation so that it uses the workspace directory, default perspective, and language settings of earlier versions.

If you follow the steps in "Using IBM Installation Manager to add or update IBM DB2 Analytics Accelerator Studio from a workstation with an internet connection" or "Using IBM Installation Manager to add or update IBM DB2 Analytics Accelerator Studio from a workstation without an internet connection" on page 82, IBM Data Studio Client 4.1 and IBM DB2 Analytics Accelerator Studio 4.1 are installed. Earlier versions of IBM DB2 Analytics Accelerator Studio, for example version 3.1, used the InstallAnywhere program for the installation. In contrast, version 4.1 installs IBM Data Studio Client 4.1 by using IBM Installation Manager and adds the IBM DB2 Analytics Accelerator functionality as a third-party extension.

When comparing IBM DB2 Analytics Accelerator Studio 4.1 with version 3.1, you will notice the following differences:

- On Windows, the default folder in the start menu is named \IBM Data Studio\ instead of \IBM DB2 Analytics Accelerator Studio\.
- The name of the default workspace is different. If you want to reuse your existing workspace for IBM DB2 Analytics Accelerator Studio, select the following directory path in the workspace launcher that is displayed when you start IBM Data Studio:

C:\Documents and Settings\<username>\IBM\ DB2 Analytics Accelerator Studio 3.1\workspace

where <username> is the name of the log-on user. This is the default workspace directory for IBM DB2 Analytics Accelerator Studio Version 3.1. You might have used a custom directory with earlier versions of the program, in which case you must change your selection accordingly.

- When you start the IBM Data Studio Client, the Accelerator perspective is not the default perspective. To activate the Accelerator perspective, select Window > Open Perspective > Other > Accelerator from the menu bar.
- IBM DB2 Analytics Accelerator Studio Version 4.1 is available in English only. However, your version of IBM Data Studio might display some interface elements in the language that you specified in the regional settings of your operating system. To avoid a mixture of elements in English and another language, configure the IBM Data Studio version that you use for IBM DB2 Analytics Accelerator for z/OS to display all interface elements in English only. To do that, follow these steps:
 - 1. Locate the eclipse.exe file in the file path of your IBM Data Studio installation.
 - 2. Start IBM Data Studio by running the following command: eclipse.exe "-nl en US"

On Windows, you can also complete these steps to the same end:

- 1. Right-click the Data Studio 4.1 Client shortcut and select Properties from the menu.
- 2. In the Properties Wizard, on the **Shortcut** tab, edit the **Target** field: At the end of the target entry, add a space character followed by "-nl en_US".
- 3. Click OK.

Updating the IBM DB2 Analytics Accelerator software, the Netezza Performance Server (NPS), the Access Server, or the replication engine

These are software components that run on the accelerator hardware, that is, the IBM PureData System for Analytics.

Related information:

http://www.ibm.com/support/docview.wss?uid=swg27039487

Installing update packages for the accelerator

IBM DB2 Analytics Accelerator for z/OS software includes specific libaries and other code to be installed on the IBM PureData System for Analytics. Updates for the Access Server, the replication engine, and the Netezza Performance Server (NPS) are delivered in separate packages.

Before you begin

Make sure that the following conditions apply:

- Make sure that IBM DB2 Analytics Accelerator for z/OS has been successfully connected to DB2 for z/OS.
- IBM UNIX System Services is installed on your z/OS data server.
- You have a user ID and a password to log on to your z/OS data server. The user ID has read access to the hierarchical file system (HFS) that is used by IBM UNIX System Services.
- The user ID is authorized to run the SYSPROC.ACCEL_UPDATE_SOFTWARE stored procedure.

About this task

The entire installation or update procedure consists of several steps:

- 1. Installing the update packages on your z/OS data server.
- 2. Transferring the update packages to the accelerator. You start this process from IBM DB2 Analytics Accelerator Studio, the remote client application of IBM DB2 Analytics Accelerator for z/OS.
- **3**. Activating the newly installed software. This step is also executed from IBM DB2 Analytics Accelerator Studio.

This topic covers step 1 only.

Procedure

- 1. Log on to your z/OS data server using a remote client, such as IBM Personal Communications.
- 2. Use SMP/E to install the PTF or APAR fix in a directory path called usr/lpp/aqt/packages in your hierarchical file system on z/OS.

Important:

The directory path usr/lpp/aqt/packages must be present, but it does not have to start in the root directory. You can create this directory path in any other directory, but then you must change the value of the environment variable AQT_INSTALL_PREFIX in the <HLQSP>.SAQTSAMP(AQTENV) data set before you transfer the packages to the accelerator in the following step. By default, AQT_INSTALL_PREFIX is set to NULL, which means that it points to a usr/lpp/aqt/packages directory path starting from the root directory. So if, for example, you want to install the packages in a directory called /mnt/software/usr/lpp/aqt/packages, you must set the AQT_INSTALL_PREFIX environment variable as follows:

AQT_INSTALL_PREFIX=/mnt/software

• The <AQT_INSTALL_PREFIX>/usr/lpp/aqt/packages directory often requires multiple gigabytes (GB) of disk space. For example, if you work with multiple accelerators in a rolling upgrade scenario, you must have multiple instances of these packages in the file system. So make sure that the directory can accommodate all the files. The *IBM DB2 Analytics Accelerator for z/OS: Program Directory* lists the exact disk space requirements.

Results

Your local usr/lpp/aqt/packages directory now contains the accelerator update package version.tar.z.

Transferring update packages for the accelerator

Transfer update packages for IBM DB2 Analytics Accelerator for z/OS, the Access Server, the replication engine, or the Netezza Performance Server (NPS) by completing the steps in this section.

Before you begin

Make sure that the following conditions apply:

- Suitable accelerator installation packages exist in the usr/lpp/aqt/packages directory on your z/OS data server. Update packages for the IBM DB2 Analytics Accelerator for z/OS software are named *version.tar.z*, where *version* is the first part of the package name, which indicates the version of the software. Update packages for the Access Server, the replication engine, or the Netezza Performance Server (NPS) are named differently.
- The value of the AQT_INSTALL_PREFIX environment variable determines where IBM DB2 Analytics Accelerator Studio looks for software update packages. If it is not set correctly, the packages cannot be located and thus cannot be transferred or installed.

Important: During the SMP/E installation or manual copying, the usr/lpp/aqt/packages was created on a volume specified by the installer. The target system must be able to access this directory. The path to the usr/lpp/aqt/packages directory does not have to start in the root directory. It might have been created in any other directory. In this case, the environment variable AQT_INSTALL_PREFIX must be set accordingly in the <HLQSP>.SAQTSAMP(AQTENV) data set, where <HLQSP> is the chosen high-level qualifier for stored-procedure libraries. By default, the AQT_INSTALL_PREFIX environment variable is set to NULL, which means that it points to a usr/lpp/aqt/packages directory path starting from the root directory. So if, for example, the packages have been copied to a directory /mnt/software/usr/lpp/ aqt/packages, the AQT_INSTALL_PREFIX environment variable must have been set as follows: AQT_INSTALL_PREFIX=/mnt/software

- The user ID has read and write access to the <AQT_INSTALL_PREFIX>/usr/lpp/aqt/packages directory, where <AQT_INSTALL_PREFIX> stands for the value of this environment variable. You specify this variable in the AQTENV data set. The AQTENV data set must be referenced by the Workload Manager (WLM) environment that was set up for the IBM DB2 Analytics Accelerator for z/OS stored procedures. Furthermore, read access is required for all subdirectories of the <AQT_INSTALL_PREFIX>/ usr/lpp/aqt/packages directory. To avoid a setup with obsolete or wrong entries, use the sample AQTENV data set that comes with IBM DB2 Analytics Accelerator for z/OS Version 4.1.0.
- When you transfer an NPS update package, it is checked whether the NPS version is compatible with the installed Netezza Firmware (FDT) and Netezza Host Platform (HPF) versions. If the NPS is incompatible, you must first update the FDT, the HPF, or both. These tasks, however, can only be completed with the help of IBM support and require a downtime of about one day. So, if the NPS update is really necessary, make sure that you plan it thoroughly in advance. For more information, follow the **Related information** link at the end of this topic.

Another caveat exists in connection with an NPS transfer: If the NPS package fails the compatibility check, the process ends prematurely without transferring any packages, no matter how many other packages were selected. It is therefore recommended that you first transfer packages for IBM DB2 Analytics Accelerator, the Access Server, or the replication engine, then activate these, and transfer and activate the NPS update in a later step.

About this task

The update packages are copied from the hierarchical file system (HFS) in z/OS to the accelerator, but are not yet activated. However, to actually use a new version, you must activate it. How to do this is described in a later topic. See the hint at the end of this one.

Installing a software update invokes the **SYSPROC.ACCEL_UPDATE_SOFTWARE** stored procedure on your data server. For information about the privileges that are required to run this procedure and further details, see the appropriate section in the *IBM DB2 Analytics Accelerator for z/OS: Stored Procedures Reference.* A link to this document is provided under **Related reference** at the end of this section.

Procedure

- 1. Start IBM DB2 Analytics Accelerator Studio.
- 2. In the Administration Explorer, select the Accelerators folder.
- 3. In the Object List Editor on the right, double-click the accelerator.
- 4. If necessary, expand the About section.

- 5. In the About section, click the Transfer updates link.
- 6. In the Transfer Software Versions window, you can see all software packages that are available in the HFS of your z/OS data server. Select the appropriate check boxes in the first column of the table to mark the packages that you want to transfer.

Attention: Make sure that you select the proper packages, that is, packages belonging to the release level that you want to upgrade to. The list in the Transfer Software Versions window might be confusing, especially if it also contains older packages. To find the correct package numbers, see the closing information for the latest program temporary fix (PTF). You find the closing information for a PTF on the support home page, in the category *Plan and install documentation* (see link under **Related information** at the end of this topic). If a PTF was shipped with a major product release, then, in general, you must transfer the packages included in the PTF rather than the packages in the base version.

7. Click **Transfer** to complete the installation.

What to do next

To save space in the SMP/E target directory of the HFS, consider deleting sub-directories that contain already transferred packages. Be careful, however, if you have more than one accelerator. Do not delete packages if these are still needed for other accelerators.

If the AQT_INSTALL_PREFIX environment variable does not point directly to the SMP/E target directories but to a copy, you might also want to delete the copy.

Related information:

"Updating other Netezza software" on page 89

Updates of the Netezza Firmware (FDT) and the Netezza Host Platform (HPF) cannot be installed with SMP/E and are therefore carried out by an IBM service engineer. Follow the steps here to open a service request (formerly called problem management record or PMR) and provide IBM support with the necessary information.

Plan and install documentation for DB2 Analytics Accelerator for z/OS

Switching between available versions of accelerator software components

To use a recently transferred software package, you must first activate that package by selecting the version that the package contains. In the case of IBM DB2 Analytics Accelerator for z/OS software, you can also switch to an older version if one of your accelerators has problems with the newest version.

