IBM Db2 Analytics Accelerator for z/OS 7.5.6

Installation Guide



Note Before you use this information and the product it supports, read the information in "Notices" on page 263.

Seventh Edition, July 2021

This edition applies to version 7.5.6 of IBM® Db2® Analytics Accelerator for z/OS® (product number 5697-DA7), and to all subsequent releases and modifications until otherwise indicated in new editions. This edition replaces SH12-7102-05. Changes to this edition are marked with a vertical bar.

© Copyright International Business Machines Corporation 2010, 2021.

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

Contents

Figures	vii
About this book	xi
What's new in version 7.5?	xiii
Chapter 1. Introduction and installation planning	
Solution overview	
Multiple database subsystems and multiple accelerators	
Product components	
Related documentation and product support	
Installation prerequisitesInstallation task flow and preparation tasks	
Installation of the IBM Integrated Analytics System	
Post-installation tasks	
1 OST IIIStatiation tasks	10
Chapter 2. Configuring TCP/IP connections	11
Chapter 3. Security considerations	13
Chapter 4. Data encryption	17
Chapter 5. Administration clients	19
Installing IBM Db2 Analytics Accelerator Studio	
Installing IBM Data Server Manager	21
Chapter 6. Enabling an existing Db2 subsystem for IBM Db2 Analytics	
Accelerator for z/OS	25
Software prerequisites for the Db2 for z/OS data server	25
Installing libraries with IBM Db2 Analytics Accelerator support	
Creating the IBM Db2 Analytics Accelerator database	
Setting ZPARMs for IBM Db2 Analytics Accelerator for z/OS	
Setting ZPARMs for IBM Db2 Analytics Accelerator in Db2 for z/OS	
Using a sample job to set ZPARMs for IBM Db2 Analytics Accelerator for z/OS	
Configuring the incremental update function	
Chapter 7. Setting up IBM Db2 Analytics Accelerator for z/OS	37
Setting up a WLM application environment for IBM Db2 Analytics Accelerator	
Setting up a WLM application environment for different product versions	
Verifying the correct setup of Db2-supplied stored procedures	
Defining WLM performance goals for IBM Db2 Analytics Accelerator for z/OS stored procedure	s45
Adjusting WLM performance goals for SYSPROC.ACCEL_LOAD_TABLES	
Setting access rights for the user who runs AQTTIJSP	
Customizing and running AQTTIJSP	
Verifying the installation of IBM Db2 Analytics Accelerator for z/OS stored procedures	50
Chapter 8. Db2 Analytics Accelerator on Z	53
Installing Db2 Analytics Accelerator on Z	

Defining an LPAR for Db2 Analytics Accelerator on Z	
Logging on to the Admin UI	
Installing and starting the appliance (single-node setup)	
Installing and starting the appliance (multi-node setup)	
Installing Db2 Analytics Accelerator on Z with GDPS failover support	
Defining an LPAR for Db2 Analytics Accelerator on Z with GDPS failover support	
Installing and starting the appliance with GDPS failover support (single-node setup)	
Shutting down and restarting a cluster (multi-node setup)	
Updating an existing configuration	
Creating a trace file (dump)	
Adding and managing users	
Checking your Db2 Analytics Accelerator on Z setup	
Replacing the SSL certificate	118
Chapter 9. Connecting IBM Db2 Analytics Accelerator for z/OS and Db2	123
Access rights	123
Access rights for power users	123
Minimum access rights	
Binding Db2 packages and granting user privileges	
Creating EXPLAIN tables	
Creating a database connection profile	
Testing the connection from your administration client	
Adding accelerators	
Obtaining the pairing code for authentication	
Completing the authentication using the Add Accelerator wizard	134
Chapter 10. Testing query acceleration	137
Defining accelerator-shadow tables	137
Loading and enabling tables	138
Running a test query	139
Chapter 11. Incremental updates	143
IBM Integrated Synchronization	
Defining a secure network port	
AT-TLS configuration	
Assigning user IDs	
Adding permissions for schema change support	
Enabling IBM Integrated Synchronization for a Db2 subsystem	
Suggestion for a high-availability setup	
Status information for error analyses	
Changing the Db2 user ID and password	
Changing the IP address and the DDF secure port	
Managing certificates	
REORG, REORG DISCARD and LOAD REPLACE behavior	
IBM InfoSphere Change Data Capture for z/OS	
Concepts and architecture	
Integration	
Continuous incremental updates	
Parallel propagation of changes	
Faulty tables processing	
Completing the installation and configuring CDC	10U 1Ω∩
Starting CDCEnabling incremental updates for a Db2 subsystem	18/
e e e e e e e e e e e e e e e e e e e	
Ctarting or stanning ingramantal undates	187
Starting or stopping incremental updates	187 190
Restarting incremental update processes	187 190 192
	187 190 192 194

Configuring the expiration behavior of WAITFORDATA queries	.197
Chapter 12. Beyond the basics	19 9
Logging on to the IBM Db2 Analytics Accelerator Console	
Using telnet to log on to the IBM Db2 Analytics Accelerator Console	
Using ssh to log on to the IBM Db2 Analytics Accelerator Console	
Setting resource limits for Db2 subsystems	
Mapping of z/OS importance levels to priorities of the accelerator database engine	
Using the call-home function	
Configuring the call-home function on the HMC	
Configuring the call-home function on the accelerator	
Enabling the call-home function	
Disabling the call-home function	
Displaying call-home settings	
Testing connectivity	
Enabling queries against system-temporal and bi-temporal tables	
Chapter 13. Installing updates	215
Installing prerequisite PTFs for Db2 for z/OS	
Updating IBM Db2 Analytics Accelerator software	
Installing update packages for the accelerator	
Transferring update packages for the accelerator	
Activating an updated accelerator software version	
Updating administration clients	
Updating IBM Db2 Analytics Accelerator Studio	
Updating IBM Data Server Manager	
Updating the IBM Db2 Analytics Accelerator stored procedures	
Updating Db2 Analytics Accelerator on Z	
Chapter 14. Removing accelerators from your Db2 for z/OS configuration	221
Cleaning up the communications database manually	
Removing the remaining database objects	
Kemoving the remaining database objects	.231
Chapter 15. Sample TCP/IP configuration	23 3
Chapter 16. Troubleshooting	2 35
Query acceleration fails after schema change	
Firewall blocks operations because crucial ports are closed	
Message DSNUTILU NOT INVOKED APF AUTHORIZED	
No CEEDUMPs	
PRIQ value too high when creating table spaces.	
The Db2 command -DIS ACCEL does not work	
Connection authorization failure (error -4214)	
Errors during ZPARM compilation	
DRDA connection does not work	
Removing orphaned system-table entries and catalog-table entries	
Package not found when running a stored procedure from IBM Db2 Analytics Accelerator Studio	
SQL code -430 from IBM Db2 Analytics Accelerator for z/OS stored procedures SQL code -471 referring to a function in the DSNAQT schema	
SYSPROC.ACCEL_LOAD_TABLES returns SQL error -471 and reason code E790002 for DSNUTILU	
Load of partitioned tables freezes during unload phase	
Chapter 17. Opening a support case	
Transferring maintenance updates	
Installing maintenance updates	.247

Appendix A. Members of SAQTSAMP	251
Appendix B. Environment variables	255
Appendix C. Disabling accelerators in a Db2 subsystem	259
Appendix D. Enabling accelerators in a Db2 subsystem	261
Notices	263
Trademarks	264
Terms and conditions for product documentation	
GDPR considerations	267
GDPR	
Why is GDPR important?	
Product Configuration for GDPR.	
How to configure the offering so that it can be used in a GDPR environment	
Data Life Cycle	
What types of data?	
Where in the process?	
For what purpose?	
Personal data used for online contact with IBM	
Data Collection	
Data Storage	
Data Access	
Data Processing	
Data Deletion	
Data Monitoring	
Responding to Data Subject Rights	
Does the offering facilitate being able to meet data subject rights?	
Glossary	273
A	
D	
L	
P	
S	
T	
V	
W	

Figures

1. Solution overview of Db2 Analytics Accelerator on Z	2
2. Solution overview of Db2 Analytics Accelerator on the IBM Integrated Analytics System	3
3. Possible connections	5
4. Installation task-flow diagram	8
5. IBM Db2 Analytics Accelerator setup	19
6. The Sign In page of IBM Data Server Manager	23
7. Assigning a medium-to-high priority service class to the SYSPROC.ACCEL_LOAD_TABLES stored procedure in the WLM ISPF application	47
8. Architecture	53
9. Failover support for a single-nodeDb2 Analytics Accelerator on Z with GDPS: schematic overview	54
10. The Partitions window of the HMC	55
11. The Customize Image Profiles window of the HMC	56
12. SSC network adapter details on the HMC	57
13. Confirmation: Operating System Messages	58
14. Browser certificate warning	59
15. The Install Software Appliance window	60
16. The Welcome page of the Admin UI	60
17. Accelerator Configuration Definition page	61
18. Accelerator Configuration Definition page	62
19. Extract of the sample configuration file (JSON format), displayed in web browser	63
20. Error message after uploading a faulty configuration file	75
21. Accelerator Configuration Definition after a successful configuration file upload	75
22. Message window showing the progress of your Db2 Analytics Accelerator on Z configuration	76

23. The Accelerator Components Health Status page is displayed after a successful configuration	76
24. Accelerator Configuration Definition page	78
25. Error message after uploading a faulty configuration file	90
26. Accelerator Configuration Definition after a successful configuration file upload	91
27. The Multiple Node Deployment Credentials window	91
28. Message window showing the progress of your Db2 Analytics Accelerator on Z configuration	92
29. The Accelerator Components Health Status page is displayed after a successful configuration	92
30. The Partitions window of the HMC	93
31. Confirmation: Operating System Messages	94
32. Accelerator Configuration Definition page	96
33. Error message after uploading a faulty configuration file	107
34. Accelerator Configuration Definition after a successful configuration file upload	107
35. Message window showing the progress of your Db2 Analytics Accelerator on Z configuration	108
36. The Accelerator Components Health Status page is displayed after a successful configuration	108
37. Shutdown from the Accelerator Components Health Status page	109
38. The Accelerator Components Health Status page	110
39. Confirmation window	110
40. The Log window of the Db2 Analytics Accelerator on Z user interface	111
41. The Dumps page	111
42. The Create Dump page	112
43. Successful user creation	113
44. List of requests	114
45. The Confirm New Account window	115
46. The Test your Network tab on the Accelerator Diagnosis page	116
47. The Db2 z/OS conn. tab on the Accelerator Diagnosis page	117

48. The System activity tab on the Accelerator Diagnosis page	117
49. The About window of Db2 Analytics Accelerator on Z user interface	119
50. The Custom SSL Certificate for Web Admin User Interface	120
51. Verifying the correct installation of the certificate	121
52. Your guide through the certificate jungle	149
53. Network connections separated by port in a DDVIPA network	159
54. Interaction diagram of the components that are involved in the incremental update process	178
55. The Configure IDAA Callhome page on the HMC	206
56. The Domain Security page on the HMC	207
57. The Installation Log viewer	222
58. The Accelerator page of the Appliance Installer	228
59. The Update page of the Appliance Installer	229
60. The Welcome page of the Secure Service Container Installer	230
61. Recommended connections between IBM Z server and IBM Integrated Analytics System (including switches)	233
62 SOL code -//30 massage window	2/11

About this book

This book provides information about the following subjects:

- Installation of IBM Db2 Analytics Accelerator support on a z/OS data server
- Maintenance of accelerators (software update, starting and stopping the container)
- Installation of Db2 for z/OS Program Temporary Fixes (PTF), which include the Db2 software that is required to integrate this product into your Db2 environment
- Configuration of the product from Db2 for z/OS (invocation of stored procedures)
- Installation of the client software (IBM Db2 Analytics Accelerator Studio)
- · Component updates

Audience

This book is intended for the following audiences:

- IBM Customer Service Centers (support personnel) installing IBM Db2 Analytics Accelerator for z/OS at the client site
- Administrators who start and stop accelerators
- 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 7.5?

IBM Db2 Analytics Accelerator for z/OS Version 7.5 is a major upgrade that provides powerful new features:

- IBM Integrated Synchronization for incremental update processing
- Multi-node installations of Db2 Analytics Accelerator on Z with increased processing power for massive workloads

Enhancements in this version

For enhancements in this or previous follow-up versions, see the What's new document.

Chapter 1. Introduction and installation planning

Read a brief product description. See which components the product consists of, what needs to be considered before the installation, which preparations are necessary, and which steps must be completed by whom.

Solution overview

IBM Db2 Analytics Accelerator for z/OS V7.5 is a bundled solution package for the acceleration of database queries. It supports two deployment options: A mainframe-only solution where the accelerator is installed in a dedicated logical partition (LPAR). This is called IBM Db2 Analytics Accelerator for z/OS deployed on IBM Z. The short name is: Db2 Analytics Accelerator on Z. The second deployment option is a solution that runs on optimized appliance hardware. It is called IBM Db2 Analytics Accelerator for z/OS Version 7.5 deployed on the IBM Integrated Analytics System (IIAS). The short name for this option is: Analytics Accelerator on Integrated Analytics System.

Both deployment options can reduce query response times dramatically.

- The mainframe-only solution (Db2 Analytics Accelerator on Z) marks the beginning of a new era: A fully integrated solution that does not require external hardware.
- IBM Db2 Analytics Accelerator on IBM Integrated Analytics System is the successor of IBM Db2 Analytics Accelerator for z/OS Version 5.1.0 (which runs on the IBM PureData® System for Analytics). The new solution runs on newly designed IBM POWER8® hardware.

Both solutions can be run individually or in combination. See <u>Figure 1 on page 2</u> and <u>Figure 2 on page 3</u> for an overview of the components involved in the setup. The overviews show examples; other, more sophisticated configurations are thinkable.

Db2 Analytics Accelerator on Z

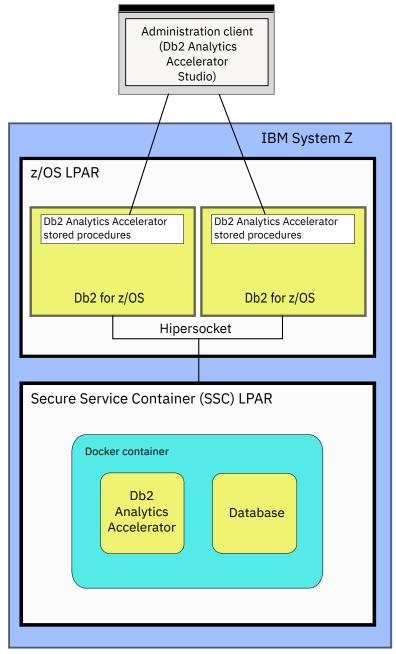


Figure 1. Solution overview of Db2 Analytics Accelerator on Z

Db2 Analytics Accelerator on IBM Integrated Analytics System

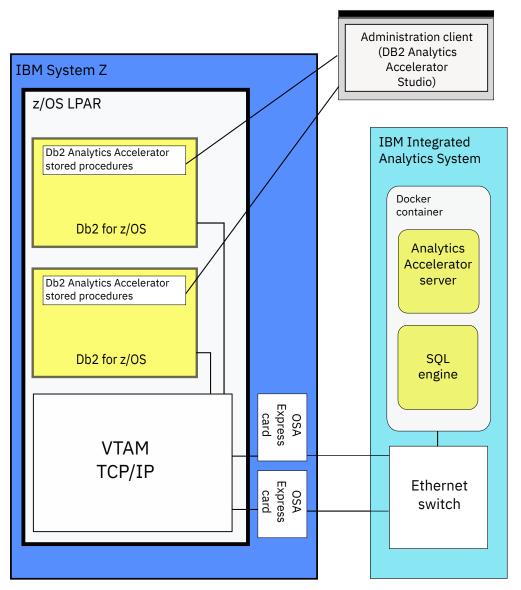


Figure 2. Solution overview of Db2 Analytics Accelerator on the IBM Integrated Analytics System

Multiple database subsystems and multiple accelerators

A single accelerator can be shared by multiple Db2 for z/OS subsystems. However, a single Db2 for z/OS subsystem can also be connected to more than one accelerator.

IBM Db2 Analytics Accelerator for z/OS supports the following subsystem configurations:

- Multiple subsystems, each of which in a separate logical partition (LPAR)
- Multiple subsystems in a common LPAR
- Multiple subsystems that make up a data sharing group (subsystems in different LPARs, on different Central Processing Complexes (CPCs))

Figure 3 on page 5 shows that Db2 subsystems can share a single accelerator as well as connect to more than just one accelerator. The leftmost box in the figure, which represents a single subsystem in a separate LPAR, is connected to two accelerators. All Db2 subsystems (including the one in the leftmost box) share the one accelerator on the left.

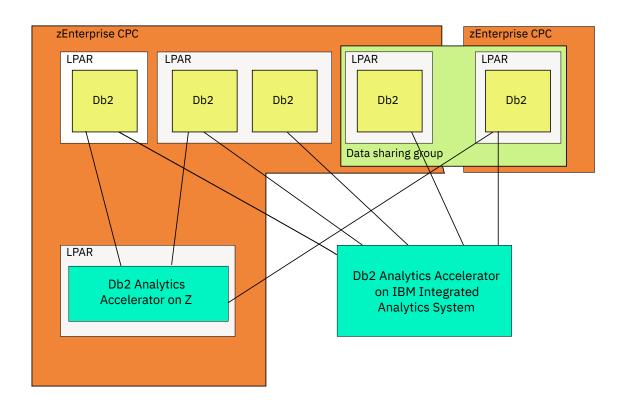


Figure 3. Possible connections

Product components

IBM Db2 Analytics Accelerator for z/OS Version 7.5 consists of the following components:

- For IBM Db2 Analytics Accelerator for z/OS on IIAS: a software package that is preinstalled on the IBM Integrated Analytics System by a service engineer
- For Db2 Analytics Accelerator on Z: a software image for the deployment on IBM Z by the customer
- IBM Db2 Analytics Accelerator stored procedures, which run on the attached IBM Z server. The stored procedures are provided on tape and are installed via SMP/E (details are described below)
- The IBM Db2 Analytics Accelerator for z/OS administration client (GUI), which runs on a workstation attached to the System z°. The following administration clients are available:

IBM Db2 Analytics Accelerator Studio

IBM Db2 Analytics Accelerator Studio is an IBM Data Studio plugin and has to be downloaded from the following page:

DB2 Analytics Accelerator Studio - Marketing Registration Services (MRS)

If you already have IBM Data Studio, install just the plugin. You get the plugin from IBM Fix Central. See:

- Prerequisites and Maintenance for IBM DB2® Analytics Accelerator for z/OS, V7
- IBM DB2 Analytics Accelerator Studio (All releases, All platforms)

IBM Data Server Manager

IBM Data Server Manager is a web application to be installed on a web server. It can be downloaded from the following page:

Data Server Manager vx.x.x.x - Marketing Registration Services (MRS)

Both administration clients are free of charge.

- IBM InfoSphere® Change Data Capture for z/OS Access Server for the incremental update function of the product on the IBM Integrated Analytics System.
- IBM InfoSphere Change Data Capture for z/OS Replication Engine for the incremental update function the product on the IBM Integrated Analytics System.

Product purchase and delivery

The product can be ordered from <u>Shopz</u>. The entire solution is available for download. The stored procedures can also be delivered on tape. The tape contains:

• IBM Db2 Analytics Accelerator for z/OS, Version 7.5.0 FMID HAQT710 (containing the stored procedures)

Related documentation and product support

You find documentation and product support in an information center, on the IBM Db2 Analytics Accelerator for z/OS support website, and also on IBM Resource Link®.

IBM Knowledge Center

You find IBM Db2 Analytics Accelerator for z/OS in the IBM Knowledge Center at:

https://www.ibm.com/support/knowledgecenter/SS4LQ8_7.5.0/com.ibm.datatools.aqt.doc/idaa_kc_welcome.html

The IBM Knowledge Center is fully accessible, meaning that it supports tools and aids for users with disabilities.

IBM Db2 Analytics Accelerator for z/OS support website

Check the IBM Db2 Analytics Accelerator for z/OS support website regularly for the latest information and updates. This site also provides links to PDF product manuals:

https://www.ibm.com/support/home/product/Y871598M07790F41/DB2_Analytics_Accelerator_for_z/OS

Support cases

If you encounter problems, you can report these problems here:

Service requests and PMRs

IBM Z documentation

Read the appropriate manual for your IBM zEnterprise® server:

- IBM z15[™] Technical Guide, SG24-8852
- IBM z14® Technical Guide, SG24-8451
- IBM z13[®] Technical Guide, SG24-8251
- IBM zEnterprise 196: System Overview, SA22-1086
- IBM zEnterprise 114: System Overview, SA22-1087
- IBM zEnterprise EC12: System Overview, SA22-1088

Installation prerequisites

Check the installation prerequisites before you install IBM Db2 Analytics Accelerator for z/OS:

 $\underline{\text{https://www.ibm.com/support/pages/prerequisites-and-maintenance-ibm-db2-analytics-accelerator-zos-v7}$

Installation task flow and preparation tasks

Complete the following tasks before IBM service personnel visits your site to install the IBM Integrated Analytics System:

About this task

The chart in this section shows the individual installation tasks that you must complete for a new installation of IBM Db2 Analytics Accelerator for z/OS. Each task is represented by a box and is assigned to the most appropriate user role. The sequence in which you must complete the tasks is indicated by the arrows and the task numbering (from A1 to D2), where capital letters are used to mark tasks that need to be carried out by a particular role. The time needed for a task is indicated on top of each box. At the bottom of a box, you find a reference to the document that tells you how to complete the task.

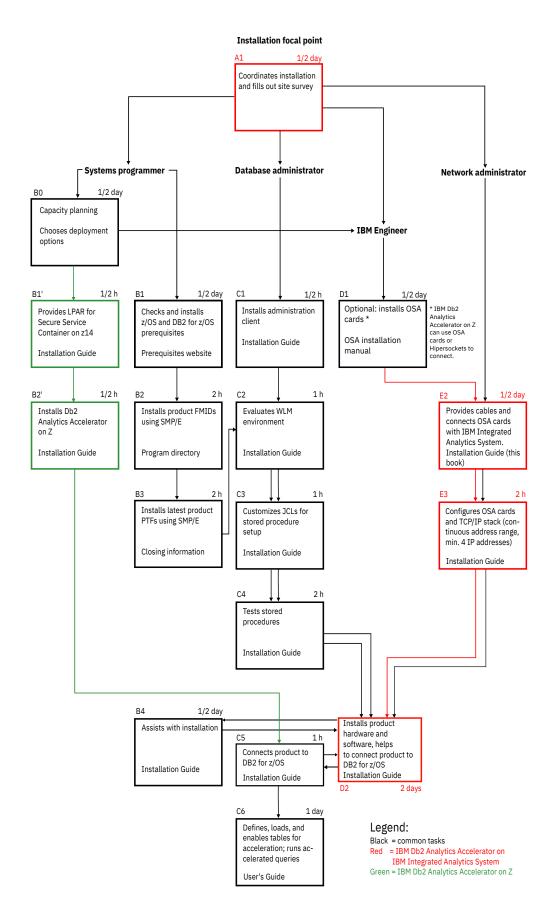


Figure 4. Installation task-flow diagram

For an update of an existing installation, refer to *Installing updates* in the *IBM Db2 Analytics Accelerator* for z/OS: Installation Guide.

Procedure

- 1. A1: The installation focal point assigns the installation tasks to the members of the IT team and fills out the site survey.
- 2. *B0*: The systems programmer does the capacity planning for IBM Z. She or he also decides which options to install: IBM Db2 Analytics Accelerator on an IBM Integrated Analytics System, Db2 Analytics Accelerator on Z, or both.
 - If the product is to be installed on an IBM Integrated Analytics System, the systems programmer gets in touch with the IBM Engineer.
- 3. *B1*: The systems programmer must check the prerequisites and ensure that all requirements are met. See "Installation prerequisites" on page 7 for more information.
- 4. C1: The database administrator installs the administration client.
 - This needs to be done up-front so that license agreements can be accepted for all products involved. IBM Db2 Analytics Accelerator Studio is the administration software of IBM Db2 Analytics Accelerator for z/OS. It is installed on a Linux or Microsoft Windows workstation that is connected to Db2 for z/OS. See: *Installing IBM Db2 Analytics Accelerator Studio* in the <u>IBM Db2 Analytics</u> Accelerator for z/OS: Installation Guide, SH12-7083.
- 5. *B1 cont. B3*: The systems programmer needs to complete the following steps so that IBM Db2 Analytics Accelerator for z/OS can be used by a Db2 subsystem on your Db2 for z/OS data server:
 - a) Apply the required z/OS and Db2 Program Temporary Fixes (PTFs).
 - b) Install the IBM Db2 Analytics Accelerator for z/OS FMIDs using SMP/E.
 - c) Install the latest PTFs for the product using SMP/E.
 PTFs might be delivered with a major product release or separately. For information on how to install the PTFs and how to handle the various components in the package, see the closing information for the PTF, which is published as a support document on the web.
- 6. *B1'* and *B2'*: If Db2 Analytics Accelerator on Z was selected as a deployment option, the systems programmer provides an LPAR for an IBM Secure Services Container on the IBM Z. In a subsequent step, the systems programmer installs Db2 Analytics Accelerator on Z.
- 7. C2-C4: The database administrator completes the following steps:
 - a) The database administrator evaluates, and, if necessary, customizes the Workload Manager (WLM) application environment.
 - b) The database administrator customizes and runs JCLs for the setup of IBM Db2 Analytics Accelerator for z/OS stored procedures.
 - For instructions, see Enabling an existing Db2 subsystem for IBM Db2 Analytics Accelerator for z/OS in the IBM Db2 Analytics Accelerator for z/OS: Installation Guide.
- 8. D1: The Central Processing Complexes (CPCs) that are supposed to interact with the IBM Integrated Analytics System must be equipped with one or two OSA-Express® cards each. Call IBM if the cards have not been ordered or delivered yet.
 - For information on how to choose the proper cards, see:
 - Network requirements for System z
- 9. *E2*: The network administrator needs to provide cables for the connections between the OSA-Express cards and the IBM Integrated Analytics System.
 - The cables must be ready to be plugged in by IBM service personnel. Therefore, the network administrator needs to run and label the cables properly.
 - **Data sharing groups:** All Db2 subsystems in the same Central Processing Complex (CPC) share the network connectivity between that CPC and the accelerator. It does not matter if these Db2 subsystems are independent, belong to the same data sharing group, or belong to different data sharing groups. Each CPC, however, must be wired individually to the accelerator.

10. E3: Although a connection to the IBM Integrated Analytics System does not exist yet, the network administrator should configure the OSA cards and the TCP/IP stack for IBM Db2 Analytics Accelerator for z/OS beforehand.

Information is found in Chapter 1: Configuring TCP/IP for connections between IBM Z and the IBM Integrated Analytics System of the IBM Db2 Analytics Accelerator for z/OS: Installation Guide.

Installation of the IBM Integrated Analytics System

D2: These tasks are completed by IBM service personnel or one of IBM's business partners, with the assistance of the systems programmer.

Procedure

- 1. The service engineer installs the hardware (IBM Integrated Analytics System).
- 2. The service engineer installs the product software on the IBM Integrated Analytics System.
- 3. The service engineer plugs in the cables to connect the IBM Integrated Analytics System with the IBM Z server.

Post-installation tasks

For new installations, the database administrator must complete a number of post-installation tasks. The systems programmer might have to assist.

Procedure

1. *C5*: The database administrator associates the product with Db2 for z/OS to authenticate IBM Db2 Analytics Accelerator for z/OS as an entitled extension of Db2 for z/OS.

If Db2 Analytics Accelerator on Z was installed, this step is carried out by the systems programmer. The product association is based on a proper authentication. The <u>IBM Db2 Analytics Accelerator for</u> z/OS: Installation Guide, SH12-7083, contains two sections that describe how to complete this task:

- Obtaining the pairing code for accelerator authentication
- Creating the authentication token using the Add New Accelerator wizard
- 2. C6: The database administrator tests the entire setup.

To do so, the database administrator needs to complete the steps in *Testing query acceleration* of the *IBM Db2 Analytics Accelerator for z/OS: Installation Guide*.

Results

IBM Db2 Analytics Accelerator for z/OS is ready for use.

What to do next

- The version of IBM Db2 Analytics Accelerator for z/OS that the service engineer has installed on the IBM Integrated Analytics System might not be the version that you want to use in production. If this is the case, first update the software. Refer to *Updating accelerator software* in the *IBM Db2 Analytics Accelerator for z/OS: Installation Guide*.
- The database administrator can start to define tables on the accelerator in order to accelerate queries.
 More detailed information can be found in the help topics or in the IBM Db2 Analytics Accelerator for z/OS: User's Guide.
- You might want to use the incremental update function so that tables on the accelerator are updated continuously. For more information, see Chapter 11, "Incremental updates," on page 143.

Chapter 2. Configuring TCP/IP connections

To transfer data between your database management system and an accelerator, you must provide several IP addresses and establish a TCP/IP connection between your IBM Z server and the IBM Integrated Analytics System or Db2 Analytics Accelerator on Z.

IBM Db2 Analytics Accelerator on IBM Integrated Analytics System

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

A private network between the IBM Z server and your IBM Integrated Analytics System is recommended, but not necessary. The number of required TCP/IP addresses depends on the configuration:

- One virtual IP address for the IBM Integrated Analytics System.
- One IP address for each processing node of the IBM Integrated Analytics System (this number varies in accordance with the model).
- One IP address for the OSA-Express card facing the IBM Integrated Analytics System.

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

Restriction: Do not use the address range of the 172.17.0.0/16 network because this network is used internally by the product.

Db2 Analytics Accelerator on Z

Db2 Analytics Accelerator on Z requires a TCP/IP network to connect Db2 for z/OS subsystems with the accelerator. This network connection can be defined through OSA Express interfaces or through Hipersockets (if the Db2 subsystem runs in an LPAR on the same Central Processing Complex).

Restriction: Do not use the address range of the 172.17.0.0/16 network because this network is used internally by the product.

Network bandwidth and speed impacts the overall solution performance. For OSA interfaces, a dedicated OSA Express card and 10 GbE connectivity with Jumbo Frames are recommended.

For Hipersockets, use an MFS size of 24 KB and an MTU size of 16 KB, for example:

CHPARM=4x (MFS 24KB, MTU 16KB)

Related reference

Sample TCP/IP configuration

In this section, you find a sample VTAM definition and a sample profile for the TCP/IP setup on your IBM Z. These samples work for both, IBM Db2 Analytics Accelerator on IBM Integrated Analytics System and Db2 Analytics Accelerator on Z.

Chapter 3. Security considerations

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 IBM Z server and the IBM Integrated Analytics System (accelerator) are connected over a private data network (PDN), which does not allow access from the internet or your organization's intranet.
- Db2 Analytics Accelerator on Z can use Hipersockets to connect to Db2 subsystems. This is a very secure option because the data never leaves the mainframe.
- 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.
- 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.
- 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 does not allow a direct remote connection. 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.

The facility to request a password is not publicly available on the internet. For your first support case requiring access to the machine, you have to submit an encrypted token to IBM, which secures the machine password. The token is stored in an IBM service database with limited access (authorized personnel only). IBM support uses the token and the serial number of the machine to generate a service password that can be used during a service session. For the next support case, you need not submit the token again. You just have to provide the machine serial number.

You can change (rotate) the secret machine password if you think it has been compromised. In this case, you have to submit the encrypted token again for your next support case.

What you can do to guarantee maximum system security

Security-related fixes and APAR information are published in the IBM Z Security Portal. To guarantee maximum system security, IBM issues the IBM Z Security Portal Security Vulnerability Bulletins. Apply the fixes recommended there. The IBM Z Security Portal is intended to help you stay current with security and system integrity fixes by providing current patch data and also provides Associated Common Vulnerability Scoring System (CVSS) V2 ratings for new APARs. For instructions on how to access the IBM Z Security Portal, see **Enterprise security > Integrity**. You find a link to the website at the end of this topic.

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.

- Use ssh rather than telnet to connect to the IBM Db2 Analytics Accelerator Console. This way, passwords and other information exchanged with the console are encrypted and not sent as plain text.
- Change the password for the IBM Db2 Analytics Accelerator Console at the first logon.

If more than a single person needs access to the console, create additional user IDs on the console, each with a different password. Do not share a single console password.

- Make sure that the hardware is located in an access-controlled area. Although direct remote connections to the accelerator are not possible, someone who knows the machine details might enter the premises and use a false identity to open a service session with IBM support.
- If available, use an access auditing solution that records accesses 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 Integrated Analytics System is encrypted. Hence there is no need to degauss or physically destroy a damaged disk after a replacement.
- 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 webbrowser 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 disclose sensitive information in the trace file or result set. Make sure to only send this data to parties with a need to know.
- To be able to trace a manipulation of accelerator-shadow tables, 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-shadow tables have been modified. Consider adding SYSIBM.USERNAMES as well.

Related reference

Setting access rights for the user who runs AQTTIJSP

The main job for the installation of the stored procedures is AQTTIJSP. You must submit this job in a later step. The ID of the user who runs this job needs SYSADM authority. Make sure that this is the case.

Access rights for power users

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 Enterprise security > Integrity > Subscription Process

Chapter 4. Data encryption

Read some information about the encryption features provided with the product and the hardware.

IBM Integrated Analytics System

The IBM Integrated Analytics System uses flash storage. The flash cards in the flash storage compartment are always encrypted using AES encryption. However, the encryption keys can currently not be changed or managed on an external key management server. The encryption of data at rest provides a safety net if a single flash card is removed. The data on the card cannot be read because it is encrypted, but is not protected against an attacker who gains access to the complete system because the keys are stored within the appliance on the IBM Integrated Analytics System. So you can safely replace flash cards in case of a failure without the risk of exposing the data on the card.

Db2 Analytics Accelerator on Z

The Secure Services Container used by Db2 Analytics Accelerator on Z employs LUKS AES encryption to encrypt all persistent storage. That is, both, the storage used by the container image and the customer-supplied storage for user data are encrypted.

Chapter 5. Administration clients

Different administration clients are available for the administration of the product: IBM Db2 Analytics Accelerator Studio and IBM Data Server Manager.

IBM Db2 Analytics Accelerator Studio is a plugin for IBM Data Studio, which allows you to administer IBM Db2 Analytics Accelerator for z/OS. IBM Data Studio is a classic workstation application, in the sense that it must be installed on all client workstations connecting to your data server.