Before you begin

- You cannot activate a software version as long as a IBM DB2 Analytics Accelerator Console window is open. Therefore, close all console windows. If you are not sure whether there are any open console windows, click **Cancel Tasks** in the relevant Accelerator view of IBM DB2 Analytics Accelerator Studio and cancel all active console tasks from the Cancel Tasks window.
- If you want to activate new versions of the Access Server and the replication engine and two DB2
 subsystems are enabled for incremental updates, disable one of these subsystems from the IBM DB2
 Analytics Accelerator Console before the activation. Otherwise, you will not be able to enable more
 than those two subsystems in the future (version 4.1 of the product allows you to enable up to ten DB2
 subsystems).

About this task

- The application of an update affects all DB2 subsystems that are connected to an accelerator.
- An update of the accelerator software does not automatically include the components for the incremental update function or the Netezza Performance Server (NPS). You must activate updates of the Access Server, the replication engine, and the Netezza Performance Server (NPS) separately.

- The activation of a different accelerator software version causes a restart of the accelerator.
- You cannot activate earlier versions of the Access Server or the replication engine. Just upgrades are possible here.

Activating a software update invokes the **SYSPROC.ACCEL_UPDATE_SOFTWARE** stored procedure on your data server. For information about the privileges that are required to run this procedure and further details, see the appropriate section in the *IBM DB2 Analytics Accelerator for z/OS: Stored Procedures Reference.* A link to this document is provided under **Related reference** at the end of this section.

Attention:

- The activation of a new version of IBM DB2 Analytics Accelerator for z/OS software components might delete the query history from the accelerator if the new version introduces changes to the query history function.
- The sequence in which you activate the different software components is important. The sequence depends on the source and on the target version. To find the proper sequence for your particular update, see Table 5 on page 81.

Procedure

- 1. In the Administration Explorer, select the Accelerators folder.
- 2. In the Object List Editor on the right, double-click the accelerator.
- **3.** In the Accelerator view, from the **Refresh** drop-down list in the upper right, select **Automatic off**. Otherwise, you might see warnings during the activation of the new software saying that the accelerator cannot be contacted.
- 4. If necessary, expand the About section.
- 5. In the **About** section, click the **Apply other software version** link.
- 6. On the first page of the Apply Software Version wizard, select the component whose active version you want to change:
 - Accelerator (IBM DB2 Analytics Accelerator for z/OS software)
 - Netezza Performance Server (NPS) (software for the Netezza hosts)
 - Access Server (Access Server component for incremental updates)
 - **Replication engine** (replication engine or Apply Agent for incremental updates)
- 7. Click Next.
- 8. On the second page of the Apply Software Version wizard, you can see all software packages on the accelerator that are currently available for the selected component. To read information about a particular version before you activate it, select the appropriate entry in the list. The information is provided in the **Details of selected version** text box at the bottom.
- 9. Activate a version by selecting the appropriate radio button in the Switch To column.
- 10. Click Finish.

Results

When the process has been completed successfully, a message similar to the following is displayed: New software version was successfully activated.

Note: In some cases, migration processes are still ongoing even though the success message was displayed, for example a restart of the Netezza database server. In a situation like this, it might take additional time until the accelerator is ready to process queries. You can check the status of an accelerator at the top of the corresponding Accelerator view. The status must be *online* for an accelerator to process queries.

What to do next

To avoid a cluttered Apply Software Version window, remove packages from the accelerator that you no longer need. For instructions, follow the link under **Related tasks**.

Related tasks:

"Removing obsolete software packages from an accelerator" on page 90 To free up space on an accelerator and reduce the number of versions that are displayed in the Apply Software Version window, you can remove packages that you no longer need.

Related information:

IBM DB2 Analytics Accelerator for z/OS: Stored Procedures Reference

Updating other Netezza software

Updates of the Netezza Firmware (FDT) and the Netezza Host Platform (HPF) cannot be installed with SMP/E and are therefore carried out by an IBM service engineer. Follow the steps here to open a service request (formerly called problem management record or PMR) and provide IBM support with the necessary information.

1. In IBM DB2 Analytics Accelerator Studio, enable tracing for the accelerator that you want to update. A trace level of DEFAULT is sufficient.

For more information, see Tracing in the IBM DB2 Analytics Accelerator Studio: User's Guide.

- 2. Save the trace information to a file.
- **3.** Open a service request at https://www.ibm.com/support/servicerequest/Home.action, preferably two weeks before the planned migration date. This gives IBM support enough time to analyze your system, coordinate maintenance actions, and update components as required.
- 4. If you plan to update the NPS, and the HPF or FDT updates are required for the NPS update, state this in the body of the PMR. Also state the target NPS version.
- 5. Attach the trace file to the service request.

What to do next

To save space in the SMP/E target directory of the HFS, consider deleting sub-directories that contain already transferred packages. Be careful, however, if you have more than one accelerator. Do not delete packages if these are still needed for other accelerators.

If the AQT_HOST_PACKAGE_DIRECTORY environment variable does not point directly to the SMP/E target directories but to a copy, you might also want to delete the copy.

Transferring Netezza Firmware (FDT) updates or Netezza Host Platform (HPF) updates

When you have submitted the service request for the update, IBM support will tell you the proper download location and the names of the packages that you have to download. After finishing the download, transfer the update packages as a preparation for the IBM service personnel, who will help you install the updates on your IBM PureData System for Analytics.

Before you begin

Netezza update packages must exist in the download directory for packages of this type in the hierarchical file system (HFS) of your z/OS data server. Update packages for the Netezza host consist of single files only. The target system must be able to access this directory.

Important: The AQT_HOST_PACKAGE_DIRECTORY environment variable points to this download directory. Unlike AQT_INSTALL_PREFIX, it specifies an absolute path (starting from the root directory).

The AQT_HOST_PACKAGE_DIRECTORY environment is also set in the <HLQSP>.SAQTSAMP(AQTENV) data set, where <HLQSP> is the chosen high-level qualifier for stored-procedure libraries. The value of this environment variable determines where IBM DB2 Analytics Accelerator Studio looks for Netezza updates. If it is not set correctly, the packages cannot be located and thus cannot be transferred to the IBM PureData System for Analytics. To avoid a setup with obsolete or wrong entries, use the sample AQTENV data set that comes with IBM DB2 Analytics Accelerator for z/OS Version 4.1.

About this task

Netezza update packages are not installed automatically. They are just transferred to a directory on the IBM PureData System for Analytics, from where you must install them manually with the help of IBM support.

Transferring a Netezza update invokes the **SYSPROC.ACCEL_UPDATE_SOFTWARE** stored procedure on your data server. For information about the privileges that are required to run this procedure and further details, see the appropriate section in the *IBM DB2 Analytics Accelerator for z/OS: Stored Procedures Reference*. A link to this document is provided under **Related reference** at the end of this section.

To start an update transfer process, follow the steps in this section.

Procedure

- 1. Start IBM DB2 Analytics Accelerator Studio.
- 2. In the Administration Explorer, select the Accelerators folder.
- 3. In the Object List Editor on the right, double-click the accelerator.
- 4. If necessary, expand the About section.
- 5. In the **About** section, click the **Transfer updates** link.
- 6. In the Transfer Software Versions window, you can see all Netezza Firmware (FDT) and Netezza Host Platform (HPF) files that are available in the HFS of your z/OS data server. Select the appropriate check boxes in the first column of the table to mark the files that you want to transfer.
- 7. Click Transfer to complete the process.

Removing obsolete software packages from an accelerator

To free up space on an accelerator and reduce the number of versions that are displayed in the Apply Software Version window, you can remove packages that you no longer need.

About this task

The removal wizard can remove all packages that have been deployed with the help of the transfer wizard, that is:

- Accelerator software packages
- Access Server packages
- Replication engine packages
- Netezza Performance Server (NPS)
- Netezza Firmware (FDT)
- Netezza Host Platform (HPF)

Note: You can only remove software packages from regular accelerators.

Procedure

- 1. Open the appropriate Accelerator view.
- 2. If necessary, expand the About section.
- 3. Click the **Remove** link in the server subsection.

- 4. In the Remove Software Versions window, select the software packages that you want to remove. To do so, select the appropriate check boxes in the first column. As a decision aid, you might want to read the information under **Details of selected version**, which is displayed for each package.
- 5. Click Remove.

What to do next

The removal wizard that is described in this topic can only remove packages from an accelerator. To free up space in the hierarchical file system (HFS) of your z/OS data server, manually delete no longer needed IBM DB2 Analytics Accelerator for z/OS or Netezza update packages from the HFS. To do so, open a UNIX System Services shell and delete the files by using the **rm** command.

Updating the z/OS components of IBM DB2 Analytics Accelerator

Read how to update the DB2 for z/OS components of IBM DB2 Analytics Accelerator for z/OS.

Before you begin

Before you can migrate to version 4.1.0, you must migrate to version 3.1.0. If you need to upgrade from version 2.1.x to 3.1.0 first, consider the following prerequisites:

- The version of the accelerator software and the stored procedures must be version 2.1.3 or later. If this is not the case, first update the accelerator software and then the stored procedures by applying PTF UK82772.
- Versions from 3.1.0 onwards require more extensive authorizations for users who load accelerator tables or move table partitions with the High Performance Storage Saver (HPSS). In addition to the authority to run the UNLOAD utility, these users must be authorized to run the RUNSTATS, REORG, and LOAD utilities.
- The authority of the user who runs the AQTTIJSP installation job for version 3.1.0 must be extended so that it includes SELECT authority for the SYSIBM.SYSTABLESPACESTATS table.

About this task

During an update, IBM DB2 Analytics Accelerator for z/OS remains operational for most of the time.

It is possible to run a version 3.1.0 or 4.1.0 accelerator if the stored procedures on a connected DB2 system are still at level 2.1.x (but not the other way around!). A mixed environment, where you have version 2.1.x stored procedures on one DB2 subsystem or member of a data sharing group and version 3.1.0 or 4.1.0 stored procedures on another, is also supported.

Important: If you first have to migrate from version 2.1.x to 3.1.0, do not activate (change the configuration to use) the updated stored procedures before updating the accelerator software to level 3.1.0. Otherwise, IBM DB2 Analytics Accelerator Studio will hang as it tries to recognize the new features in version 3.1.0.

Procedure

- 1. To update the libraries with IBM DB2 Analytics Accelerator for z/OS support in DB2 for z/OS (including the stored procedures), use the apply function in SMP/E and follow the ++HOLD actions in the program temporary fixes (PTFs).
- 2. If you intend to accelerate static SQL queries, verify that the table
- SYSACCEL.SYSACCELERATEDPACKAGES and related indexes exist in the relevant DB2 subsystems.
- I If not, create this table and the related indexes in each connected subsystem by running the DB2 for
- z/OS job DSNTIJAS. You find instructions in the section *Creating the IBM DB2 Analytics Accelerator for*
- *z/OS database*. Follow the link under **Related tasks** tasks at the end of this topic.