IBM Data Server Manager is a multi-purpose web application, which can also be used to administer IBM Db2 Analytics Accelerator for z/OS. Because it is a web application that comes with its own web server, it is sufficient to install it on a single computer within your company network. As long as client workstations can connect to this computer over the network, they can use IBM Data Server Manager for the administration of IBM Db2 Analytics Accelerator for z/OS.

IBM Data Server Manager is easier to maintain because updates must be installed in one place only, and need not be rolled out to a number of client workstations.

Instead of using one of the supported administration clients, you can also control and operate the product by using Db2 commands and the stored procedures that come with the product.

Launching a function from an administration client causes the client to connect to your data server, which in turn calls a stored procedure that executes the corresponding command on the accelerator. The accelerator is thus always controlled and maintained indirectly via the data server.

In most cases, you do not have to change existing database applications to use IBM Db2 Analytics Accelerator. If an accelerator has been properly deployed and tables have been added to the accelerator, queries that come from database applications are usually routed automatically from your data server to the accelerator provided that this results in shorter response times.

See the following diagram:

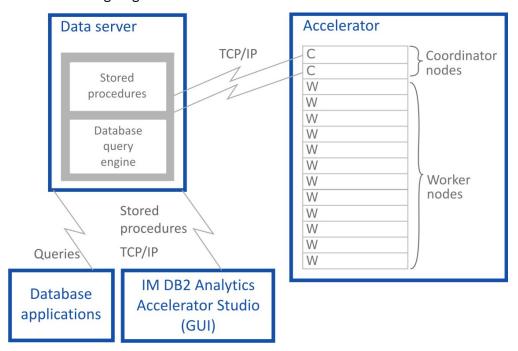


Figure 5. IBM Db2 Analytics Accelerator setup

Installing IBM Db2 Analytics Accelerator Studio

Follow the steps in this section to install the administration client, IBM Db2 Analytics Accelerator Studio.

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.

If you already have an installation of IBM Data Studio, and just need the IBM Db2 Analytics Accelerator Studio plug-in, you can also download and add the plug-in by using the update function of IBM Installation Manager. To do this, follow the link to Updating IBM Db2 Analytics Accelerator Studio at the end of this topic.



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

- 1. From a workstation with a direct internet connection, download the installation package, which is found at:
 - Db2 Analytics Accelerator Studio
- 2. Type your IBM ID in the appropriate field and click **Continue**. If you do not have an IBM ID, click Create IBMid 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.

Note: Sometimes, you are asked to provide information such as your first name, last name, email address and your preferred business contact method although you have already signed in with your IBM ID. In that case, just provide the requested information again to proceed.

- 3. On the following page, enter the password of your IBM ID and click Log in.
- 4. On the following page, under **License**, select **I agree** to agree to the terms and conditions.
- 5. Click I Confirm.
- 6. You can use either of these methods do download the installation package:
 - Download using Download Director
 - Download using http

For the http method, click the appropriate tab (the tab for the Download Director is in front, so it does not have to be clicked).

- 7. Click **Download now** and wait for the download to finish.
- 8. 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.

- 9. Depending on your choice in step <u>"8" on page 20</u>, extract the package on the download machine, or let your users extract it on their local machines.
- 10. Let your users start one of the installation programs, that is, **launchpad.exe** or

launchpad_win_noadmin64.exe

The **launchpad.exe** program allows you to launch an administrative installation and a restricted installation. The **launchpad_win_noadmin64.exe** program only allows you to launch a restricted installation and only works on 64-bit Windows 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.

- 11. After selecting an installation method, IBM Installation Manager starts and lists the installable packages. Your users must select the following packages:
 - IBM Data Studio client
 - IBM Db2 Analytics Accelerator Studio

If any of these are already installed, they will be grayed out or not listed. Sometimes, IBM Installation Manager will offer to install a newer version.

12. Let your users complete the IBM Installation Manager wizard.

Related tasks

Updating IBM Db2 Analytics Accelerator Studio

Follow the instructions here to download and install an upgrade of IBM Db2 Analytics Accelerator Studio.

Installing IBM Data Server Manager

Download and install IBM Data Server Manager on a computer within the network of your organization.

Before you begin

IBM Data Server Manager is available for the following operating systems and computers:

- Linux on x86 (Intel) computers
- Linux on IBM pSeries (POWER®) systems
- · Linux on IBM Z
- IBM AIX® systems
- Microsoft Windows 64-bit systems

About this task

Procedure

1. On a computer with a direct internet connection IBM Data Server Manager, open a web browser and go to the following page

Marketing Registration Services (MRS).

2. Type your IBM ID in the appropriate field and click **Continue**.

If you do not have an IBM ID, click **Create IBMid** 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.

Note: Sometimes, you are asked to provide information such as your first name, last name, email address and your preferred business contact method although you have already signed in with your IBM ID. In that case, just provide the requested information again to proceed.

- 3. On the following page, enter the password of your IBM I and click Sign in.
- 4. On the **Data Server Manager v2.x.x.x** page, answer the following questions by selecting an answer from the drop-down lists:
 - a. What is your current reason for downloading?
 - b. How will you be involved in the decision making process for your projects?
- 5. Scroll down to the **License** section, and select **I** agree to agree to the terms and conditions.
- 6. Click I Confirm.
- 7. You can use either of these methods do download the installation package:
 - Download using Download Director
 - Download using http

For the http method, click the appropriate tab (the tab for the Download Director is in front, so it does not have to be clicked).

- 8. Click **Download now** next to the appropriate package and wait for the download to finish.
- 9. If necessary, transfer the package to the computer, virtual machine, or partition on which you want to install IBM Data Server Manager.
- 10. Extract the package contents. This is not required on Windows machines.
- 11. Proceed as follows:
 - On AIX or Linux:
 - a. Open a terminal window.
 - b. Change to the directory that contains the extracted package content.
 - c. Run ./setup.sh.
 - On a Windows machine, double-click the package name. This will start an installation wizard.
- 12. Follow the instructions of the installation program or wizard.

The program or wizard will suggest default network ports. You can change these according to your needs. You will also be prompted to specify a logon user ID and a password.

13. Wait until the installation program or wizard completes and the server has been started.

What to do next

Check if you can log on to IBM Data Server Manager:

- 1. On the computer on which IBM Data Server Manager has been installed, start a web browser.
- 2. Enter the following URL in the address field:

https://localhost:<port>/console

where <port> is the port number you chose for https connections during the installation. By default, this is port 11081.

- 3. Your browser might display a security warning. Accept the warning and continue.
- 4. The **Sign In** page of IBM Data Server Manager should be displayed:

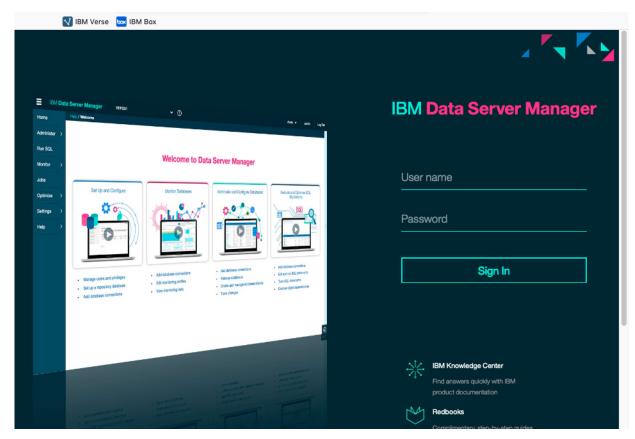


Figure 6. The **Sign In** page of IBM Data Server Manager

- 5. In the appropriate fields, type the user name and password you provided in step "12" on page 22.
- 6. Click Sign In.

Related tasks

Updating IBM Data Server Manager

Follow the instructions here to download and install an upgrade of IBM Data Server Manager.

Chapter 6. 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 for z/OS data server

Before you begin with the installation, make sure that you meet the software requirements for the Db2 for z/OS data server.

For more information, refer to the following website:

Prerequisites and Maintenance for IBM Db2 Analytics Accelerator for z/OS, V7

Installing libraries with IBM Db2 Analytics Accelerator support

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

Before you begin

Make sure that Db2 for z/OS is at the proper maintenance level.

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. To set up the stored procedures and related libraries for the product, use the SMP/E *Apply* function to install the following packages:

FMID HAOT710

Includes the stored procedures of IBM Db2 Analytics Accelerator.



Attention: An SMP/E installation of HAQT710 will delete all previous versions of IBM Db2 Analytics Accelerator, for example HAQT510..To retain older installations, install HAQT710 into a separate SMP/E environment.

FMID JAQT711

Includes the license library for IBM Db2 Analytics Accelerator on an IBM Integrated Analytics System.

FMID JAQT712

Includes the license library for Db2 Analytics Accelerator on Z.

FMID HCHCA21

Includes IBM InfoSphere Change Data Capture for z/OS components

- 2. 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.

Related information

Installing or upgrading InfoSphere CDC

Configuring the incremental update function

IBM Db2 Analytics Accelerator for z/OS offers two implementation options for the handling of incremental updates: IBM Integrated Synchronization and IBM InfoSphere Change Data Capture for z/OS (CDC). If you want to use CDC, you are advised to configure 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 task before you create the IBM Db2 Analytics Accelerator stored procedures by following the steps in "Customizing and running AQTTIJSP" on page 49.

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.

Note: You can also set, change, or override most of the ZPARMS that are discussed here in one of the following ways:

- Online, during a query session
- In the BIND options for packages (static SQL)
- · As part of connection properties

Setting ZPARMs for IBM Db2 Analytics Accelerator in Db2 for z/OS

In Db2 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.

AUTO

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 archived by the High-Performance Storage Saver (HPSS) is searched when a query is executed.

YES

Archived data is searched.

NO

Archived 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.

QUERY_ACCEL_OPTIONS = NONE | TRUE

The following options are currently available for this ZPARM.

NONE (default)

Means that values cannot be set.

TRUE

Brings up the DSNTIP8A panel, which allows you to set one or more of the values in the range from 1 to 11 (separated by a comma):

```
------ ACCELERATION OPTIONS
DSNTIP8A
                    Acceleration Options Panel
Allow queries to run on an accelerator server that contain: 1 MBCS EBCDIC ===> NO NO or YES
                                         ===> NO
                                                     NO or YES
 2 INSERT FROM SELECT
                                         ===> NO
                                                     NO or YES
 3 BYTE-BASED STRING FUNCTIONS ===> NO
4 DD/MM/YYYY LOCAL DATE FORMAT ===> NO
5 SYSTEM TIME TEMPORAL QUERIES ===> NO
                                                     NO or YES
                                                     NO or YES
                                                   NO or YES
 5 SYSTEM_TIME TEMPORAL QUERIES
 6 TS COLS WITH PRECISION OF 12
                                         ===> NO
                                                     NO or YES
 7 YYYYMMDD
              LOCAL DATE FORMAT
                                         ===> NO
                                                      NO or YES
                                    ===> NO
 8 FAVOR NEW ACCELERATOR_TYPE
                                                     NO or YES
```

```
9 ENABLE UNCERTAINTY COST ESTIMATION ===> NO
                                               NO or YES
10 BALANCE WORKLOAD BETWEEN VERSIONS ===> NO
                                               NO or YES
11 USE ONLY NEW ACCELERATOR_TYPE
                                              NO or YES
PRESS: ENTER to continue UP/DOWN to scroll
                                              RETURN to exit
```

1 MBCS EBCDIC

NO

Queries that include data encoded in a multibyte-character-set EBCDIC encoding scheme are blocked from accelerated query processing.

Queries that include data encoded in multibyte EBCDIC are not blocked from accelerated query processing, although accelerators use a different encoding scheme (Unicode, UTF-8) for the same data. Consequently, a query processed by Db2 for z/OS might return a result set that is different from the result set that the accelerator returns for the same query.

The reason for result differences lies in the different collating sequences used for data on the accelerator and for data in Db2. This difference affects data ordering and the results of range predicates. Therefore, differences are likely to occur if table columns include a mixture of uppercase characters, lowercase characters, numeric characters, or national language characters, and if the SQL codes contains range predicates or ordering statements for the columns including these characters.

2 INSERT FROM SELECT

NO

For an INSERT operation that includes a SELECT statement, the SELECT portion of the statement is not routed to an accelerator, that is, the entire statement is executed in Db2 for z/OS.

YES

For an INSERT operation that includes a SELECT statement, this option causes the SELECT portion of the statement to be routed to an accelerator and the INSERT operation to be performed by Db2 for z/OS.

Notes:

- It is not necessary to set this option to YES for in-database transformation with accelerator-only tables because transactions on this type of table can only be executed on an accelerator. The processing location is automatically determined by the type of the table.
- 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 BYTE-BASED STRING FUNCTIONS

NO

Queries that contain built-in functions for which Db2 processes each byte of the input string, rather than each character of the input string, are not accelerated. Such queries are processed by Db2 for z/OS.

YES

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 and IBM 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 DD/MM/YYYY LOCAL DATE FORMAT

NO

Queries containing expressions with a DATE data type in a LOCAL format are not accelerated.

YES

Allows the acceleration of queries that use expressions with a DATE data type in a LOCAL format.

IBM Db2 Analytics Accelerator 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 define the specific dd/mm/yyyy date format. If the LOCAL format is not defined as dd/mm/yyyy, the query might fail with SQLCODE -4742 and reason code 19 or return unpredictable results.

Option 4 cannot be used if option 7 is specified because these options are mutually exclusive.

The installation panel or bind option specifies LOCAL as the format for dates and if dd/mm/yyyy is used as the representation of the LOCAL format. If dd/mm/yyyy is not used as the date representation of the LOCAL format, 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.

5 SYSTEM_TIME TEMPORAL QUERIES

NO

Queries against queries against system-temporal and bi-temporal tables are not accelerated.

YES

Allows you to run accelerated queries against system-temporal and bi-temporal tables. When this value is set, timestamp columns with a precision of 12, which are used in the PERIOD definition of system-temporal intervals, are included when such a table is defined on an accelerator. The inclusion entails the truncation of the timestamp 12 values in these columns to values with precision 6. As a result, these columns can be referenced by an accelerated query.

6 TS COLS WITH PRECISION OF 12

NO

Queries that contain timestamp columns with a precision of 12 are not accelerated.

YES

Allows the truncation of timestamp values with precision 12 to a precision of 6. The database engine of the product can only handle timestamps with a precision up to 6. Base tables in DB2 10 for z/OS or higher might contain values with a precision up to 12. Hence an upfront conversion (truncation) of such values is required before these are loaded into an accelerator-shadow table. An accelerated query would fail if longer values were encountered.

You need not set this value for queries against system-temporal or bi-temporal tables because the use of option 5 already includes a truncation of timestamp values in these tables. However, option 6 must be set if you want to run accelerated queries against timestamp columns with precision 12 in other tables. This includes, business-temporal tables that use such columns.

7 YYYYMMDD LOCAL DATE FORMAT

NO

Queries containing expressions with a DATE data type in a LOCAL format are not accelerated.

YES

Same as option 4, but for LOCAL Db2 DATE values in the yyyymmdd format (with two space characters at the end).

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 define the specific YYYYMMDD date format. If the LOCAL format is not defined as YYYYMMDD, the query might fail with SOLCODE -4742 and reason code 19 or return unpredictable results.

Option 7 cannot be used if option 4 is enabled because these options are mutually exclusive.

8 FAVOR NEW ACCELERATOR_TYPE

NO

Favors the older accelerator (version 5.x.x) in an environment that contains version 5 and version 7 accelerators. That is, queries eligible for query acceleration on both types of accelerators are always sent to the version 5 accelerator.

YES

Favors the newer accelerator (version 7.x.x) in an environment that contains version 5 and version 7 accelerators. That is, queries eligible for query acceleration on both types of accelerators are always sent to the version 7 accelerator.

9 ENABLE UNCERTAINTY COST ESTIMATION

NO

Uncertainty cost consideration is disabled for accelerated gueries.

Enables uncertainty cost consideration for accelerated queries.

10 BALANCE WORKLOAD BETWEEN VERSIONS

NO

In an environment with version 5 and version 7 accelerators, the individual workload on each accelerator will not be taken into account when the routing decision is made. That is, a query will be sent to one of the available accelerators at random or according to the setting of the options 8 or 10 on this panel.

YES

Bases the routing decision on the individual workload of each accelerator in an environment with version 5 and version 7 accelerators. That is, the accelerator with the lowest workload will be selected.

You can also influence workload balancing by using the CURRENT ACCELERATOR special register. For example, if this option is set to YES, and the CURRENT ACCELERATOR special register specifies a particular accelerator, the workload will be routed to that server only if it is available. This holds true even if accelerators with a lower workload are available.

11 USE ONLY NEW ACCELERATOR TYPE

In an environment with version 5, version 7, and even newer accelerators, EXPLAIN processing, accelerator modeling, and query execution is available for all accelerators regardless of their version.

YES

Reserves EXPLAIN processing, accelerator modeling, and query execution exclusively for the newest accelerator type in environments where version 5, version 7, or later accelerator types coexist.

Important:

- You can change the value of QUERY_ACCEL_OPTIONS online.
- If you use IBM Db2 Analytics Accelerator 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.

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 ===> NO NO, COMMAND, or AUTO
2 GET ACCEL ARCHIVE ===> NO NO 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 WITH FAILBACK
4 = ELIGIBLE
5 = ALL

PRESS: ENTER to continue RETURN to exit HELP for more information
```

DSNTIP8A is a pop-up 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

Subsystem parameters that are not on installation panels

Using a sample job to set ZPARMs for IBM Db2 Analytics Accelerator for z/OS

Instead of setting the ZPARMs directly in Db2 for z/OS, you can customize and run a sample job to the same end.

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.

AUTO

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 archived by the High-Performance Storage Saver (HPSS) is searched when a query is executed.

YES

Archived data is searched.

NO

Archived 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 OUERY ACCELERATION ZPARM provides the default setting for the CURRENT OUERY 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.

QUERY ACCEL OPTIONS = NONE | TRUE

The following options are currently available for this ZPARM.

NONE (default)

Means that values cannot be set.

TRUE

Brings up the DSNTIP8A panel, which allows you to set one or more of the values in the range from 1 to 11 (separated by a comma):

```
----- ACCELERATION OPTIONS --
           Acceleration Options Panel
DSNTIP8A
===>
Allow queries to run on an accelerator server that contain:
 1 MBCS EBCDIC
                                   ===> NO
                                             NO or YES
2 INSERT FROM SELECT
                                   ===> NO
                                             NO or YES
```

```
3 BYTE-BASED STRING FUNCTIONS
                                     ===> NO
                                                NO or YES
 4 DD/MM/YYYY LOCAL DATE FORMAT
                                     ===> NO
                                                NO or YES
5 SYSTEM_TIME TEMPORAL QUERIES
                                      ===> NO
                                                NO or YES
6 TS COLS WITH PRECISION OF 12
                                                NO or YES
                                      ===> NO
                                      ===> NO
                                                NO or YES
7 YYYYMMDD
             LOCAL DATE FORMAT
8 FAVOR NEW ACCELERATOR TYPE
                                     ===> NO
                                                NO or YES
 9 ENABLE UNCERTAINTY COST ESTIMATION ===> NO
                                                NO or YES
10 BALANCE WORKLOAD BETWEEN VERSIONS ===> NO
                                                NO or YES
11 USE ONLY NEW ACCELERATOR_TYPE
                                     ===> NO
                                                NO or YES
PRESS: ENTER to continue UP/DOWN to scroll
                                               RETURN to exit
```

1 MBCS EBCDIC

NO

Queries that include data encoded in a multibyte-character-set EBCDIC encoding scheme are blocked from accelerated query processing.

YES

Queries that include data encoded in multibyte EBCDIC are not blocked from accelerated query processing, although accelerators use a different encoding scheme (Unicode, UTF-8) for the same data. Consequently, a query processed by Db2 for z/OS might return a result set that is different from the result set that the accelerator returns for the same query.

The reason for result differences lies in the different collating sequences used for data on the accelerator and for data in Db2. This difference affects data ordering and the results of range predicates. Therefore, differences are likely to occur if table columns include a mixture of uppercase characters, lowercase characters, numeric characters, or national language characters, and if the SQL codes contains range predicates or ordering statements for the columns including these characters.

2 INSERT FROM SELECT

NO

For an INSERT operation that includes a SELECT statement, the SELECT portion of the statement is not routed to an accelerator, that is, the entire statement is executed in Db2 for z/OS.

YES

For an INSERT operation that includes a SELECT statement, this option causes the SELECT portion of the statement to be routed to an accelerator and the INSERT operation to be performed by Db2 for z/OS.

Notes:

- It is not necessary to set this option to YES for in-database transformation with accelerator-only tables because transactions on this type of table can only be executed on an accelerator. The processing location is automatically determined by the type of the table.
- 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 BYTE-BASED STRING FUNCTIONS

NO

Queries that contain built-in functions for which Db2 processes each byte of the input string, rather than each character of the input string, are not accelerated. Such queries are processed by Db2 for z/OS.

YES

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 and IBM 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 DD/MM/YYYY LOCAL DATE FORMAT

NO

Queries containing expressions with a DATE data type in a LOCAL format are not accelerated.

YES

Allows the acceleration of queries that use expressions with a DATE data type in a LOCAL format.

IBM Db2 Analytics Accelerator 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 define the specific dd/mm/yyyy date format. If the LOCAL format is not defined as dd/mm/yyyy, the query might fail with SQLCODE -4742 and reason code 19 or return unpredictable results.

Option 4 cannot be used if option 7 is specified because these options are mutually exclusive.

The installation panel or bind option specifies LOCAL as the format for dates and if dd/mm/ yyyy is used as the representation of the LOCAL format. If dd/mm/yyyy is not used as the date representation of the LOCAL format, 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.

5 SYSTEM_TIME TEMPORAL QUERIES

NO

Queries against queries against system-temporal and bi-temporal tables are not accelerated.

YES

Allows you to run accelerated queries against system-temporal and bi-temporal tables. When this value is set, timestamp columns with a precision of 12, which are used in the PERIOD definition of system-temporal intervals, are included when such a table is defined on an accelerator. The inclusion entails the truncation of the timestamp 12 values in these columns to values with precision 6. As a result, these columns can be referenced by an accelerated query.

6 TS COLS WITH PRECISION OF 12

NO

Queries that contain timestamp columns with a precision of 12 are not accelerated.

YES

Allows the truncation of timestamp values with precision 12 to a precision of 6. The database engine of the product can only handle timestamps with a precision up to 6. Base tables in DB2 10 for z/OS or higher might contain values with a precision up to 12. Hence an upfront conversion (truncation) of such values is required before these are loaded into an accelerator-shadow table. An accelerated query would fail if longer values were encountered.

You need not set this value for queries against system-temporal or bi-temporal tables because the use of option 5 already includes a truncation of timestamp values in these tables. However, option 6 must be set if you want to run accelerated queries against timestamp columns with precision 12 in other tables. This includes, business-temporal tables that use such columns.

7 YYYYMMDD LOCAL DATE FORMAT

NO

Queries containing expressions with a DATE data type in a LOCAL format are not accelerated.

YES

Same as option 4, but for LOCAL Db2 DATE values in the yyyymmdd format (with two space characters at the end).

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 define the specific YYYYMMDD date format. If the LOCAL format is not defined as YYYYMMDD , the query might fail with SQLCODE -4742 and reason code 19 or return unpredictable results.

Option 7 cannot be used if option 4 is enabled because these options are mutually exclusive.

8 FAVOR NEW ACCELERATOR_TYPE

NO

Favors the older accelerator (version 5.x.x) in an environment that contains version 5 and version 7 accelerators. That is, queries eligible for query acceleration on both types of accelerators are always sent to the version 5 accelerator.

YES

Favors the newer accelerator (version 7.x.x) in an environment that contains version 5 and version 7 accelerators. That is, queries eligible for query acceleration on both types of accelerators are always sent to the version 7 accelerator.

9 ENABLE UNCERTAINTY COST ESTIMATION

NO

Uncertainty cost consideration is disabled for accelerated queries.

YES

Enables uncertainty cost consideration for accelerated queries.

10 BALANCE WORKLOAD BETWEEN VERSIONS

NO

In an environment with version 5 and version 7 accelerators, the individual workload on each accelerator will not be taken into account when the routing decision is made. That is, a query will be sent to one of the available accelerators at random or according to the setting of the options 8 or 10 on this panel.

YES

Bases the routing decision on the individual workload of each accelerator in an environment with version 5 and version 7 accelerators. That is, the accelerator with the lowest workload will be selected.

You can also influence workload balancing by using the CURRENT ACCELERATOR special register. For example, if this option is set to YES, and the CURRENT ACCELERATOR special register specifies a particular accelerator, the workload will be routed to that server only if it is available. This holds true even if accelerators with a lower workload are available.

11 USE ONLY NEW ACCELERATOR_TYPE

NO

In an environment with version 5, version 7, and even newer accelerators, EXPLAIN processing, accelerator modeling, and query execution is available for all accelerators regardless of their version.

YES

Reserves EXPLAIN processing, accelerator modeling, and query execution exclusively for the newest accelerator type in environments where version 5, version 7, or later accelerator types coexist.

Important:

- You can change the value of QUERY ACCEL OPTIONS online.
- If you use IBM Db2 Analytics Accelerator 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.

Procedure

- 1. Add or change the listed parameters of the DSN6SPRM macro, which is used to provide input to the DSNTIJUZ sample job in your working libraries.
- 2. Stop Db2. This is not required if you just want to change the values of QUERY_ACCEL_OPTIONS.
- 3. Submit DSNTIJUZ to assemble the new ZPARM load member.
- 4. Restart Db2 so that the changes can take effect. If just QUERY_ACCEL_OPTIONS has been changed, you do not have to restart Db2. Instead, run the following Db2 command:

SET SYSPARM LOAD

5. Run sample queries or a test load to verify the proper functioning of Db2.

Related information

CURRENT QUERY ACCELERATION

Subsystem parameters that are not on installation panels

Configuring the incremental update function

IBM Db2 Analytics Accelerator for z/OS offers two implementation options for the handling of incremental updates: IBM Integrated Synchronization and IBM InfoSphere Change Data Capture for z/OS (CDC). If you want to use CDC, you are advised to configure 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

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

Incremental updates

The incremental update function of IBM Db2 Analytics Accelerator for z/OS allows you to update accelerator-shadow tables continually. Changes to the data in original Db2 for z/OS tables are thus propagated to the corresponding target tables with a high frequency and just a brief delay. This way, query results from an accelerator are always extracted from recent, close-to-real-time data.

Chapter 7. 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

Follow the steps here to set up a suitable Workload Manager (WLM) application environment for the stored procedures of the product.

Before you begin

Make sure that the required Db2 program temporary fixes (PTF) for IBM Db2 Analytics Accelerator for z/OS support have been installed.

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.

<HLQBASE>

HLQ for your Db2 libraries

<HLQSP>

HLQ for the IBM Db2 Analytics Accelerator stored-procedure libraries

<HLOSP>.SAOTLICI

Contains information about the deployment options that you are entitled to use:

AQTLICP

A member for IBM Db2 Analytics Accelerator on an IBM Integrated Analytics System

AOTLICZ

A member for Db2 Analytics Accelerator on Z

<HLQDb2SSN>

HLQ for Db2 subsystem-specific libraries

<HLQACTIVE>

A suggested HLQ for copies of <HLQSP>. To work on copies is recommended because these are independent of the original libraries under SMP/E control. If needed, you can create new copies from the unchanged originals.

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.

<HLQSP>.SAQTDBRM

Contains database request modules (DBRMs) that must be bound to Db2.

<HLOSP>.SAOTMOD

Contains shared libraries and load modules for the stored procedures.

Important: This data set must be APF-authorized if you want to use any of these products or features:

• IBM Db2 Analytics Accelerator *Loader*

Procedure

I

- 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 stored procedures. Use the following properties:

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 N
       JCL FOR RUNNING THE WLM-ESTABLISHED STORED PROCEDURES
       ADDRESS SPACE
                  -- THE MVS REGION SIZE FOR THE ADDRESS SPACE.
//*
          RGN
//*
//*
//*
//*
          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.
          Db2VERS -- Db2-VERSION (I.E. V110)
                     SET BY APPLICATION ENVIRONMENT
//*
//*
                         DSNWLMV9 ==> V110
       The user ID that is used to start the task must have
//*
       read access to the <HLQs> in the STEPLIB statement
//********************************
//DSNWLMV9 PROC RGN=0K, APPLENV=DSNWLMV9, 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=<HLQACTIVE>.SAQTMOD
//SYSTSPRT DD
                 SYSOUT=A
//CEEDUMP DD
                SYSOUT=H
//0UT1
          DD
                SYSOUT=A
//UTPRINT DD
                 SYSOUT=A
//DSSPRINT DD
                SYSOUT=A
//SYSPRINT DD
                 SYSOUT=A
                 DSN=<HLQACTIVE>.SAQTSAMP(AQTENV),DISP=SHR
//AQTENV
          DD
//AQTDEF6 DD
                 DSN=<HLQACTIVE>.SAQTSAMP(AQTDEF6),DISP=SHR
```

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 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 stored procedures call the following Db2 for z/OS stored procedures:
 - SYSPROC.ADMIN_INFO_SYSPARM
 - SYSPROC.DSNUTILU

In addition, the Db2 stored procedure SYSPROC.ADMIN_COMMAND_DB2 is called by the AQTTIJSP installation job and by your administration client.

Run these Db2-supplied stored procedures in a separate WLM application environments that is not shared with the IBM Db2 Analytics Accelerator 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 the IBM Db2 Analytics Accelerator stored procedures are supposed to use. To do so, add the following statements to the procedure that starts the address space:

```
//SYSTCPD DD DISP=SHR,DSN=<TCPIP.DATA file>
//TCPIP EXEC PGM=BPXTCAFF,PARM=TCPIP
```

Note: The SYSTCPD card might not be a sufficient stack selector for outbound connections. In fact, the decisive route and stack selector is the common inet routing table, which determines which route in which stack is used. You can restrict the choice of an application or batch job to one particular stack in the inet routing table by adding the BPXTCAFF call. It depends on the network configuration whether this is necessary.

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 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 sample AQTENV data set member that came with this version of the product. Do not re-customize an older version of AQTENV because fundamental settings might have changed and using an old configuration might lead to problems.
- The value of NUMTCB must be in the range between 15 and 30 (inclusively), so as not to exceed the 31-bit memory limits of the stored procedures.
- 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

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

• If you are planning to use an external program that comes with a set of callback functions, add the name of the dynamic link library (DLL) with the callback resources to the //STEPLIB section as shown. The DLL you specify must have been compiled as a 31-bit XPLINK application.

```
//STEPLIB DD DISP=SHR,DSN=<HLQDb2SSN>.SDSNEXIT
// DD DISP=SHR,DSN=<HLQBASE>.SDSNLOAD
// DD DISP=SHR,DSN=<HLQBASE>.SDSNLOD2
// DD DISP=SHR,DSN=<HLQACTIVE>.SAQTMOD
// DD DISP=SHR,DSN=<THIRD.PARTY.CALLBACK>
```

where <THIRD. PARTY. CALLBACK> is the name of the DLL.

At present, external programs can be called by the SYSPROC.ACCEL_LOAD_TABLES stored procedure. This stored procedure allows you to address callback interfaces from certain XML input elements, so that callback functions can be started before, when, and after the execution of the external program. An important aim of such functions is a better runtime stability during program execution. For example, newer versions of the IBM Db2 Analytics Accelerator *Loader* provide callback functions for the following purposes:

- To check the version compatibility of the external program
- To set the proper interception point, at which the external program interrupts the regular load procedure (call to Db2 Unload Utility) and takes over
- To end the execution of the external program in case of errors

For more information about the callback interfaces, see the description of the $table_load_specification$ parameter. You find it in the section on the SYSPROC.ACCEL_LOAD_TABLES stored procedure in the IBM Db2 Analytics Accelerator for z/OS: Stored Procedures Reference.

To write your own callback interfaces, you can use the sample headers in the appendix of the *Stored Procedures Reference*. The header can be used for applications in the C or C++ programming language.

Related tasks

Installing libraries with IBM Db2 Analytics Accelerator support

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

Installing IBM Db2 Analytics Accelerator Studio

Follow the steps in this section to install the administration client, IBM Db2 Analytics Accelerator Studio.

Related reference

Environment variables

The job control language (JCL) for the configuration of the Workload Manager (WLM) environment for IBM Db2 Analytics Accelerator 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

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

DB2 11 for z/OS: DB2-supplied stored procedures and user-defined functions

Setting up a WLM application environment for different product versions

To run IBM Db2 Analytics Accelerator for z/OS Version 7.x.x (Db2 Analytics Accelerator on Z or IBM Db2 Analytics Accelerator on IBM Integrated Analytics System) and an older version side-by-side, like IBM

Db2 Analytics Accelerator for z/OS Version 5.1.0 PTF-2 or later, the Workload Manager (WLM) application environment must be set up to accommodate two sets of stored procedures.

Before you begin

Make sure that the following components have been installed:

- Db2 for z/OS with the required Db2 program temporary fixes (PTFs) for both product features
- Db2 command line processor for calling and verifying the WLM environment from the local environment

Note: The Db2 command line processor is a Java[™] 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.

<HLQBASE>

HLQ for your Db2 libraries

<HLQSP>

HLQ for the IBM Db2 Analytics Accelerator stored-procedure libraries

<HLQDb2SSN>

HLQ for Db2 subsystem-specific libraries

<HLQXML4C1>

,HLQ for the XML toolkit. The XML toolkit is only required for product version 5. Therefore, you only need it if you want to run product version 5 and product version 7 in the same WLM environment. It is not needed if you want to run different version 7 accelerators, such as version 7.1.9 and version 7.5.0, in the same WLM environment.

<HLQACTIVEV7>

A suggested HLQ for copies of the IBM Db2 Analytics Accelerator for z/OS Version 7.x.x libraries. To work on copies is recommended because these are independent of the original libraries under SMP/E control. If needed, you can create new copies from the unchanged originals.

<HLQACTIVEV5>

A suggested HLQ for copies of the libraries for IBM Db2 Analytics Accelerator for z/OS Version 5.1.x.

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

<HLOSP>.SAOTSAMP

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.

<HLQSP>.SAQTDBRM

Contains database request modules (DBRMs) that must be bound to Db2.

<HLQSP>.SAQTMOD

Contains shared libraries and load modules for the stored procedures.

Important: This data set must be APF-authorized if you want to use any of these products or features:

- IBM Db2 Analytics Accelerator Loader
- RACF to control federated access

Procedure

1. Copy the <HLQSP>.SAQTMOD load-module data-set as <HLQACTIVEV7>.SAQTMOD for reference in your WLM procedure. Also copy the <HLQSP>.SAQTSAMP(AQTENV) member as

<HLQACTIVEV7>. 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 stored procedures. Use the following properties:

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.

Important: The STEPLIB DD statement for the IBM Db2 Analytics Accelerator V7.1 libraries ([...] <HLQACTIVEV7>.SAQTMOD) must come before the statement for the IBM Db2 Analytics Accelerator V5.1 libraries ([...] <HLQACTIVEV5>.SAQTMOD).

```
//********************
///* PROCEDURE NAME = DSNWLM
//*
       JCL FOR RUNNING THE WLM-ESTABLISHED STORED PROCEDURES
       ADDRESS SPACE
//*
//*
//*
                     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 PROC RGN=0K,APPLENV=DSNWLMV9,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=<HLQACTIVEV7>.SAQTMOD
DD DISP=SHR,DSN=<HLQACTIVEV5>.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=<HLQACTIVEV7>.SAQTSAMP(AQTENV),DISP=SHR
```

Important:

- If you are migrating from an earlier version of IBM Db2 Analytics Accelerator for z/OS, make sure to delete the reference to the AQTOSR data set from this template or procedure. The AQTOSR data set is no longer used.
- 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 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 stored procedures call the following Db2 for z/OS stored procedures:
 - SYSPROC.ADMIN INFO SYSPARM
 - SYSPROC.DSNUTILU

In addition, the Db2 stored procedure SYSPROC.ADMIN_COMMAND_Db2 is called by the AQTTIJSP installation job and by your administration client.

Run these Db2-supplied stored procedures in a separate WLM application environments that is not shared with the IBM Db2 Analytics Accelerator 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 the IBM Db2 Analytics Accelerator stored procedures are supposed to use. To do so, add the following statements to the procedure that starts the address space:

```
//SYSTCPD DD DISP=SHR,DSN=<TCPIP.DATA file>
//TCPIP EXEC PGM=BPXTCAFF,PARM=TCPIP
```

Note: The SYSTCPD card might not be a sufficient stack selector for outbound connections. In fact, the decisive route and stack selector is the common inet routing table, which determines which route in which stack is used. You can restrict the choice of an application or batch job to one particular stack in the inet routing table by adding the BPXTCAFF call. It depends on the network configuration whether this is necessary.

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 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 sample AQTENV data set that came with this
 version of the product. Do not re-customize an older version of AQTENV because fundamental
 settings might have changed and using an old configuration might lead to problems.
- The value of NUMTCB must be in the range between 15 and 30 (inclusively), so as not to exceed the 31-bit memory limits of the stored procedures.
- 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 your administration client only works if the value of AQT_HOST_PACKAGE_DIRECTORY has been 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 your administration client:

```
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
```

• If you are planning to use an external program that comes with a set of callback functions, add the name of the dynamic link library (DLL) with the callback resources to the //STEPLIB section as shown. The DLL you specify must have been compiled as a 31-bit XPLINK application.

```
//STEPLIB DD DISP=SHR,DSN=<HLQDb2SSN>.SDSNEXIT
// DD DISP=SHR,DSN=<HLQBASE>.SDSNLOAD
// DD DISP=SHR,DSN=<HLQBASE>.SDSNLOD2
// DD DISP=SHR,DSN=<HLQACTIVE>.SAQTMOD
// DD DISP=SHR,DSN=<THIRD.PARTY.CALLBACK>
...
```

where <THIRD.PARTY.CALLBACK> is the name of the DLL.

At present, external programs can be called by the SYSPROC.ACCEL_LOAD_TABLES stored procedure. This stored procedure allows you to address callback interfaces from certain XML input elements, so that callback functions can be started before, when, and after the execution of the external program. An important aim of such functions is a better runtime stability during program execution. For example, newer versions of the IBM Db2 Analytics Accelerator *Loader* provide callback functions for the following purposes:

- To check the version compatibility of the external program
- To set the proper interception point, at which the external program interrupts the regular load procedure (call to Db2 Unload Utility) and takes over
- To end the execution of the external program in case of errors

For more information about the callback interfaces, see the description of the $table_load_specification$ parameter. You find it in the section on the SYSPROC.ACCEL_LOAD_TABLES stored procedure in the IBM Db2 Analytics Accelerator for z/OS: Stored Procedures Reference.

To write your own callback interfaces, you can use the sample headers in the appendix of the *Stored Procedures Reference*. The header can be used for applications in the C or C++ programming language.

Related tasks

<u>Installing libraries with IBM Db2 Analytics Accelerator support</u>

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

Installing IBM Db2 Analytics Accelerator Studio

Follow the steps in this section to install the administration client, IBM Db2 Analytics Accelerator Studio.

Related reference

Environment variables

The job control language (JCL) for the configuration of the Workload Manager (WLM) environment for IBM Db2 Analytics Accelerator 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

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

DB2 11 for z/OS: DB2-supplied stored procedures and user-defined functions

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 stored procedures. Verify that this and a few other requirements are met by following the steps here.

Procedure

- 1. 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. In the start-up JCL job for the WLM environment for SYSPROC.DSNUTILU, use the MNSPAS parameter. Set the parameter to a value that matches the maximum number of table partitions that are to be loaded in parallel. Set the value at least to 10, but to no more than 50.
 - Mind that the number of parallel load operations (simultaneous runs of the SYSPROC.ACCEL_LOAD_TABLES stored procedure) and the maximum number of simultaneous unload operations per stored procedure call (determined by the AQT_MAX_UNLOAD_PARALLEL environment variable) influence the total number of parallel DSNUTILU calls.
- 4. Verify that all WLM environments include the following libraries in their STEPLIB statements:

```
//STEPLIB DD DISP=SHR,DSN=<HLQDb2SSN>.SDSNEXIT
// DD DISP=SHR,DSN=<HLQBASE>.SDSNLOAD
// DD DISP=SHR,DSN=<HLQBASE>.SDSNLOD2
```

Related information

DB2 11 for z/OS: DB2-supplied stored procedures and user-defined functions

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 Setting performance objectives for distributed workloads by using z/OS Workload Manager in the IBM Knowledge Center.

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 your administration client, 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.

```
Subsystem-Type Xref Notes Options Help
Command ===>
                                                       SCROLL ===>
Subsystem Type . : DDF
                              Fold qualifier names?
                                                        (Y or N)
Description . . . Distributed requests
                                        M=Move
Action codes: A=After
                           C=Copy
                                                  I=Insert rule
               B=Before
                          D=Delete row R=Repeat IS=Insert Sub-rule
                                                          More ===>
               --Qualifier--
                                                   ----Class-----
Action
         Type
                    Name
                            Start
                                                Service
                                                            Report
                                      DEFAULTS: DDF
     1
        UΙ
                                                DDF
                     ACCEL L*
                                                DDFLOAD
```

Figure 7. 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

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 main job for the installation of the stored procedures is AQTTIJSP. You must submit this job in a later step. The ID of the user who runs this job needs SYSADM authority. Make sure that this is the case.

If the user who is supposed to run the job is denied SYSADM authority for some reason, the user can try to run the job with the following privileges:

- 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 the 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 all SYSIBM catalog tables (SYSIBM.*):
 - SYSIBM.SYSCOLUMNS
 - SYSIBM.SYSCOPY
 - SYSIBM.SYSDATATYPES
 - SYSIBM.SYSDUMMY1
 - SYSIBM.SYSFIELDS
 - SYSIBM.SYSFOREIGNKEYS
 - SYSIBM.SYSINDEXES
 - SYSIBM.SYSINDEXPART
 - SYSIBM.SYSKEYCOLUSE
 - SYSIBM.SYSKEYS
 - SYSIBM.SYSRELS
 - SYSIBM.SYSTABCONST
 - SYSIBM.SYSTABLEPART
 - SYSIBM.SYSTABLES
 - SYSIBM.SYSTABLESPACE
 - SYSIBM.SYSTABLESPACESTATS
- UPDATE on the following table:
 - 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

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.
- Make sure that the Workload Manager (WLM) application environment has been started. You can check this by running the following command from the TSO or ISPF command line:

```
DISPLAY WLM, APPLENV=<env_name>
```

where <env_name> is the name of the WLM application environment that you configured for the IBM Db2 Analytics Accelerator stored procedures.

To enter this command from the ISPF primary menu, follow these steps:

- 1. Enter SD for System Display and Search Facility.
- 2. At the COMMAND INPUT prompt, enter /.
- 3. On the **System Command Extension** panel, type the **DISPLAY WLM, APPLENV=<env_name>** command next to one of the ===> prompts.
- 4. Press Enter.

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.

Important: The hyphen (-) is not a placeholder. Therefore, do not replace it with the name of a WLM environment in clauses like 'WLMENV(-)'. Replacing a hyphen in this way leads to SQL error -628 during the execution of the AQTTRIN job step. The actual placeholders are strings that start and end with an exclamation mark (!).

- 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 AQTTIJSP according to the instructions in the job.
 - If you want to install new versions of the database objects, keep the value INSTALL for the MODE parameter in job step AQTTRIN.

Important: If you want to run different product versions in parallel (coexistence), such as product version 5 and product version 7, create the stored procedures for both versions by using the INSTALL option. That is, run a new stored procedure installation even if one of the products is already installed. Do not use REINSTALL, as this will lead to problems.

• If you just want to see the SQL statements that would be executed, but do not want to run these statements, change the value of the MODE parameter to INSTALL-PREVIEW or REINSTALL-PREVIEW. The JCL is then copied to the data set that is specified in the JCLOUT DD statement, but not executed.

then the INSTALL option for MODE should be used to create the V7 procs in coexiszemce to V5

• If IBM support requests a trace file of the installation program, uncomment the following lines or blocks of code in the JCL:

AQTTRACE

Uncomment this step to allocate a trace data set.

AQTTRACE DD

Uncomment this statement in the AQTTRIN step to activate tracing.

- 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

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

An installation verification step is part of running the AQTTIJSP job. The following steps are relevant only if you encountered errors during the execution of AQTTIJSP. To verify the installation after the surfacing of AQTTIJSP errors, customize the AQTSJI00 job member (JCL) and run it.

Before you begin

Make sure that the user who will run AQTSJI00 has the privileges to read the following tables:

- Db2 catalog tables
- SYSACCEL.SYSACCELERATORS
- SYSACCEL.SYSACCELERATEDTABLES

About this task

To verify the installation of the stored procedures, you do not need a connection to IBM Db2 Analytics Accelerator for z/OS. 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.

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 AOTSJI00 JCL by running SAOTSAMP(AOTSJI00).

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.

Related reference

Members of SAQTSAMP

The SAQTSAMP data set contains various samples for installing, configuring or running the product's stored procedures. The following table lists all these members and provides brief descriptions of their functions.

Chapter 8. Db2 Analytics Accelerator on Z

Db2 Analytics Accelerator on Z is the product feature that runs entirely on a mainframe computer. Unlike IBM Db2 Analytics Accelerator on an IBM Integrated Analytics System, it does not require external hardware. This chapter deals with the definition of a logical partition (LPAR) for and the configuration of Db2 Analytics Accelerator on Z.

Architecture

Figure 8 on page 53 shows what major components the product consists of and how it fits into an existing IBM Z environment. The block that represents Db2 Analytics Accelerator on Z has a pink shading to set it off from the existing IBM Z environment, which has a light blue shading.

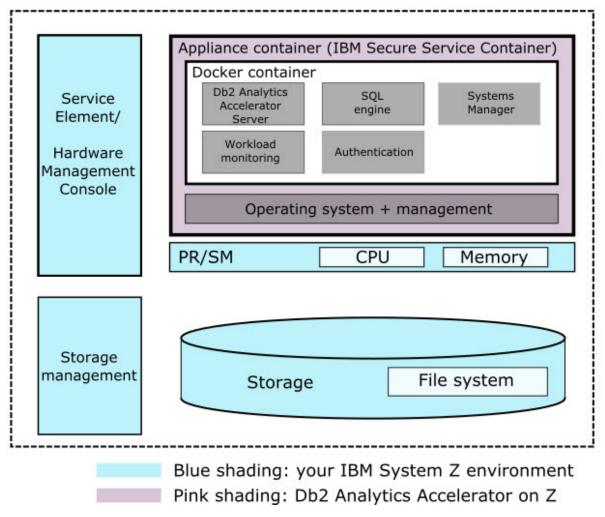


Figure 8. Architecture

Ways to set up Db2 Analytics Accelerator on Z

You have the following options when it comes to setting up Db2 Analytics Accelerator on Z:

- Setup on a single dedicated LPAR (single-node setup)
- Single-node setup in a Geographically Dispersed Parallel Sysplex (GDPS®) for failover support
- Setup on multiple dedicated LPARs (multi-node setup)

• Multi-node setup in a GDPS

A single-node installation is sufficient in many cases. A multi-node installation is for large-scale systems with a capacity of 30 Integrated Facilities for Linux (IFLs) or more. For systems like this, the distribution of work to different processing nodes (LPARs) leads to a better performance.

To address failover situations, both, single-node and multi-node installations of Db2 Analytics Accelerator on Z, can be integrated into a GDPS environment.

Important: Currently, Db2 Analytics Accelerator on Z supports GDPS Metro and GDPS Global.

<u>Figure 9 on page 54</u> shows two sites, (CEC A and CEC B), with an active and a passive accelerator (IDAA-on-Z LPAR A and LPAR B).

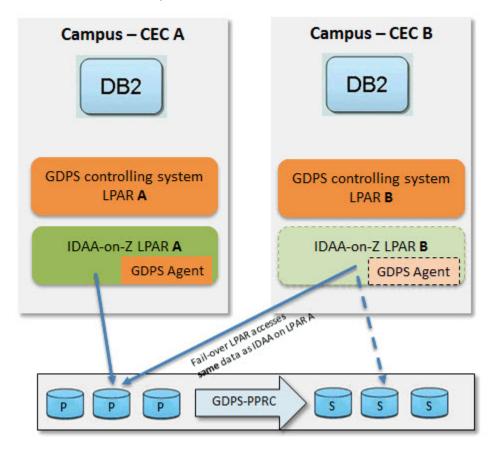


Figure 9. Failover support for a single-nodeDb2 Analytics Accelerator on Z with GDPS: schematic overview

If the active accelerator fails, GDPS detects this issue and activates the standby accelerator, which uses the same disks as the previously active accelerator. If the entire site fails, a storage swap will be triggered so that the standby accelerator uses the secondary set of disks. This type of setup requires some planning with regard to the use of networks and storage, which typically vary per site. You also have to make connectivity information available to the GDPS controlling systems. Because the network configuration for a GDPS setup is different from a setup without GDPS, it is important to decide early whether you want to use GDPS or not.

For more details on GDPS concepts and on the configuration of IBM Db2 Analytics Accelerator in this environment, refer to the GDPS Metro documentation.

Installing Db2 Analytics Accelerator on Z

Db2 Analytics Accelerator on Z can be installed with a single processing node or with multiple processing nodes that share the accelerator workload.

About this task

A single-node installation is sufficient in many cases. It requires a single Secure Service Container (SSC) LPAR.

A multi-node installation requires exactly six dedicated Secure Service Container (SSC) LPARs, which must be organized in an LPAR group (cluster). A multi-node installation is recommended if your IBM Z has 30 Integrated Facilities for Linux (IFLs) or more. Tests have shown that the distribution of large workloads to different LPARs leads to a better performance. It is also a good idea to start with a multi-node setup if you expect your workload to grow considerably.



Attention: A single-node installation cannot be converted to a multi-node installation later, or vice versa. This requires a new installation.

Defining an LPAR for Db2 Analytics Accelerator on Z

This list of steps describes how to create a single, dedicated LPAR for Db2 Analytics Accelerator on Z from the Hardware Management Console (HMC). For a multi-node setup, create six LPARs by repeating these steps.

Procedure

- 1. Open the Hardware Management Console (HMC) of your IBM Z.
- 2. Go to the Partitions tab.
- 3. Right-click a not yet activated partition (LPAR) and select **Operational Customization > Customize/ Delete Activation Profiles**.

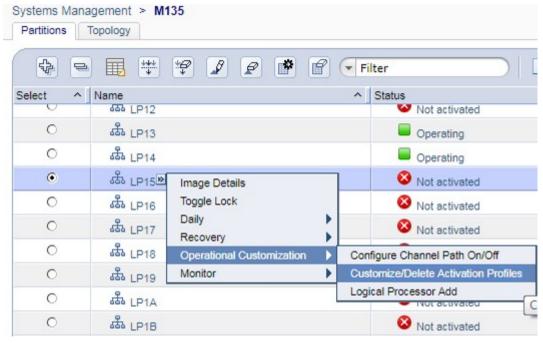


Figure 10. The Partitions window of the HMC

4. In the **General** section of the **Customize Image Profiles** window, select **Mode** > **SSC**. You might want to add a description for the profile in the **Description** field.

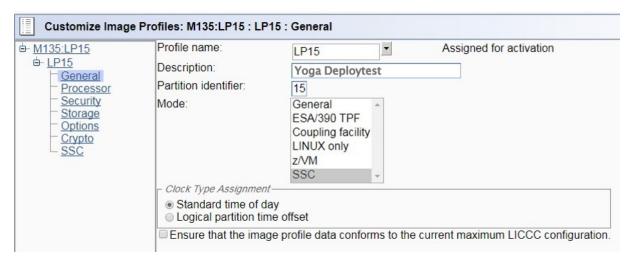


Figure 11. The Customize Image Profiles window of the HMC

- 5. In the navigation tree on the left, click **SSC** to go to that section.
- 6. In the Configure Management Network window, provide the following details:

Option	Description
CHPID	According to your environment, for example 10
Port	0
VLAN ID	Any available ID.
	Important: In a multi-node setup, all nodes must have the same VLAN ID.
IP address type	Static IPv4
IP address	A valid, available IP address in the management network.
	Important: In a multi-node setup, all IP addresses must be in the same subnet (see description of Mask field).
Mask	A subnet mask that is valid in the given environment, for example 22.
	Important: In a multi-node setup, all nodes must have the same subnet mask.

Restriction: Do not use the address range of the 172.17.0.0/16 network because this network is used internally by the product.

7. Click OK.

The **SSC** section looks as follows:

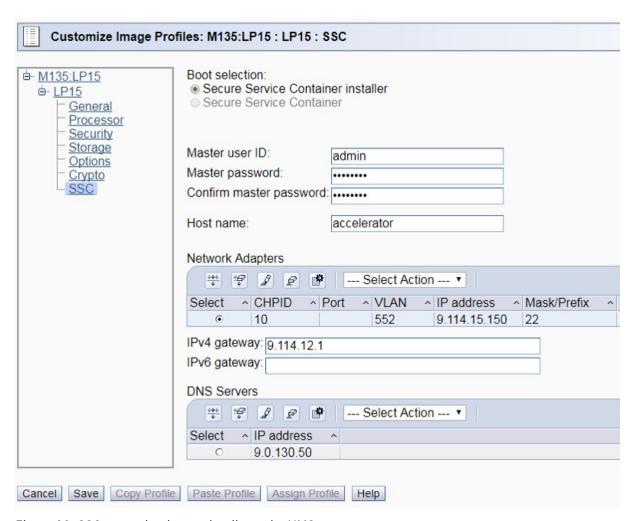


Figure 12. SSC network adapter details on the HMC

Note: The screen capture above shows a value in the IPv4 gateway field. This is the IP address of the default gateway (router) for access from and to other sub-nets. In most cases, this gateway cannot be used to access Db2 for z/OS subsystems in different sub- nets. For this purpose, use the static_routes option in the configuration file, which is introduced in the upcoming chapters.



Attention: Do not define more than one network adapter for the SSC LPAR. If more than one definition is present, your Db2 Analytics Accelerator on Z installation will fail eventually.

8. Provide a **Master user ID** and a **Master password**.

You will need this later when you start the Appliance Installer to configure and start the appliance (accelerator).

If you are configuring multiple LPARs for a multi-node environment, also refer to Figure 27 on page 91 to better understand the role of the LPAR passwords.

9. Click Save.

You return to the Partitions tab.

10. Right-click the LPAR again. This time, select **Daily** > **Activate**.

The system confirms the activation by displaying a window labeled Operating System Messages:

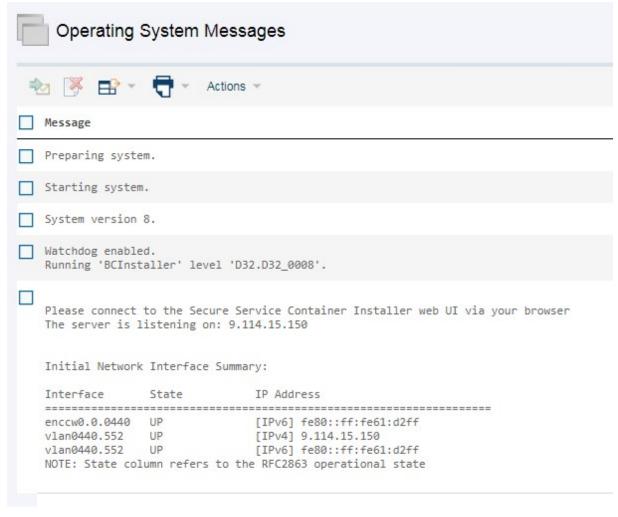


Figure 13. Confirmation: Operating System Messages

Logging on to the Admin UI

To complete your Db2 Analytics Accelerator on Z configuration, you must log on to the administration user interface (Admin UI). The logon steps are the same for all configurations.

About this task

Restriction: For the following steps, you need Mozilla Firefox or Google Chrome. Other browsers are not supported.

Procedure

- 1. Enter the IPv4 address of the Secure Service Container (SSC) LPAR (see Figure 12 on page 57).
- On the Login page, log in with the Master user ID and Master password that you specified as you defined the SSC LPAR. (see Figure 12 on page 57)

At the very first logon, your browser might display a certificate warning like the following:



Your connection is not secure

The owner of 9.152.151.197 has configured their website improperly. To protect your information from being stolen, Firefox has not connected to this website.

Learn more...

Go Back	Advanced
Report errors like this to help Mozilla identify and block mal	icious sites
9.152.151.197 uses an invalid security certificate.	
The certificate is not trusted because it is self-signed.	
Error code: SEC_ERROR_UNKNOWN_ISSUER	
Add Exception	

Figure 14. Browser certificate warning

This is because valid certificates require a trusted issuer (certificate authority or CA) and must have been defined before the domain is accessed. The Admin UI uses a non-trusted, self-signed certificate because specific customer domains are not known at the time the certificate is being created.

- 3. Click **Add Exception** (or whatever this button might be labeled in your browser) to accept the certificate and permit the connection.
- 4. If you use the Admin UI for the first time, you see the **License Agreement Page**. Click **I Agree** to continue.
- 5. On the **Install Software Appliance** page, upload the package (compressed file) that contains the appliance.

Note: The package is several GB in size, so the upload might take considerable time if the network between the source and the target is limited (for example, if the network is accessed from a remote location).

You also have to make sure that the target disk is large enough. Additional space is needed when the appliance is configured after the upload. Make sure that you meet the requirements, depending on your chosen setup type.

Single-node (with or without GDPS):

See Table 1 on page 61.

Multi-node (with or without GDPS):

See the Before you begin section.

Click **Browse** to navigate to the download location of the package file.

6. Select the target disk from the **Target Disk on Server** drop-down list.

The disk must be reachable from the previously defined LPAR.

Your page should look similar to the following:

Install Software Applian	Ce
To use a Software Appliance you can upload a Software Appliance.	in image file from the local machine to a target disk on the server or attach a disk with an already installed
Upload image to target disk	
Attach existing disk	
Local Installation Image	7.5.2.0-Information_Manage
	Image Details
	Name: Db2 Analytics Accelerator for z/OS Version: 7.5.2 Description: Db2 Analytics Accelerator Appliance Size: 6.78GB
Target Disk on Server	Device Type FICON DASD FCP
	Disk 0.0.9c09 (3390/0e)
y.	
	Cancel Apply
	₹ Sec
	Ser

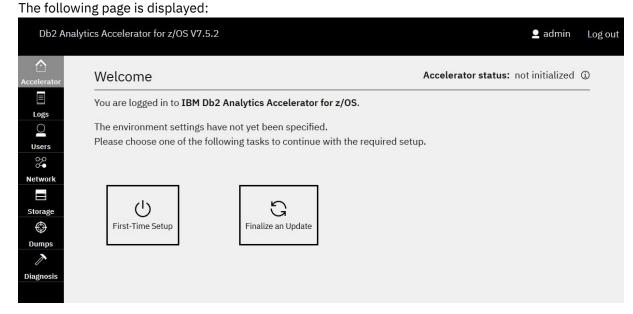
Figure 15. The Install Software Appliance window

- 7. When finished, click **Apply**.
- 8. In the **Confirm Appliance Installation** window, leave the **Reboot automatically ...** check box selected and click **Yes**.

Note: During the reboot, the browser might show 401 authorization errors. These errors can be ignored.

After a successful completion of the installation, you see the **Login** page of the appliance user interface.

9. Enter your **Master user ID** and **Master password** once more to log in.



- Figure 16. The **Welcome** page of the Admin UI
- 10. On the Welcome page, click First-Time Setup.

You see the following page:

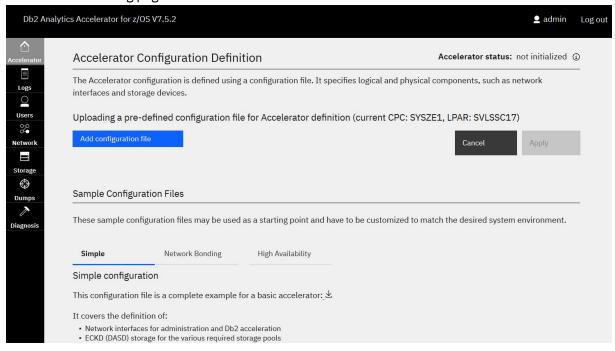


Figure 17. Accelerator Configuration Definition page

What to do next

Continue with the steps for your particular configuration. You can choose between the following setup types:

- Single-node setup
- Multi-node setup
- · Single-node setup with GDPS failover support
- Multi-node setup with GDPS failover support

Installing and starting the appliance (single-node setup)

Follow these steps after the accelerator (SSC) LPAR definition for a single-site setup of Db2 Analytics Accelerator on Z.

Before you begin

Storage must be provided for various storage pools. <u>Table 1 on page 61</u> provides an overview and minimum size recommendations.

Table 1. Recommended pools and pool sizes					
Pool	JSON key	Description	Size		
Appliance data pool	"data_devices"	Accelerator database including accelerator-shadow tables, accelerator-only tables, temporary database space, and temporary result space.	> 200 GB plus 90 percent of uncompressed Db2 for z/OS data (including accelerator-only tables) plus 20 percent of the accelerator (SSC) LPAR memory.		

Table 1. Recommended pools and pool sizes (continued)					
Pool	JSON key	Description	Size		
Appliance operation	"boot_device"	Size of the appliance image, which was defined and used during the image file upload step ("Install Software Appliance").	> 40 GB Important: This must be a single disk.		
Appliance runtime	"runtime_devices"	Temporary storage space required for the appliance runtime (container execution environment)	> 80 GB		

About this task

Restriction: For the following steps, you need Mozilla Firefox or Google Chrome. Other browsers are not supported.

Procedure

- 1. Define a dedicated Secure Service Container (SSC) LPAR as described in "Defining an LPAR for Db2 Analytics Accelerator on Z" on page 55.
- 2. Log on to the Admin UI and proceed to the **Welcome** page. For a description, see "Logging on to the Admin UI" on page 58.
- 3. On the **Welcome** page, click **First-Time Setup**.

You see the following page:

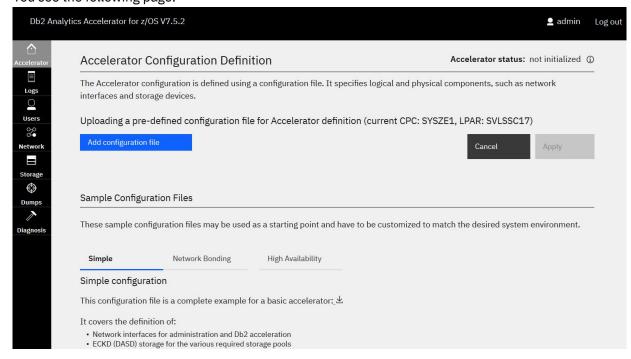


Figure 18. Accelerator Configuration Definition page

4. Starting with product version 7.1.9, all configuration settings are made by uploading a configuration file in JavaScript Object Notation (JSON) format.

A sample configuration file is provided at the bottom of the **Welcome** page. Under the heading **Sample Configuration Files**:

a. Click the **Simple** tab.



Note: Depending on your settings, some browsers will display the JSON file rather than download it. In that case, click **File** > **Save Page As** and save the file with an extension of . json.

The sample configuration file is stored in the Download folder of your web browser.

Important: Compared with previous product releases, the configuration file delivered with product version 7.5.3 and later versions has changed considerably.

When you open this file in your web browser, it looks as follows:

```
JSON
        Raw Data
                    Headers
Save Copy Collapse All Expand All Trilter JSON
 version:
                                 "7.5.3"
 accelerator_name:
                                 "igor01"
 accelerator_description:
                                 "SAMPLE-single-node"
                                 "single-node"
 accelerator_type:
                                 "10.1.1.101/24"
 db2_pairing_ipv4:
metwork_interface_bindings:
                                 "activation-profile"
    mgmt_nw:
    db2_nw:
                                 "osa@Ap@"
runtime environments:

→ 0:
                                 "CPC001"
      cpc_name:
                                 "IGOR001"
      lpar_name:
    metwork_interfaces:
       w 0:
                                 "osa@Ap@"
           name:
           device:
                                 "0.0.0a00"
           port:
v storage_environments:
  ₹ 0:
    ▼ boot_device:
```

Figure 19. Extract of the sample configuration file (JSON format), displayed in web browser

- 5. Open this file in a text editor of your choice and modify the settings according to your needs.
 - An editor capable of validating JSON files is recommended because the configuration file must be valid JSON. If it cannot be parsed correctly, you will run into errors. Valid JSON means:
 - Quotes are required around attribute values, even if these are plain numbers.
 - Colons must be used to separate attribute names from their values.
 - Object definitions consisting of key/value pairs must be enclosed in curly braces.
 - Arrays or lists must be enclosed in square brackets.

For your reference, take a look at the code of the sample configuration file:

```
"version": "7.5.6",
"accelerator_name": "igor01",
"accelerator_description": "SAMPLE-single-node",
"accelerator_type": "single-node",
"db2_pairing_ipv4": "10.1.1.101/24",
```

```
"network_interface_bindings": {
       "db2_nw": "osa0Ap0",
"mgmt_nw": "activation-profile"
   },
"runtime_environments": [
          "cpc_name": "CPC001",
"lpar_name": "IGOR001"
          "network_interfaces": [
                "device": "0.0.0a00",
"name": "osa0Ap0",
"port": "0"
             3
         ]
      }
    "storage_environments": [
          "boot_device": {
    "type": "dasd",
    "device": "0.0.5e29"
          "runtime_devices": {
    "type": "dasd",
              "devices": [
"0.0.5e25"
                 "0.0.5e26"
          data_devices": {
 "type": "dasd",
              "devices": [
                 "0.0.5e14",
                    "0.0.5e80"
                    "0.0.5e8f"
             ]
          3
      },
          "boot_device": {
    "type": "dasd",
    "device": "0.0.1b11"
          "runtime_devices": {
   "type": "dasd",
   "devices": [
       "0.0.1625",
                 "0.0.1b26"
             ]
          },
"data_devices": {
  "type": "dasd",
  "initials": [
              "0.0.1d00",
"0.0.1d0f"
             ]
         }
      }
  ]
3
```

Most attributes in the configuration file are required.

"version" (required)

The version of the accelerator. The version in the configuration file must match the version of the accelerator exactly. Do not change the value so that it deviates from the version of the accelerator.

"accelerator_name" (required)

The name of the accelerator. This attribute is used to identify the accelerator in the Admin UI and in log and trace files written to the logs/dumps/traces/panels directory. You can change the name online. A reset or restart is not required.

"accelerator_description" (optional)

Optional text description. You might want to add some information about the accelerator, especially if it is helpful for an administrator. You can change the description online. A reset or restart is not required.

"accelerator_type" (required)

As the name suggests, the type of the accelerator. Set this attribute to the value "single-node". This means that the runtime and storage environments you define later contain definitions for a single LPAR and a single storage environment only. A single-node accelerator can use up to 40 Integrated Facilities for Linux (IFLs) and 3 TB of memory.

Important: Once this parameter has set for and used by an accelerator, it cannot be changed anymore. If you want to switch to a multi-node installation, you must provide an entirely new configuration. Also, if you switch from one mode to the other, the existing data is not migrated or preserved.

"db2_pairing_ipv4" (required)

The IP address used to pair your Db2 subsystem with the specified accelerator. This IP address uniquely identifies the accelerator and is used by Db2 for z/OS to connect to the accelerator.

You can change the value online, that is, without a reset or restart. However, since the pairing address is stored in the Db2 for z/OS catalog during a pairing, you also have to (manually) correct the value in the catalog. Otherwise, your accelerator will not work anymore.

You can specify a netmask as part of the IPv4 address, like /24 for a subnet with 256 addresses. For example:

```
"db2_pairing_ipv4": "10.101.31.172/24"
```

This specifies the IP address 10.101.31.172 as the identifier of a subnet that comprises the address range from 10.101.31.0 to 10.101.31.255.

"temp_working_space" (optional)

The accelerator needs temporary storage for extensive sort operations that cannot be executed exclusively in the system memory. Increase the size if certain operations cannot be completed because the temporary storage and the system memory do not suffice. For example:

You can set the "temp_working_space" parameter to the following values:

"automatic"

This is the default, that is, the value used if you omit the parameter. This setting results in the creation of a database-managed (DMS) table space. That is, the size is controlled by the Db2 for z/OS database manager. The size will be the smaller of the following two values:

- 80 percent of the LPAR memory
- 50 percent of the free space of the data pool

"unlimited"

This setting will use all available space if necessary. However, it is not recommended because it might lead to a disk overflow, which causes subsequent errors.

"none"

Just a very small temporary workspace is used. A query that includes an extensive sort which leads to a memory overflow will fail immediately.

Fixed size

A DMS table space (see "automatic") with a size you determine. For example:

```
"temp_working_space": "500 GB"
```

The specified size is reserved for the temporary work space and serves as a size limit at the same time.