- **3**. For the configuration, revisit the sections that are listed at the end of this topic. Consider the following:
 - a. When you set up the WLM application environment:
 - 1) Reuse the existing WLM JCL procedure for the version 3.1.x stored procedures.
 - 2) Edit the JCL and change the high-level qualifier (HLQ) in the following statements so that the production HLQ for version 4.1.0 data sets is used:
 - DD <HLQACTIVE>.SAQTMOD
 - AQTENV DD <HLQACTIVE>.SAQTSAMP(AQTENV)
 - **3)** Append the following line to the same JCL:
 - AQTOSR DD DSN=<HLQACTIVE>.SAQTOSR(AQTOSR),DISP=SHR
 - b. Copy the customized settings from your old version 3.1.0 AQTENV production member (the member that you previously referred to in the AQTENV DD statement of the JCL for the WLM setup) to the new AQTENV data set member for version 4.1.0. Do not refresh your WLM application environment yet.

Important: Do not reuse the AQTENV data set member for version 3.1.0 because this member is incompatible with version 4.1.0.

- c. For an upgrade installation, customize the AQTTIJSP job as documented and run just the necessary job steps:
 - 1) AQTTICO4
 - 2) AQTTIND
 - 3) AQTTIGR
 - 4) If you use DB2 11 for z/OS, additionally execute the steps AQTTINDA and AQTTIGRA.
- d. Refresh the WLM application environment for the stored procedures. This activates stored procedure code for version 4.1.0. You might want to verify this by starting IBM DB2 Analytics Accelerator Studio. If the stored procedure code was updated successfully, the graphical user interface shows the new version 4.1.0 features.
- e. Free the version 3.1.0 packages, that is, packages whose name starts with SYSACCEL.AQT03.

Related tasks:

"Creating the IBM DB2 Analytics Accelerator database" on page 18

Customize and submit the DSNTIJAS job to create the database and tables for IBM DB2 Analytics Accelerator for z/OS in DB2 for z/OS.

"Setting up a WLM application environment for IBM DB2 Analytics Accelerator for z/OS stored procedures" on page 27

Follow the steps in this section to set up a suitable Workload Manager (WLM) application environment.

"Verifying the correct setup of DB2-supplied stored procedures" on page 31

The DB2 for z/OS stored procedures SYSPROC.ADMIN_INFO_SYSPARM, SYSPROC.DSNUTILU, and SYSPROC.ADMIN_COMMAND_DB2 must run in different Workload Manager (WLM) environments that are separate from the one used by the IBM DB2 Analytics Accelerator for z/OS stored procedures. Verify that this and a few other requirements are met by following the steps here.

"Customizing and running AQTTIJSP" on page 35

Customize the AQTTIJSP job member (JCL) for the installation of IBM DB2 Analytics Accelerator stored procedures before you run the job.

"Verifying the installation of IBM DB2 Analytics Accelerator for z/OS stored procedures" on page 36 Verifying the installation of IBM DB2 Analytics Accelerator for z/OS stored procedures does not require a connection to IBM DB2 Analytics Accelerator for z/OS on an IBM PureData System for Analytics.

Related reference:

"Setting access rights for the user who runs AQTTIJSP" on page 34

The AQTTIJSP installation job, which you need to run in a later step, requires a user ID with sufficient DB2 access rights. These access rights are listed here.

Running different versions of IBM DB2 Analytics Accelerator for z/OS at the same time

In connection with a DB2 data sharing group, you can run IBM DB2 Analytics Accelerator for z/OS Version 2.1.x and later versions at the same time. Read how to create an environment that allows coexistence.

About this task

When you migrate the z/OS components of IBM DB2 Analytics Accelerator - including the stored procedures - from version 2.1.x to a later version in a DB2 10 for z/OS data sharing group, you can update one or more members to the later version, while other members stay at version 2.1.x. The coexistence of different versions is also possible if the members reside in the same logical partition (LPAR). In the latter case, the members running different versions share the same Workload Manager (WLM) application environment, which means that the same JCL procedure is used to define this environment. You must therefore configure this JCL procedure accordingly.

Procedure

- Identify a DB2 10 member (A) to be migrated from version 2.1.x to a later version, while at least one other member (B) of the same data sharing group continues to run IBM DB2 Analytics Accelerator for z/OS Version 2.1.x.
- 2. Check whether you must use the same WLM JCL procedure for versions 2.1.x and the later version. This is the case if member B runs in the same LPAR as member A, or if your JCL procedure library, for example USER.PROCLIB, is shared throughout the sysplex. Depending on the result of this check:
 - If you must use the same WLM JCL procedure for both versions, continue with step 3.
 - If you can use different WLM JCL procedures, create a second procedure with the same name for the newer stored procedures of member A, and follow steps 3a on page 92 through 3d on page 92 for member A. Skip step 3e on page 92 because the version 2.1.x stored procedure packages are still required for coexistence.

If the WLM JCL procedure is shared by the different members, use a data set and alias naming convention similar to the one described in *Special considerations for WLM-managed stored procedures in coexistence* of the DB2 10 for z/OS information center. See the link at the end of this topic.

- **3**. In step 3a on page 92 of section, choose a string for <HLQACTIVE> that includes the **DB2SSN** JCL parameter. This allows release-specific data set names for the relevant members of the data sharing group.
- 4. In step 3b on page 92, do not include the **DB2SSN** parameter in the AQTOSR DD statement. Use only the name that this library has in the newer version. This is because the AQTOSR library does not exist in version 2.1.x.
- 5. Use aliases to redirect the data-set names of both installed versions to member-specific data-set names.

Example

If member A is named DB2A, member B is named DB2B, AQT is the high-level-qualifier, and you want to migrate member A from version 2.1.x to a later version, proceed as follows:

1. Update the WLM JCL procedure as shown:

//DSNWLMVA PROC RGN=0K, APPLENV= DSNWLMVA, NUMTCB=15,

- //IEFPROC EXEC PGM=DSNX9WLM,REGION=&RGN,TIME=NOLIMIT,
- // PARM='&DB2SSN,&NUMTCB,&APPLENV'
- //STEPLIB DD DISP=SHR,DSN=<HLQDB2SSN>.SDSNEXIT
- // DD DISP=SHR,DSN=<HLQBASE>.SDSNLOAD
- // DD DISP=SHR,DSN=<HLQBASE>.SDSNLOD2
- // DD DISP=SHR,DSN=AQT.&DB2SSN..SAQTMOD
- // DD DISP=SHR,DSN=<HLQXML4C1.10>.SIXML0D1

//SYSTSPRT DD SYSOUT=A
//CEEDUMP DD SYSOUT=H
//OUT1 DD SYSOUT=A
//UTPRINT DD SYSOUT=A
//DSSPRINT DD SYSOUT=A
//SYSPRINT DD SYSOUT=A
//AQTENV DD DSN=AQT.&DB2SSN..SAQTSAMP(AQTENV),DISP=SHR
//AQTOSR DD DSN=AQT.DB2A.SAQTOSR(AQTOSR),DISP=SHR

2. Create the following aliases for the later version and member A:

AQT.DB2A.SAQTMOD	*ALIAS
AQT.DB2A.SAQTSAMP	*ALIAS
AQT.DB2A.AQTOSR	*ALIAS

3. Create the following aliases for version 2.1.x and member B:

AQT.DB2B.SAQTMOD	*ALIAS
AQT.DB2B.SAQTSAMP	*ALIAS

Related tasks:

"Setting up a WLM application environment for IBM DB2 Analytics Accelerator for z/OS stored procedures" on page 27

Follow the steps in this section to set up a suitable Workload Manager (WLM) application environment.

Related information:

Special considerations for WLM-managed stored procedures in coexistence
Chapter 12. Troubleshooting

In the following sections, you find descriptions of known IBM DB2 Analytics Accelerator for z/OS problems. The author and the development team have tried to provide a solution wherever possible. However, a solution might be unavailable because the cause of a problem cannot be clearly identified. This is mostly the case if multiple causes can lead to the same symptom. It can also be that a solution has not yet been found. In such cases, contact IBM support.

Message AQT10200I and SQL code -206 when running stored procedures after migration

Stored procedure calls end with message AQT10200I and SQLCODE= -206 after a migration of IBM DB2 Analytics Accelerator for z/OS from version 2.1.x to a later version.

Symptoms

The following message is displayed:

AQT10200I - The OPEN CURSOR operation failed. Error information: "DSNT408I SQLCODE = -206, ERROR: ARCHIVE IS NOT VALID IN THE CONTEXT WHERE IT IS USED"

Causes

Required program temporary fixes (PTFs) for DB2 for z/OS have not been installed or the respective ++HOLD migration actions have not been applied yet.

Resolving the problem

Follow the steps in section *Installing prerequisite PTFs for DB2 10 for z/OS* of the *IBM DB2 Analytics Accelerator for z/OS: Installation Guide*.

Related tasks

"Installing prerequisite PTFs for DB2 10 for z/OS or DB2 11 for z/OS" on page $82\,$

When you update your existing installation, first apply the prerequisite program temporary fixes (PTFs) for DB2 10 for z/OS or DB2 11 for z/OS.

Message DSNUTILU NOT INVOKED APF AUTHORIZED

You run the DSNUTILU stored procedure and receive a reply message that says DSNUTILU NOT INVOKED APF AUTHORIZED.

Symptoms

You cannot run the DSNUTILU stored procedure successfully.

Causes

The stored procedure is not APF-authorized.

Resolving the problem

Make sure that all libraries used by the WLM application environment for DB2 stored procedures are APF-authorized in the STEPLIB statement of the startup JCL procedure.

Error during bind step for stored procedure or function

You receive an error message as you try to bind an IBM DB2 Analytics Accelerator for z/OS stored procedure or function.

Symptoms

The error message is similar to this one: BIND PACKAGE(SYSACCEL) MEMBER(AQT02TRC) ACTION(REPLACE) ISO(CS) CURRENT..... EBIND) KEEPDYNAMIC(NO) LIBRARY('DB291.ISAO.SAQTDBRM') DSNX100I § BIND SQL WARNING USING HACA02 AUTHORITY PLAN=(NOT APPLICABLE) DBRM=AQT02TRC STATEMENT=154 FUNCTION DSNAQT.ACCEL_READFILE IS NOT DEFINED <<<<

Causes

Mismatch between parameter lists.

This error occurs if a new version of an IBM DB2 Analytics Accelerator for z/OS stored procedure has a modified parameter list compared with the old version. In the example, the versions of the module AQT02TRC and the FUNCTION DSNAQT.ACCEL_READFILE interface do not match.

Resolving the problem

Contact IBM support.

No CEEDUMPs

You cannot find a CEEDUMP although a signal was caught while you were running an IBM DB2 Analytics Accelerator for z/OS stored procedure.

Symptoms

A signal was caught while running an IBM DB2 Analytics Accelerator for z/OS stored procedure, but a CEEDUMP has apparently not been produced.

Causes

Missing configuration settings.

Resolving the problem

To analyze problems that cause the stored procedure to end abnormally or stop with a signal, a CEEDUMP is required. The JCL startup procedure for the Workload Manager (WLM) application environment of the IBM DB2 Analytics Accelerator for z/OS stored procedure allows you to specify a location for a CEEDUMP.

- Make sure that the RUN OPTION 'TRAP(0N)' is set for the stored procedures. This option is set for all IBM DB2 Analytics Accelerator for z/OS stored procedures by default, but it could have been overwritten when the AQTTIJSP job member in the SAQTSAMP data set was customized at installation time. It could also be that the setting was changed by an ALTER PROCEDURE run.
- Make sure that the JCL startup procedure for the Workload Manager (WLM) application environment of the IBM DB2 Analytics Accelerator for z/OS stored procedure contains a DD statement for CEEDUMP.