A change of the setting requires a restart of the accelerator. Changes do not take effect before the restart. If you use the setting "automatic", the size of the DMS table space is re-adjusted

dynamically after each start of the accelerator, according to the memory and disk usage rates that are encountered.

If the free storage space has shrunk during the time from one restart to the next, a smaller temporary workspace is provided after the latest restart. If you were close to 100 percent storage usage before a restart, the temporary workspace might be reduced so much during the restart that certain jobs cannot be executed anymore. In that case, consider adding storage to the accelerator. This can be done while your system is online.

Important: Because the self-adjusting size mechanism does not exist for older versions, the setting "temp_working_space": "unlimited" is added during a configuration upgrade to version 7.5.6 or later. If you export an upgraded JSON configuration, it will contain this setting, rather than "temp_working_space": "automatic". If you re-apply the exported configuration, the unlimited temporary workspace might lead to a disk overflow, as indicated in the description of this value.

"network_interface_bindings" (required)

In the network interface bindings section, you map physical interfaces to network names that can be specified in a runtime environment. The settings in this section cannot be changed online. If you need to change them, you must restart the Db2 Analytics Accelerator on Z accelerator.

"mgmt_nw"

This network, which is used by the Admin UI and other support interfaces, is defined by the HMC activation profile of the accelerator (SSC) LPAR. It is not part of the Db2 Analytics Accelerator on Z configuration. Therefore, use the attribute value "activation-profile". Note that the name of the management network might change if someone updates the activation profile of the accelerator (SSC) LPAR on the HMC.

"db2 nw"

This network name points to the IP address of your Db2 subsystem (counterpart of the "db2_pairing_ipv4"). It is used during the pairing process, and all network traffic between your Db2 subsystem and the accelerator will run through this interface.

The attribute value must be the same as one of the "name: " attributes in your "network_interfaces" definitions further down in the configuration file. The value must be an alphanumeric character string no longer than 8 characters. Compare this with the sample code.

In the network interfaces: section, which is described below, you find the details of all networks, including the network used for the pairing process. In the example, this is the network device with the ID 0.0.4b00.

"gdps_nw"

This network interface is used for GDPS failover support. The value is the interface name of the alternative data network in a failover scenario. You can omit this value if you do not use GDPS.

An example of the "network_interface_bindings" block:

```
"network_interface_bindings": {
   "db2_nw": "my_db2_network",
   "gdps_nw": "my_gdps_network",
   "mgmt_nw": "activation-profile"
```

Notes:

- Multiple interfaces can use the same physical connection. For example, in a GDPS setup, the interface name specified as the value of "gdps_nw" might be the same as the value of "db2_nw". In this case, the same physical connection would be used by both interfaces, and only one network definition would be required in the "network_interfaces" block further down.
- Do not use the value "activation-profile" for any network interface other than the "mgmt_nw" because the activation profile might change in the HMC. Such a change might have

repercussions if a network that refers to the activation profile is used for other purposes. The change might make your network definition unusable.

"runtime_environments" (required)

This block defines the network interfaces and other LPAR-specific settings of the accelerator (SSC LPAR). Each of these LPARs is identified by the CPC name and the LPAR name. A set of networks must be defined for each LPAR. This is usually the Db2 network, and, optionally, an additional network to reach the GDPS server. The definition of a runtime environment is also required for GDPS setups because GDPS controls the storage configuration, but not the configuration of accelerator LPARs. Specify the following attributes to identify an accelerator (SSC) LPAR:

"cpc_name"

The name of the CPC as defined on the Hardware Management Console (HMC).

You cannot change this name online. If you must change it because the CPC name changes or because you want to move the accelerator to a different CPC, first create an additional runtime environment that contains the new value. Then shut down the accelerator, change the old value and reactivate the accelerator.

"lpar_name"

The name of the accelerator (SSC) LPAR as defined on the HMC.

You cannot change this name online. If you must change it because the LPAR name changes or because you want to move the accelerator to a different LPAR, first create an additional runtime environment that contains the new value. Then shut down the accelerator, change the old value and reactivate the accelerator.

Continue with the network interfaces for the accelerator (SSC) LPAR:

"network_interfaces" (required)

This keyword defines the physical network interfaces (OSA or hipersocket) that are used by a runtime environment. Each network name defined in the "network_interface_bindings" section must be mapped to a physical network interface in the corresponding runtime environment. You can change all of the definitions in this block online. The following attributes must be specified for each physical network:

"name" (required)

The name of the network interface in the "network_interface_bindings" section.

"ipv4"

The IPv4 address and the subnet used for an interface, for example:

```
10.4.1.101/16
```

Important: The IPv4 address that the "db2_nw" network uses is specified as the "db2_pairing_ipv4" address. Therefore, do not specify an additional "ipv4" address for the "db2_nw" network.

"device" (required)

The identifier of an OSA-Express card or a Hipersocket.

"port" (optional)

The network port to be used. If this value is omitted, the port number defaults to "0". Most OSA cards have a single physical port "0".

"vlan" (optional)

If a virtual LAN (VLAN) has been defined for the accelerator (SSC) LPAR and you want to use this VLAN as an interface for Db2 Analytics Accelerator on Z, you can specify the VLAN name here.

Example:

```
"device": "0.0.4b00",
            "name": "my_db2_network",
"port": "1",
"vlan": "600"
         }
     ]
  3
]
```

In this example, you find a network interface definition for an accelerator (SSC) LPAR named SVLSSCOC. This LPAR runs in a CPC named SYSZE1. The LPAR has a network device 0.0.4b00 (an OSA-Express card or a Hipersocket). The name that Db2 Analytics Accelerator on Z uses for this network device is "my_db2_network".

"static_routes" (optional)

This option is used to define additional network routes for an interface.

If the IP address of a Db2 for z/OS LPAR or the GDPS keys LPAR is in a different subnet than the IP address assigned to the accelerator, an additional route definition is needed to establish the connection. An additional static route also helps to avoid undesired network traffic through a default gateway, which might have been defined in the HMC activation profile of the accelerator (SSC) LPAR.

"ipv4"

The IPv4 address or the IPv4 address and subnet of the target network.

"via"

The IPv4 address of the routing device.

Example: The accelerator's pairing IP address is 10.20.1.33/24 and there are two Db2 for z/OS LPARs with the IP addresses 10.1.1.47/24 and 10.1.1.48/24.

One or more gateways connect both subnets. One gateway is accessed through IP address 10.20.1.1, the other through 10.1.1.1.

To allow traffic from one network to the other, the TCPIP.PROFILE definition in z/OS defines a route to 10.20.1.0/24, which uses the gateway 10.20.1.1. The accelerator uses the following configuration to enable traffic to the 10.1.1.0 network using the corresponding gateway at 10.1.1.1:

```
"accelerator_name": "S1",
"db2_pairing_ipv4": "10.20.1.33/24",
"network_interface_bindings": {
   "db2_nw": "db2_nw_connection",
"mgmt_nw": "activation-profile"
},
"runtime_environments": [
      "network_interfaces": [
            "name": "db2_nw_connection",
            "device": "0.0.0440",
"vlan": "552",
             "static_routes": [ { "ipv4": "10.1.1.0/24", "via": "10.20.1.1" } ]
      ]
   3
],
```

This way, all traffic to an IPv4 address that starts with 10.1.1 uses the OSA-Express card with device ID 0.0.0440 via gateway 10.20.1.1. All network traffic between the accelerator and destinations in the 10.1.1.0/24 subnet is thus bound to that OSA device.

"bond_settings" (optional)

This attribute allows you to define several network cards (OSA-Express cards) as a single device. Bonding is usually employed in a high-availability setup, as the remaining network cards in the setup can take over if one network card fails. It is also possible to run all available network cards simultaneously.

Example:

In this example, two OSA cards (devices 0a00 and 1b00) are combined to one bonding device called "db2_nw_connection". The device works in "active-backup" mode, meaning that at any time, just one of the network cards is active. The other card takes over when the active card fails.

You can alternatively specify "mode": "802.3ad", in which case all network cards of the device will be active at the same time. "802.3ad" stands for the IEEE 802.3ad link aggregation mode.

In 802.3ad mode, you need at least two physical devices. Specify these in the same way as you specify the devices for active-backup mode. That is, use a "slaves" list as shown in the previous example.

"options" (optional)

It is not necessary to specify "options" for "bond_settings". If the options are omitted, default values are used. Whether options apply to a particular setup depends on the selected mode ("active-backup" or "802.3ad"). For a detailed description of these options, see *Chapter 7. Configure Network Bonding* in the *Red Hat® Enterprise Linux 7: Networking Guide*. A link is provided at the end of this topic.

Restriction: Currently, you cannot specify just a subset of the available options. You either have to specify no options at all, in which case default values are used, or specify all options pertaining to a particular mode.

All of the following values can be changed online.

"primary"

Valid in *active-backup* mode only. The first physical device to be used. This is "0.0.0a00" according to the previous example. The primary device is the first of the bonding interfaces. It will be used as the active device unless it fails.

"primary-reselect": "always"

Valid in *active-backup* mode only. Determines how the active physical device is selected after a failure. Specify "always", which means that an attempt will be made to make the first physical device (labeled "primary") active again.

Other allowed options are "better", which means that the fastest device will be used as the active device, or "failure", which means that the active physical device is only changed if the currently active device fails.

"failover-MAC": "none"

Valid in *active-backup* mode only. Allows you set all physical devices to the same MAC address or determine these addresses according to a policy. Specify the value

"none", which means that the same MAC address will be used for all physical devices.

"no-gratuitous-ARPs": "0"

Valid in *active-backup* mode only. Determines the number of peer notifications after a failover event. Specify "0", which means no notifications. This option corresponds to the num_grat_arp or num_unsol_na option in the *Red Hat Enterprise Linux 7: Networking Guide*.

"transmit-hash-policy": "layer2"

Valid in 802.3ad mode only. Selects a policy according to which the MAC addresses of the physical devices are determined. Specify layer2, which means that traffic to a particular network peer is assigned to the same network device, which is determined solely by its MAC address. Other allowed options are "layer3+4" and "layer2+3". The option "layer3+4" means that multiple network devices can be used to reach a single network peer even if a single network connection does not span multiple network devices. The option "layer2+3" is similar to "layer 2", but the network device is selected by its IP address in addition to its MAC address.

Note: In the *Red Hat Enterprise Linux 7: Networking Guide*, this option is called "xmit-hash-policy"

"LACP-rate": "slow" | "fast"

Valid in 802.3ad mode only. The rate at which physical devices transmit Link Aggregation Control Protocol Data Units (LACPDUs). Specify "slow", which means every 30 seconds, or "fast", which means every 1 second.

"link-monitoring": "MII"

Selects the method to be used for monitoring the physical device's ability to carry network traffic. Select "MII", which stands for *media-independent interface*. With this setting, the driver, the MII register, or the ethtool can be queried for monitoring information about a physical device. Alternatively, you can specify "ARP" to use the ARP monitor.

"monitoring-frequency": "100"

The time interval that passes between two monitoring events. It is an integer value that stands for milliseconds. Use a value of "100".

"link-up-delay": "0"

Delay that needs to pass before network traffic is sent to a physical device after link monitoring has reported the device to be up. Specify "0", which means *no delay*.

"link-down-delay": "0"

Delay that needs to pass before network traffic is routed to the failover device after link monitoring has reported the failure of the previously active device. Specify "0", which means no delay.

Example (active-backup mode):

```
"options": {
   "primary": "0.0.0a00",
   "primary-reselect": "always",
   "failover-MAC": "none",
   "no-gratuitous-ARPs": "0",
   "link-monitoring": "MII",
   "monitoring-frequency": "100",
   "link-up-delay": "0",
   "link-down-delay": "0",
}
```

Example (802.3ad mode):

```
"options": {
    "LACP-rate": "slow",
    "transmit-hash-policy": "layer2",
    "link-monitoring": "MII",
    "monitoring-frequency": "100",
    "link-up-delay": "0",
```

```
"link-down-delay": "0"
}
```

"zfcp_devices" (required if ZFCP drives are used)

The FICON® Express ports of your devices must be listed in the "runtime_environments" section. This is required for ZFCP storage devices only, as the ports of ECKD devices are handled by the firmware of the CPC. For ZFCP devices, however, you must list the port names. See the following example:

The accelerator uses multiple paths, that is, it tries to use all specified ports. For that reason, a list of different ports rather than just one port can increase the performance.

You can use a single port identifier for a ZFCP device or for an ECKD device, but not for both.

You can change the ZFCP port names online.

"storage_environments" (required)

This block lists all storage devices, that is, disks or disk enclosures. At a minimum, this block contains the name of the storage on which the accelerator is initially deployed. You can define additional storage for mirroring or failover purposes. Note, however, that the accelerator does not manage the mirroring or copying of data to additional devices. The "storage_environments" section combines the "primary_storage" and "storage_maps" sections found in configuration files of earlier releases. During the first-time deployment, these devices are formatted, which means that the existing data on these devices is erased.

If you use the GDPS integration, GDPS will manage the storage environments. In this case, you need to define the "primary_storage" environment only.

The accelerator uses three categories of storage: the boot device, the runtime data pool, and the data pool for operative data. Each storage device or storage pool can use DASD (ECKD) or ZFCP (SCSI) devices. The mixing of different device types is not supported. However, the size of individual devices in a pool is not restricted.

Important: You can add devices to a storage pool at any time while the accelerator is online. New devices are automatically integrated into a storage pool. This does not require a reset or a restart. However, you cannot remove devices from a storage pool in this way. A removal requires an entirely new installation.

You must define these storage devices by using the following attributes in the configuration file:

"boot_device" (required)

The boot device contains the software image that is written by the Secure Service Container (SSC) installer. The accelerator will be started from this device. It is also the target device for uploading the SSC installer image before the initial deployment or before an update. The boot device must be a single device with at least 40 GB net storage capacity. For example:

```
"boot_device": {
   "type": "dasd",
   "device": "0.0.5e29"
}
```

```
"boot_device": {
    "type": "zfcp",
    "udid": "6005076241bb50242000000000000008e"
```

"runtime_devices" (required)

The runtime storage is used by the accelerator software for internal processing. It does not contain user data and its size is fixed. The size does not depend on the amount of user data processed by the accelerator. Specify a list of devices with a total net capacity of at least 80 GB.During normal operation, the utilization rate should not exceed 80 percent. If it does exceed 80 percent most of the time, consider adding devices. For example:

```
"runtime_devices": {
   "type": "dasd",
"devices": [
    "0.0.998c"
```

The mixing of different device types in a single list, that is, DASD (ECKD) and ZFCP (SCSI) is not supported.

"data_devices" (required)

The data storage is used to store the user data and temporary data of the accelerator (table data). It is typically the largest storage area of your entire configuration. Its size is determined by the amount of data the accelerator has to handle. During normal operation, the utilization rate should not exceed 80 percent. If it does exceed 80 percent most of the time, consider adding devices. For example:

```
"data_devices": {
  "type": "dasd",
"devices": [
     "0.0.9c00"
     "0.0.9c01"
     "0.0.9c02"
3
```

"type" (required)

This is the type of storage to be used (disk type). Possible values are "dasd" for extended count key data (ECKD) volumes and "zfcp" for Small Computer System Interface (SCSI) volumes. You must specify the type for each device category (that is, the boot device, the runtime device, and the data device).

Important:

- It is not possible to mix ECKD and SCSI devices in a single device category or device pool.
- If you use DASD (ECKD) storage, HyperPAV aliases are strongly recommended because they increase the processing speed.
- DASD (ECKD) devices are formatted by the **dasdfmt** program of the Linux operating system on the accelerator. This can take a long time, sometimes even hours for large devices or storage pools. HyperPAV aliases also help speed up formatting. Therefore, define HyperPAV aliases also in your initial JSON configuration file.
- You can use DASD (ECKD) devices of different sizes in a single pool.
- You can use ZFCP devices of different sizes in a single pool. However, the use of devices smaller than 5 TB and larger than 10 TB is not recommended.
- For ZFCP devices, use Ficon Express (FE) ports because these ports guarantee a much better performance and availability. FE ports are defined in the runtime environment.
- Although the adding of devices is supported while the accelerator is online, the type cannot be changed after the initialization of the storage pool. To change the type, you must remove the accelerator and reinstall it.

Example (ECKD or "dasd"):

```
"type": "dasd",
"devices": [
    "0.0.9c00",
    "0.0.9c01",
    "0.0.9c02"
]
```

Example (SCSI or "zfcp"):

```
"type": "zfcp",
"udids": [
   "0c984712545423523614b8d812345632",
   "0c0a0b5c9d15555a545545b46456c4d6"
]
```

"device" or "devices" (required)

This attribute is used to list the devices by their names or identifiers. You must specify a device or a list of devices for each device category (that is, the boot device, the runtime device, and the data device).

ECKD example:

ZFCP example:

Note: To change the ID of the boot device, a few extra steps are required:

- a. Copy the current storage environment in your JSON file. That is, create a duplicate block in the file.
- b. Change the ID in the copied block.
- c. Restart the accelerator.
- d. Remove the old storage environment from the JSON file when the accelerator is online again.

HyperPAV aliases:

HyperPAV aliases can be used in the following ways:

- · Automatic HyperPAV aliases
- · Explicitly listed HyperPAV aliases

In automatic HyperPAV mode, all HyperPAV alias devices that are visible to an LPAR and that are connected to the same control-unit image (LCU) are used automatically for that LPAR. To enable the automatic HyperPAV mode, you must add a definition to the storage environments section in the JSON configuration file.

Important:

- HyperPAV aliases can be used with DASD (ECKD) storage only.
- Make sure that only the volumes and HyperPAV aliases you want to use on a particular LPAR are
 visible to that LPAR. This is even more important if you use automatic alias devices because in that
 case, your accelerator (SSC) LPAR has to sift through all visible devices just to determine and activate
 the alias devices.

Example:

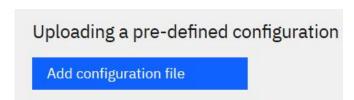
Mind that in "auto" mode, the system uses all available HyperPAV alias devices.

To use just a subset of the available devices, it is preferable to use explicitly listed HyperPAV aliases. as in the following example:

In this particular case, the system uses a range of HyperPAV aliases from 0.1.4000 to 0.1.4007 plus a single HyperPAV alias with the ID 0.1.1234.

6. When you're finished with your configuration file, upload it to the Admin UI.

On the **Accelerator Configuration Definition** page, under the heading **Uploading Configuration File**, click the upload button (see also Figure 18 on page 62):



If something is wrong with the file you uploaded, an error message is displayed on the page:

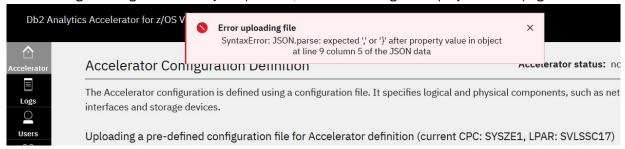


Figure 20. Error message after uploading a faulty configuration file

7. If errors occurred, fix these and repeat the upload (steps "5" on page 63 and "6" on page 74).

If no errors occurred, the Accelerator Configuration Definition page shows the settings of your configuration file in a folder structure. You can expand the folders to display the settings they contain by clicking the plus signs in front of the folder icons.

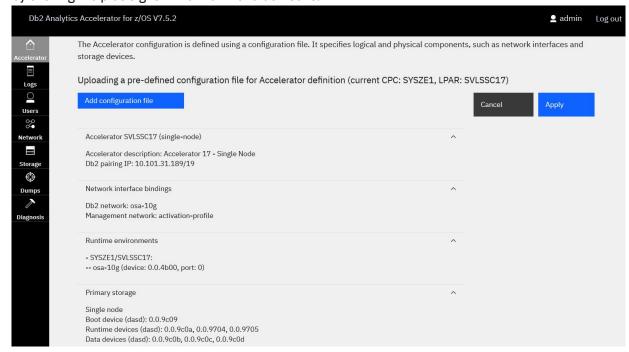


Figure 21. Accelerator Configuration Definition after a successful configuration file upload

8. Click Apply.

You see a message window indicating that the configuration is in progress. During that time, you can follow the line of events if you click the Logs tab. The page does not refresh automatically, but you can click the **Refresh** button in the upper right corner.



Figure 22. Message window showing the progress of your Db2 Analytics Accelerator on Z configuration

Results

When these processes have finished, the **Accelerator Components Health Status** page is displayed automatically. The page should now give you the following information:

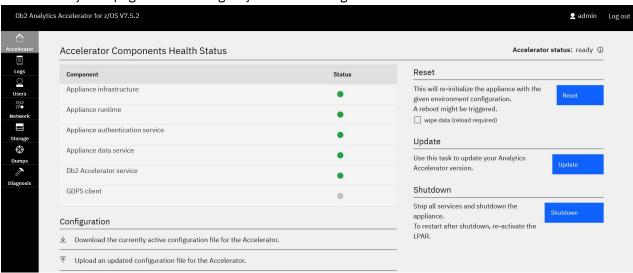


Figure 23. The Accelerator Components Health Status page is displayed after a successful configuration

The message Accelerator status: ready on the top right indicates that all installation steps have been completed and that components have been started for Db2 Analytics Accelerator on Z.

Installing and starting the appliance (multi-node setup)

Follow these steps after the accelerator (SSC) LPAR definition for the multi-node setup of Db2 Analytics Accelerator on Z.

Before you begin

The following requirements must be met for a multi-node installation:

- Exactly six accelerator (SSC) LPARs, organized in an LPAR group (cluster)
- At least 1.5 TB of main memory.

- At least 7 network connection points that can be reached from the outside of the Hipersocket network (through OSA-Express cards).
- At least 18 individual disk drives (FICON (ECKD) or FCP (SCSI):
 - 40 GB or more are required for each of the six boot disks in the cluster.
 - 80 GB or more are required for each runtime disk in the cluster. Six or a multiple of six runtime disks can be specified.
 - Six or a multiple of six disks can be used for the data pool. Each data disk needs as much storage as the data pool of a single-node installation.
- An absolute limit must be set for the number of the Integrated Facilities for Linux (IFLs) that the cluster can use, for example 40 IFLs. The minimum is 30 IFLs; the maximum is 190 IFLs.
- The accelerator (SSC) LPARs must belong to the same Central Processing Complex (CPC).
- The accelerator (SSC) LPARs must be connected by a Hipersocket network.
- Every accelerator (SSC) LPAR in the cluster must be connected to a management network.
- One accelerator (SSC) LPAR is used for the head node.

The head node is the controlling node. Externally, it communicates with the networks outside of the cluster. It is paired with one or more Db2 subsystems; it connects to the management network defined in the HMC activation profile, and optionally connects to GDPS servers. Internally, the head node communicates with the data nodes.

The head node must have access to 30 shared IFLs. The initial processing weight of these IFLs needs to be set to 100 percent in the HMC activation profile. Furthermore, the head node requires at least 256 GB of main memory.

• The other accelerator (SSC) LPARs are used for the data nodes.

The data nodes mostly communicate with the head node and with each other. Communication hardly leaves the Hipersocket network. Data nodes require a connection to a management network, but this connection is only used initially, when software is transferred from the head node, and for the collection of trace data.

The five data nodes claim 25 percent of the processing capacity of the shared IFLs. The initial processing weight of the IFLs must be set to 10 percent. The data nodes require between 256 GB and 4 TB of main memory.

About this task

Restriction: For the following steps, you need Mozilla Firefox or Google Chrome. Other browsers are not supported.

Procedure

- 1. Define six dedicated Secure Service Container (SSC) LPARs as described in <u>"Defining an LPAR for Db2</u> Analytics Accelerator on Z" on page 55.
- 2. Log on to the Admin UI and proceed to the **Welcome** page.
 - For a description, see "Logging on to the Admin UI" on page 58.
- 3. On the **Welcome** page, click **First-Time Setup**.

You see the following page:

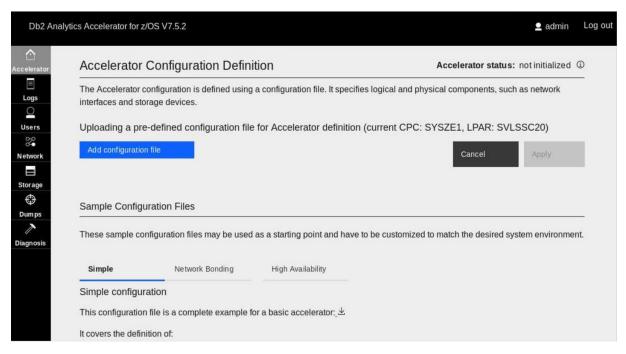


Figure 24. Accelerator Configuration Definition page

4. Starting with product version 7.1.9, all configuration settings are made by uploading a configuration file in JavaScript Object Notation (JSON) format.

Important: For a multi-node setup, a sample configuration file is not provided. This is because the setup is complex, and a sample file cannot cover all possibilities without becoming very confusing. Contact IBM support for help with the creation of a configuration file in JSON format.

- 5. Create a JSON file in a text editor of your choice and include settings as shown in the following steps. An editor capable of validating JSON files is recommended because the configuration file must be valid JSON. If it cannot be parsed correctly, you will run into errors. Valid JSON means:
 - Quotes are required around attribute values, even if these are plain numbers.
 - Colons must be used to separate attribute names from their values.
 - Object definitions consisting of key/value pairs must be enclosed in curly braces.
 - · Arrays or lists must be enclosed in brackets.

Most of the following attributes are required.

"version" (required)

The version of the accelerator. The version in the configuration file must match the version of the accelerator exactly. Do not change the value so that it deviates from the version of the accelerator.

"accelerator_name" (required)

The name of the accelerator. This attribute is used to identify the accelerator in the Admin UI and in log and trace files written to the logs/dumps/traces/panels directory. You can change the name online. A reset or restart is not required.

"accelerator_description" (optional)

Optional text description. You might want to add some information about the accelerator, especially if it is helpful for an administrator. You can change the description online. A reset or restart is not required.

"accelerator_type" (required)

As the name suggests, the type of the accelerator. Set this attribute to the value "multi-node". This means that the runtime and storage environments you define later contain definitions for six LPARs and six storage environments. A single processing node of a multi-node accelerator can use up to 30 Integrated Facilities for Linux (IFLs) and 1.5 TB of memory.

Important: Once this parameter has set for and used by an accelerator, it cannot be changed anymore. If you want to switch to a single-node installation, you must provide an entirely new configuration. Also, if you switch from one mode to the other, the existing data is not migrated or preserved.

"db2_pairing_ipv4" (required)

The IP address used to pair your Db2 subsystem with the head node of the specified accelerator. This IP address uniquely identifies the head node and is used by Db2 for z/OS to connect to the accelerator.

You can change the value online, that is, without a reset or restart. However, since the pairing address is stored in the Db2 for z/OS catalog during a pairing, you also have to (manually) correct the value in the catalog. Otherwise, your accelerator will not work anymore.

You can specify a netmask as part of the IPv4 address, like /24 for a subnet with 254 usable addresses. For example:

```
"db2_pairing_ipv4": "10.108.16.184/24"
```

This specifies the IP address 10.108.16.184 as the identifier of a subnet that comprises the address range from 10.108.16.1 to 10.108.16.254.

Tip: All IP addresses in the configuration file can point to a subnet.

"temp_working_space" (optional)

The accelerator needs temporary storage for extensive sort operations that cannot be executed exclusively in the system memory. Increase the size if certain operations cannot be completed because the temporary storage and the system memory do not suffice. For example:

You can set the "temp_working_space" parameter to the following values:

"automatic"

This is the default, that is, the value used if you omit the parameter. This setting results in the creation of a database-managed (DMS) table space. That is, the size is controlled by the Db2 for z/OS database manager. The size will be the smaller of the following two values:

- 80 percent of the LPAR memory
- 50 percent of the free space of the data pool

"unlimited"

This setting will use all available space if necessary. However, it is not recommended because it might lead to a disk overflow, which causes subsequent errors.

"none"

Just a very small temporary workspace is used. A query that includes an extensive sort which leads to a memory overflow will fail immediately.

Fixed size

A DMS table space (see "automatic") with a size you determine. For example:

```
"temp_working_space": "500 GB"
```

The specified size is reserved for the temporary work space and serves as a size limit at the same time.

A change of the setting requires a restart of the accelerator. Changes do not take effect before the restart. If you use the setting "automatic", the size of the DMS table space is re-adjusted dynamically after each start of the accelerator, according to the memory and disk usage rates that are encountered.

If the free storage space has shrunk during the time from one restart to the next, a smaller temporary workspace is provided after the latest restart. If you were close to 100 percent storage usage before a restart, the temporary workspace might be reduced so much during the restart that certain jobs cannot be executed anymore. In that case, consider adding storage to the accelerator. This can be done while your system is online.

Important: Because the self-adjusting size mechanism does not exist for older versions, the setting "temp_working_space": "unlimited" is added during a configuration upgrade to version 7.5.6 or later. If you export an upgraded JSON configuration, it will contain this setting, rather than "temp_working_space": "automatic". If you re-apply the exported configuration, the unlimited temporary workspace might lead to a disk overflow, as indicated in the description of this value.

"network_interface_bindings" (required)

In the network interface bindings section, you map physical interfaces to network names that can be specified in a runtime environment. The settings in this section cannot be changed online. If you need to change them, you must restart the Db2 Analytics Accelerator on Z accelerator.

"mgmt nw"

This network, which is used by the Admin UI and other support interfaces, is defined by the HMC activation profile of the accelerator (SSC) LPAR. It is not part of the Db2 Analytics Accelerator on Z configuration. Therefore, use the attribute value "activation-profile". Note that the name of the management network might change if someone updates the activation profile of the accelerator (SSC) LPAR on the HMC.

"db2 nw"

This network name points to the IP address of your Db2 subsystem (counterpart of the "db2_pairing_ipv4"). It is used during the pairing process, and all network traffic between your Db2 subsystem and the accelerator will run through this interface.

The attribute value must be the same as one of the "name: " attributes in your "network_interfaces" definitions further down in the configuration file. The value must be an alphanumeric character string no longer than 8 characters. Compare this with the sample code.

In the network_interfaces: section, which is described below, you find the details of all networks, including the network used for the pairing process. In the example, this is the network device with the ID 0.0.4b00.

"cluster nw"

The name of the Hipersocket network that connects your nodes.

The attribute value must be the same as one of the "name: " attributes in your "network interfaces" definitions further down in the configuration file.

In the network_interfaces: section, which is described below, you find the details of all networks, including the Hipersocket network. In the example, this is the network device with the ID 0.0.7f00.

An example of the "network_interface_bindings" block:

```
"network_interface_bindings": {
  "db2_nw": "my_db2_network",
"cluster_nw": "my_hiper",
"mgmt_nw": "activation-profile"
```

Note: Do not use the value "activation-profile" for any network interface other than the "mgmt_nw": because the activation profile might change in the HMC. Such a change might have repercussions if a network that refers to the activation profile is used for other purposes. The change might make your network definition unusable.

"runtime_environments" (required)

This block defines the network interfaces and other LPAR-specific settings of the accelerator (SSC LPARs). Each LPAR is identified by the CPC name and the LPAR name. A set of networks must be defined for each accelerator (SSC) LPAR. This is usually the Db2 network, and the Hipersocket network of your cluster. Specify the following attributes to identify an accelerator (SSC) LPAR:

The name of the CPC as defined on the Hardware Management Console (HMC).

You cannot change this name online. If you must change it because the CPC name changes or because you want to move the accelerator to a different CPC, first create an additional runtime environment that contains the new value. Then shut down the accelerator, change the old value and reactivate the accelerator.

"head"

The definition of the head node. For example:

"lpar_name"

The name of an accelerator (SSC) LPAR as defined on the HMC.

You cannot change an "lpar_name" online. If you must change it because the LPAR name changes or because you want to use different LPARs for the accelerator, first create an additional runtime environment that contains the new values. Then shut down the accelerator, change the old values and reactivate the accelerator.

Continue with the network interfaces for the accelerator (SSC) LPAR:

"network_interfaces" (required)

This keyword defines the physical network interfaces (OSA or hipersocket) that are used by a runtime environment. Each network name defined in the

"network_interface_bindings" section must be mapped to a physical network interface in the corresponding runtime environment. You can change all of the definitions in this block online. The following attributes must be specified for each physical network:

"name" (required)

The name of the network interface in the "network_interface_bindings" section

"ipv4"

The IPv4 address and the subnet used for an interface, for example:

```
10.4.1.101/16
```

Important: The IPv4 address that the "db2_nw" network uses is specified as the "db2_pairing_ipv4" address. Therefore, do not specify an additional "ipv4" address for the "db2_nw" network.

"device" (required)

The identifier of an OSA-Express card or a Hipersocket.

"port" (optional)

The network port to be used. If this value is omitted, the port number defaults to "0". Most OSA cards have a single physical port "0".

"vlan" (optional)

If a virtual LAN (VLAN) has been defined for the accelerator (SSC) LPAR and you want to use this VLAN as an interface for Db2 Analytics Accelerator on Z, you can specify the VLAN name here.

Example:

"static_routes" (optional)

This option is used to define additional network routes for an interface.

If the IP address of a Db2 for z/OS LPAR or the GDPS keys LPAR is in a different subnet than the IP address assigned to the accelerator, an additional route definition is needed to establish the connection. An additional static route also helps to avoid undesired network traffic through a default gateway, which might have been defined in the HMC activation profile of the accelerator (SSC) LPAR.

"ipv4"

The IPv4 address or the IPv4 address and subnet of the target network.

"via"

The IPv4 address of the routing device.

Example: The accelerator's pairing IP address is 10.20.1.33/24 and there are two Db2 for z/OS LPARs with the IP addresses 10.1.1.47/24 and 10.1.1.48/24.

One or more gateways connect both subnets. One gateway is accessed through IP address 10.20.1.1, the other through 10.1.1.1.

To allow traffic from one network to the other, the TCPIP.PROFILE definition in z/OS defines a route to 10.20.1.0/24, which uses the gateway 10.20.1.1. The accelerator uses the following configuration to enable traffic to the 10.1.1.0 network using the corresponding gateway at 10.1.1.1:

This way, all traffic to an IPv4 address that starts with 10.1.1 uses the OSA-Express card with device ID 0.0.0440 via gateway 10.20.1.1. All network traffic between

the accelerator and destinations in the 10.1.1.0/24 subnet is thus bound to that OSA device.

"bond_settings" (optional)

This attribute allows you to define several network cards (OSA-Express cards) as a single device. Bonding is usually employed in a high-availability setup, as the remaining network cards in the setup can take over if one network card fails. It is also possible to run all available network cards simultaneously.

Example:

In this example, two OSA cards (devices 0a00 and 1b00) are combined to one bonding device called "db2_nw_connection". The device works in "active-backup" mode, meaning that at any time, just one of the network cards is active. The other card takes over when the active card fails.

You can alternatively specify "mode": "802.3ad", in which case all network cards of the device will be active at the same time. "802.3ad" stands for the IEEE 802.3ad link aggregation mode.

In 802.3ad mode, you need at least two physical devices. Specify these in the same way as you specify the devices for active-backup mode. That is, use a "slaves" list as shown in the previous example.

"options" (optional)

It is not necessary to specify "options" for "bond_settings". If the options are omitted, default values are used. Whether options apply to a particular setup depends on the selected mode ("active-backup" or "802.3ad"). For a detailed description of these options, see *Chapter 7. Configure Network Bonding* in the *Red Hat Enterprise Linux 7: Networking Guide*. A link is provided at the end of this topic.

Restriction: Currently, you cannot specify just a subset of the available options. You either have to specify no options at all, in which case default values are used, or specify all options pertaining to a particular mode.

All of the following values can be changed online.

"primary"

Valid in *active-backup* mode only. The first physical device to be used. This is "0.0.0a00" according to the previous example. The primary device is the first of the bonding interfaces. It will be used as the active device unless it fails.

"primary-reselect": "always"

Valid in *active-backup* mode only. Determines how the active physical device is selected after a failure. Specify "always", which means that an attempt will be made to make the first physical device (labeled "primary") active again.

Other allowed options are "better", which means that the fastest device will be used as the active device, or "failure", which means that the active physical device is only changed if the currently active device fails.

"failover-MAC": "none"

Valid in *active-backup* mode only. Allows you set all physical devices to the same MAC address or determine these addresses according to a policy. Specify the value "none", which means that the same MAC address will be used for all physical devices.

"no-gratuitous-ARPs": "0"

Valid in *active-backup* mode only. Determines the number of peer notifications after a failover event. Specify "0", which means no notifications. This option corresponds to the num_grat_arp or num_unsol_na option in the *Red Hat Enterprise Linux 7: Networking Guide*.

"transmit-hash-policy": "layer2"

Valid in 802.3ad mode only. Selects a policy according to which the MAC addresses of the physical devices are determined. Specify layer2, which means that traffic to a particular network peer is assigned to the same network device, which is determined solely by its MAC address. Other allowed options are "layer3+4" and "layer2+3". The option "layer3+4" means that multiple network devices can be used to reach a single network peer even if a single network connection does not span multiple network devices. The option "layer2+3" is similar to "layer 2", but the network device is selected by its IP address in addition to its MAC address.

Note: In the *Red Hat Enterprise Linux 7: Networking Guide*, this option is called "xmit-hash-policy"

"LACP-rate": "slow" | "fast"

Valid in 802.3ad mode only. The rate at which physical devices transmit Link Aggregation Control Protocol Data Units (LACPDUs). Specify "slow", which means every 30 seconds, or "fast", which means every 1 second.

"link-monitoring": "MII"

Selects the method to be used for monitoring the physical device's ability to carry network traffic. Select "MII", which stands for *media-independent* interface. With this setting, the driver, the MII register, or the ethtool can be queried for monitoring information about a physical device. Alternatively, you can specify "ARP" to use the ARP monitor.

"monitoring-frequency": "100"

The time interval that passes between two monitoring events. It is an integer value that stands for milliseconds. Use a value of "100".

"link-up-delay": "0"

Delay that needs to pass before network traffic is sent to a physical device after link monitoring has reported the device to be up. Specify "0", which means *no delay*.

"link-down-delay": "0"

Delay that needs to pass before network traffic is routed to the failover device after link monitoring has reported the failure of the previously active device. Specify "0", which means *no delay*.

Example (active-backup mode):

```
"options": {
   "primary": "0.0.0a00",
   "primary-reselect": "always",
   "failover-MAC": "none",
   "no-gratuitous-ARPs": "0",
   "link-monitoring": "MII",
   "monitoring-frequency": "100",
   "link-up-delay": "0",
```

```
"link-down-delay": "0"
}
```

Example (802.3ad mode):

```
"options": {
    "LACP-rate": "slow",
    "transmit-hash-policy": "layer2",
    "link-monitoring": "MII",
    "monitoring-frequency": "100",
    "link-up-delay": "0",
    "link-down-delay": "0"
}
```

"zfcp_devices" (required if ZFCP drives are used)

The FICON Express ports of your devices must be listed in the "runtime_environments" section. This is required for ZFCP storage devices only, as the ports of ECKD devices are handled by the firmware of the CPC. For ZFCP devices, however, you must list the port names. See the following example:

The accelerator uses multiple paths, that is, it tries to use all specified ports. For that reason, a list of different ports rather than just one port can increase the performance.

You can use a single port identifier for a ZFCP device or for an ECKD device, but not for both.

You can change the ZFCP port names online.

"data1"

The definition of the first data node. It has only one network interface: the Hipersocket network of the cluster.

Example:

Specify the other data nodes in the same manner, that is, create the sections "data2", "data4", and "data5"

Complete example:

```
"network_interfaces": [
              "name": "my_hiper",
"device": "0.0.7f00",
"ipv4": "172.84.0.180/23"
               "name": "my_db2_network",
"device": "0.0.4b00"
       ]
   },
"data1": {
   "lpar_name": "LPAR1"
   "ork interfaces"
        "network_interfaces": [
              "name": "my_hiper",
"device": "0.0.7f00",
"ipv4": "172.84.0.181/23"
       ]
   },
"data2": {
    "lpar_name": "LPAR2"
    "-work interfaces"
        "network interfaces": [
              "name": "my_hiper",
"device": "0.0.7f00",
"ipv4": "172.84.0.182/23"
           3
       ]
   },
"data3": {
        "lpar_name": "LPAR3"
        "network_interfaces": [
          "name": "my_hiper",
  "device": "0.0.7f00",
  "ipv4": "172.84.0.183/23"
       ]
   },
"data4": {
        "lpar_name": "LPAR4"
        "network interfaces": [
              "name": "my_hiper",
"device": "0.0.7f00",
"ipv4": "172.84.0.184/23"
           3
       ]
   },
"data5": {
        "lpar_name": "LPAR5"
        "network interfaces": [
              "name": "my_hiper",
"device": "0.0.7f00",
"ipv4": "172.84.0.185/23"
           7
       ]
    3
}
```

"storage_environments" (required)

This block lists all storage devices, that is, disks or disk enclosures. At a minimum, this block contains the name of the storage on which the accelerator is initially deployed. You can define additional storage for mirroring or failover purposes. Note, however, that the accelerator does not manage the mirroring or copying of data to additional devices. The "storage_environments" section combines the "primary_storage" and "storage_maps" sections found in configuration files of earlier releases. During the first-time deployment, these devices are formatted, which means that the existing data on these devices is erased.

The accelerator uses three categories of storage: the boot device, the runtime data pool, and the data pool for operative data. Each storage device or storage pool can use DASD (ECKD) or ZFCP

(SCSI) devices. The mixing of different device types is not supported. However, the size of individual devices in a pool is not restricted.

Important: You can add devices to a storage pool at any time while the accelerator is online. New devices are automatically integrated into a storage pool. This does not require a reset or a restart. However, you cannot remove devices from a storage pool in this way. A removal requires an entirely new installation.

"head" (required)

This section lists the storage devices for the head node.

Each node uses three categories of storage: the boot device, the runtime data pool, and the data pool for operative data. You must define these storage devices by using the following attributes in the configuration file:

"boot_device" (required)

A boot device contains the software image that is written by the Secure Service Container (SSC) installer. The cluster will be started from the boot devices of its nodes. The boot devices are also the target devices for uploading the SSC installer image before the initial deployment or before an update. A boot device must be a single device with at least 40 GB net storage capacity.

"runtime_devices" (required)

A runtime device is used by the Docker container that runs the accelerator software on each node. It does not contain user data and its size is fixed. The size does not depend on the amount of user data processed by the accelerator Specify a list of devices with a total net capacity of at least 80 GB for each node .During normal operation, the utilization rate of a single device should not exceed 80 percent. If it does exceed 80 percent most of the time, consider adding devices.

"data_devices" (required)

A data device is used to store user data and the temporary data of the accelerator (table data). It is typically the largest storage area of an individual node. Its size is determined by the amount of data that the node has to handle. During normal operation, the utilization rate of a single device should not exceed 80 percent. If it does exceed 80 percent most of the time, consider adding devices.

"type" (required)

This is the type of storage to be used (disk type). Possible values are "dasd" for extended count key data (ECKD) volumes and "zfcp" for Small Computer System Interface (SCSI) volumes. You must specify the type for each device category (that is, the boot device, the runtime device, and the data device).

Important:

- It is not possible to mix ECKD and SCSI devices in a single device category or device pool.
- If you use DASD (ECKD) storage, HyperPAV aliases are strongly recommended because they increase the processing speed.
- DASD (ECKD) devices are formatted by the dasdfmt program of the Linux operating system on the accelerator. This can take a long time, sometimes even hours for large devices or storage pools. HyperPAV aliases also help speed up formatting. Therefore, define HyperPAV aliases also in your initial JSON configuration file.
- You can use DASD (ECKD) devices of different sizes in a single pool.
- You can use ZFCP devices of different sizes in a single pool. However, the use of devices smaller than 5 TB and larger than 10 TB is not recommended.
- For ZFCP devices, use Ficon Express (FE) ports because these ports guarantee a much better performance and availability. FE ports are defined in the runtime environment.

 Although the adding of devices is supported while the accelerator is online, the type cannot be changed after the initialization of the storage pool. To change the type, you must remove the accelerator and reinstall it.

Example (ECKD or "dasd"):

```
"type": "dasd",
"devices": [
  "0.0.9c00",
  "0.0.9c01",
```

Example (SCSI or "zfcp"):

```
"type": "zfcp",
"udids": [
   "0c984712545423523614b8d812345632",
"0c0a0b5c9d15555a545545b46456c4d6"
```

"device" or "devices" (required)

This attribute is used to list the devices by their names or identifiers. You must specify a device or a list of devices for each device category (that is, the boot device, the runtime device, and the data device).

Example:

```
"storage_environments": [
  {
    "head": {
       "boot device": {
          "type": "dasd"
          "device": "0.0.9b11"
       },
"runtime_devices": {
  "type": "dasd",
  "devices": [
            "0.0.9b12"
            "0.0.9b13"
         ]
        data_devices": {
          "type": "dasd",
"devices": [
            [ "0.0.9c00", "0.0.9cff" ]
         1
       3
```

Note: To change the ID of a boot device, a few extra steps are required:

- a. Copy the current storage environment in your JSON file. That is, create a duplicate block in the file.
- b. Change the ID in the copied block.
- c. Restart the accelerator.
- d. Remove the old storage environment from the JSON file when the accelerator is online again.

"data1" ... "data5"

Your data nodes. Specify a "boot_device", "runtime_devices", and "data_devices" for each of your data nodes in the same way as for the head node ("storage_environments").

Complete storage definition (partly shortened):

```
"storage_environments": [
   "head": {
      "boot device": {
```

```
type": "dasd",
                 "device": "0.0.9b11"
            },
"runtime_devices": {
  "type": "dasd",
  "devices": [
     "0.0.9b12",
     "0.0.9b13"
                     "0.0.9b12",
            },
"data_devices": {
    " "dasd",
                 "type": "dasd",
"devices": [
                    [ "0.0.9c00", "0.0.9cff" ]
            }
         "data1": {
    "boot_device": ...,
    "data_devices": ...,
    "runtime_devices": ...
       },
"data2": {
   "boot_device": ...,
   "data_devices": ...,
   "----+ime_devices": .
             "runtime_devices": ...
        "data3": {
    "boot_device": ...,
    "data_devices": ...,
    "runtime_devices": ...
        },
"data4": {
            "boot_device": ...,
"data_devices": ...,
             "runtime_devices": ...
        },
"data5": {
            "boot_device": ...,
"data_devices": ...,
"runtime_devices": ...
    }
]
```

HyperPAV aliases:

HyperPAV aliases can be used in the following ways:

- Automatic HyperPAV aliases
- Explicitly listed HyperPAV aliases

In automatic HyperPAV mode, all HyperPAV alias devices that are visible to an LPAR and that are connected to the same control-unit image (LCU) are used automatically for that LPAR. To enable the automatic HyperPAV mode, you must add a definition to the storage environments section in the JSON configuration file.

Important:

- HyperPAV aliases can be used with DASD (ECKD) storage only.
- Make sure that only the volumes and HyperPAV aliases you want to use on a particular LPAR are
 visible to that LPAR. This is even more important if you use automatic alias devices because in that
 case, your accelerator (SSC) LPAR has to sift through all visible devices just to determine and
 activate the alias devices.

Example:

```
},
"data_devices": {
 hyperpav": "auto"
```

Mind that in "auto" mode, the system uses all available HyperPAV alias devices.

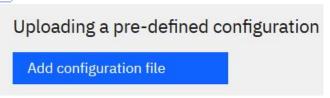
To use just a subset of the available devices, it is preferable to use explicitly listed HyperPAV aliases. as in the following example:

```
"hyperpav": [
     "0.1.4000"
"0.1.4007"
    0.1.1234"
]
```

In this particular case, the system uses a range of HyperPAV aliases from 0.1.4000 to 0.1.4007 plus a single HyperPAV alias with the ID 0.1.1234.

6. When you're finished with your configuration file, upload it to the Admin UI.

On the Accelerator Configuration Definition page, under the heading Uploading Configuration File, click the upload button (see also "Installing and starting the appliance (multi-node setup)" on page 76):



If something is wrong with the file you uploaded, an error message is displayed on the page:

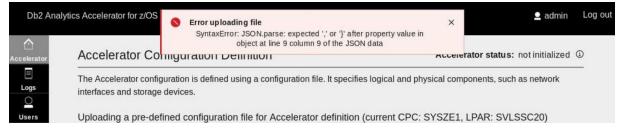


Figure 25. Error message after uploading a faulty configuration file

7. If errors occurred, fix these and repeat the upload (steps "5" on page 78 and "6" on page 90).

If no errors occurred, the Accelerator Configuration Definition page shows the settings of your configuration file in a folder structure. You can expand the folders to display the settings they contain by clicking the plus signs in front of the folder icons.

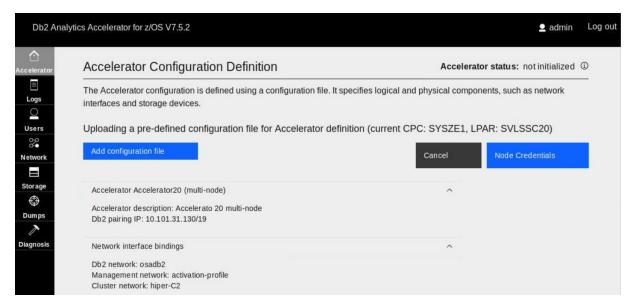


Figure 26. Accelerator Configuration Definition after a successful configuration file upload

- 8. Click Node Credentials.
- In the Multiple Node Deployment Credentials window, enter the IP addresses of the accelerator's data-node LPARs. Also provide the user ID and the password of each accelerator (SSC) LPAR, as defined by the HMC.

This is required for the initial deployment only. That is, the cluster configuration will continue to work if the user IDs or the passwords change. The HMC credentials of the data-node LPARs are not used after the initial setup.

See "Installing and starting the appliance (multi-node setup)" on page 76.

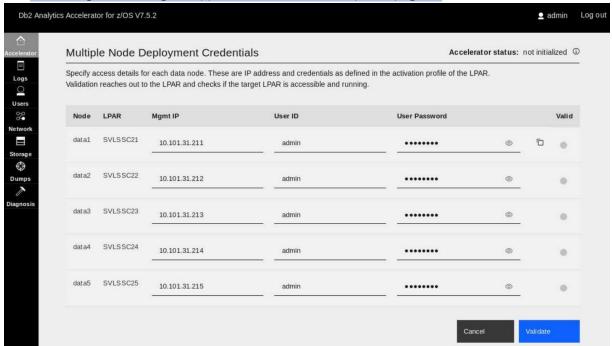


Figure 27. The Multiple Node Deployment Credentials window

10. Click Validate.

You see a message window indicating that the configuration is in progress:

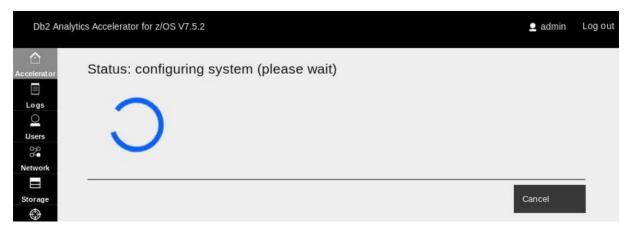


Figure 28. Message window showing the progress of your Db2 Analytics Accelerator on Z configuration

Results

When these processes have finished, the **Accelerator Components Health Status** page is displayed automatically. The page should now give you the following information:

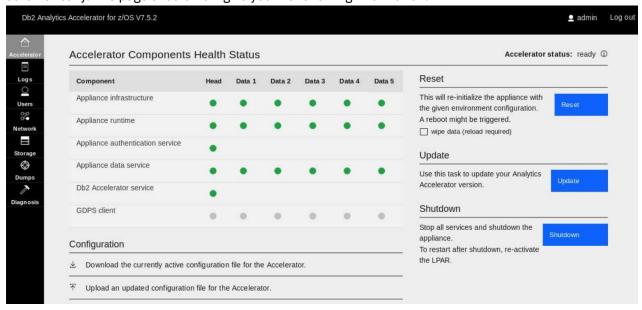


Figure 29. The Accelerator Components Health Status page is displayed after a successful configuration

The message Accelerator status: ready on the top right indicates that all installation steps have been completed and that components have been started for Db2 Analytics Accelerator on Z.

Installing Db2 Analytics Accelerator on Z with GDPS failover support

You can use GDPS to install Db2 Analytics Accelerator on Z in a high-availability environment with failover support.

Before you begin

GDPS stands for Geographically Dispersed Parallel Sysplex[®]. It is a multi-site disaster recovery technology for IBM Z systems: If a system at one site ceases to function, another instance of that system at another site will take over. This technology can also be used for an installation of Db2 Analytics Accelerator on Z. To follow the instructions here, a GDPS infrastructure must already exist.

GDPS is a separately priced feature, a license of which is not part of the IBM Db2 Analytics Accelerator product license. If you have not used GDPS before and need a license, contact IBM support. The support team can also provide you with or direct you to specific GDPS documentation.

About this task

You have to install the Db2 Analytics Accelerator on Z only once, in the primary LPAR. The backup LPAR uses the same disks as the primary LPAR, and therefore, no additional installation is required for the backup LPAR.

In addition, you need to specify some of the properties of your GDPS servers in the JSON configuration file for Db2 Analytics Accelerator on Z.

Defining an LPAR for Db2 Analytics Accelerator on Z with GDPS failover support

This section describes how to create a single, dedicated LPAR for an installation of Db2 Analytics Accelerator on Z with GDPS failover support. However, depending on your setup (single-node or multi-node), you need two or six of these LPARs. In a single-node setup, one is the primary, active LPAR, and the other one is the backup site for the failover case. Repeat this sequence of steps for each LPAR you need to create.

Procedure

- 1. Open the Hardware Management Console (HMC) of your IBM Z.
- 2. Go to System Management > Machine > LPARs.
- 3. Select the **Partitions** tab.
- 4. Right-click the (not yet activated) partition that you want to use and select **Operational Customization > Customize/Delete Activation Profile**.

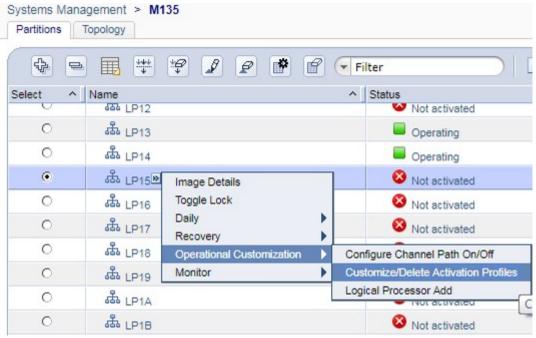


Figure 30. The Partitions window of the HMC

- 5. In the navigation pane on the left of the **Customize Image Profiles** window, select **SSC**.
- 6. In the right pane, click **Reset User Settings**.
- 7. Enter a Master user ID and a Master password.
- 8. Confirm the master password by retyping it in the appropriate field.
- 9. Click Apply.

You return to the SSC entry page.

10. Click Reset Network Settings.

The Network Adapters page opens.

11. In the **Host name** field, specify a host name for the network connection.

Important: In a GDPS setup, the host name for the primary LPAR and the backup LPAR must be equal.

- 12. Click the appropriate icon or **Select Action** underneath and specify exactly one single valid IP address.
- 13. Click Save.

You return to the **Partitions** tab.



Attention: Do not define more than one network adapter for the SSC LPAR. If more than one definition is present, your Db2 Analytics Accelerator on Z installation will fail eventually.

14. Right-click the LPAR again. This time, select **Daily** > **Activate**.

The system confirms the activation by displaying a window labeled **Operating System Messages**:

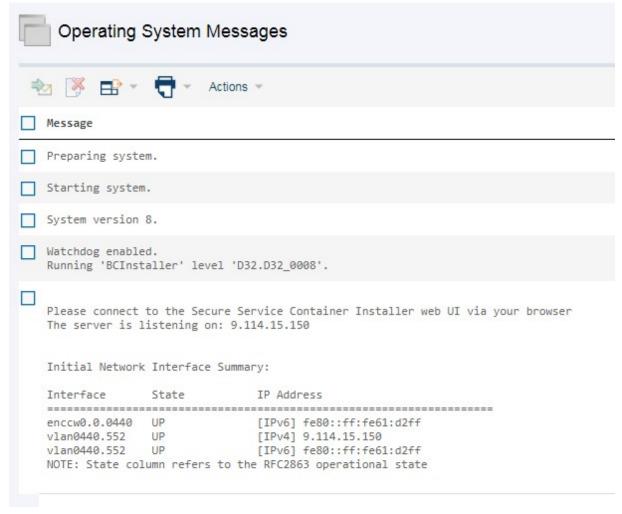


Figure 31. Confirmation: Operating System Messages

Installing and starting the appliance with GDPS failover support (single-node setup)

This installation is similar to the single-site setup, but requires an additional accelerator (SSC) LPAR, an existing GDPS environment and a few extra configuration steps in the **Admin UI**.

Before you begin

Storage must be provided for various storage pools. <u>Table 1 on page 61</u> provides an overview and minimum size recommendations.

Table 2. Recommended pools and pool sizes			
Pool	JSON key	Description	Size
Appliance data pool	"data_devices"	Accelerator database including accelerator-shadow tables, accelerator-only tables, temporary database space, and temporary result space.	> 200 GB plus 90 percent of uncompressed Db2 for z/OS data (including accelerator-only tables) plus 20 percent of the accelerator (SSC) LPAR memory.
Appliance operation	"boot_device"	Size of the appliance image, which was defined and used during the image file upload step ("Install Software Appliance").	> 40 GB
			Important: This must be a single disk.
Appliance runtime	"runtime_devices"	Temporary storage space required for the appliance runtime (container execution environment)	> 80 GB

About this task

Restriction: For the following steps, you need Mozilla Firefox or Google Chrome. Other browsers are not supported.

Procedure

- 1. Define two dedicated Secure Service Container (SSC) LPARs as described in "Defining an LPAR for Db2 Analytics Accelerator on Z with GDPS failover support" on page 93.
- Log on to the Admin UI and proceed to the Welcome page.
 For a description, see "Logging on to the Admin UI" on page 58.
- 3. Click First-Time Setup.
- 4. When finished, click **Apply**.

The process needs between 1 and 2 minutes to complete. During that time, the **Apply** button is grayed out.

5. On the **Welcome** page, click **First-Time Setup**. You see the following page:

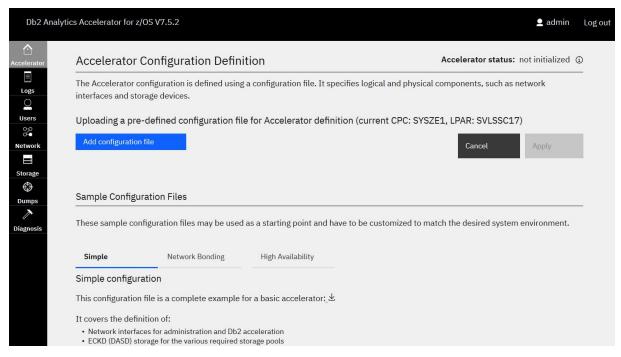


Figure 32. Accelerator Configuration Definition page

6. Starting with product version 7.1.9, all configuration settings are made by uploading a configuration file in JavaScript Object Notation (JSON) format.

A sample configuration file is provided at the bottom of the **Welcome** page. Under the heading Sample Configuration Files:

- a. Click the **High Availability / GDPS** tab.
- b. Click the download button (

The sample configuration file is stored in the Download folder of your web browser.

Important: Compared with previous product releases, the configuration file delivered with product version 7.5.3 and later versions has changed considerably.

- 7. Open this file in a text editor of your choice and modify the settings according to your needs.
 - An editor capable of validating JSON files is recommended because the configuration file must be valid JSON. If it cannot be parsed correctly, you will run into errors. Valid JSON means:
 - Quotes are required around attribute values, even if these are plain numbers.
 - Colons must be used to separate attribute names from their values.
 - Object definitions consisting of key/value pairs must be enclosed in curly braces.
 - · Arrays or lists must be enclosed in brackets.

For your reference, take a look at the code of the sample GDPS configuration file:

```
"version": "7.5.6",
"accelerator_name": "gdpse2e"
"accelerator_description": "SAMPLE-single-node with GDPS support", "accelerator_type": "single-node", "db2_pairing_ipv4": "10.109.36.184/24",
"network_interface_bindings": {
   "db2_nw": "db2_nw_connection",
   "gdps_nw": "gdps_nw_connection",
   "mgmt_nw": "activation-profile"
},
"runtime_environments": [
        "cpc_name": "SYSTEM1",
        "lpar_name": "LPAR1",
```

```
"network_interfaces": [
                "name": "db2_nw_connection",
"device": "0.0.4b00",
"port": "0"
                "name": "gdps_nw_connection",
"device": "0.0.4206",
"ipv4": "10.101.31.81/24",
"port": "0"
         ]
      },
{
         "cpc_name": "SYSTEM2",
"lpar_name": "LPAR2",
          "network_interfaces": [
                "name": "db2_nw_connection",
"device": "0.0.4200",
"port": "0"
                "name": "gdps_nw_connection",
"device": "0.0.4203",
"ipv4": "10.101.31.82/24"
         ]
      }
   ],
"storage_environments": [
         "boot_device": {
    "type": "dasd",
    "device": "0.0.1b11",
    "storage_uid": "0x7779823794723"
           runtime_devices": {
             "type": "dasd",
             "devices":
                "0.0.1b25"
                "0.0.1b26"
             ],
"storage_uid": "0x7779823794812"
           data_devices": {
             "type": "dasd",
"devices": [
                "0.0.1c14"
                [ "0.0.1d00", "0.0.1d0f" ]
             ],
"storage_uid": "0x7779823794812"
         3
      }
   ]
3
```

Note: The "storage_uid": values are required for GDPS Metro 4-site only.

Most attributes in the configuration file are required.

"version" (required)

The version of the accelerator. The version in the configuration file must match the version of the accelerator exactly. Do not change the value so that it deviates from the version of the accelerator.

"accelerator_name" (required)

The name of the accelerator. This attribute is used to identify the accelerator in the user interface. It reoccurs in the dump/trace/log output.

"accelerator_description" (optional)

Optional text description. You might want to add some information about the accelerator.

"accelerator_type" (required)

As the name suggests, the type of the accelerator. Currently, only one type is supported. Set this attribute to the value "single-node".

"db2_pairing_ipv4" (required)

The IP address used to pair your Db2 subsystem with the specified accelerator. This IP address uniquely identifies the accelerator and is used by Db2 for z/OS to connect to the accelerator. You cannot change it without dropping the accelerator and creating a new pairing. It always has to be the same address, no matter in which Central Processing Complex (CPC) or LPAR the accelerator is started.

You can specify a netmask as part of the IPv4 address, like /24 for a subnet with 256 addresses. For example:

```
"db2_pairing_ipv4": "10.109.36.184/24"
```

This specifies the IP address 10.109.36.184 as the identifier of a subnet that comprises the address range from 10.109.36.0 to 10.109.36.255.

Tip: All IP addresses in the configuration file can point to a subnet.

"network interface bindings" (required)

You must bind your Db2 Analytics Accelerator on Z configuration to certain network interfaces.

"mgmt nw"

This network, which is used by the Admin UI and other support interfaces, is defined by the HMC activation profile of the accelerator (SSC) LPAR. It is not part of the Db2 Analytics Accelerator on Z configuration. Therefore, use the attribute value "activation-profile". Note that the name of the management network might change if someone updates the activation profile of the accelerator (SSC) LPAR on the HMC.

"db2_nw"

This network name points to the IP address of your Db2 subsystem (counterpart of the "db2 pairing ipv4"). It is used during the pairing process, and all network traffic between your Db2 subsystem and the accelerator will run through this interface.

The attribute value must be the same as one of the "name: " attributes in your "network_interfaces" definitions further down in the configuration file. The value must be an alphanumeric character string no longer than 8 characters. Compare this with the sample GDPS configuration file.

In the network_interfaces: section, which is described below, you find the details of all networks, including the network used for the pairing process. In the example, this is the network device with the ID 0.0.4b00.

"gdps_nw"

This network interface is used for GDPS failover support. It is specifies the network used to contact the GDPS servers (see keyword gdps_servers).

An example of the "network_interface_bindings" block:

```
"network_interface_bindings": {
       "mgmt_nw": "activation-profile",
"db2_nw": "db2_nw_connection",
"gdps_nw": "gdps_nw_connection"
```

Notes:

- Multiple interfaces can use the same physical connection. For example, in a GDPS setup, the interface name specified as the value of "gdps_nw": might be the same as the value of "db2_nw":. In this case, the same physical connection would be used by both interfaces, and only one network definition would be required in the "network interfaces" block further down.
- Do not use the value "activation-profile" for any network interface other than the "mgmt nw" because the activation profile might change in the HMC. Such a change might have repercussions if a network that refers to the activation profile is used for other purposes. The change might make your network definition unusable.

"runtime_environments" (required)

This block defines the network interfaces and other LPAR-specific settings of the accelerator (SSC LPAR) involved in the GDPS setup. You still have to define a runtime environment because GDPS controls the storage configuration, but not the configuration of accelerator LPARs. Each LPAR is identified by the CPC name and the LPAR name, and a set of networks must be defined for each accelerator (SSC) LPAR. This is usually the Db2 network, and an additional network to reach the GDPS server. Specify the following attributes to identify an accelerator (SSC) LPAR:

"cpc name"

The name of the CPC.

"lpar_name"

The name of the accelerator (SSC) LPAR.

Continue with the network interfaces for the accelerator (SSC) LPAR:

"network_interfaces" (required)

This keyword introduces the block of network interface definitions for the accelerator (SSC) LPAR. The following attributes must be specified for a single network:

"name" (required)

The name of the network interface. You can choose it freely.

"ipv4" (required)

The IP address to be used for connections to the failover environment.

"device" (required)

The identifier of an OSA-Express card or a Hipersocket.

"port" (optional)

The network port to be used. If this value is omitted, the port number defaults to "0".

Example:

```
"runtime_environments": [
          "cpc_name": "SYSTEM1",
"lpar_name": "LPAR1",
           "network_interfaces": [
                 "name": "db2_nw_connection",
"device": "0.0.4b00",
"port": "0"
                 "name": "gdps_nw_connection",
"ipv4": "10.109.36.184/24",
"device": "0.0.4206",
"port": "0"
              3
          ]
          "cpc_name": "SYSTEM2",
"lpar_name": "LPAR2",
           "network_interfaces": [
                 "name": "db2_nw_connection",
                 "device": "0.0.4b00",
"port": "0"
              },
                 "name": "gdps_nw_connection",
"ipv4": "10.101.36.185/24",
"device": "0.0.4206",
"port": "0"
             3
          ]
      3
```

In this example, you find network interface definitions for two accelerator (SSC) LPARs named LPAR1 and LPAR2. These LPARs run in CPCs named SYSTEM1 and SYSTEM2. The primary

LPAR in the GDPS setup uses the network device 0.0.4b00 (an OSA-Express card or a Hipersocket). The name that Db2 Analytics Accelerator on Z uses for this network device is "db2_nw_connection".

The failover LPAR uses the network device 0.0.4206. Its IP address is specified, so that the device can be connected to in case of a failover.

"static_routes" (optional)

This option is used to define additional network routes for an interface.

If the IP address of a Db2 for z/OS LPAR or the GDPS keys LPAR is in a different subnet than the IP address assigned to the accelerator, an additional route definition is needed to establish the connection. An additional static route also helps to avoid undesired network traffic through a default gateway, which might have been defined in the HMC activation profile of the accelerator (SSC) LPAR.

"ipv4"

The IPv4 address or the IPv4 address and subnet of the target network.

"via"

The IPv4 address of the routing device.

Example: The accelerator's pairing IP address is 10.20.1.33/24 and there are two Db2 for z/OS LPARs with the IP addresses 10.1.1.47/24 and 10.1.1.48/24.

One or more gateways connect both subnets. One gateway is accessed through IP address 10.20.1.1, the other through 10.1.1.1.

To allow traffic from one network to the other, the TCPIP.PROFILE definition in z/OS defines a route to 10.20.1.0/24, which uses the gateway 10.20.1.1. The accelerator uses the following configuration to enable traffic to the 10.1.1.0 network using the corresponding gateway at 10.1.1.1:

This way, all traffic to an IPv4 address that starts with 10.1.1 uses the OSA-Express card with device ID 0.0.0440 via gateway 10.20.1.1. All network traffic between the accelerator and destinations in the 10.1.1.0/24 subnet is thus bound to that OSA device.

"bond_settings" (optional)

This attributes allows you to define several network cards (OSA-Express cards) as a single device. Bonding is usually employed in a high-availability setup, as the remaining network cards in the setup can take over if one network card fails. It is also possible to always run all available network cards simultaneously.

Example:

In this example, two OSA cards (devices 0a00 and 1b00) are combined to one bonding device called "db2_nw_connection". The device works in "active-backup" mode, meaning that at any time, just one of network cards is active. The other card takes over when the active card fails.