PRIQ value too high when creating table spaces

During table-space creation, you receive a message saying that the PRIQ value was exceeded.

Symptoms

The problem occurs if you want to create a table and a table space must be created before this. The following error message might be displayed: IDC32211 CONSTANT '16777216' NOT WITHIN VALUE RANGE

Causes

Program Temporary Fix (PTF) UK43901 is missing.

Resolving the problem

Install the PTF or use smaller sizes.

The DB2 command -DIS ACCEL does not work

You cannot run the DB2 for z/OS command -DIS ACCEL.

Symptoms

The -DIS ACCEL command does not return the expected accelerator information.

Resolving the problem

- 1. Verify that the DB2 ZPARMs are configured as described in the section "Installing libraries with IBM DB2 Analytics Accelerator for z/OS support" on page 17*Installing DB2 libraries with IBM DB2 Analytics Accelerator for z/OS support* of the *IBM DB2 Analytics Accelerator for z/OS: Installation Guide*.
- 2. Verify that the IBM DB2 Analytics Accelerator for z/OS libraries are part of the STEPLIB statement.

Connection authorization failure (error -4214)

You receive a message about a connection authorization failure with error code -4214.

Symptoms

The full message is similar to this one: [jcc][t4][2010][11246][4.7.89] Connection authorization failure occurred. Reason: Local security service non-retryable error. ERRORCODE=-4214, SQLSTATE=28000

Causes

No access to the DB2 subsystem, although z/OS can be accessed.

Diagnosing the problem

The error is probably due to a missing RACF authorization.

Resolving the problem

Check the RACF security settings.

Errors during ZPARM compilation

You receive error messages during ZPARM compilation.

Symptoms

You receive the following error messages: X00370097 VOLTDEVT=SYSDA, X00370098 XLKUPDLT=NO, X00370099 ZOSMETRICS=NO ** ASMA017W Undefined keyword parameter; default to positional, including keyword - DSN6S/OPTIOWGT ** ASMA017W Undefined keyword parameter; default to positional, including keyword - DSN6S/OPTJBPR ** ASMA017W Undefined keyword parameter; default to positional, including keyword - DSN6S/ZOSMETRICS ** ASMA435I Record 3165 in DB291.ISAO.LAB.SDSNMACS(DSN6SPRM) on volume: TSMS08 57+* Avoid overflow DK153 000000 00000 00814 60+DSN6SPRM CSECT CSECT name 02-DSNDS

Causes

These errors are caused by the following parameters, which are unknown because they have been introduced by Program Temporary Fixes (PTFs) that are newer than those included in the base level of the libraries with IBM DB2 Analytics Accelerator for z/OS support:

- OPTIOWGT
- OPTJBPR
- ZOSMETRICS

Resolving the problem

Update the libraries with IBM DB2 Analytics Accelerator for z/OS support.

DRDA connection does not work

You can ping the accelerators, but you cannot establish a distributed relational database access (DRDA) connection between your database management system and the IBM PureData System for Analytics.

Symptoms

DB2 commands or IBM DB2 Analytics Accelerator for z/OS stored procedures cannot establish a TCP/IP connection with the accelerator. Running the SYSPROC.ACCEL_TEST_CONNECTION stored procedure might expose this symptom.

Resolving the problem

Make sure that the distributed data facility (DDF) of DB2 for z/OS uses the same TCP/IP stack as the ping program.

Removing orphaned system-table entries and catalog-table entries

If an IBM PureData System for Analytics was physically removed, but the accelerator was not removed before by running the *Remove accelerator* function from IBM DB2 Analytics Accelerator Studio or by running the SYSPROC.ACCEL_REMOVE_ACCELERATOR stored procedure, you find invalid entries in a number of DB2 for z/OS catalog tables and system tables.

Symptoms

You find invalid entries in the following tables:

- SYSIBM.IPNAMES
- SYSIBM.USERNAMES
- SYSIBM.LOCATIONS
- SYSACCEL.SYSACCELERATORS
- SYSACCEL.SYSACCELERATEDTABLES

Resolving the problem

- 1. Disable the accelerator in the DB2 subsystem, by using the -DIS ACCEL command or the appropriate function in IBM DB2 Analytics Accelerator Studio.
- 2. To remove the invalid entries from the system tables, run the following SQL commands in the order indicated:
 - a. DELETE FROM SYSIBM.IPNAMES WHERE LINKNAME=(SELECT LINKNAME FROM SYSIBM.LOCATIONS WHERE LOCATION=(SELECT LOCATION FROM SYSACCEL.SYSACCELERATORS WHERE ACCELERATORNAME=<acceleratorName>));
 - b. DELETE FROM SYSIBM.USERNAMES
 WHERE LINKNAME=(SELECT LINKNAME FROM SYSIBM.LOCATIONS
 WHERE LOCATION=(SELECT LOCATION FROM SYSACCEL.SYSACCELERATORS
 WHERE ACCELERATORNAME=<acceleratorName>));
 - c. DELETE FROM SYSIBM.LOCATIONS
 WHERE LOCATION=(SELECT LOCATION FROM SYSACCEL.SYSACCELERATORS
 WHERE ACCELERATORNAME=<acceleratorName>);
 - d. DELETE FROM SYSACCEL.SYSACCELERATORS
 WHERE ACCELERATORNAME=<acceleratorName>;
 - e. DELETE FROM SYSACCEL.SYSACCELERATEDTABLES WHERE ACCELERATORNAME=<acceleratorName>;

Related tasks

Appendix E, "Disabling accelerators in a DB2 subsystem," on page 117 It is recommended that you disable an accelerator before you activate an IBM DB2 Analytics Accelerator for z/OS software update or take actions to solve a problem.

Package not found when running a stored procedure from IBM DB2 Analytics Accelerator Studio

You receive a message saying that a package was not found when you try to run an IBM DB2 Analytics Accelerator for z/OS stored procedure from IBM DB2 Analytics Accelerator Studio.

Symptoms

You receive a message similar to this one: SQL0805N Package "<location>.NULLID.SYSSTAT.5359534C564C3031" was not found. SQLSTATE=51002

Causes

The package has not been bound due to an IBM DB2 Analytics Accelerator Studio installation error.

Resolving the problem

Bind the package manually. The following methods can be used:

- From the DB2 command-line client:
 - 1. cd <DB2-client-install-folder>\bnd

where <DB2-client-install-folder> is the fully qualified path to the installation folder of the DB2 command-line client.

- db2 connect to <database-name> where <database-name> is the database to which the stored procedure belongs.
- 3. db2 bind @db2cli.lst grant public
- Using the DB2Binder utility from a Windows command-prompt:
 - 1. Adjust the following command as needed. Then press the Enter key.

```
cd /d <idaa-studio-install-directory>
\plugins\com.ibm.datatools.db2_2.1.403.v20120228_2105\driver
```

where <idaa-studio-install-directory> is the drive and installation directory of IBM DB2 Analytics Accelerator Studio on your local workstation, for example C:\Program Files\IBM\ IBM DB2 Analytics Accelerator Studio 2.1. The full name of the com.ibm.datatools.db2_ directory changes with each new driver. So make sure that you choose the correct directory.

2. Enter:

```
..\..\jre\bin\java -cp db2jcc4.jar;
db2jcc4_license_cisuz.jar;
db2jcc4_license_cu.jar com.ibm.db2.jcc.DB2Binder
-url jdbc:db2://<server>:<port>/<location>
-user <user-id> -password <password>
```

where

```
<server>
```

Is the host name of the DB2 data server

```
<port>
```

Is the port on which the DB2 data server listens to JDBC requests

<location>

Is the unique name of the database server. An application uses the location name to access a DB2 database server. A database alias can be used to override the location name when accessing a remote server.

<user-id>

Is a user ID with the privilege of running the DB2 Binder utility

<password>

Is the password belonging to <user-id>

Packet delivery failure

The stored procedures can establish a connection to the IBM PureData System for Analytics, but sometimes, the packet delivery fails, especially during loads.

Symptoms

You notice communication timeout errors between stored procedures and IBM DB2 Analytics Accelerator for z/OS. Usually, a stored procedure sends a request and does not receive an answer. The problem is prevalent in SYSPROC.ACCEL_CONTROL_ACCELERATOR calls using the **<getInfo>** parameter if a reply requires the sending of data packages bigger than 1.5 KB.

Causes

In most cases, this is caused by network components that do not support jumbo frames.

Resolving the problem

- 1. Make sure that all network components, especially switches, support jumbo frames. Sometimes, you must configure the switches to support jumbo frames.
- **2**. To verify that your configuration works, click the **Test connection** button in the Add Accelerator window.

SQL code -430 from IBM DB2 Analytics Accelerator for z/OS stored procedures

A stored procedure of IBM DB2 Analytics Accelerator for z/OS ends abnormally, and you receive an error message with SQL code -430.

Symptoms

IBM DB2 Analytics Accelerator for z/OS stored procedures end abnormally. In IBM DB2 Analytics Accelerator Studio, the Administration Explorer returns a message window like this one:

😑 Prob	lem Occurred	
8	'DSE Model Refresher' has encount PROCEDURE SYSPROC.ACCEL_CONTROL_ACCE (SPECIFIC NAME SYSPROC.ACCEL_CONTROL_ACCE ABNORMALLY TERMINATED. SQLCO SQLSTATE=38503, DRIVER=4.7.89	ered a problem. LERATOR) HAS DDE=-430,
	ОК	<< Details
PROCEDI	JRE SYSPROC.ACCEL_CONTROL_AC JRE SYSPROC.ACCEL_CONTROL_AC	CELERATOR (SF CELERATOR (SF
<	Ш	>

Figure 4. SQL code -430 message window

Causes

This is probably a configuration problem.

Diagnosing the problem

- 1. Verify that the IBM DB2 Analytics Accelerator for z/OS stored procedures run in a separate Workload Manager (WLM) environment. Each application environment must be set up according to the instructions in *Setting up a WLM application environment for IBM DB2 Analytics Accelerator for z/OS stored procedures*. In particular, the ADMIN_INFO_SYSPARM and DSNUTILU stored procedures must run in different WLM environments and NUMTCB must be set to the correct value.
- 2. Rerun the AQTSJI00 and AQTSJI01 jobs to verify the correct installation of the IBM DB2 Analytics Accelerator for z/OS stored procedures.

SQL code -471 referring to a function in the DSNAQT schema

You receive a message with SQLCODE= -471 and reason code 00E79002 when trying to run the internal function DSNAQT.ACCEL_READFILE.