You can alternatively specify "mode": "802.3ad", in which case all network cards of the device will be active at the same time. "802.3ad" stands for the IEEE 802.3ad link aggregation mode.

In 802.3ad mode, you need at least two physical devices. Specify these in the same way as you specify the devices for active-backup mode. That is, use a "slaves" list as shown in the example.

"options" (optional)

It is not necessary to specify "options" for "bond_settings". If the options are omitted, default values are used. Whether options apply to a particular setup depends on the selected mode ("active-backup" or "802.3ad"). For a detailed description of these options, see *Chapter 7. Configure Network Bonding* in the *Red Hat Enterprise Linux 7: Networking Guide*. A link is provided at the end of this topic.

Restriction: Currently, you cannot specify just a subset of the available options. You either have to specify no options at all, in which case default values are used, or specify all options pertaining to a particular mode.

All of the following values can be changed online.

"primary"

Valid in *active-backup* mode only. The first physical device to be used. This is "0.0.0a00" according to the previous example. The primary device is the first of the bonding interfaces. It will be used as the active device unless it fails.

"primary-reselect": "always"

Valid in *active-backup* mode only. Determines how the active physical device is selected after a failure. Specify "always", which means that an attempt will be made to make the first physical device (labeled "primary") active again.

Other allowed options are "better", which means that the fastest device will be used as the active device, or "failure", which means that the active physical device is only changed if the currently active device fails.

"failover-MAC": "none"

Valid in *active-backup* mode only. Allows you set all physical devices to the same MAC address or determine these addresses according to a policy. Specify the value "none", which means that the same MAC address will be used for all physical devices.

"no-gratuitous-ARPs": "0"

Valid in *active-backup* mode only. Determines the number of peer notifications after a failover event. Specify "0", which means no notifications. This option corresponds to the num_grat_arp or num_unsol_na option in the *Red Hat Enterprise Linux 7: Networking Guide*.

"transmit-hash-policy": "layer2"

Valid in 802.3ad mode only. Selects a policy according to which the MAC addresses of the physical devices are determined. Specify layer2, which means that traffic to a particular network peer is assigned to the same network device, which is determined solely by its MAC address. Other allowed options are "layer3+4" and "layer2+3". The option "layer3+4" means that multiple network devices can be used to reach a single network peer even if a single network connection does not span multiple network devices. The option "layer2+3" is similar to "layer 2", but the network device is selected by its IP address in addition to its MAC address.

Note: In the *Red Hat Enterprise Linux 7: Networking Guide*, this option is called "xmit-hash-policy"

"LACP-rate": "slow" | "fast"

Valid in 802.3ad mode only. The rate at which physical devices transmit Link Aggregation Control Protocol Data Units (LACPDUs). Specify "slow", which means every 30 seconds, or "fast", which means every 1 second.

"link-monitoring": "MII"

Selects the method to be used for monitoring the physical device's ability to carry network traffic. Select "MII", which stands for *media-independent interface*. With this setting, the driver, the MII register, or the ethtool can be queried for monitoring information about a physical device. Alternatively, you can specify "ARP" to use the ARP monitor.

"monitoring-frequency": "100"

The time interval that passes between two monitoring events. It is an integer value that stands for milliseconds. Use a value of "100".

"link-up-delay": "0"

Delay that needs to pass before network traffic is sent to a physical device after link monitoring has reported the device to be up. Specify "0", which means no delay.

"link-down-delay": "0"

Delay that needs to pass before network traffic is routed to the failover device after link monitoring has reported the failure of the previously active device. Specify "0", which means *no delay*.

Example (active-backup mode):

```
"options": {
    "primary": "0.0.0a00",
    "primary-reselect": "always",
    "failover-MAC": "none",
    "no-gratuitous-ARPs": "0",
    "link-monitoring": "MII",
    "monitoring-frequency": "100",
    "link-up-delay": "0",
    "link-down-delay": "0",
}
```

Example (802.3ad mode):

```
"options": {
  "LACP-rate": "slow",
  "transmit-hash-policy": "layer2",
  "link-monitoring": "MII",
  "monitoring-frequency": "100",
  "link-up-delay": "0",
  "link-down-delay": "0"
}
```

"zfcp_devices" (required if ZFCP drives are used)

The FICON Express ports of your devices must be listed in the "runtime_environments" section. This is required for ZFCP storage devices only, as the ports of ECKD devices are

handled by the firmware of the CPC. For ZFCP devices, however, you must list the port names. See the following example:

The accelerator uses multiple paths, that is, it tries to use all specified ports. For that reason, a list of different ports rather than just one port can increase the performance.

You can use a single port identifier for a ZFCP device or for an ECKD device, but not for both.

You can change the ZFCP port names online.

"storage_envrionments" (required)

I

This block lists all storage devices, that is, disks or disk enclosures. At a minimum, this block contains the name of the storage on which the accelerator is initially deployed. You can define additional storage for mirroring or failover purposes. The "storage_environments" section combines the "primary_storage" and "storage_maps" sections found in configuration files of earlier releases. During the first-time deployment, these devices are formatted, which means that the existing data on these devices is erased.

The accelerator uses three categories of storage: the boot device, the runtime data pool, and the data pool for operative data. Each storage device or storage pool can use DASD (ECKD) or ZFCP (SCSI) devices. The mixing of different device types is not supported. However, the size of individual devices in a pool is not restricted.

"boot_device" (required)

The boot device contains the software image that is written by the Secure Service Container (SSC) installer. The accelerator will be started from this device. It is also the target device for uploading the SSC installer image before the initial deployment or before an update. The boot device must be a single device with at least 40 GB net storage capacity. For example:

```
"boot_device": {
   "type": "dasd",
   "device": "0.0.5e29"
}
```

or

"runtime_devices" (required)

The runtime storage is used by the accelerator software for internal processing. It does not contain user data and its size is fixed. The size does not depend on the amount of user data processed by the accelerator. Specify a list of devices with a total net capacity of at least 80 GB.During normal operation, the utilization rate should not exceed 80 percent. If it does exceed 80 percent most of the time, consider adding devices. For example:

```
"runtime_devices": {
  "type": "dasd",
  "devices": [
```

```
"0.0.998c"
]
}
```

The mixing of different device types in a single list, that is, DASD (ECKD) and ZFCP (SCSI) is not supported.

"data_devices" (required)

The data storage is used to store the user data and temporary data of the accelerator (table data). It is typically the largest storage area of your entire configuration. Its size is determined by the amount of data the accelerator has to handle. During normal operation, the utilization rate should not exceed 80 percent. If it does exceed 80 percent most of the time, consider adding devices. For example:

```
"data_devices": {
    "type": "dasd",
    "devices": [
        "0.0.9c00",
        "0.0.9c01",
        "0.0.9c02"
    ]
}
```

"type" (required)

This is the type of storage to be used (disk type). Possible values are "dasd" for extended count key data (ECKD) volumes and "zfcp" for Small Computer System Interface (SCSI) volumes. You must specify the type for each device category (that is, the boot device, the runtime device, and the data device).

Important:

- It is not possible to mix ECKD and SCSI devices in a single device category or device pool.
- If you use DASD (ECKD) storage, HyperPAV aliases are strongly recommended because they increase the processing speed.
- DASD (ECKD) devices are formatted by the **dasdfmt** program of the Linux operating system on the accelerator. This can take a long time, sometimes even hours for large devices or storage pools. HyperPAV aliases also help speed up formatting. Therefore, define HyperPAV aliases also in your initial JSON configuration file.
- You can use DASD (ECKD) devices of different sizes in a single pool.
- You can use ZFCP devices of different sizes in a single pool. However, the use of devices smaller than 5 TB and larger than 10 TB is not recommended.
- For ZFCP devices, use Ficon Express (FE) ports because these ports guarantee a much better performance and availability. FE ports are defined in the runtime environment.
- Although the adding of devices is supported while the accelerator is online, the type cannot be changed after the initialization of the storage pool. To change the type, you must remove the accelerator and reinstall it.

Example (ECKD or "dasd"):

```
"type": "dasd",
"devices": [
    "0.0.9c00",
    "0.0.9c01",
    "0.0.9c02"
]
```

Example (SCSI or "zfcp"):

```
"type": "zfcp",
"udids": [
   "0c984712545423523614b8d812345632",
   "0c0a0b5c9d15555a545545b46456c4d6"
]
```

"device" or "devices" (required)

This attribute is used to list the devices by their names or identifiers. You must specify a device or a list of devices for each device category (that is, the boot device, the runtime device, and the data device).

Example:

"gdps_mode" (required)

Set this to "true".

"gdps_configuration" (required)

This block contains the details of your GDPS setup.

"gdps_servers" (required)

A list of the GDPS K-systems (servers) to be used in your setup. Give each server a name, and specify the IP address and the port to be used.

Example:

```
"gdps_mode": "true";
"gdps_configuration": {
    "gdps_servers": {
        "ipv4": "10.2.1.11",
        "port": "1020"
        },
        "server2": {
              "ipv4": "10.2.2.11",
              "port": "1030"
        }
}
```

In this example, two K-systems are specified, which is a common setup for a *GDPS Metro Mirror* or a *GDPS Metro Global GM 3-site* configuration. A configuration for a *GDPS Metro Global GM 4-site* typically requires four K-systems (GDPS servers), specified by the names server1, server2, server3 and server4.

For a GDPS Metro Mirror 2-site or a GDPS Metro Global GM 3-site configuration, specify the GDPS K-systems as follows:

Site 1 of the production region:	server1
Site 2 of the production region :	server2

In contrast, for a GDPS Metro Global GM 4-site configuration, where both regions are equivalent production regions, you must specify the GDPS K-systems as follows:

Site 1 of the installation region:	server1
Site 2 of the installation region:	server2

Site 1 of the non-installation region:	server3
Site 2 of the non-installation region:	server4

"storage_options" (optional)

You can use this optional attribute to specify storage-related options pertaining to your GDPS environment, such as HyperPAV aliases. Mind that HyperPAV aliases can only be used for DASD storage.

Example:

Note: The "storage_uid": values are required for GDPS Metro 4-site only.

The following options are related to HyperPAVs. These options control the HyperPAV handling for your GDPS environment.

"hyperpav-default"

Set this attribute to the default settings for HyperPAV aliases. Possible values:

"none"

HyperPAV aliases are not used. This is the default setting.

"auto"

All HyperPAV aliases are used.

"hyperpav-explicit"

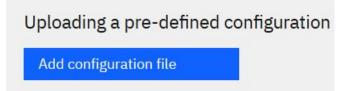
Set this attribute to use just a subset of the existing HyperPAV aliases. For example:

```
"aliases": "auto"
}
```

You see HyperPAV alias settings for two storage environments. The first block of entries specifies a list of alias device IDs explicitly. These are to be used by the environment with the UID "0x7779823794723". In the second block, you find settings for a storage environment with the UID "0x7779823784103". For this environment, the default setting "auto" has been chosen. That is, the treatment of HyperPAV aliases will be the same as if the setting "hyperpav-default": "auto" was used. All other, unlisted storage environments will use the setting "hyperpav-default": "none".

8. When you're finished with your configuration file, upload it to the Admin UI.

On the **Accelerator Configuration Definition** page, under the heading **Uploading Configuration File**, click the upload button (see also Figure 18 on page 62):



If something is wrong with the file you uploaded, an error message is displayed on the page:

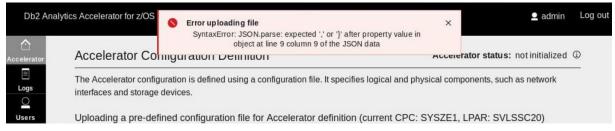


Figure 33. Error message after uploading a faulty configuration file

9. If errors occurred, fix these and repeat the upload (steps "Installing and starting the appliance with GDPS failover support (single-node setup)" on page 95 and "Installing and starting the appliance with GDPS failover support (single-node setup)" on page 95).

If no errors occurred, the **Accelerator Configuration Definition** page shows the settings of your configuration file in a folder structure. You can expand the folders to display the settings they contain by clicking the plus signs in front of the folder icons.

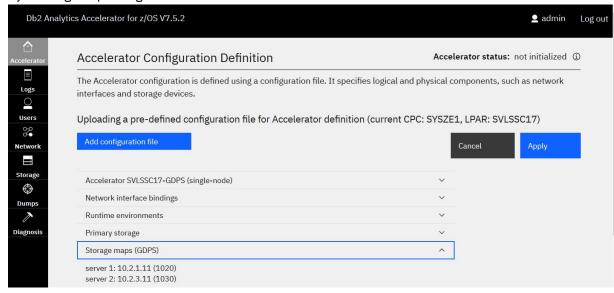


Figure 34. Accelerator Configuration Definition after a successful configuration file upload 10. Click **Apply**.

You see a message window indicating that the configuration is in progress:

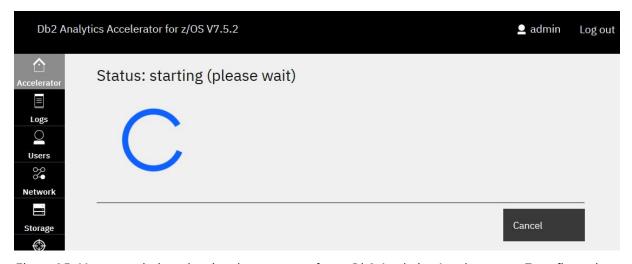


Figure 35. Message window showing the progress of your Db2 Analytics Accelerator on Z configuration

Results

When these processes have finished, the **Accelerator Components Health Status** page is displayed automatically. The page should now give you the following information:

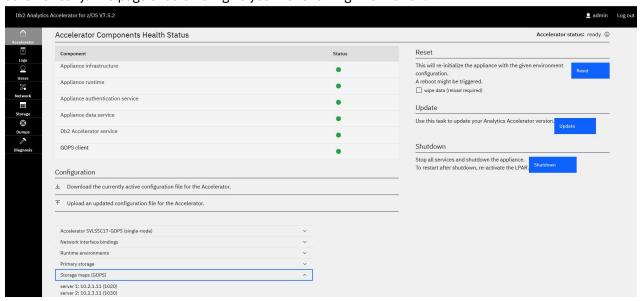


Figure 36. The Accelerator Components Health Status page is displayed after a successful configuration

The message Accelerator status: ready on the top right indicates that all installation steps have been completed and that components have been started for Db2 Analytics Accelerator on Z.

Shutting down and restarting a cluster (multi-node setup)

You might want to shut down a cluster (multi-node installation) or restart it after a shutdown or failure.

About this task

On IBM Z hardware, an individual LPAR does not fail because one of the processor cores is out of order. In fact, the operating system on the LPAR does not even notice the event. From that perspective, it seemed unnecessary to develop a system that allows users to restart individual nodes. Hence you can only restart an entire cluster.

The entire cluster must also be restarted if an administrator shut down an individual node from the Hardware Management Console (HMC). It is not possible to run a cluster with the remaining nodes.

If a cluster went offline as the result of a failure rather than a regular shutdown, the restart takes longer because recovery work needs to be completed first.

Procedure

- Enter the IPv4 address of the Secure Service Container (SSC) LPAR in a web browser (see <u>Figure 12 on</u> page 57).
- On the **Login** page of the Admin UI, log in with the **Master user ID** and **Master password** that you specified as you defined the SSC LPAR. (see Figure 12 on page 57)
 - You see the Accelerator Components Health Status page.
- To shut down a cluster, click the **Shutdown** button on the **Accelerator Components Health Status** page of the Db2 Analytics Accelerator on Z user interface.

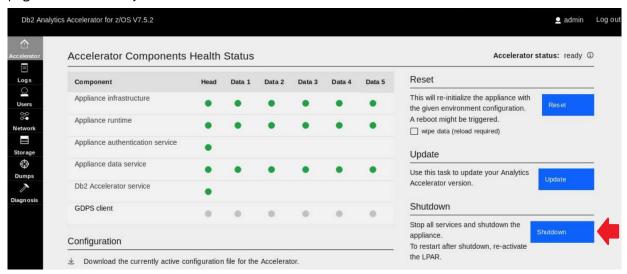


Figure 37. Shutdown from the Accelerator Components Health Status page

 To restart a cluster, re-activate all the nodes that belong to the cluster, that is, the head node and the five data nodes.

You re-activate all the nodes by re-activating the LPARs from the HMC.

Updating an existing configuration

You can easily replace an existing Db2 Analytics Accelerator on Z configuration by uploading a modified configuration file.

About this task

You can add storage without a downtime of the accelerator (appliance). If you change the configuration of the Db2 network ("db2_nw"), you must restart the accelerator. Other changes to the network configuration are applied dynamically, so that a restart is unnecessary.

Restriction: For the following steps, you need Mozilla Firefox or Google Chrome. Other browsers are not supported.

Procedure

- 1. Enter the IPv4 address of the Secure Service Container (SSC) LPAR in a web browser (see <u>Figure 12 on</u> page 57).
- 2. On the **Login** page of the Admin UI, log in with the **Master user ID** and **Master password** that you specified as you defined the SSC LPAR. (see Figure 12 on page 57)

You see the Accelerator Components Health Status page.

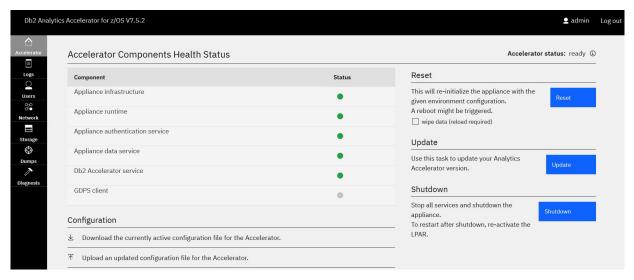


Figure 38. The Accelerator Components Health Status page

3. It is recommended that you use the existing configuration file as the basis for your modifications. You

download this file by clicking Download the currently active configuration file for the Accelerator. Navigate to the file location of your choice and save the file.

Important: Compared with previous product releases, the syntax of the configuration file delivered with product version 7.5.3 and later versions has changed considerably. To upgrade from product version 7.5.2 or an earlier version, or make changes after an upgrade from these versions, always download the latest sample configuration file to guide you through the configuration process. The following chapters show the latest configuration syntax for your setup:

- Single-node
- Multi-node
- GDPS
- 4. Once your updated configuration file is ready for deployment, click Upload an updated configuration file for the Accelerator.
- 5. Navigate to the location of the file, select it, and click **OK**.

You see a confirmation window like this:



Figure 39. **Confirmation** window

6. Click OK.

Results

The accelerator continues to work. If you have added storage in your configuration file, the corresponding disk drive is formatted.

What to do next

Mind that you have to restart the accelerator (appliance) if you have changed the network configuration.

You can check the deployment process if you click **Log** in the navigator on the left of the Db2 Analytics Accelerator on Z user interface. The **Log** window lists recent events, as in the following example:

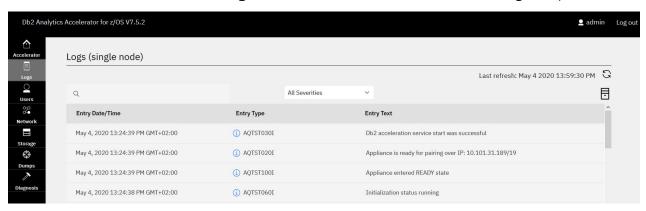


Figure 40. The Log window of the Db2 Analytics Accelerator on Z user interface

Creating a trace file (dump)

If you encounter installation problems or other problems with your Db2 Analytics Accelerator on Z, you might want to contact IBM support for help. In many cases, IBM support will ask you for a trace file that includes a snapshot of the data processed by the Db2 Analytics Accelerator on Z. On the administration interface, such a trace file is called a dump. This section describes how to create a dump.

Procedure

- 1. Enter the IPv4 address of the Secure Service Container (SSC) LPAR (see Figure 12 on page 57)
- 2. On the **Login** page of the Admin UI, log in with the **Master user ID** and **Master password** that you specified as you defined the SSC LPAR. (see Figure 12 on page 57)
- 3. In the navigator on the left, click **Dumps** to show the **Dumps** page.

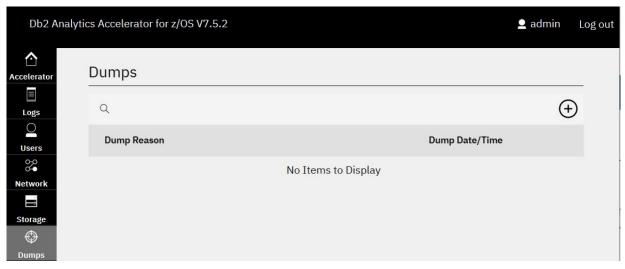


Figure 41. The Dumps page



You see the following page:

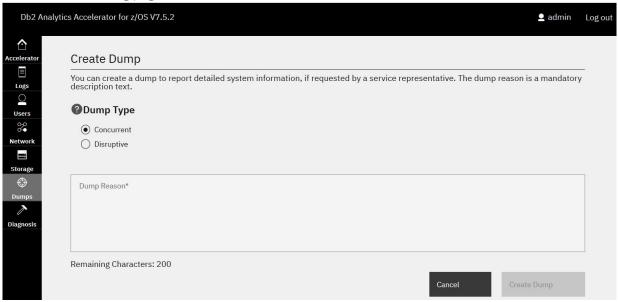


Figure 42. The Create Dump page

5. Select the **Dump Type**.

Concurrent

Leaves the Db2 Analytics Accelerator on Z accelerator active (online) as you create the dump.

Disruptive

Stops the Db2 Analytics Accelerator on Z accelerator before the dump is created.

- 6. Enter a description in the text box with the placeholder text Dump Reason. A description is mandatory.
- 7. Click Create Dump.

Results

The dump or trace file is written to the configured downloads folder of your web browser. You can send this file to IBM support.

Adding and managing users

The secure service container (SSC) LPAR, which is used to host your Db2 Analytics Accelerator on Z, allows you to define and give access to additional users. Read the following instructions to add additional users to your accelerator LPAR, remove them, or reset their passwords.

About this task

The SSC LPAR does in fact have an entire user management system, which includes users, groups, and roles. However, this system supports just one role at this time. It is the administrator or admin role. At the moment, a user who maintains a Db2 Analytics Accelerator on Z requires this role; a role with fewer privileges does not suffice. Since roles other than admin do not exist, it does not make much sense to create groups that elevate the privileges of otherwise underprivileged users. For that reason, group creation is not covered by this topic.

Procedure

- · To add a user:
 - a) Log on to the Admin UI.

b) In the navigator on the left, click Users.

The **Users** window has several tabs. The tabs give you direct access to the different functional areas of user management: users, groups, roles, and password requests.

- c) Select the **Users** tab.
- d) Click \bigoplus on the upper right of the window.
- e) Enter the user's name in the User name field.
- f) Click Next.
- g) Skip group assignment by clicking Next.
- h) On the **Assign Roles** tab, select the **admin** check box.
- i) Click Create.

The successful user creation is confirmed in a window that contains a HTTPS link for the new user.

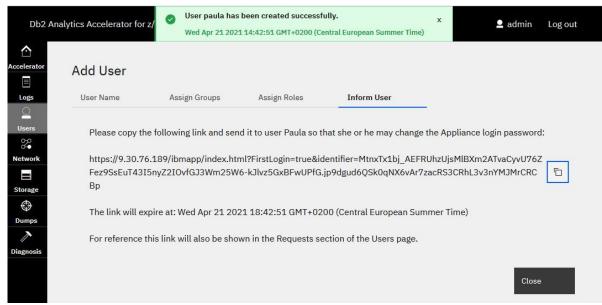


Figure 43. Successful user creation

- j) Copy the link and send it to the new user. The user can then set a password and log on to the **Admin UI**. The link expires after a certain time.
- To reset a user's password:

The following steps can be completed by an administrator or by the user:

- a) Open the Admin UI logon page in a browser and click Forgot your password?.
- b) Enter the user's name in the **User ID** field.
- c) Click Reset.

This generates a password reset request, and you are taken to the logon page.

To complete the following steps, you must be an administrator who can still log on.

- a) Log on to the Admin UI.
- b) In the navigator on the left, click **Users**.
- c) Select the **Request** tab.

You see a list of password reset requests. The number of open requests is indicated on the tab and in the navigator on the left.

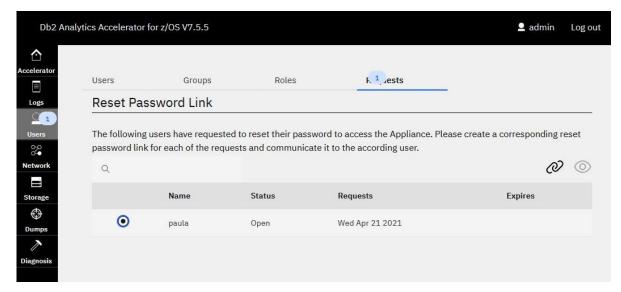


Figure 44. List of requests

- d) Select the request by clicking the appropriate radio button on the left.
- e) Click on the upper right of the window.

A small window with a link is displayed, like the link in Figure 43 on page 113.

f) Close the small window. The **Status** column of the request list now shows **Link created**.



The icon is next to chain icon you clicked before.

The small window with the link opens again. Copy the link and send it to yourself or to the user who needs to reset the password.

Depending on where you sent the link, the following steps are taken by yourself as the administrator, or by the user who needs to reset the password.

a) When the link is followed, it takes you or the user to the **Confirm New Account** window of the **Admin UI**.

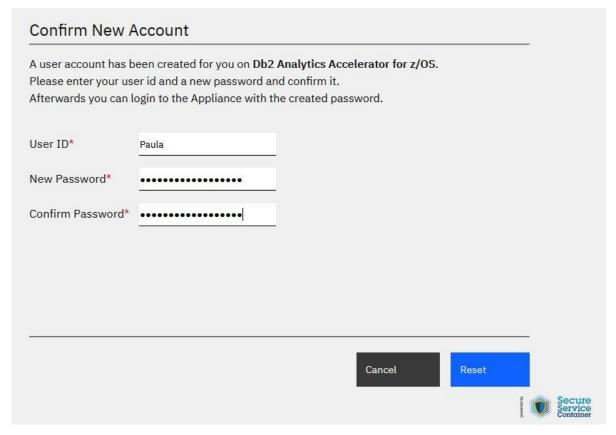


Figure 45. The Confirm New Account window

- b) You or the user must fill in the fields and click **Reset**.
- c) You or the user is taken to the logon page of the **Admin UI**, where the new credentials can be used.
- To remove a user:
 - a) Log on to the Admin UI.
 - b) Select the **Users** tab.
 - c) In the list of **Configured Users**, select the user you want to remove.
 - d) Click 🛅
 - e) In the **Delete User** confirmation window, click **Delete**.

Checking your Db2 Analytics Accelerator on Z setup

From the **Diagnosis** window of the Admin UI, you can start several report and verification tools to check the functioning of your setup.

About this task

A Db2 Analytics Accelerator on Zcommunicates with various other components over your configured network interface.

The connection to the Db2 for z/OS is particularly critical with regard to bandwidth, availability, latency.

Using the tools in the **Diagnosis** window, you can check connections to IP addresses, connections to the Db2 network, and collect system report data, which might be useful for solving a customer case.

Procedure

1. Log on to the Admin UI.

- 2. In the navigator on the left, click Diagnosis.
- 3. The **Diagnosis** window has several tabs. Each tab gives you access to a verification or report tool.
 - · Test your Network tab

On the **Test your Network** tab, you can enter a target IP address in the **IP Address** field can click **Check** to verify whether the target IP address can be reached and whether any packet losses occurred. The result of this ping command is shown in a box underneath.

A second box further down gives the details of the network connection that was checked, that is, the network interface that was used, the source IP address of the accelerator, and the routes to the target IP address. This information might be helpful in detecting network issues during an initial configuration.

To write the results of a check to a text file and download that file, you can click the **Results** button.

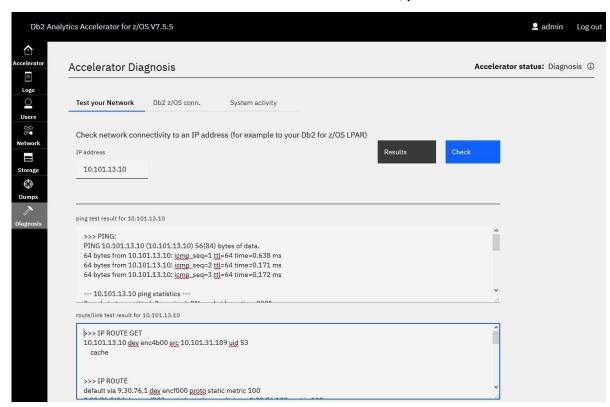


Figure 46. The **Test your Network** tab on the **Accelerator Diagnosis** page

Db2 z/OS conn. tab

When a Db2 subsystem has been connected to your accelerator, the subsystem and the accelerator communicate and exchange data over certain ports.

If you click the **Retrieve** button on the **Db2 z/OS conn.** tab, the IP addresses involved in this communication are reported in the table underneath. The IP address of the Db2 for z/OS LPAR is identified in the first column, and details of this and the other network endpoints are given in the second column.

The **Route check** that you find reported in the third column is a brief test ensuring that the Db2 connection does not use any other than the intended network interface, such as the network of the activation profile.

Finally, the IP address of the Db2 subsystem LPAR is pinged, and a maximum transmission unit (MTU) size check is carried out. This way, you can check whether the size of transferred packages matches the required size. The results of this check are displayed in the fourth column.

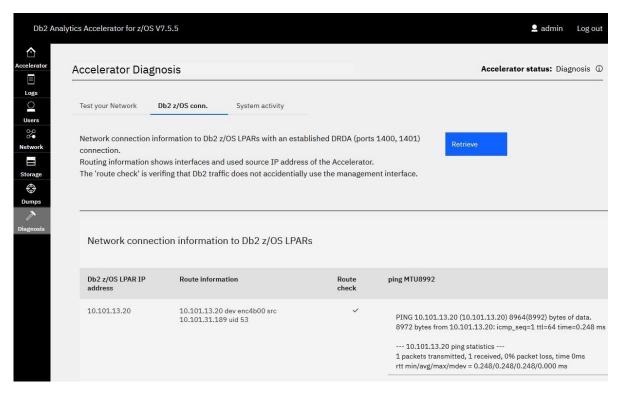


Figure 47. The **Db2 z/OS conn.** tab on the **Accelerator Diagnosis** page

· System activity tab

On this tab, you can capture the load activities, CPU utilization, memory utilization, storage device utilization, and network utilization of your accelerator.

In the **Number of minutes to run** field, enter an integer number. This number represents the number of minutes the report tool will be capturing data for. Then click **Start**.

The results of the capturing process are written to a file. This file is listed in a table at the bottom. The table displays the name of the file, the capturing start, and the capturing end times. It also offers controls to delete or download the file.

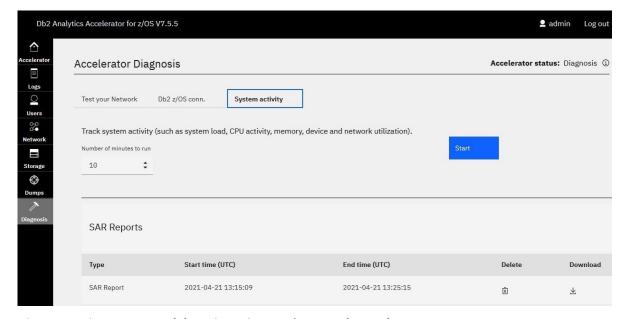


Figure 48. The **System activity** tab on the **Accelerator Diagnosis** page

Replacing the SSL certificate

By default, Db2 Analytics Accelerator on Z uses a self-signed certificate for secure web traffic (https://). This certificate is provided by IBM. To meet security standards that might exist for your organization, you can replace this certificate with a custom secure-sockets-layer (SSL) certificate. This custom certificate can be signed by your organization or by an external certificate authority (CA).

Before you begin

A custom certificate for Db2 Analytics Accelerator on Z must be in PEM format. That is, you need two files in PEM format:

• A file with an extension of key that contains the private key. This is a text file that starts with the entry:

```
"----BEGIN PRIVATE KEY---- "
```

A file with an extension of crt that contains the certificate. This is a text file that contains one or more
of the following entries:

```
"----BEGIN CERTIFICATE----- "
```

You might have a certificate file in PKCS#7, PKCS#12 or DER format. If so, then you must convert the certificate to the PEM format first. A tool like OpenSSL provides you with a conversion function to do this.

Your certificate file might also include the private key. Again, you can use a tool like OpenSSL to extract the private key (in an unencrypted format) from the certificate and save it to a separate file.

Procedure

- 1. Enter the IPv4 address of the Secure Service Container (SSC) LPAR in a web browser (see <u>Figure 12 on page 57</u>)
- 2. On the **Login** page of the Admin UI, log in with the **Master user ID** and **Master password** that you specified as you defined the SSC LPAR. (see Figure 12 on page 57)
- 3. Click the information icon () in the upper right of the window. The icon is available in all windows of the user interface.

The **About** window is displayed:

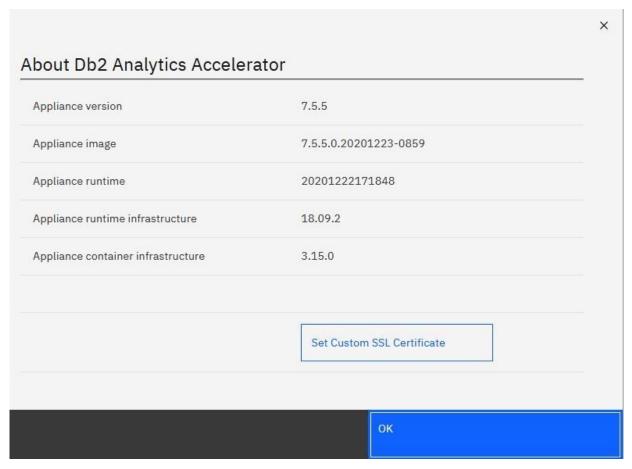


Figure 49. The About window of Db2 Analytics Accelerator on Z user interface

4. In the **About** window, click **Set Custom SSL Certificate**.

You see the following window:

ou may replace the shipped IBM self-signed certificate for a domain name you assigned to your Accelerator's LPA here are several formats for SSL certificates. You need the private key file (*.key) and a certificate file (*.you have a file in PFX format (*.pfx or *.pk12), you have les.	AR. *.crt) in PEM format.
Private key file:	Upload *.key
Certificate file for your LPAR's domain name:	Upload *.crt or *.cer
Delete custom SSL certificate (Restore original certificate).	Restore

Figure 50. The Custom SSL Certificate for Web Admin User Interface

- 5. Click **Upload *.key** and follow the instructions to upload and install your private key (.key) file.
- 6. After returning to the **Custom SSL Certificate for Web Admin User Interface**, click **Upload *.crt or *.cer** and follow the instructions to upload and install your certificate (.crt) file.
- 7. Click **OK** and refresh the browser window.
- 8. Verify that the new certificate has been installed correctly: Click the padlock icon to the left of your browser's address field.