Symptoms

You see an error message that is similar to the following:

DSNT408I SQLCODE = -471, ERROR: INVOCATION OF FUNCTION OR PROCEDURE DSNAQT.ACCEL_LIST_SOFTWARE FAILED DUE TO REASON 00E79002 DSNT418I SQLSTATE = 55023 SQLSTATE RETURN CODE DSNT415I SQLERRP = DSNX9GPL SQL PROCEDURE DETECTING ERROR DSNT416I SQLERRD = -40 0 0 -1 0 0 SQL DIAGNOSTIC INFORMATION DSNT416I SQLERRD = X'FFFFFD8' X'00000000' X'00000000' X'FFFFFFF' X'00000000' X'00000000' SQL DIAGNOSTIC INFORMATION

Causes

In addition to externally published stored procedures in the SYSPROC schema, IBM DB2 Analytics Accelerator for z/OS employs user-defined functions internally. These functions are included in the DSNAQT schema and serve the following purposes:

- Reading temporary trace files from IBM DB2 Analytics Accelerator for z/OS stored procedures (function DSNAQT.ACCEL_READFILE)
- Checking available software versions of IBM DB2 Analytics Accelerator Studio (function DSNAQT.ACCEL_GETVERSION)

If such a function was stopped for some reason (for example by a database administrator), the end-user who has triggered the execution of the internal function sees a DB2 error -471 00E79002.

Resolving the problem

1. From your DB2 subsystem, check the function status by using the following command:

-DIS FUNCTION SPECIFIC (DSNAQT.*)

 If the function is not in the state STARTED, start it with the following command: -STA FUNCTION SPECIFIC (DSNQAT.*)

The trace information that was supposed to be delivered to the caller of the stored procedure in a DB2 result set has been kept in its temporary location (/tmp by default). Thus you can still transfer this information after restarting the stopped function.

SYSPROC.ACCEL_LOAD_TABLES returns SQL error -471 and reason code E790002 for DSNUTILU

You try to load accelerator tables, but the SYSPROC.ACCEL_LOAD_TABLES stored procedure returns SQL error -471 and reason code E790002. This error is related to the DB2 stored procedure SYSPROC.DSNUTILU.

Symptoms

The load performance is poor and load processes do not run to completion.

Causes

Wrong setup of the Workload Manager (WLM) application environment for DSNUTILU

Resolving the problem

See "Adjusting WLM performance goals for SYSPROC.ACCEL_LOAD_TABLES" on page 32.

Load of partitioned tables freezes during unload phase

You can load unpartitioned tables without problems, but the process stalls when you try to load tables with more than one partition.

Symptoms

The DB2 command -DISPLAY UTIL(*) shows two or more UNLOAD utility processes that appear to be active, but only with a few unloaded rows. The number of displayed unloaded rows does not change when you run the -DISPLAY UTIL(*) command repeatedly.

Causes

The DB2 stored procedure SYSIBM.DSNUTILU has been started more than once in the same address space.

Resolving the problem

1. Make sure that NUMTCB is set to 1.

You can change the value in the JCL for the setup of the Workload Manager (WLM) application environment (part of SYSIBM.DSNUTILU) so that it is set when you run the SYSIBM.DSNUTILU stored procedure, or set it as a start parameter in the definition of the WLM application environment.

Important: A NUMTCB start parameter in the definition of the application environment overrides an equivalent setting in the JCL.

2. Make sure that the address space for DSNUTILU is managed by the WLM and not limited to a single instance per system or per sysplex.

Message SYSACCELERATORS DOES NOT HAVE PROPER COLUMN DEFINITIONS

The message SYSACCELERATORS DOES NOT HAVE PROPER COLUMN DEFINITIONS is returned by the -start ACCEL command.

Symptoms

The complete message text that is displayed is similar to the following screen output:

-D8Q0 START ACCEL(SAPACCEL) DSNX892I -D8Q0 DSNX8CTG DSNACCEL TABLE 829 SYSACCELERATORS DOES NOT HAVE PROPER COLUMN DEFINITIONS, COLUMN ACCELERATORIPNAME IS MISSING OR INCORRECTLY DEFINED DSNX800I -D8Q0 DSNX8STA ACCELERATOR FUNCTION IS NOT AVAILABLE DSNX810I -D8Q0 DSNX8CMD START ACCEL FOLLOWS -DSN9022I -D8Q0 DSNX8CMD '-START ACCEL' NORMAL COMPLETION

Causes

You added IBM DB2 Analytics Accelerator for z/OS to a DB2 V9.1 for z/OS subsystem, but forgot to set the ACCEL_LEVEL ZPARM to the correct value. The ACCEL_LEVEL parameter specifies the version of the accelerator that DB2 for z/OS is supposed to use.

Resolving the problem

Define the ACCEL_LEVEL ZPARM in the DSN6SPRM macro. Values:

V1 (default)

Instructs DB2 for z/OS to use version 1 of IBM DB2 Analytics Accelerator for z/OS.

V2 Instructs DB2 for z/OS to use version 2 of IBM DB2 Analytics Accelerator for z/OS.

Notes:

- You cannot change the ACCEL_LEVEL ZPARM online.
- The ACCEL_LEVEL ZPARM is deprecated in DB2 Version 9.1 for z/OS.
- This information does not apply to DB2 10 for z/OS. DB2 10 for z/OS does not support version 1 of IBM DB2 Analytics Accelerator for z/OS, so the parameter does no longer exist.

Repeated IBM InfoSphere Change Data Capture for z/OS errors after abnormal end of DB2 for z/OS

IBM InfoSphere Change Data Capture for z/OS repeatedly reports errors after an abnormal end of DB2 for z/OS.

Symptoms

An error message similar to the following is displayed:

CHC9210E (CDCSZA1) DB2 CAF IFI-READS request has failed. RetturnCode=X'00000100', ReasonCode=X'00F30018'

Causes

If incremental updates are configured, and DB2 for z/OS ends unexpectedly, IBM InfoSphere Change Data Capture for z/OS is not notified of the shutdown. IBM InfoSphere Change Data Capture for z/OS still "expects" to find the instance of DB2 that was previously running. When DB2 for z/OS is restarted, IBM InfoSphere Change Data Capture for z/OS tries to renew the connection to the known DB2 instance. That is, it uses the session configuration or parameters of the previously stalled connection. These will never be correct, so the error message is displayed after each retry.

Resolving the problem

Restart the IBM InfoSphere Change Data Capture for z/OS address space.

Appendix A. Sample TCP/IP configuration

In this section, you find a sample VTAM definition and a sample profile for the TCP/IP setup on your System z.

Schematic overview

The following figures give a schematic overview of a working setup. A detailed description of this setup can be found in *Chapter 5: QDIO mode for z/OS* of the *OSA-Express Implementation Guide*. To read this guide, click the **Related information** link at the end of this topic.



Figure 5. Recommended connections between IBM zEnterprise[®] server and IBM PureData System for Analytics (including switches)

Important: The switches that you use must support jumbo frames (a frame size of 8992). With some switches, you must enable support for this type of frames in the switch configuration.

Sample VTAM definition

The following sample VTAM definition is for the setup in Figure 5.

**				
* GIGABIT ETHERNET OSA EXPRESS *				
* VIAM NODE :	IPAMNZA *			
**				
TRLE0 VBUIL	D TYPE=TRL			
IPANZA0 TRLE	LNCTL=MPC,READ=E100, WRITE=E101,DATAPATH=E102, PORTNAME=GBE100,MPCLEVEL=QDIO			
IPANZA1 TRLE	LNCTL=MPC,READ=E300, WRITE=E301,DATAPATH=E302, PORTNAME=GBE300,MPCLEVEL=QDIO			

The following profile uses the names introduced in the VTAM definition.

Sample TCP/IP profile (PROFILE.TCPIP data set)

```
The following sample TCP/IP profile also refers to the setup in Figure 5 on page
105.
IPCONFIG MULTIPATH PERCONNECTION
; If you do not use a switch, use NOMULTIPATH instead
; STATIC VIPA DEFINITIONS for IBM DB2 Analytics Accelerator network
DEVICE VIPA1 VIRTUAL 0
LINK VIPA1L VIRTUAL 0 VIPA1
; Interfaces to IBM DB2 Analytics Accelerator
;Ten Gigabit Interface Definition for Portname GBE100
   INTERFACE TENGBEE1 DEFINE IPAQENET
    PORTNAME GBE100
   IPADDR 10.1.0.137/24
   MTU 8992
    INBPERF DYNAMIC PRIR
    VMAC ROUTEALL
    SOURCEVIPAINT VIPA1L
;Ten Gigabit Interface Definition for Portname GBE300
   INTERFACE TENGBEE3 DEFINE IPAQENET
   PORTNAME GBE300
    IPADDR 10.1.0.237/24
   MTU 8992
    INBPERF DYNAMIC PRIR
    VMAC ROUTEALL
   SOURCEVIPAINT VIPA1L
;
HOME
    10.1.0.37
                   VIPA1L
BEGINRoutes
; Direct Routes - Routes that are directly connected to my interfaces
; Destination Subnet Mask First Hop Link Name Packet Size ;
                      255.255.255.0 = TENGBEE1 mtu 8992
255.255.255.0 = TENGBEE3 mtu 8992
 ROUTE 10.1.0.0
 ROUTE 10.1.0.0
ENDRoutes
;
; In case a global SOURCEVIPA or TCPSTACKSOURCEVIPA address is defined,
; the source IP address for requests targeting the accelerator network
; needs to be explicitly defined as the VIPA for the accelerator
; network on z/OS.
; Otherwise, communication and tools (such as ssh) addressing the private network
; will fail because the requestor address cannot be reached.
SRCIP
 DESTINATION
                     10.1.0.0/24 10.1.0.37
ENDSRCIP
 START TENGBEE1 ; 10 GBE to Netezza
 START TENGBEE3 ; 10 GBE to Netezza
;
Related information:
Http://www.redbooks.ibm.com/redbooks/pdfs/sg245948.pdf
http://www.ibm.com/support/docview.wss?uid=swg27023654
http://www.ibm.com/support/docview.wss?uid=swg27024236
```

Appendix B. Members of SAQTSAMP

The SAQTSAMP data set contains various samples for installing, configuring or running the IBM DB2 Analytics Accelerator for z/OS stored procedures. The following table lists all these members and provides brief descriptions of their functions.