Make sure you find this or similar information in the window that opens:

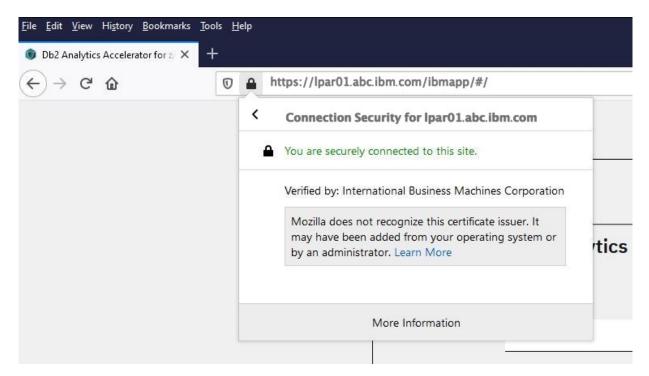


Figure 51. Verifying the correct installation of the certificate

Chapter 9. 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.

Access rights

The various IBM Db2 Analytics Accelerator for z/OS components require different authorizations. Which authorizations a user requires depends on the role that this user has to play. However, in nearly all cases, the rights that users need go beyond the obvious. Unfortunately, this information unit cannot list the required authorizations for all conceivable roles. It therefore restricts itself to two sets of authorizations, one that restricts access to the absolute minimum, and one that gives a user total access to all components (power user).

Access rights for power users

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 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
- Authorization to run the ACCESS DB command on the databases that the tables reside in (needed to refresh Db2 real-time statistics)

Required z/OS power-user access rights

A power user requires the following access rights in RACF and in the z/OS UNIX file system (zFS):

- An OMVS segment is required for the user ID.
- Write access to the /tmp directory (UNIX System Services pipes are created in this directory).
- Read access to /usr/lpp/IBM/aqt/V7R1M0 and all of its subdirectories.

Db2 for z/OS power-user authorizations for administration clients

You might want to enable your power user to run administration clients. 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-shadow tables and the calculation of the overall table size.

Minimum access rights

A user might only be required to view information about an accelerator, the accelerator-shadow tables on the accelerator and their status, or the query history. Such a user needs only minimal access rights. See which access rights such a user requires.

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 the SYSIBM.USERNAMES.NEWAUTHID column.

Required minimum authorizations in Db2 for z/OS

A read-only user requires the following authorizations in Db2 for z/OS:

- EXECUTE on the SYSACCEL.* packages
- EXECUTE on the following functions:
 - DSNAQT.ACCEL_READFILE3
 - DSNAQT.ACCEL_GETVERSION
 - DSNAQT.ACCEL_CONTROL_GETACCELERATORINFO
- EXECUTE on the SYSPROC.ADMIN INFO SYSPARM stored procedure
- MONITOR1 privilege (needed to call the SYSPROC.ADMIN INFO SYSPARM stored procedure internally)
- · DISPLAY privilege

Minimum Db2 for z/OS authorizations for administration clients

The recommended minimum set of authorizations for administration clients is:

- EXECUTE on the following IBM Db2 Analytics Accelerator stored procedures:
 - SYSPROC.ACCEL_TEST_CONNECTION

- SYSPROC.ACCEL_CONTROL_ACCELERATOR
- SYSPROC.ACCEL_GET_QUERY_DETAILS2
- SYSPROC.ACCEL_GET_QUERY_DETAILS
- SYSPROC.ACCEL_GET_QUERIES2
- SYSPROC.ACCEL_GET_QUERIES
- SYSPROC.ACCEL_GET_TABLES_INFO
- SYSPROC.ACCEL_GET_TABLES_DETAILS
- SELECT on the DSNAQT.ACCEL_NAMES view. This privilege is required for the enumeration of accelerators.
- SELECT on the SYSACCEL.SYSACCELERATEDTABLES table.
- SELECT on the following catalog tables of the database management system:
 - SYSIBM.SYSCOLUMNS
 - SYSIBM.SYSCONTROLS
 - SYSIBM.SYSDATABASE
 - SYSIBM.SYSDUMMY1
 - SYSIBM.SYSRELS
 - SYSIBM.SYSTABLEPART
 - SYSIBM.SYSTABLES
 - SYSIBM.SYSVIEWS

Read access to these tables is required for 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, an administration client, 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 administration clients 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 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

- 1. To create and bind the Db2 packages, follow the instructions on this website:
 - https://www.ibm.com/support/knowledgecenter/SS62YD_4.1.1/com.ibm.datatools.qrytune.configothers.doc/topics/enabledb2zfromclient_ds.html
- 2. To grant the EXECUTE privilege to other users, proceed as follows:
 - a) In the Data Source Explorer in the IBM Data Studio client, expand the connection to the subsystem.
 Right-click the icon for the subsystem and select Analyze and Tune > Configure for Tuning > Advanced Configuration and Privilege Management.
 - 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 your administration client, 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 other administration clients in the same way.

Procedure

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

https://www.ibm.com/support/knowledgecenter/SS62YD_4.1.1/com.ibm.datatools.qrytune.configothers.doc/topics/db2z_configure_ds.html

In addition to these, you need an EXPLAIN table called DSN_QUERYINFO_TABLE. This table is described here:

https://www.ibm.com/support/knowledgecenter/SSEPEK_10.0.0/com.ibm.db2z10.doc.perf/src/tpc/db2z_dsnqueryinfotable.dita

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 Accelerator** wizard. However, you can only start this wizard after connecting to a Db2 subsystem.

About this task

No matter which administration client you use, 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** or by selecting the profile from the drop-down list of the top banner in IBM Data Server Manager.

Procedure

- 1. Start IBM Db2 Analytics Accelerator Studio or IBM Data Server Manager.
- 2. If the **Welcome** screen is displayed, close it.
- 3. Open the wizard or window to add a database connection:

IBM Db2 Analytics Accelerator Studio	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 .
IBM Data Server Manager	 a. Click in the upper left. b. Select SETTINGS > Manage Connections. c. Select Add > Add a database.

4. Give the database connection profile a name:

IBM Db2 Analytics Accelerator Studio	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.
IBM Data Server Manager	Type the name in the Database connection name field.

5. Select the database management system:

IBM Db2 Analytics Accelerator Studio	From the Select a database manager list, select Db2 for z/OS .
	Make sure that in the JDBC driver drop-down list, IBM Data Server Driver for JDBC and SQLJ (JDBC 4.0) Default is selected.
IBM Data Server Manager	From the Data server type drop-down list, select Db2 for z/OS .

6. Specify the **Location** name of the Db2 subsystem. This is the name of the Db2 subsystem in the Db2 system table SYSIBM.LOCATIONS.

IBM Db2 Analytics Accelerator Studio	Type the name in the Location field.
IBM Data Server Manager	Type the name in the Location field.

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.

7. Specify the host name or IP address of the computer on which the database server is located:

IBM Db2 Analytics Accelerator Studio	Type the host name or IP address in the Host field.
IBM Data Server Manager	Type the host name or IP address in the Host name field.

8. Specify the port number for connections to the database server:

IBM Db2 Analytics Accelerator Studio	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.
IBM Data Server Manager	In the Port number field, type the number of the port that was configured during the installation of IBM Data Server Manager. By default, this is 11511 for secure (https://) connections.

9. Select specific options of your administration client:

IBM Db2 Analytics Accelerator Studio	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.
IBM Data Server Manager	Select:
	a. Enable operation to allow a user (ID) to administer, tune, and query the connected Db2 subsystem.
	b. Enable data collection to allow a user (ID) to collect data from the connected Db2 subsystem.
	Both options are required for successful operation of IBM Db2 Analytics Accelerator. You specify the corresponding user IDs on the next tab.

- 10. If you use IBM Data Server Manager, click the **Credential** tab.
- 11. Specify the user ID of the logon user or operation user:

IBM Db2 Analytics Accelerator Studio	In the User name field, type the user ID that you want to use to log on to the database server.
IBM Data Server Manager	In the group box labeled Operation credentials , you find a User ID field. Type the ID of the logon user in this field.

Note that you can only use your administration client successfully if the logon user has sufficient rights to run the stored procedures that are called by the administration client 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.

12. Specify the password belonging to the logon user ID or operation user ID:

IBM Db2 Analytics Accelerator Studio	 In the Password field, type the password of the logon user. 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. • Leave the Default schema field blank.
IBM Data Server Manager	a. Under Operation credentials , in the Password field, type the password of the logon user.

	b. To save the logon password so that you do not need to provide it each time you connect to the same Db2 subsystem, select Save credentials to repository .
--	---

13. Specify a secondary user ID (optional for IBM Db2 Analytics Accelerator Studio), or the data collection user ID for IBM Data Server Manager:

IBM Db2 Analytics Accelerator Studio	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.
IBM Data Server Manager	Under Data collecting credentials , in the User ID field, type the ID of the data collection user (this user is related to the Enable data collection option in step "Creating a database connection profile" on page 126).

14. Test the connection or connections:

IBM Db2 Analytics Accelerator Studio	Click Test Connection to check if you can log on to the database server.
IBM Data Server Manager	a. In the Operation credentials group box, click the Test Connection button to check if your operation user can log on to the Db subsystem.
	b. In the Data collecting credentials group box, click the Test Connection button to check if your data collection user can log on to the Db subsystem.

15. Save your database connection profile:

IBM Db2 Analytics Accelerator Studio	Click Finish .
IBM Data Server Manager	Click OK .

Results

After creating the profile, your administration client 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 current view of your administration client (more than one minute), you can set a filter to limit the number of schemas to be loaded:

IBM Db2 Analytics Accelerator Studio	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.
IBM Data Server Manager	1. Select > ADMINISTER > Schemas.
	You can use the Filter by Name field to quickly filter the list of schemas using ordinary characters and wildcard characters (% for strings and _ for individual characters). For example, entering A% will display just the schemas whose names start with the letter A.
	Another way is to use the filter editor:
	a. Click the funnel icon in the upper right.
	b. In the Filter window, you can define many types of filter rules. For example, to exclude all schemas whose names start with the characters BLU:
	i) Leave the first drop-down list unchanged.
	ii) From the second drop-down list, select does not contain.
	iii) In the text entry field, type BLU.
	iv) Click Apply . The Filter window closes, and the list of schemas is refreshed. Notice that the funnel icon is completely filled with color when a filter has been applied.

Related reference

Access rights for power users

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 your administration client

Follow the steps in this section to see if you can connect to a Db2 subsystem from your administration client.

Procedure

Depending on your administration client, use one of the following approaches:

IBM Db2 Analytics Accelerator Studio	In the Administration Explorer , double-click the icon representing the Db2 subsystem. For example:
IBM Data Server Manager	IBM Data Server Manager tries to reconnect to the previously connected Db2 subsystem by default. If the reconnection fails, or you want to connect to a different subsystem, you can select the appropriate connection name from the drop-down list on the top banner. For example:
	■ IBM Data Server Manager DB0ABC12 • ①
	Help / Welcome srv1.abc.xy2.com DB0ABC12 € Connections
	Alternatively:
	a. Click and select SETTINGS > Manage Connections. b. Select the box in front of a listed connection.

Results

You see the following results if the connection was successful:

IBM Db2 Analytics Accelerator Studio	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 .
IBM Data Server Manager	The views on the main menu, under the ADMINISTER submenu, are populated.

Examples

1	After a successful connection, you can see the following folder hierarchy in the Administration
	Explorer:

	STLEC1 (DB2 for z/OS V9.1 (New-Function Mode)) Accelerators Aliases Application Objects Auxiliary Tables Constraints Databases Indexes MQTs Schemas Sequences Storage Groups Synonyms Tables Table Spaces	
IBM Data Server Manager	Tables in the connected Db2 subsystem are	
	displayed after selecting > ADMINISTER > Tables:	
	IBM Data Server Manage	er
	Administer / Tables	
	Create	
	TABLE SCHEMA	
	BITT TC0_IU	
	BITT_HISTORY TC0_IU	
I		

Adding accelerators

Adding an accelerator to your configuration is a two-step process. First, you must obtain a pairing code from the IBM Db2 Analytics Accelerator Console. Second, you enter this paring code along with the IP address and the name of the accelerator in the **Add Accelerator** wizard of your administration client.

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 the administration client must coordinate their steps.

Procedure

1. Ask the network administrator or the person who did the TCP/IP setup for the IP address of the accelerator. Make a note of this information. You need to enter it as you complete the steps that follow.

For IBM Db2 Analytics Accelerator on an IBM Integrated Analytics System, this is the virtual IP or wall IP address.

For Db2 Analytics Accelerator on Z, this is the IP address of the network that you labeled DB2 in the **Appliance Installer**.

2. Log on to the IBM Db2 Analytics Accelerator Console by using **telnet** or **ssh**. The preferred method is **ssh**.

For more information:

- "Using telnet to log on to the IBM Db2 Analytics Accelerator Console" on page 199
- "Using ssh to log on to the IBM Db2 Analytics Accelerator Console" on page 200
- 3. Press the Pause key, then Enter to display the following screen:

4. Type 1 and press Enter to display the submenu:

- 5. Type 1 and press Enter:
- 6. 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 the administration client, 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 : 203.0.113.8
   Port : 1400
   Port (AT-TLS) : 11400
   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.

- 7. Make a note of the following information on the **console**:
 - · Pairing code
 - IP address
 - Port (for unencrypted network communication). Use this port if you are not sure.
 - Port (AT-TLS) (for encrypted network communication).

The use of encryption requires extra configuration steps on the accelerator and on the participating z/OS LPARs.

- 8. Press Enter to return to the main menu of the **console**.
- 9. 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 specify the IP address, the 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 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 your administration client.
- 2. Go to the folder or page that lists the accelerators:

IBM Db2 Analytics Accelerator Studio	Select the Accelerators folder in the Administration Explorer .
IBM Data Server Manager	a. Click . b. Select > ADMINISTER > Accelerators.

3. Start the Add Accelerator wizard:

IBM Db2 Analytics Accelerator Studio	 a. On the menu bar of the Object List Editor, click the downward-pointing arrow next to the green plus sign. b. From the drop-down menu, select Add Accelerator.
IBM Data Server Manager	Click Add.

4. 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.



Attention:

- If you want to connect more than one Db2 subsystem to the same accelerator, you must use a different location name for each pairing. For each IP address and port combination, the assignment to an accelerator must be logically unique. If you disregard this rule and use the same location name, you will lose the connection to the previously paired Db2 subsystem, and the data on the accelerator will be deleted. This mechanism was implemented on purpose because it ensures that a newly added accelerator is always "clean".
- Reuse an accelerator name only if you are absolutely sure that the accelerator was stopped before it was removed.

Background: When an accelerator is paired with Db2 for z/OS, Db2 starts a heartbeat thread to monitor vital accelerator functions. This heartbeat thread is bound to the IP address of the accelerator. It continues to run if you remove the accelerator without stopping it before. The continuing thread is canceled only when you shut down Db2.

Your new accelerator will take over the still running heartbeat thread of the removed accelerator if all of the following conditions are true:

- a. You did not stop the accelerator before the removal.
- b. You did not shut down and restart Db2 for z/OS after the removal.
- c. You used the name of the removed accelerator for the new accelerator.

If thereafter problems with the new accelerator occur, and you need to check the heartbeat information, you might fail to notice that this information does not pertain to the accelerator you are scrutinizing.

- 5. In the **Pairing code** field, type the pairing code.
- 6. In the **IP address** field, type the IP address of the accelerator.
- 7. Specify the correct port:

IBM Db2 Analytics Accelerator Studio	In the Port field, type 1400. This is the fixed port for network communication between the z/OS data server and the accelerator.
IBM Data Server Manager	The Port field is prefilled with port number 1400. Leave this as is, unless you know for sure that a different port has been configured.

- 8. Click **Test Connection** to check whether the accelerator with the given address can be connected to.
- 9. Click OK.

A connection test is carried out.

Related tasks

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.

Related reference

Access rights for power users

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

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

Chapter 10. 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

Defining accelerator-shadow tables

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

Procedure

1. Open the list of accelerators:

IBM Db2 Analytics Accelerator Studio	In the Administration Explorer , go to the Accelerators folder.
IBM Data Server Manager	a. Click in the upper left. b. Select ADMINISTER > Accelerators.

2. Select the accelerator that you want to add tables to:

IBM Db2 Analytics Accelerator Studio	a. In the Object List Editor , double-click the accelerator name. b. In the Accelerator view, click Add on the toolbar, which is located above the list of
IBM Data Server Manager	tables. a. Select a row in the list of accelerators. b. Click the Manage button on the toolbar. c. Click the Add button on the toolbar.

- 3. In the **Add Tables** or **Select Tables to Add** wizard, select the tables that are referenced in your test query.
- 4. Confirm by closing the **Add Tables** or **Select Tables to Add** wizard.

IBM Db2 Analytics Accelerator Studio	Click Finish.
IBM Data Server Manager	Click OK .

Results

The newly created accelerator-shadow tables are shown in the **Accelerator** view of IBM Db2 Analytics Accelerator Studio or in the list of tables of IBM Data Server Manager.

Loading and enabling tables

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

About this task

Loading or updating the data in accelerator-shadow tables for query acceleration requires your administration client 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. Select the accelerator that holds tables referenced by your test query:

IBM Db2 Analytics Accelerator Studio	a. In the Administration Explorer, select the Accelerators folder. b. In the Object List Editor, double-click the accelerator containing the tables that you want to load.
IBM Data Server Manager	 a. Click in the upper left. b. Select ADMINISTER > Accelerators. c. In the list of accelerators, select the accelerator containing the tables you want to load. d. Click Manage.

2. You see a list of the tables on the accelerator. Select the accelerator-shadow tables that you want to load

To see the tables, you might have to expand the schema nodes first, by clicking the plus sign in front of a schema name. Selecting an entire schema will select all accelerator-shadow tables belonging to that schema for loading.

3. Click the **Load** button on the toolbar.

The Load Tables window opens.

4. Select the appropriate tables or partitions.

IBM Db2 Analytics Accelerator Studio	Select the appropriate check boxes in front of the table or partition names, or use the buttons on the right:
	Select All Selects all tables and partitions that are shown in the Load Tables window.
	Deselect All Deselects all tables and partitions.
	Select Recommended Selects all accelerator-shadow tables and partitions for which a load or reload is recommended in the Load Recommended column.

	Expand All Expands all table nodes so that the partitions come into view.
	Collapse All Collapses all table nodes so that the table partitions are hidden from view.
	Next Unknown Jumps to and selects the next accelerator- shadow table or partition with a status of Unknown - no statistics in the Load Recommended column. This might prove useful if the list is crowded.
IBM Data Server Manager	Select the appropriate check boxes in front of the table or partition names.
	To quickly select just the tables and partitions for which a load is recommended, click the Select Recommended button.

- 5. Do not change value of the **Lock DB2 tables while loading** drop-down list. Leave the setting **None**.
- 6. Leave After the load, enable acceleration for disabled tables selected.
- 7. 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. In IBM Db2 Analytics Accelerator Studio, the status of the accelerator is shown on top of the **Accelerator** view. In IBM Data Server Manager, it is shown in the lower half of the window after selecting the accelerator in 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. Open the SQL editor:

IBM Db2 Analytics Accelerator Studio	a. On the toolbar of the Administration Explorer , click the downward-pointing arrow next to the New button. b. Select New SQL Script from the menu.
IBM Data Server Manager	a. Click in the upper left.

b. Select RUN SQL .
Note: This option is available only if an active database connection has been selected from the drop-down list on the top banner.

2. Type or paste your SQL statement in the text area:

IBM Db2 Analytics Accelerator Studio	In the blank space of the Script <x>.sql workspace that opens on the upper right, type the following statement on the first line:</x>
	SET CURRENT QUERY ACCELERATION = ENABLE;
IBM Data Server Manager	In the text area that opens in the upper half of the window, type the following statement on the first line:
	SET CURRENT QUERY ACCELERATION = ENABLE;

- 3. Complete your query. Type a simple SELECT statement, such as SELECT_SUM(QUANTITY) FROM SALES.FIGURES_FACT.
- 4. Run your SQL script:

IBM Db2 Analytics Accelerator Studio	a. Click D b. Click Finish .
	The SQL Results pane in the lower left shows you whether the query ran successfully.
IBM Data Server Manager	Select Run > Run all from the toolbar.

5. To check whether the accelerator was used, follow these steps:

IBM Db2 Analytics Accelerator Studio	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.
IBM Data Server Manager	a. Click in the upper left.
	b. Select ADMINISTER > Accelerators.
	c. In the list of accelerators, select the accelerator the query was submitted to.
	d. Click Manage .
	e. Click the Queries tab.
	i .

Results

You only see queries that were run on the selected accelerator. In-house Db2 queries are not listed. Your test query should appear at the top of the table or list. 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 for z/OS: User's Guide*. You find a link under **Related information** at the end of this topic.

Chapter 11. Incremental updates

The incremental update function of IBM Db2 Analytics Accelerator for z/OS allows you to update accelerator-shadow tables continually. Changes to the data in original Db2 for z/OS tables are thus propagated to the corresponding target tables with a high frequency and just a brief delay. This way, query results from an accelerator are always extracted from recent, close-to-real-time data.

The product offers the following options to implement incremental updates:

- · IBM Integrated Synchronization
- IBM InfoSphere Change Data Capture for z/OS (CDC)

If you chose CDC as your implementation option, you already installed the necessary CDC components by completing the steps in "Installing libraries with IBM Db2 Analytics Accelerator support" on page 25. 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.

Incremental update processing requires a setup on the mainframe as well as on the accelerator. For change processing on the mainframe, the following feature or product exists:

IBM Integrated Synchronization

IBM Integrated Synchronization is a built-in product feature that makes it easy to set up incremental updates. It does not require IBM InfoSphere Change Data Capture for z/OS (CDC), which is, although bundled with IBM Db2 Analytics Accelerator, an extra product. IBM Integrated Synchronization therefore requires fewer customization and activation steps.

In addition, IBM Integrated Synchronization has further advantages:

- Reduced complexity.
- · Updates are processed very fast.
- Reduced CPU consumption on the mainframe because most of the processing is done on the client side (accelerator).
- On the mainframe, the workload can be handled by IBM z Systems® Integrated Information Processors (zIIPs).
- Speedier support case processing through better support channel integration

Compatibility

You need not remove existing CDC implementations as long as they serve different Db2 subsystems. A single accelerator can be the common replication target of different Db2 subsystems using either CDC or IBM Integrated Synchronization. However, this type of setup is not recommended.

Feature prerequisites

To use IBM Integrated Synchronization, you need:

- Db2 12 for z/OS with the latest fix for APAR PH06628 (PTF UI63356 at the time of this writing) installed, running at function level V12R1M500. You can retrieve a list of all recommended fixes from the Prerequisites and Maintenance website.
- Distributed data facility (DDF) with a secure port, configured for network encryption through AT-TLS
- For data sharing groups: ability to always connect to the same member using one of the following options:
 - A member-specific secure port that has a local alias
 - The high-availability setup for data sharing groups by means of DDVIPA and the TIMEDAFFINITY option. See "Suggestion for a high-availability setup" on page 159.

• IBM Db2 Analytics Accelerator for z/OS Version 7.5.0 or later.

Prerequisites for tables and hints

- Tables to be included must have a unique constraint (primary key or primary index). If such a key does not exist in the accelerator-shadow table or cannot be determined, you must redefine the table and specify such a key. You cannot do this from IBM Db2 Analytics Accelerator Studio; you must run the corresponding stored procedure (SYSPROC.ACCEL_ADD_TABLES) with a special option. The columns that you choose for the key must contain unique values or form such values when they are combined. The values of the selected key columns will then be written to a hidden column in the accelerator-shadow table. This column, called *informational unique constraint*, then serves as the primary key or primary index. For more information, see the entry about the table_specifications parameter of the SYSPROC.ACCEL_ADD_TABLES stored procedure in the IBM Db2 Analytics Accelerator for z/OS: Stored Procedures Reference.
- 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-shadow tables.

- The Db2 source tables of your accelerator-shadow tables have an attribute named DATA CAPTURE. 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.
 - The DATA CAPTURE attribute is set by an ALTER TABLE statement, which is run as part of the SYSPROC.ACCEL_SET_TABLES_REPLICATION stored procedure. However, the attribute can only be set successfully if the ID of the user who runs the stored procedure has ALTER TABLE authorization. If this is not the case, a database administrator must set the attribute for all tables in Db2 for z/OS.
- You might have to reload or even remove tables from an accelerator after the execution of an ALTER TABLE or ALTER TABLESPACE statement in Db2 for z/OS. For more information, see What to do after changing a table or table space in Db2 for z/OSin the IBM Db2 Analytics Accelerator for z/OS: User's Guide.

Intervention-free handling of faulty tables

A table in error is suspended from the replication process. This means updates are no longer applied and therefore the table data might be outdated.

To prevent queries from running against suspended tables, query acceleration must be disabled for these tables. This is done automatically, and each such event is reported as a warning in the event viewer for incremental updates. An intervention on your part is not required.

All that is required as a preparation is that a few extra privileges are granted. If IBM Integrated Synchronization is already installed, have a look at the steps "4" on page 155 and "6" on page 155 in the section "Assigning user IDs" on page 154. These steps describe how to grant the required privileges.

Important:

- If the required privileges are missing, IBM Integrated Synchronization cannot be enabled for a Db2 subsystem.
- For already enabled Db2 subsystems, missing privileges lead to error messages in the event viewer for incremental updates. These messages will state that query acceleration could not be disabled.

Schema changes

Schema changes are supported in the sense that you need not remove, redefine, re-enable replication for, and finally reload an accelerator-shadow table after a schema change to an original Db2 for z/OS table. The schema change will be detected by the IBM Integrated Synchronization function, and will, if necessary, initiate a modification of the target accelerator-shadow table so that incremental updates can continue.

While the schema change support includes the adding of all sorts of columns (even column types not supported by IBM Db2 Analytics Accelerator), a few restrictions exist none the less:

- Data-definition-language (DDL) modifications that effect a schema change cannot be written to the spill queue. That is, if a manual load operation suspends an incremental update that includes a schema change, the affected table will be set to the *Error* state, and incremental updates will not continue for this table (no draining of the spill queue).
- Schema changes are not supported if key columns are affected by the change.
- Schema changes are not supported for tables of the following types:
 - Tables for which federated access has been configured
 - Accelerator-only tables
 - Tables that have been archived by the high-performance storage saver (HPSS)
- Schema changes are not supported for columns of the following data types:
 - Distinct types (user-defined types based on existing, built-in Db2 for z/OS data types)
 - Large object types (BLOB, CLOB, or DBCLOB)
- Mixed DDL modifications cannot be handled, only modifications of the same type can. That is, an ALTER TABLE statement that adds two columns to a table is allowed, but not a statement that performs two different operations on a table, such as ADD COLUMN and SET DATA TYPE.
- DROP COLUMN schema changes are not supported.
- If you run a WAITFORDATA query on a table right after changing the table schema (ALTER TABLE), the query might end with SQLCODE -4742 and reason code 14 because the schema change had not been processed on the accelerator at the time the query started. Avoid running a WAITFORDATA query immediately after a schema change.
- After 255 updates of the table version (data row format version), replication is disabled for a table. That is, replication is disabled after 255 schema changes if the version is updated after each change.
- You cannot add columns of data type ROWID.

Note: The schema change feature works best if the workload is low. Remember that tables are modified structurally. A high number of queries running at the same time might lead to lock situations and timeouts for both, the queries and the schema-modifying DDL statements. As a consequence, tables are removed from the replication process, and queries end with an error.

Restrictions

- It is not possible to enable tables if the table name or the schema name contains GB18030 characters (Simplified Chinese) of Unicode plane 2 (U+20000-U+2FFFF: Supplementary Ideographic Plane). Trying to do so results in an error.
- To be able to load or reload replication-enabled tables, replication must be started. Loading or reloading replication-enabled tables does not work if replication is stopped. In the unlikely situation that a replication-enabled table prevents a replication start, proceed as follows:
 - 1. Disable replication for this table.
 - 2. Remove the table from the accelerator.
 - 3. Restart replication.
 - 4. Redefine (re-add) he table to the accelerator.
 - 5. Re-enable replication for this table.

6. Load the table.

Related concepts

Suggestion for a high-availability setup

A high-availability setup ensures that your IBM Integrated Synchronization process remains intact in case of a network connection or similar failure. Several ways for such a setup are thinkable. The setup suggested here is based on a DDVIPA network and uses a Db2 data sharing group. It is relatively easy to implement and, once set up, does not require any manual interverention on your part because the maintenance is fully automatic.

Defining a secure network port

The IBM Integrated Synchronization feature uses the distributed data facility (DDF) to connect to Db2 for z/OS. Connections must be encrypted because the log records that feed the incremental update stream might contain sensitive data. Unencrypted connections will be refused. The first step on the way to encrypted connections is to check whether a secure port exists, and to define one if it doesn't.

Procedure

1. To check if a secure port exists, run the **-DIS DDF** command from TSO.

In this example, the command output shows that the secure port (SECPORT) is set to 15111:

```
RESPONSE=XYZ1
DSNL080I -DB11 DSNLTDDF DISPLAY DDF REPORT FOLLOWS:
DSNL081I STATUS=STARTD
DSNL082I LOCATION LUNAME GENERICLU
DSNL083I LOCDB11 NATIVE.IPWADB11 -NONE
DSNL084I TCPPORT=12511 SECPORT=15111 RESPORT=15011 IPNAME=-NONE
```

If the value of SECPORT is 0, you must define a secure port.

- 2. To define a secure port, use one of the following methods:
 - Set the ZPARM DRDA SECURE PORT accordingly on the DSNTIP5 panel.
 - Specify the secure port in the boot strap data set (BSDS) by using the DSNJU003 change log inventory utility. This method requires a restart of Db2 for z/OS because the BSDS can only be updated while Db2 is down. For more information, follow the appropriate link at the bottom.
- 3. To work with data sharing groups, IBM Integrated Synchronization requires a stable connection to the log reader task of the member that the session was started from.

That is, for incremental update processing, the connection must always go to the same member. To achieve this, use one of the following options:

Option 1:

This is a full data sharing setup that uses DDVIPA with the TIMEDAFFINITY option. With this type of setup, a dedicated SECPORT is assigned to the Db2 members and an exclusive location alias for IBM Integrated Synchronization is defined on each member. This means that all members have the same SECPORT.

- a. Define a dedicated location alias and a secure port (SECPORT) for IBM Integrated Synchronization on all data sharing members. If you already use a SECPORT for other workloads, choose a different SECPORT for IBM Integrated Synchronization.
- b. Set up high availability for IBM Integrated Synchronization as described in "Suggestion for a high-availability setup" on page 159 and start the location alias on all data sharing members that participate in the high-availability setup.

Option 2:

In a setup like this, the data sharing setup uses a DDVIPA network, but the TIMEDAFFINITY option is not set. A dedicated SECPORT is assigned to the Db2 members and an exclusive location alias for IBM Integrated Synchronization is defined on each member. This means that all members have the same SECPORT.

However, the location alias is started on a single member only. To shift the workload from a member A to a member B, you must stop the location alias on member A and then start the alias on member B.

- a. Define a dedicated location alias and a secure port (SECPORT) for IBM Integrated Synchronization on all data sharing members. If you already use a SECPORT for other workloads, choose a different SECPORT for IBM Integrated Synchronization.
- b. Start the location alias on the member that provides the Db2 log data for IBM Integrated Synchronization. Make sure that the location alias is started on this member only. If this member or the network fails, or if you have to shut down the member for maintenance purposes, the location alias can be started on another member so that IBM Integrated Synchronization can continue to work.

With a configuration like this, all members used for connections can share the same DDVIPA IP address. This makes the handling and the setup of SSL certificates easier because a single, shared certificate can be used.

To define a secure port for every member that is supposed to serve as a connection endpoint for IBM Integrated Synchronization, you have the following options:

- Using the -MODIFY DDF command. This is the recommended method because you need not stop Db2 for z/OS.
- Defining the location aliases in the BSDS. If you prefer this method, follow the links at the end.

The following example shows how to use the **-MODIFY DDF** command to define and start a location alias for a single member DB2A that listens on secure port 15011:

```
-DB2A MODIFY DDF ALIAS(DB2AMBR1) ADD
-DSNL300I -DB2A DSNLTMDF MODIFY DDF REPORT FOLLOWS:
-DSNL302I ALIAS DB2AMBR1 IS SET TO ADD
-DSNL301I DSNLTMDF MODIFY DDF REPORT COMPLETE

-DB2A MODIFY DDF ALIAS(DB2AMBR1) SECPORT(15011)
-DSNL300I -DB2A DSNLTMDF MODIFY DDF REPORT FOLLOWS:
-DSNL302I ALIAS DB2AMBR1 IS SET TO SECPORT 15011
-DSNL301I DSNLTMDF MODIFY DDF REPORT COMPLETE

-DB2A MODIFY DDF ALIAS(DB2AMBR1) START
-DSNL300I -DB2A DSNLTMDF MODIFY DDF REPORT FOLLOWS:
-DSNL302I ALIAS DB2AMBR1 IS SET TO START
-DSNL302I ALIAS DB2AMBR1 IS SET TO START
-DSNL301I DSNLTMDF MODIFY DDF REPORT COMPLETE

-DSNL301I DSNLTMDF MODIFY DDF REPORT COMPLETE
-DSNL314I -DB2A DSNLILNR THE ALIAS DB2AMBR1 IS STARTED
```

Running **-DIS DDF** generates the following screen output:

```
-DIS DDF
-DSNL080I -DB2A DSNLTDDF DISPLAY DDF REPORT FOLLOWS:
-DSNL081I STATUS=STARTD
-DSNL082I LOCATION LUNAME GENERICLU
-DSNL083I LOCDB2 NATIVE.APP2DB2A -NONE
-DSNL084I TCPPORT=446 SECPORT=12000 RESPORT=5001 IPNAME=-NONE
...
-DSNL087I ALIAS PORT SECPORT STATUS
-DSNL088I DB2AMBR1 0 15011 STARTD
...
-DSNL099I DSNLTDDF DISPLAY DDF REPORT COMPLETE
```

Related information

DSNJU003 (change log inventory)
-MODIFY DDF command
Configuring the Db2 server for SSL

AT-TLS configuration

IBM Integrated Synchronization uses the z/OS distributed data facility (DDF). Network connections to the DDF must be encrypted using the secure socket layer (SSL) encryption standard. These encrypted

connections are not processed by Db2 for z/OS, but by a TCP/IP component of the z/OS operating system, which is called AT-TLS.