Member name	Function
AQTENV	Default environment variable settings for the IBM DB2 Analytics Accelerator for z/OS stored procedures.
AQTNOTE	Member that contains legal information.
AQTSCALL	Sample application program in the C programming language with embedded SQL CALL statements for calling IBM DB2 Analytics Accelerator for z/OS stored procedures.
AQTSCI01	DB2 command line processor script that calls the following stored procedures in the order indicated:1. SYSPROC.ACCEL_REMOVE_TABLES
	 (returns only the version of the stored procedure) 2. SYSPROC.ADMIN_INFO_SYSPARM 3. SYSPROC.ADMIN_COMMAND_DB2(-DIS ACCEL) 4. SYSPROC.ADMIN_COMMAND_DB2(-DIS GROUP) 5. SYSPROC.DSNUTILU(UNLOAD)

Table 6. Members of the SAQTSAMP data set

Table 6. Members of the SAQTSAMP data set (continued)

Member name	Function
AQTSCI02	DB2 command line processor script that calls the following stored procedures and DB2 commands in the order indicated: 1. SYSPROC.ACCEL_ADD_ACCELERATOR for the initial setup
	 Running this stored procedure requires a valid pairing code. You must therefore edit AQTSCI02 so that a valid pairing code is provided when SYSPROC.ACCEL_ADD_ACCELERATOR is invoked. Important: In the sample script, the call of this stored procedure has been commented out. To activate the call, you must uncomment the corresponding line. SYSPROC.ACCEL_TEST_CONNECTION SYSPROC.ACCEL_ADD_TABLES SYSPROC.ACCEL_LOAD_TABLES SYSPROC.ACCEL_LOAD_TABLES
	To enforce a reload of the tables and enable these for query acceleration:
	1. SYSPROCACCEL LOAD TABLES
	2. SYSPROC.ACCEL_SET_TABLES_ACCELERATION(ON)
	 To execute the query: START -ACCEL Query execution SYSPROC.ACCEL_GET_QUERIES STOP -ACCEL To check various other functions: SYSPROC.ACCEL_ALTER_TABLES SYSPROC.ACCEL_REMOVE_TALBES SYSPROC.ACCEL_CONTROL_ACCELERATOR (including the various subfunctions) SYSPROC.ACCEL_REMOVE_ACCELERATOR Important: In the sample script, the call of this stored procedure has been commented out. To activate the call, you must uncomment the corresponding line.
	Note: This script requires customization.
AQTSJI00	JCL that collects information about required IBM DB2 Analytics Accelerator for z/OS databases and tables.
AQTSJI01	JCL that calls the DB2 command line processor for verifying IBM DB2 Analytics Accelerator for z/OS stored procedures without an accelerator.
AQTSJI02	JCL that calls all verification steps one-by-one. For a successful completion, a running and connected accelerator is required.
AQTSJI03	JCL that compiles, links, and invokes the AQTSCALL sample program, which calls IBM DB2 Analytics Accelerator for z/OS stored procedures. For a successful completion of the job, a running accelerator is required.

Table 6. Members of the SAQTSAMP data set (continued)

Member name	Function
AQTSJI04	JCL for a manual recovery of a partition that was moved with the High Performance Storage Saver (HPSS). With the help of this JCL, you can recover a moved partition if the accelerator that holds the data is unavailable. No changes will be made on the accelerator. The JCL uses the image copies in DB2 for z/OS for the recovery. Important: When the accelerator returns to the <i>online</i> state at a later point-in-time, the data might not be in sync anymore with the data in DB2 for z/OS. This might result in different query results. To solve this problem, restore all moved partitions to DB2 for z/OS, and delete the corresponding table from the accelerator when the accelerator is back online.
AQTSSCHK	 UNIX System Services shell script that checks the output of DB2 command line processor scripts containing IBM DB2 Analytics Accelerator for z/OS stored procedure calls. If BPXBATCH is used to call this script, as in the AQTSJI02 member, the return code of this job step is determined by the MESSAGE output parameters of the IBM DB2 Analytics Accelerator for z/OS stored procedures. Divide this return code by 256 to determine the severity of an error: 0 All stored procedures that were called returned information messages 4 Warning (return code 1024) 8 Error (return code 2048) 12 Severe (return code 3072)
AQTSSCPY	UNIX System Services shell script that copies SAQTSAMP members to the /tmp/ivp directory in the hierarchical file system. Note: DB2 command line processor scripts cannot be run from a data set.
AQTSXADD	XML definition of a simple table specification (<tablespecifications> element). This XML code can be used as input for the SYSPROC.ACCEL_ADD_TABLES stored procedure.</tablespecifications>
AQTSXALT	XML definition of a table specification as input for the SYSPROC.ACCEL_ALTER_TABLES stored procedure.
AQTSXCN0	XML code for the deletion of collected trace data. The XML code is used as the value of the COMMAND parameter of the SYSPROC.ACCEL_CONTROL_ACCELERATOR stored procedure.
AQTSXCN1	XML code for the retrieval of accelerator status information. The XML code is used as the value of the COMMAND parameter of the SYSPROC.ACCEL_CONTROL_ACCELERATOR stored procedure.
AQTSXCN2	XML code for the configuration of accelerator tracing. The XML code is used as the value of the COMMAND parameter of the SYSPROC.ACCEL_CONTROL_ACCELERATOR stored procedure.
AQTSXCN3	XML code for the collection of trace data. The XML code is used as the value of the COMMAND parameter of the SYSPROC.ACCEL_CONTROL_ACCELERATOR stored procedure.
AQTSXCN4	XML code that lists the active tasks on the accelerator. The XML code is used as the value of the COMMAND parameter of the SYSPROC.ACCEL_CONTROL_ACCELERATOR stored procedure.
AQTSXIM0	XML code that returns just the version of an IBM DB2 Analytics Accelerator for z/OS stored procedure without executing its actual function. The XML code is used as the value of the MESSAGE input parameter.

Table 6. Members of the SAQTSAMP data set (continued)

Member name	Function
AQTSXQHI	XML definition that serves as an input value for the QUERY_SELECTION parameter of the SYSPROC.ACCEL_GET_QUERIES stored procedure.
AQTSXSD1	All XML schema definitions (xsd files) for the input and output parameters of the IBM DB2 Analytics Accelerator for z/OS stored procedures.
AQTSXTCO	XML definition that serves as an input value for the DIAGNOSTIC_INPUT parameter of the SYSPROC.ACCEL_TEST_CONNECTION stored procedure.
AQTSXTS0	XML definition of a table set as input for various stored procedures.
AQTSXTSL	XML definition that serves as an input value for the TABLE_LOAD_SPECIFICATION parameter of the SYSPROC.ACCEL_LOAD_TABLES stored procedure.
AQTSXTSU	XML definition for an update of selected table partitions. The definition serves as an input value for the TABLE_LOAD_SPECIFICATION parameter of the SYSPROC.ACCEL_LOAD_TABLES stored procedure.
AQTTIJSP	JCL for the installation of the IBM DB2 Analytics Accelerator for z/OS stored procedures.

Appendix C. Environment variables

The job control language (JCL) for the configuration of the Workload Manager (WLM) environment for IBM DB2 Analytics Accelerator for z/OS stored procedures contains a data definition (DD) "AQTENV". This data definition includes a data set in which environment variables are defined. These variables control the behavior of some stored procedures.

Important:

• When editing the AQTENV data set with an ISPF editor, make sure not to use the NUM ON option. Otherwise the line numbers of the columns from 72 to 80 become part of the variable values. Also, do not insert blanks before or after the equals sign and make sure that you do not have trailing blanks at the end of the line.

The AQTENV data set is made available to the stored procedures by the RUN OPTION 'ENVAR("_CEE_ENVFILE_S=DD:AQTENV")', which is set in the CREATE PROCEDURE statement for each procedure. If a line in the AQTENV data set matches the pattern *NAME=VALUE*, the environment variable *NAME* is set to *VALUE*. For more information, follow the **Related information** link at the end.

- Make sure that the permissions for the AQTENV data set include read access for all users who execute stored procedures. Otherwise, the environment variable settings do not take effect and corresponding error messages are written to the system log.
- After modifying settings in the AQTENV data set, refresh the WLM environment so that the changes can take effect.
- Do not reuse the AQTENV data set member that was delivered with version 3.1.0 because this member is incompatible with IBM DB2 Analytics Accelerator for z/OS Version 4.1.0.

When upgrading to a new release, always use the AQTENV sample that is provided with the new version as a template for your own AQTENV data set. This ensures that you do not set obsolete environment variables and specify all mandatory new variables that were introduced with the new release.

New in version 4.1.0 PTF-2: You can set or override environment variables temporarily on a per-call basis. Settings will be valid only for the duration of a stored-procedure call. For more information, follow the link to the *IBM DB2 Analytics Accelerator for z/OS: Stored Procedures Reference* at the end of this topic.

AQT_ARCHIVE_COPY1

Specifies the name of the first image-copy data-set to be created when the SYSPROC.ACCEL_ARCHIVE_TABLES stored procedure moves partition data from DB2 for z/OS to an accelerator. For example:

AQT_ARCHIVE_COPY1 = &USERID..&DB..&TS..P&PART..&UNIQ.

where

&USERID.

ID of the user who runs SYSPROC.ACCEL_ARCHIVE_TABLES

&DB.

Name of the database that a partition resides in

Т

L

L

Т

&TS.

Name of the table space that the partition resides in

&PART.

Identifier of the (physical) partition. The letter P in the example is a text or string constant used as a prefix. This is required because &PART. resolves to a numeric value like 00001, and this not valid for qualifiers in a data-set name.

&UNIQ.

Causes the creation of a unique identifier

The template specification in the example could, for instance, result in the following image copy name:

BCKE.V4L1.BCKERTSE.CKRANGE3.P00001.D72R4KHN

- All template variables that are documented for the DB2 COPY utility can be used, with the exception of &SEQ (&SQ), &LIST (&LI), and &DSNUM.
- The chosen variables must ensure the uniqueness of image-copy data-set names. It is therefore recommended that you use at least the &PART. and &UNIQ. template variables.
- Templates must resolve to valid z/OS data set names.
- The template data-set names that you use must have been mapped to suitable data classes in the DFSMS.

Important: The AQT_ARCHIVECOPY_HLQ environment variable that was used with earlier versions of the HPSS is deprecated. Remove it from the AQTENV data set.

AQT_ARCHIVE_COPY2

Specifies the name of the second image-copy data-set to be created when the SYSPROC.ACCEL_ARCHIVE_TABLES stored procedure moves partition data from DB2 for z/OS to an accelerator. For more details, see AQT_ARCHIVE_COPY1.

AQT_ARCHIVE_RECOVERYCOPY1

Specifies the name of the third image-copy data-set to be created when the SYSPROC.ACCEL_ARCHIVE_TABLES stored procedure moves partition data from DB2 for z/OS to an accelerator. For more details, see AQT_ARCHIVE_COPY1.

AQT_ARCHIVE_RECOVERYCOPY2

Specifies the name of the fourth image-copy data-set to be created when the SYSPROC.ACCEL_ARCHIVE_TABLES stored procedure moves partition data from DB2 for z/OS to an accelerator. For more details, see AQT_ARCHIVE_COPY1.

AQT_ENABLE_MULTIPLE_ENCODING

Enables the SYSPROC.ACCEL_ADD_TABLES stored procedure to add both, EBCDIC and UNICODE tables, to the same accelerator, provided that these tables are located in the same DB2 subsystem.

The setting takes effect only if it is the first table that is added to an accelerator or if the accelerator already contains EBCDIC tables. You cannot add EBCDIC tables if the accelerator already contains UNICODE tables; in this case, you can only define tables of both types if you first remove all tables.

You cannot define ASCII tables and UNICODE or EBCDIC tables on the same accelerator.

Note: In version 4 PTF-1 of the product, enabling tables with UNICODE and EBCDIC encoding might result in sorting differences between DB2 for z/OS and the accelerator for certain queries against UNICODE tables. Such differences can occur if the queries contain explicit cast statements to CHAR or VARCHAR data types. These problems disappear when you install version 4 PTF-2.

AQT_HOST_PACKAGE_DIRECTORY

Specifies the directory in the hierarchical file system (HFS) of your z/OS system in which those update packages for the Netezza host are stored that are not installed with SMP/E. Currently, upgrades for the Netezza Host Platform (HPF) and the Netezza Firmware (FDT) cannot be installed with SMP/E. Instead, packages of this type must be downloaded via FTP to the specified directory, so that you can transfer these to the Netezza host. To install these packages, you must open a service request. See *Updating other Netezza software* in the *IBM DB2 Analytics Accelerator for z/OS: User's Guide* for more information.For more information, follow the link to *Updating other Netezza software* at the end of this topic.

AQT_INSTALL_PREFIX

1

I

L

T

L

I

Points to the installation directory that was specified during the SMP/E installation. The software packages are always installed in the usr/lpp/aqt/packages directory. This directory path is created in the specified installation directory. To successfully complete the software installation from IBM DB2 Analytics Accelerator Studio, AQT_INSTALL_PREFIX must point to the correct installation directory.

For example, if you set AQT INSTALL PREFIX=/mnt/software, the package

By default, AQT_INSTALL_PREFIX is not set (its value is an empty string). This means that IBM DB2 Analytics Accelerator Studio expects to find an installation path /usr/lpp/aqt/packages starting from the root directory. For the installation to succeed in this case, the packages must have been

AQT_MAX_RETRIES_DSNUTILU

put in the root directory by SMP/E.

Determines how many times the SYSPROC.ACCEL_LOAD_TABLES stored procedure tries to call DSNUTILU if the attempt failed in the first place with SQL code -471 and reason code 00E79002 (timeout error). The default value is 2.

If partition or table data is being moved by the

location is /mnt/software/usr/lpp/aqt/packages.

SYSPROC.ACCEL_ARCHIVE_TABLES stored procedure, AQT_MAX_RETRIES_DSNUTILU also controls the number of pruning attempts. Mostly, an attempt to prune the data fails because an exclusive lock on the data cannot be obtained during the timeout period that is set by the UTMOUT subsystem parameter. For more information, follow the appropriate link under **Related information** at the end of this topic.

AQT_MAX_UNLOAD_IN_PARALLEL

The maximum number of parallel DSNUTILU invocations used by the SYSPROC.ACCEL_LOAD_TABLES, SYSPROC.ACCEL_ARCHIVE_TABLES, and SYSPROC.ACCEL_RESTORE_ARCHIVE_TABLES stored procedures when loading data from a partitioned table. Increasing the value leads to a better performance, provided that enough processors are available to

handle additional parallel processes. Note also that increasing the value of AQT_MAX_UNLOAD_IN_PARALLEL to more than 8 might not increase the throughput any further. The default value is 4.

Important: If you increase the value of

AQT_MAX_UNLOAD_IN_PARALLEL, you must, in most cases, also increase the NUMTCB value of the WLM application environment for the IBM DB2 Analytics Accelerator for z/OS stored procedures. During the initial setup, the NUMTCB value is set to 15. Increase this value according to the following formula:

<value of NUMTCB> > = 3 × <value of AQT_MAX_UNLOAD_IN_PARALLEL> + 1

Thus, if you increase the value of AQT_MAX_UNLOAD_IN_PARALLEL to 6, you must set NUMTCB to 19 at least.

AQT_SECONDS_BEFORE_RETRY_DSNUTILU

Specifies the interval in seconds between DSNUTILU calls (retries). By default, this interval is set to 60 seconds.

AQT_SORTDEVT

T

1

Т

Specifies the device type to use when sorting temporary data sets dynamically. This setting is used by the sort program when anIBM DB2 Analytics Accelerator for z/OS stored procedure sends a request to a DB2 Utility that includes a sort job like this.

AQT_UTILITY_TMP (new in version 4.1.0 PTF-2)

Specifies a template for the generation of unique temporary data-set names to be used by DB2 utilities. For example, several IBM DB2 Analytics Accelerator stored procedures invoke DB2 utilities, which in turn create temporary data sets with the user ID of the stored-procedure caller as the high-level qualifier. If you do not want this to happen, you can set AQT_UTILITY_TMP to generate a different name. The syntax is: AQT_UTILITY_TMP = &USERID..AQT.&UNIQ.

The specification needs to resolve to valid and unique data-set names. To ensure uniqueness, always make the &UNIQ. variable part of the template value. The string AQT is a constant. For a description of the DB2 template variables, see the entry for AQT_ARCHIVE_COPY1 further up in this list.

Related information:

IBM DB2 Analytics Accelerator for z/OS: Stored Procedures Reference

Z/OS V1R11.0 XL C/C++ Programming Guide (SC09-4765-10)

E Specifying how long utilities wait for resources

IBM DB2 Analytics Accelerator for z/OS: User's Guide

Appendix D. Using the DB2 command line processor from UNIX System Services

To start a batch job, you can use the DB2 command line processor from UNIX System Services. In this case, the z/OS BPXBATCH utility is used to invoke the command line processor.

Before you begin

Make sure that the following conditions apply:

- The DB2 libraries for IBM DB2 Analytics Accelerator for z/OS have been installed.
- The IBM DB2 Analytics Accelerator for z/OS stored procedures have been installed.
- The DB2 command line processor has been installed.

Note: The DB2 command line processor is a Java application that requires IBM Data Server drivers for JDBC.

 The accelerator has been successfully added to the DB2 subsystem from IBM DB2 Analytics Accelerator Studio.

About this task

To create a suitable batch job with minimum effort, you can copy and customize one or more of the DB2 command-line scripts and UNIX System Services shell scripts that are included in the SAQTSAMP data set.

Procedure

 Create a separate DB2 command line processor properties-file for the user who will run the batch job. Make sure that the user uses this properties file, for example by adding the CLPPROPERTIESFILE environment variable to the .profile of this user. Remove read access to this file for everybody else. This way, you can safely add a connection alias for the accelerator network connection to the properties file, including the password. Use the following syntax to add this information to the properties file:

<connection_alias>=<db2host>:<db2port>/<db2location>,<uid>,<password>

where

<db2host>

Is the host name or IP address of the System z server on which DB2 runs

<db2port>

Is the port for network connections between the DB2 command-line client and the DB2 host. To identify the DB2 port, run the following DB2 command:

- <ssid> DIS DDF

where <ssid> is the DB2 subsystem ID.

<db2location>

Is the location of the DB2 subsystem that is supposed to interact with an accelerator.

<uid>

Is the ID of the user running the batch job

<password>

The password of the user running the batch job

- 2. Copy the following DB2 command-line scripts or UNIX System Services shell scripts and JCLs. The scripts are included in the SAQTSAMP data set, which is delivered with IBM DB2 Analytics Accelerator for z/OS.
 - AQTSCI02
 - AQTSJI02
 - AQTSSCHK
 - AQTSSCPY
 - Sample XML input strings depending on the stored procedures that you want to test. For example, AQTSXTCO is a sample XML input string for the SYSPROC.ACCEL_TEST_CONNECTION stored procedure. The SAQTSAMP data set contains such sample XML input strings for most IBM DB2 Analytics Accelerator for z/OS stored procedures.

For a description, follow the **Related reference** link at the end.

- 3. Customize the following scripts according to your needs:
 - AQTSCI02
 - AQTSJI02
 - AQTSCHK

Note: AQTSSCHK checks the number of successful stored procedure calls. If you change the number of calls in AQTSCI02, you must also change it in AQTSJI02 because AQTSJI02 passes this number to AQTSCHK as the number of calls to check for.

- Members containing sample XML input strings, such as AQTSXTCO.
- 4. Submit AQTSJI02. AQTSJI02 uses AQTSSCPY to copy AQTSCI02 and various XML files to the /tmp/ivp directory in the hierarchical file system, starts the command line processor with the copied files as input, and finally runs AQTSSCHK to verify the output.

Related tasks:

"Testing the stored procedures" on page 48 Run a few more tests to check whether the stored procedures behave as expected.

Related reference:

Appendix B, "Members of SAQTSAMP," on page 107

The SAQTSAMP data set contains various samples for installing, configuring or running the IBM DB2 Analytics Accelerator for z/OS stored procedures. The following table lists all these members and provides brief descriptions of their functions.

Appendix E. Disabling accelerators in a DB2 subsystem

It is recommended that you disable an accelerator before you activate an IBM DB2 Analytics Accelerator for z/OS software update or take actions to solve a problem.

Procedure

Use one of the following methods (DB2 command or IBM DB2 Analytics Accelerator Studio):

- In DB2 for z/OS, enter -stop ACCEL <name> where <name> is the name of the accelerator
- In IBM DB2 Analytics Accelerator Studio:
 - 1. Go to the Administration Explorer.
 - 2. Select the Accelerators folder.
 - 3. In the Object List Editor on the right, select the accelerator.
 - 4. Complete one of the following steps:
 - Click on the accelerator to disable the selected accelerator without canceling running queries. The accelerator status first changes to *Stopping*, then to *Stopped*. During the *Stopping* phase, running queries are completed.
 - Click the button on the right of the button and select Force from the menu to disable the selected accelerator and cancel all running queries. The status of the accelerator changes to *Stopped* immediately.

Note: Disabling does not make an accelerator unusable or remove it from the configuration. It just deactivates it. To re-enable an accelerator, click the

button on the toolbar.

Appendix F. Enabling accelerators in a DB2 subsystem

If you have disabled an accelerator before activating an accelerator update or solving a problem, (re-)enable the accelerator so that it becomes operational again.

Procedure

Use one of the following methods (DB2 command or IBM DB2 Analytics Accelerator Studio):

- In DB2 for z/OS, enter -start ACCEL <name> where <name> is the name of the accelerator.
- In IBM DB2 Analytics Accelerator Studio:
 - 1. Go to the Administration Explorer.
 - 2. Select the Accelerators folder.
 - 3. In the Object List Editor on the right, select the accelerator.
 - 4. Click **Start** on top of the Object List Editor. The status of the accelerator changes from *Stopped* to *Online*.

Glossary

This glossary includes terms and definitions related to the installation of IBM DB2 Analytics Accelerator for z/OS.

The following cross-references are used in this glossary:

- *See* refers you from a term to a preferred synonym, or from an acronym or abbreviation to the defined full form.
- See also refers you to a related or contrasting term.

To view glossaries for other IBM products, go to www.ibm.com/software/ globalization/terminology (opens in new window).

Α

access plan graph

A visual representation of a query that shows the database objects that are accessed by the query and the order in which this is done.

APF See authorized program facility (APF).

authorized program facility (APF)

In a z/OS environment, a facility that permits the identification of programs that are authorized to use restricted functions.

D

DDF See distributed data facility (DDF).

distributed data facility (DDF)

A set of DB2 for z/OS components through which DB2 for z/OS communicates with another relational database management system.

L

lock A means of preventing uncommitted changes made by one application process from being perceived by another application process and for preventing one application process from updating data that is being accessed by another process. A lock ensures the integrity of data by preventing concurrent users from accessing inconsistent data.

Μ

manifest

An XML file that contains metadata and procedural instructions for an IBM DB2 Analytics Accelerator for z/OS update package, such as the version of the package, update information (which versions are updated to which newer version), and parameters for internal package processing.

plan (also execution plan)

A Netezza file with an extension of pln that shows in which order the tables in a query are scanned.

port-forwarding

A networking mechanism that allows Secure Shell access to a host in a private network from the outside.

Ρ

Secure Shell (SSH)

A UNIX-based command interface and protocol for securely getting access to a remote computer.

S-FTP See SSH File Transfer Protocol.

SSH File Transfer Protocol

A network protocol that provides the ability to transfer files securely over any reliable data stream.

SSH tunnel

A secure and encrypted path through a network.

SSH See Secure Shell (SSH).

systems programmer

A programmer who plans, maintains, and controls the use of an operating system with the aim of improving the overall productivity of an installation.

table space

- A logical unit of storage in a database. In DB2 for z/OS, a table space is a page set and can contain one or more tables.
- In Netezza terminology, a logical collection of extents that are assigned to a table. A table space contains all the disk space that is allocated to a given table or table fragment and includes pages allocated to data and to indexes, pages that store TEXT or BYTE data in the dbspace, and bitmap pages that track page use within the extents.

throughput

A measure of the amount of information transmitted over a network in a given period of time. Throughput is generally measured in bits per second (bps), kilobits per second (Kbps), or megabits per second (Mbps).

V

Т

virtual IP address

An IP address that is shared among multiple domain names or multiple servers. Virtual IP addressing enables one IP address to be used either when insufficient IP addresses are available or as a means to balance traffic to multiple servers.

wall IP address

A Netezza term for the virtual IP address that is used to contact the Netezza system over the network. A floating IP address that is automatically assigned to the active Netezza host. Using the wall IP address from the z/OS data server, the active Netezza host can always be reached. This is transparent behavior in case of a failover. See also virtual IP address.

WLM application environment

A z/OS Workload Manager attribute that is associated with one or more procedures. The WLM application environment determines the address space in which a given procedure runs.

Index

A

accelerator 73, 79 accelerators disabling 117 enabling 119 updating software 84 access DB2 subsystem 42 rights AQTTIJSP 34 for IBM DB2 Analytics Accelerator for z/OS stored procedures 39 RACF 39 Access Server activation 67 update packages 66, 85 updates installing 85 activation of software updates 87 Access Server 67 replication engine 67 allocation 73, 79 application environment 27 AQTENV 111 AQTSCI02 48, 50 AQTSJI01 customization 38 running 38 AQTSJI02 48 customization 50 running 51 AQTSXTCO 50 AQTTIJSP access rights 34 running 35 assigning accelerator resources to 73, 79 auditing 9 authentication 45 authorizations 39 automatic update of table data 57

В

batch jobs 115 binding DB2 packages 40

С

CDC configuring 25, 61 installing 17 starting as z/OS started task 65 classification of workloads DDF 33 DSNUTILU 33 SYSPROC.ACCEL_LOAD_TABLES 33 z/OS 34 CLPPROPERTIESFILE 115 command line processor 115 connection alias 48, 115 profile 42 testing 44 credentials 45 customization AQTSCI02 50 AQTSJI01 38 AQTTIJSP 35

D

damaged disk 9 data loss 57 data sharing group different versions of IBM DB2 Analytics Accelerator for z/OS 93 database connection profile 42 DB2 authorizations for stored procedures 39 DB2 for z/OS access from IBM DB2 Analytics Accelerator Studio 40 access from IBM Optim Query Tuner 40 command-line processor 115 connection to IBM DB2 Analytics Accelerator for z/OS 39 disabling accelerators 117 enabling accelerators 119 DB2 subsystem 73, 79 DB2 Unload Utility 33 DB2-supplied stored procedures setup verification 31 DDF workload classification 33 disabling accelerators in DB2 for z/OS 117 double-byte EBCDIC (graphic) 18, 22 drain 73, 79 DSNTIJAS 18 DSNTIJUZ 22 DSNUTILU performance goals 33

Ε

```
EBCDIC
GRAPHIC 18, 22
MIXED 18, 22
enabling
accelerators 119
existing DB2 subsystem 17
INSERT FROM SELECT 18, 22
mixed and double-byte EBCDIC 18, 22
Optim Query Tuner 40
environment variables 111
AQT_HOST_PACKAGE_DIRECTORY 89
AQT_INSTALL_PREFIX 66, 85
```

G

GB18030 characters 57 glossary 121

Η

HPSS partitions after migration 5

IBM DB2 Analytics Accelerator for z/OS connection testing 44 connection to DB2 for z/OS 39 Console 45 migrating 81 update packages 85 updates installing 81, 85 libraries 91 stored procedures 91 WLM performance goals for stored procedures 32 IBM DB2 Analytics Accelerator Studio installation 13, 14 language settings 83 perspective in IBM Data Studio 83 workspace directory 83 IBM InfoSphere Change Data Capture for z/OS 25 configuring 61 installing 17 starting as z/OS started task 65 IBM Installation Manager IBM Data Studio 13, 14 IBM DB2 Analytics Accelerator Studio 13, 14 incremental updates 57 configuring 25 log cache 64 log cache size 64 INSERT FROM SELECT 18, 22 installation DSNTIJAS 18 from Installation Manager 13, 14 IBM DB2 Analytics Accelerator database 18 IBM DB2 Analytics Accelerator Studio 13, 14 testing 53 updates 81 verification 53

J

JCLs 107

L

language settings, IBM DB2 Analytics Accelerator Studio 83 location name 47 log cache size, incremental updates 64 loss of data 57

Μ

migrating from version 2.1.x to 3.10 1 from version 3.1.0 to 4.1.0 1 HPSS partitions 4 moved partitions 4 migration 81 mixed EBCDIC 18, 22 moved partitions after migration 5

Ν

Netezza Firmware (FDT) transferring updates 89 updating 89 Netezza Host Platform (HPF) transferring updates 89 updating 89 Netezza Performance Server (NPS) 85 update packages 85 updates installing 85 Netezza priorities 75 NPS *See* Netezza server software

0

Optim Query Tuner 40

Ρ

pairing code 45 performance goals DSNUTILU 33 SYSPROC.ACCEL_LOAD_TABLES 33 WLM 32 perspective, IBM DB2 Analytics Accelerator Studio 83 power user 39 prerequisites DB2 data server 17 IBM DB2 Analytics Accelerator Studio 13, 14 priority genstats 76 groom 76 Netezza 75 of queries 75 table loading 76 WLM 75 private data network (PDN) 9

Q

query acceleration testingdefining tables 53loading and enabling tables 53running a test query 53

R

RACF access for IBM DB2 Analytics Accelerator for z/OS stored procedures 39 replication engine activation 67 update packages 66, 85 updates installing 85 requirements for DB2 data server 17 resource 73, 79 resource allocation 73, 79 running AQTSJI01 38 AQTTIJSP 35

S

sample job members 107 SAQTSAMP 107 schema filters 42 security 9 sensitive information 9 service password 9 setting ZPARMs 22 ZPARMs, DB2 11 for z/OS 18 shared accelerator 73, 79 sharing credentials 45 size of log cache, incremental updates 64 SMP/E 27, 66, 85, 91 software prerequisites for DB2 data server 17 required PTFs DB2 10 for z/OS 82 DB2 11 for z/OS 82 version activation 67, 87 switching 67, 87 starting CDC 65 IBM InfoSphere Change Data Capture for z/OS 65 stored procedures 39 clean up SYSACCEL tables after testing accelerator 51 DB2-supplied 31 installation verification 36 SAQTSAMP(AQTSJI00) 37 SAQTSAMP(AQTSJI01) 38 testing 48 AQTSJI02 50 AQTSXTCO 50 switching between software versions 87 Access Server 67 replication engine 67 SYSPROC.ACCEL_LOAD_TABLES 33

T

TCP/IP configuration 7 sample configuration 105 testing query acceleration 53 stored procedures 48 tracing sensitive information 9 transferring update Access Server 66, 85 IBM DB2 Analytics Accelerator for z/OS software 85 Netezza Firmware (FDT) 89 Netezza Host Platform (HPF) 89 replication engine 66, 85 troubleshooting 95

U

update different versions of IBM DB2 Analytics Accelerator for z/OS 93 required PTFs DB2 10 for z/OS 82 DB2 11 for z/OS 82 updating Access Server 84 credentials 45 from version 2.1.x to 3.10 1 from version 3.1.0 to 4.1.0 1 IBM DB2 Analytics Accelerator software 84 IBM DB2 Analytics Accelerator Studio from IBM Installation Manager 82 IBM DB2 Analytics Accelerator Studio from IBM Installation Manager (internet connection) 83 IBM DB2 Analytics Accelerator Studio from IBM Installation Manager (no internet connection) 82 Netezza Firmware (FDT) 89 Netezza Host Platform (HPF) 89 Netezza Performance Server (NPS) 84 replication engine 84 updating table data continuously 57 user IDs 42

W

wall IP address 7 WLM application environment 27, 111 importance level 75 performance goals for IBM DB2 Analytics Accelerator for z/OS stored procedures 32 for SYSPROC.ACCEL_LOAD_TABLES 33 workload classification DDF 33 DSNUTILU 33 SYSPROC.ACCEL_LOAD_TABLES 33 z/OS 34 Workload Manager See WLM workspace directory, IBM DB2 Analytics Accelerator Studio 83

Ζ

ZPARMs 18, 22

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing IBM Corporation North Castle Drive Armonk, NY 10504-1785 U.S.A.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

Intellectual Property Licensing Legal and Intellectual Property Law IBM Japan Ltd. 19-21, Nihonbashi-Hakozakicho, Chuo-ku Tokyo 103-8510, Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk. IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Deutschland GmbH Dept. M358 IBM-Allee 1 71139 Ehningen Germany

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

All IBM prices shown are IBM's suggested retail prices, are current and are subject to change without notice. Dealer prices may vary.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating
platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

Each copy or any portion of these sample programs or any derivative work, must include a copyright notice as follows:

© (your company name) (year). Portions of this code are derived from IBM Corp. Sample Programs.

© Copyright IBM Corp._enter the year or years_.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

Trademarks

IBM, the IBM logo, and ibm.com[®] are trademarks of International Business Machines Corporation, registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at "Copyright and trademark information" at http://www.ibm.com/legal/copytrade.shtml.

Intel, Intel Iogo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.



Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft and Windows are trademarks of Microsoft Corporation in the United States, other countries, or both.

Netezza is a registered trademark of Netezza Corporation, an IBM Company.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Readers' Comments — We'd Like to Hear from You

IBM DB2 Analytics Accelerator for z/OS Installation Guide Version 4.1.0

Publication No. SH12-7038-01

We appreciate your comments about this publication. Please comment on specific errors or omissions, accuracy, organization, subject matter, or completeness of this book. The comments you send should pertain to only the information in this manual or product and the way in which the information is presented.

For technical questions and information about products and prices, please contact your IBM branch office, your IBM business partner, or your authorized remarketer.

When you send comments to IBM, you grant IBM a nonexclusive right to use or distribute your comments in any way it believes appropriate without incurring any obligation to you. IBM or any other organizations will only use the personal information that you supply to contact you about the issues that you state on this form.

Comments:

Thank you for your support.

Submit your comments using one of these channels:

- Send your comments to the address on the reverse side of this form.
- · Send your comments via email to: swsdid@de.ibm.com

If you would like a response from IBM, please fill in the following information:

Name

Address

Company or Organization

Phone No.

Email address



Cut or Fold Along Line





Product Number: 5697-DAB

SH12-7038-01