To use AT-TLS, the following prerequisites must be met:

- The IBM Encryption Facility for z/OS (ISCF) must be installed.
- Transparent Transport Layer Security (TTLS) must be set as the standard to be used in the configuration statement of your TCPCONFIG data set (at the highest level).
- The policy agent (PAGENT) must be started.
- A server certificate is required. This certificate must be added to a dedicated RACF key ring.

The instructions in the following sections presuppose that ISCF, TCP/IP, and the policy agent are already set up and running.

Tip: You can complete the first four of the following steps (from *Modifying the RACF profile* to *Exporting and transferring the signer certificate*) by customizing and running the sample JCL job AQTSSLSA. This sample job is a member of the SAQTSAMP data set, which comes with the stored procedures of the product.

Related information

DB2 for z/OS: Configuring TLS/SSL for Secure Client/Server Communications

Modifying the RACF profile

Add definitions for the key ring and access permissions to the RACF profile as shown. The profile meant here is the RACF profile in the z/OS system or LPAR that contains the Db2 subsystem or data sharing group you want to configure IBM Integrated Synchronization for.

Example

The following example shows a JCL that creates the required RACF objects and grants the required permissions for the user who runs the Db2 started tasks. The name of this user is DB2USER. Adapt this example according to your needs.

```
//CERTAUTH JOB ,
// MSGLEVEL=(1,1),REGION=0M
//CRTCRT EXEC PGM=IKJEFT01
//SYSTSPRT DD SYSOUT=*
//SYSUADS DD DSN=SYS1.UADS,DISP=SHR
//SYSLBC DD DSN=SYS1.BRODCAST,DISP=SHR
//SYSTSIN DD *
SETROPTS CLASSACT(DIGTCERT DIGTRING)
RDEFINE FACILITY IRR.DIGTCERT.LISTRING UACC(NONE)
RDEFINE FACILITY IRR.DIGTCERT.LIST UACC(NONE)
PERMIT IRR.DIGTCERT.LIST CLASS(FACILITY) ID(DB2USER) ACCESS(CONTROL)
PERMIT IRR.DIGTCERT.LISTRING CLASS(FACILITY) ID(DB2USER) ACCESS(READ)
SETR RACLIST (DIGTRING) REFRESH
SETR RACLIST (DIGTCERT) REFRESH
SETR RACLIST (FACILITY) REFRESH
```

Tip: You can add these commands and the commands in the sections that follow to the same JCL, so that finally, only one job needs to be submitted.

Creating certificates

To continue with the setup, you need a signer certificate and a server certificate. Use the signer certificate to sign the server certificate. The signer certificate can be a certificate from an external certificate authority (CA), an internal CA, or a self-signed certificate.

About this task

SSL connections require a certificate that identifies the server and that can be used to encrypt the network connections to this server.

In the z/OS operating system, such certificates are managed by RACF, and are stored in a RACF object called a key ring. You must specify this key ring in your TCPCONFIG data set, so that it can be used by the AT-TLS component.

The following diagram contains a topology map, which details where the certificates are located, where they go, and how components interact with these certificates.

The diagram shows an IBM Integrated Synchronization setup for a standalone Db2 subsystem *Db2 A*. There is another standalone Db2 subsystem *Db2 B* in the same LPAR. This subsystem is not connected to an accelerator.

The diagram also includes a certificate for encryption of data in motion (Acc.cert), which is not required for IBM Integrated Synchronization. However, it has been included to show its location along with the other certificates within the same RACF infrastructure.

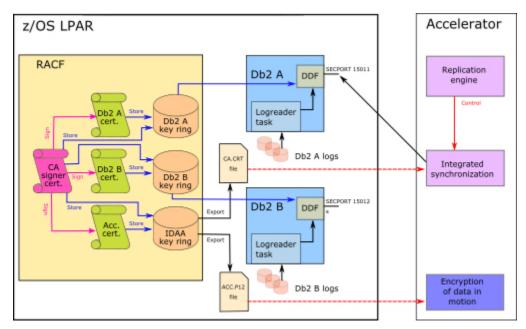


Figure 52. Your guide through the certificate jungle

The certificate that you use can be a certificate issued by an external or an internal certificate authority (CA). Certificates issued by known external CAs have the advantage that they can be validated correctly by connecting clients without completing any prerequisite steps, provided that the CA is known to the client application. Since clients often use predefined libraries that are distributed freely, this is, in many cases, not a problem.

Libraries of internal CAs can usually be linked by client applications within an organization. This also allows you to tap resources of an established infrastructure.

However, if you do not have access to certificates of an external or internal CA, you can create a so-called self-signed certificate. Such certificates require an import on the client side (the accelerator) because the client application needs to "know" whether the certificate and the server identity behind it can be trusted.

The steps in the upcoming sections provide you with all the information needed to create a self-signed certificate and to set up a working AT-TLS connection using this certificate.

If you already have a certificate that you can reuse for this purpose, you can skip the steps that deal with the certificate creation and just adapt the steps in which the correct certificate must be specified.

Procedure

 In accordance with the following example, add RACF commands to a JCL job that finally create a signer certificate. This certificate is used to sign the server certificate that you create in a later step. If you already have a certificate that can be used for this purpose (from an external or an internal CA), you can skip this step.

```
RACDCERT CERTAUTH -
GENCERT -
SUBJECTSDN(OU('DB2 SERVER CA') -
O('IBM') -
C('USA')) -
NOTAFTER(DATE(2030-12-31)) -
WITHLABEL('DB2 SERVER CA') -
KEYUSAGE(CERTSIGN)
```

The signer certificate created by example bears the label 'DB2 SERVER CA'. RACF uses this label to reference the certificate. You must specify the name of this label in a few upcoming configuration steps. Modify the label name to make its purpose easily recognizable within your organization.

- 2. Add commands that will create the server certificate and sign it with the signer certificate from step "1" on page 149.
 - As the owner of this certificate, specify the user ID that starts the DDF started task (DB2USER in the example below).
 - As the canonical name (CN), specify the DNS host name to contact your Db2 subsystem. For a data sharing group, this would be the host name of the group rather than a member host name.

Here is an example:

```
RACDCERT ID(DB2USER) -
DELETE(LABEL('DB2ASERVER CERTIFICATE'))
RACDCERT ID(DB2USER) -
GENCERT -
SUBJECTSDN(CN('DB2A') -
O('IBM') -
C('USA')) -
ALTNAME(IP(10.9.3.7)) -
NOTAFTER(DATE(2030-12-31)) -
WITHLABEL('DB2ASERVER CERTIFICATE') -
SIGNWITH(CERTAUTH LABEL('DB2 SERVER CA'))
SETR RACLIST (DIGTRING) REFRESH
SETR RACLIST (DIGTCERT) REFRESH
SETR RACLIST (FACILITY) REFRESH
```

Important: The canonical name (CN) in the server certificate is not verified when Db2 for z/OS and IBM Db2 Analytics Accelerator are connected. Any host name will be accepted if the certificate was signed by the CA whose public key was downloaded to the accelerator.

Creating the key ring

The next step is to create a key ring and store both, the signer certificate and the server certificate, in that key ring.

About this task

In most cases, you will not be able to reuse an existing key ring because the server certificate must be the default certificate to be returned by the key ring. This might interfere with the requirements of other encrypted connections.

The same key ring can be used by multiple Db2 subsystems under the condition that all subsystems use the same user ID for the DDF started task. If access by different user IDs is required, then access must be managed correctly by RACF methods.

Procedure

Add RACF commands to a JCL job as shown in the following example in order to create a key ring in RACF and store the certificates.

Tip: Make the user ID that runs the DDF started task the owner of the key ring. This way, you can omit a number of access authorizations.

Use the following example as a reference:

Exporting and transferring the signer certificate

Export and transfer the signer certificate to import it into the trusted keystore of the target accelerator in a later step.

Procedure

1. Add RACF commands to a JCL job that will export the signer certificate to a data set in DER format, as shown in the following example.

Tip: You can also complete this and the next step by customizing and running the sample JCL job AQTSSLSA. This sample job is a member of the SAQTSAMP data set, which comes with the stored procedures of the product.

Example:

```
RACDCERT CERTAUTH -

EXPORT(LABEL('DB2 SERVER CA')) -

DSN('<HLQ>.CERT.DB2.CERTAUTH.CRT') -
FORMAT(CERTDER)
```

where <HLQ> stands for the high-level-qualifier of the data set.

2. Before you can download the signer certificate to the accelerator, you must copy the DER data set to the z/FS file system of your IBM Z.

The location must be the directory that the AQT_HOST_PACKAGE_DIRECTORY environment variable points to. Environment variables for IBM Db2 Analytics Accelerator are set in the AQTENV data set. Make sure to use the *binary* transfer mode when you copy the data set to the z/FS. For example:

```
OPUT '<HLQ>.CERT.DB2.CERTAUTH.CRT' '/usr/lpp/aqt/packages/CERTAUTH.crt' binary
```

Your certificate can then be transferred to the accelerator.

- 3. Start IBM Db2 Analytics Accelerator Studio from a client workstation. Open the **Accelerator** view of the accelerator that you want to use as an endpoint of your IBM Integrated Synchronization setup.
- 4. Click **Encryption details** in the header of the **Accelerator** view (at the top).
- 5. In the Encryption Details window, click Transfer new certificates

The **Transfer Certificates** window opens.

- 6. In the **Transfer Certificates** window, under the heading **Available certificates**, select the certificate file and click **Transfer certificate from System Z**.
- 7. Select the copied certificate file in the **Select Certificates** window and click **Add selected**.

You return to the Transfer Certificates window.

- 8. Click Transfer.
- 9. Confirm the message in the **Transfer File** window by clicking **OK**.

Defining TTLS rules

Follow the examples in the steps below to define TTLS rules. A TTLS rule is required for each connection that is used by the IBM Integrated Synchronization process. TTLS rules are picked up by the Policy Agent

(PAGENT). They specify the port numbers, network directions, key rings (certificates), and the encryption algorithms a connection can use.

Procedure

- 1. Rules for the Policy Agent are defined in the main TCP/IP configuration file. So you have to find this file first.
 - a) Open the JCL job for the TCPIP started task.
 - b) Locate the PROFILE DD card in this job.
 - c) Check if AT-TLS is enabled by looking at the TCPCONFIG statement. It needs to contain the following entry:

```
TCPCONFIG
TTLS ; Enable AT-TLS Function
.
.
```

- d) If AT-TLS has been enabled, you can identify the relevant TCP/IP configuration file by checking the Policy Agent started task.
- e) Locate the PARM option in the JCL job for the Policy Agent started task.

 It refers to a configuration file, which in turn refers to the AT-TLS configuration file by means of the TTLSConfig statement:

```
TTLSConfig //'SYS1.TCPPARMS(TTLS)'
```

So in this case, the name of the AT-TLS configuration file is SYS1.TCPPARMS(TTLS).

2. Open the AT-TLS configuration file identified in the previous step and add the required settings for IBM Integrated Synchronization.

Connections to the distributed data facility (DDF) for IBM Integrated Synchronization must be encrypted. So each connection used for this purpose requires a TTLS rule. Connections are identified by their IP addresses, port numbers, and started task names.

The AT-TLS configuration file contains many blocks of settings, which are identified by a unique name. Each unique name can be used to refer to a shared set of settings from other rules. You can add new definitions to the end of the configuration file. Just make sure that the chosen identifiers (names) are unique.

a) A connection for IBM Integrated Synchronization between a single Db2 subsystem and an accelerator requires a TTLS rule as shown in the following example:

This rule refers to a standalone Db2 subsystem with the name DB11. The DDF listens to incoming connections on secure port 15111. The name of the Db2 subsystem is reflected in the address space name of the DDF, which is DB11DIST.

Important: A rule for a data sharing group is slightly different. However, most of the steps in this section also apply to the rule for a data sharing group. See the next section for instructions on how to make the necessary changes for a data sharing group.

The rule contains two references to other definition blocks within the AT-TLS configuration file:

TTLSGroupActionRef Db2SslGroup

• TTLSEnvironmentActionRef DB11SslEnv

This is a common way of referring to TTLS configuration options, which allows a reuse of the same options in other rules by just specifying the unique name of the block or section.

b) A suitable definition block for the TTLSGroupActionRef looks as follows:

This definition block enables TTLS and sets a specific trace option.

c) The other reference, TTLSEnvironmentActionRef points to a definition block that specifies the key ring to be accessed by the connection, the role in the TTLS handshaking process, and the encryption algorithms that can be used by the connection. A suitable definition block for this reference looks like this:

The key ring referred to is DB11KEYRING; the handshake role must be set to Server. The encryption algorithms are defined in a separate block, so again, you find a reference here. This reference points to a block named Db2SslCipherParms.

d) The following example shows what the Db2Ss1CipherParms block might look like:

- 3. Save your AT-TLS configuration file when your rule definitions are complete.
- 4. To let the changes take effect, refresh the address space of the Policy Agent from the console:

```
/f PAGENT, REFRESH
```

Defining a TTLS rule for a data sharing group

As indicated before in this document, the setup for a data sharing group is slightly different because IBM Integrated Synchronization requires dedicated connections to individual members.

About this task

In addition, you need to ensure that the same members are contacted each time a connection is created. This does not work in a setup with group IP addresses and shared secure ports. Because a secure port is needed for each member involved, you must use a port group in your TTLS rules rather than a (single) local port range.

Procedure

- 1. Open the TCP/IP configuration file identified in step <u>"1" on page 152</u> of the previous previous to add or change a rule for a data sharing group.
- 2. Add or change a TTLS rule as shown in the following example:

```
TTLSRule DB2ARule

LocalPortGroupRef
JobName
Direction
TIDSGroupActionRef
TTLSEnvironmentActionRef

PortGroup DB2ASecurePorts

PortRange
Port 12000
PortRange
Port 15011
Port 15011
}
```

If you compare this with the example in step "2.a" on page 152, you can see that the LocalPortRange parameter has been dropped. Instead, you find a reference to a PortGroup defined outside of the main rule. This port group specifies two ports on the member DB2A: one is the common port used by all data sharing members, the other is the private port available to that particular member only.

Important: Server certificates must be created for each LPAR in which your data sharing members reside because the host names of the LPARs are used to identify the servers.

If all your members reside in the same LPAR, they can share a single server certificate and also the same RACF key ring, provided that all members are also started with the same user ID.

If your members are in different LPARs, you must reference different key rings in each LPAR, and use separate TTLS rules that must be defined in the AT-TLS configuration file of each LPAR.

Assigning user IDs

To connect to Db2 for z/OS, the IBM Integrated Synchronization component on the accelerator requires a user ID for authentication. A user ID is also needed to verify authorizations for access to the IBM Integrated Synchronization API and to the log records.

About this task

Although you can use any existing user ID for that purpose, a dedicated user ID is recommended. In the upcoming examples, the user ID *DB2SYNC* is used.

Procedure

1. Create a dedicated user ID for IBM Integrated Synchronization in RACF.

Use password authentication or PassTicket authentication for this user.

Often, a new PassTicket cannot be generated as fast as required by consecutive logon requests. So if you use PassTicket authentication, make sure to add the APPLDATA ('NO REPLAY PROTECTION') clause to the RACF definition. That is, submit an authentication command that looks like this:

```
RDEFINE PTKTDATA <applName> SSIGNON(KEYMASKED(<key>)) APPLDATA('NO REPLAY PROTECTION')
```

where

<applName>

Is the name of the application

<key>

Is your PassTicket key

2. Add a DSNR profile.

This profile controls the access to the IBM Integrated Synchronization API. The user ID created in step "1" on page 154 requires read access to this profile. Ownership must be assigned to the user ID that starts the Db2 subsystem or data sharing member that you want to connect to. The profile name must consist of the Db2 subsystem or member name followed by a dot and the keyword ACCEL.

For example:

```
DB2A.ACCEL
```

The RACF commands you need to run to create this profile and grant the required read access are (in accordance with the names previously introduced):

```
RDEFINE DSNR (DB2A.ACCEL) OWNER(DB2USER) UACC(NONE)
PERMIT DB2A.ACCEL CLASS(DSNR) ID(DB2SYNC) ACCESS(READ)
```

3. The user ID created in step <u>"1"</u> on page 154 also needs read access to the DB2A.DIST profile.

Most likely, this .DIST profile already exists because it controls all remote access to Db2. If it does not yet exist, define it in the same way as the DB2A.ACCEL profile.

Example:

```
RDEFINE DSNR (DB2A.DIST) OWNER(DB2USER) UACC(NONE)
PERMIT DB2A.DIST CLASS(DSNR) ID(DB2SYNC) ACCESS(READ)
```

4. To access the log records, the DB2SYNC user ID needs MONITOR2 authorization.

Grant this authorization by using the following commands:

```
GRANT MONITOR2 TO DB2SYNC;
COMMIT;
```

5. The DB2SYNC user ID also requires permission to run the SYSPROC.ACCEL_SET_TABLES_ACCELERATION and SYSPROC.ACCEL_REMOVE_TABLES stored procedures.

For details, see the descriptions of these stored procedures in the *IBM Db2 Analytics Accelerator for z/OS: Stored Procedures Reference*. The permissions are listed under the heading *Authorizations for z/OS*.

6. In addition, the DB2SYNC user ID requires read access to the AQTDEF6 and AQTENV data set members of your IBM Db2 Analytics Accelerator installation. The name of the data set that these members belong to depends on your product configuration. Grant the read privilege to the RACF user profile that has access to these members.

This will probably not be one of the profiles already mentioned, but yet another RACF profile. On a generic basis, the command for this reads:

```
PERMIT <profile name> ID(DB2SYNC) ACCESS(READ)
```

where <profile name > is the name of the identified RACF profile.

7. Define an OMVS segment in RACF for the DB2SYNC user ID.

Adding permissions for schema change support

Set permissions as described to allow an interruption-free handling of schema changes in Db2 for z/OS tables.

About this task

In earlier versions of IBM Integrated Synchronization, the replication cycle broke when the structure of a Db2 for z/OS table changed (schema change), for example, because a column had been added to the table. The table was suspended from the replication process and the table state was changed to *Error*. Queries that included this table were no longer accelerated. To restore incremental update and query acceleration capabilities, you had to remove, redefine, re-enable replication for, and reload the related

accelerator-shadow table. This is no longer necessary. For more information, see <u>"Schema changes" on page 145.</u>

However, to process tables with schema changes correctly, metadata of the IBM Integrated Synchronization function must be updated, and this requires the querying of system tables. To allow these queries, you must extend the privileges of the dedicated user ID for IBM Integrated Synchronization.

Procedure

- 1. Grant the dedicated user ID for IBM Integrated Synchronization (*DB2SYNC* in the examples of this guide) the SELECT privilege on the following system tables:
 - SYSIBM.SYSTABLES
 - SYSIBM.SYSCOLUMNS
 - SYSIBM.SYSDUMMY1

To set these privileges for the sample user ID *DB2SYNC*:

```
GRANT SELECT ON SYSIBM.SYSTABLES TO DB2SYNC;
GRANT SELECT ON SYSIBM.SYSCOLUMNS TO DB2SYNC;
GRANT SELECT ON SYSIBM.SYSDUMMY1 TO DB2SYNC;
COMMIT;
```

2. Grant the dedicated user ID for IBM Integrated Synchronization the SELECT and UPDATE privileges on the SYSACCEL.SYSACCELERATEDTABLES table.

To set this privilege for the sample user ID DB2SYNC:

```
GRANT SELECT, UPDATE ON SYSACCEL.SYSACCELERATEDTABLES TO DB2SYNC; COMMIT;
```

Enabling IBM Integrated Synchronization for a Db2 subsystem

Follow the steps in this topic to enable IBM Integrated Synchronization for a single Db2 subsystem or data sharing member from the IBM Db2 Analytics Accelerator Console. Repeat the process for other subsystems or members 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

If you use IBM Db2 Analytics Accelerator on IBM Integrated Analytics System, you can enable IBM Integrated Synchronization for up to 10 Db2 subsystems or data sharing groups that are connected to one IBM Integrated Analytics System.

If you use Db2 Analytics Accelerator on Z, the number of Db2 subsystems or data sharing groups that can be enabled for IBM Integrated Synchronization varies with the number of processors (Integrated Facilities for Linux or IFLs) and the amount of memory configured for the accelerator LPAR. For example, if 10 IFLs and 320 GB of memory have been allocated to your accelerator LPAR, you can enable up to 5 Db2 subsystems or data sharing groups for incremental updates.

Procedure

1. Ask the network administrator or the person who did the TCP/IP setup for the IP address of the accelerator. Make a note of this information. You need to enter it as you complete the steps that follow.

For IBM Db2 Analytics Accelerator on an IBM Integrated Analytics System, this is the virtual IP or wall IP address.

For Db2 Analytics Accelerator on Z, this is the IP address of the network that you labeled DB2 in the **Appliance Installer**.

2. Log on to the IBM Db2 Analytics Accelerator Console by using **telnet** or **ssh**. The preferred method is **ssh**.

For more information:

- "Using telnet to log on to the IBM Db2 Analytics Accelerator Console" on page 199
- "Using ssh to log on to the IBM Db2 Analytics Accelerator Console" on page 200
- 3. Press the Pause key, then Enter to display the following screen:

4. Type 2 and press Enter to display the submenu:

```
main -> 2

You have the following options:

(0) - Go back one level
(1) - Enable incremental updates
(2) - Disable incremental updates
(3) - Restart incremental update processes
(4) - Include or exclude tables not enabled for incremental updates
    in WAITFORDATA queries
(Default 0) >
```

- 5. Type 1 and press Enter.
- 6. 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 in front of a listed Db2 subsystem and press Enter. For your reference: The name of the Db2 subsystem used in many examples is DB2A. Compare "Defining a secure network port" on page 146 for example.
 - b) Available incremental update technologies:

```
    Integrated Synchronization
    IBM InfoSphere Change Data Capture for z/OS
    Select the incremental update technology ('' or '0' to exit):
```

Type 1 for IBM Integrated Synchronization and press Enter.

- c) Enter Db2 z/OS group ip ('' or '0' to exit):

 Enter the IP address of the logical partition (LPAR) on which the selected Db2 subsystem or data sharing member resides, for example 10.9.3.7.
- d) Enter Db2 z/OS secure DDF port: (Default 446) >
 Enter the secure port that you defined for the distributed data facility (DDF) of the selected subsystem or data sharing member by following the steps in "Defining a secure network port" on page 146, for example 15111.

e) Press 'y' to use passticket authentication. Passticket authentication is recommended but requires to configure passtickets on z/OS side. (y/n): [y]:

Enter y or n, depending on the chosen authentication method for the dedicated IBM Integrated Synchronization user ID (DB2SYNC in the example).

- f) Enter the DB2 UserID for replication ('' or '0' to exit):

 This is the dedicated user ID for IBM Integrated Synchronization that you created by following the steps in "Assigning user IDs" on page 154 (DB2SYNC if you followed the example). Type this ID and press Enter.
- g) Enter the password (in TSO, use PF3 to hide input): Enter the password belonging to the user ID entered in step "6.f" on page 158.
- h) Enter the password again to confirm (in TSO, use PF3 to hide input): Repeat the password and press Enter to confirm it.
- i) Follow the appropriate list of steps according to your product version:

7.5.4 and earlier

i) Select the certificate that identifies the DB2 subsystem:

Import the signer certificate that you exported by following the steps in <u>"Exporting and transferring the signer certificate" on page 151</u>. If this certificate has been exported and transferred to the accelerator, it is listed as a selectable choice:

Certificate Name: CERTAUTH
issuer= OU=DB2 SERVER CA,O=IBM,C=USA
subject= OU=DB2 SERVER CA,O=IBM,C=USA
notBefore=May 16 07:00:00 2019 GMT
notAfter=Dec 31 06:59:59 2030 GMT
<No Alias>

Certificates available: 1 : CERTAUTH

Select a certificate or enter 0 to go
back: (Default 0) >

Type the number in front of the certificate name under Certificates available: and press Enter.

ii) Use certificate 'CERTAUTH' (y/n)
[n]:

Confirm your certificate choice. Type y and press Enter.

7.5.5 and later

i) Press 'y' to download the certificates from <ip- address>:<port>. (y/n) (Default y):

Type y and press Enter to confirm. The certificates (signer and server certificate) are displayed so that you can check them.

Examples:

Certificate Name: ABCXYZ1-certificate-1 issuer= OU=IDAA DB2 SERVER CA, O=IBM, L=SVL, ST=CA, C=USA subject= CN=abc.org.ibm.com,OU=DWB1,O=SVLDC11,C= notBefore=Aug 14 07:00:00 2019 GMT notAfter=Jan 1 06:59:59 2031 GMT <No Alias> Fingerprint=61:24:98:1D:DB:1D:2B:F6:6C: 4C:E9:48:2B:83:8D:83 Certificate Name: ABCXYZ1-certificate-2 issuer= OU=IDAA DB2 SERVER CA,O=IBM,L=SVL,ST=CA,C=USA subject= OU=IDAA DB2 SERVER CA, O=IBM, L=SVL, ST=CA, C=USA notBefore=Jul 23 07:00:00 2019 GMT notAfter=Jan 1 06:59:59 2031 GMT <No Alias> MD5 Fingerprint=66:38:C8:56:BB:F4:1C:08:73: 77:85:D7:59:66:58:6F ii) Check these certificates

carefully. If you enter 'y', they will be installed to identify the Db2 server <ip-address>:<port> as a trusted access point for replication. (y/n) (Default y):

When you are finished checking the certificates, type y and press Enter.

j) Press 'y' to configure Db2 subsystem 'ABCXYZ1' for replication:

Type y, then press Enter.

If you completed this last step successfully, the word done. is displayed on the screen.

- k) Press the Enter key repeatedly until you reach the main menu of the IBM Db2 Analytics Accelerator Console.
- 7. Type x and press Enter to exit the console.

What to do next

Having configured a Db2 subsystem for IBM Integrated Synchronization, 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 your administration client.

Suggestion for a high-availability setup

A high-availability setup ensures that your IBM Integrated Synchronization process remains intact in case of a network connection or similar failure. Several ways for such a setup are thinkable. The setup suggested here is based on a DDVIPA network and uses a Db2 data sharing group. It is relatively easy to implement and, once set up, does not require any manual interverention on your part because the maintenance is fully automatic.

The following prerequisites must be met for the suggested setup:

- A DDVIPA network between the z/OS LPAR used by Db2 for z/OS and your accelerators. DDVIPA is short for distributed dynamic virtual IP addressing.
- A Db2 for z/OS data sharing group used as the source for your IBM Integrated Synchronization process.

In general, a DDVIPA network is used to evenly distribute incoming requests to the available Db2 data sharing members. As connections might vary from request to request, such a network seems unsuitable for IBM Integrated Synchronization at a first glance because IBM Integrated Synchronization requires a permanent (unchanging) network connection between a Db2 data sharing member and an accelerator.

However, the distribution channels in a DDVIPA network are separated by the ports that are used. Each IP distribution channel is bound to a specific port number. It is possible to define the same secure port for data sharing members within the same group. See the following diagram:

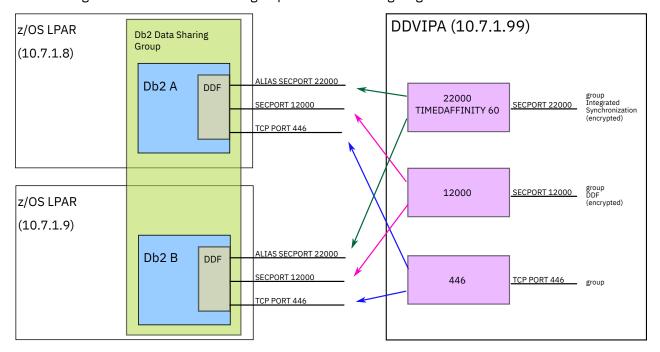


Figure 53. Network connections separated by port in a DDVIPA network

IBM Integrated Synchronization uses a timeout value of 60 seconds for log reader tasks. If no new requests are made during this period, Db2 for z/OS ends the active session automatically, and resources

are cleaned up. Therefore, accelerators make new requests within this period to keep the log reader task active. A TIMEDAFFINITY value of 60 seconds, a length equal to that of the timeout period for log reader tasks, ensures that the DDVIPA network will route these requests to the same data sharing member, as long as the member and the connected accelerator remain up and running.

If a data sharing member ceases to listen for requests on the assigned port, the DDVIPA network ignores the TIMEDAFFINITY parameter and starts a new distribution process to find a working member that accepts connections at the specified port address. Because the port number is the same as for the old connection, and the data in both members is identical, the incremental update process can continue without interruption. Once the new connection is established, the accelerator will remain connected to the second member until this new connection fails for some reason.

If, on the other hand, the accelerator stops making requests, and restarts at a time exceeding the timeout period of 60 seconds, it will likely connect to a different data sharing member. Usually, this occurs if incremental updates were suspended, in which case a new session is started anyway.

The suggested setup also supports multiple accelerators that use IBM Integrated Synchronization. Each accelerator can be assigned to a different data sharing member, or to the same member depending on the other distribution options. The accelerators remain connected to their assigned data sharing members, independent of each other.

When a failover to a different data sharing member occurs, the administrator will be notified of this event because several error messages are issued when the new session starts. The incremental updates will remain active despite the failover.

A planned maintenance isn't very complicated either. You just have stop the location alias of the member that you want to shut down. As long as other members are active that use the same port in their location aliases, requests will be redirected to these other members.

Setting up the network for data sharing groups

If just a single data sharing member receives IBM Integrated Synchronization requests at any given time, you need not change your DDVIPA configuration. However, if you have more members, you must add the TIMEDAFFINITY parameter to the relevant VIPADISTRIBUTE statements.

Procedure

- 1. Open your TCP/IP configuration data set in a suitable editor.
- 2. Add TIMEDAFFNITY 60 to all VIPADISTRIBUTE statements that define connections for IBM Integrated Synchronization.
- 3. Save the data set.

Example

See the following sample extract to guide you through the setup, or copy it to adjust and complete it. If you copy the extract, replace the placeholders with a real netmask, IP addresses, and port numbers. Also delete the ellipses (...):

```
<ddvipa-netmask> <ddvipa-ip-address>
VIPADISTRIBUTE DEFINE TIMEDAFFINITY 60 <ddvipa-ip-address>
                     <db2-integrated-synchronization-secport>
               DESTIP <db2-member-ip-address-1>
                      <db2-member-ip-address-2>
                      <db2-member-ip-address-n>
VIPABACKUP
               1 <ddvipa-ip-address>
VIPADISTRIBUTE DEFINE TIMEDAFFINITY 60 <ddvipa-ip-address>
               PORT
                     <db2-integrated-synchronization-secport>
               DESTIP <db2-member-ip-address-1>
                      <db2-member-ip-address-2>
                      <db2-member-ip-address-n>
              2 <ddvipa-ip-address>
VTPABACKUP
VIPADISTRIBUTE DEFINE TIMEDAFFINITY 60 <ddvipa-ip-address>
               PORT <db2-integrated-synchronization-secport>
```

where:

<ddvipa-netmask>

Is the netmask that applies to the interface on which the DVIPA listens for incoming requests.

<ddvipa-ip-address

Is the IP address the DVIPA uses to listen for incoming requests and from where it distributes the requests to active data sharing members.

<db2-integrated-synchronization-secport>

Is the dedicated secure port (number) defined in the location alias.

<db2-member-ip-address-1> ... <db2-member-ip-address-n>

Are the IP addresses of the data sharing members for which a location alias has been defined and that accept IBM Integrated Synchronization requests.

Defining a location alias for dedicated secure ports

Db2 for z/OS location aliases have individual connection properties, such as a dedicated secure port for remote network communication. This makes a location alias an ideal vehicle for providing a dedicated secure port for IBM Integrated Synchronization.

About this task

A location alias is independent of an already existing secure port (SECPORT) definition, except that it cannot have the same secure port number. In fact, you need not define a regular DDF secure port for your IBM Integrated Synchronization setup. It is sufficient to create a location alias with a secure port of its own. In this case, the general access to Db2 for z/OS would not be achieved by using secure (encrypted) connections.

You create a location alias by running a sequence of -MODIFY DDF ALIAS Db2 commands. The commands can be submitted from the command-line or added to the bootstrap data set (BSDS). Adding the commands to the BSDS is not recommended because it requires you to stop and restart Db2 to let the changes take effect. Submitting the command from the command-line, on the contrary, will activate the location alias immediately. The Db2 for z/OS Knowledge Center collection describes both methods (see links at the bottom).

Procedure

Create a location alias by running the following commands (preferably from the Db2 command-line):

```
a) -MODIFY DDF ALIAS(<alias-name>) ADD
```

b) -MODIFY DDF ALIAS(<alias-name>) SECPORT(<db2-subsystem-secport>)

c) -MODIFY DDF ALIAS(<alias-name>) START

where:

<alias-name>

Is the name you must give the location alias. It can be up to 16 characters long.

<db2-subsystem-secport>

Is the dedicated secure port number. The number must be unique on the LPAR or in the Sysplex, so that the accelerator can unequivocally identify the correct port.

Important: Location aliases are bound to a specific Db2 subsystem. In data sharing groups, they must be defined for each member that participates in the IBM Integrated Synchronization process.

Related information

-MODIFY DDF (Db2 command)
DSNJU003 (change log inventory)

Managing secure port availability with location aliases

Location aliases can be started and stopped individually for a Db2 subsystem or data sharing group. This allows you to control the availability of your Db2 data sources for IBM Integrated Synchronization.

About this task

You can easily make a Db2 subsystem or data sharing group unavailable for IBM Integrated Synchronization by stopping its location alias, for example to carry out a planned maintenance job. Other network traffic targeted at the general Db2 secure port would not be affected by this.

Procedure

• To prevent incoming synchronization requests from reaching a Db2 subsystem or data sharing group, replace <alias-name> in the following command with the real name of the location alias and submit the command from the Db2 command-line:

```
-MODIFY DDF ALIAS(<alias-name>) STOP
```

To reverse the previous action, submit the same command with the start parameter:

```
-MODIFY DDF ALIAS(<alias-name>) START
```

Status information for error analyses

When IBM Integrated Synchronization has been configured, your Db2 for z/OS system externalizes additional status information. You can use this information for a verification of your setup or for error analyses.

Health reporting

When IBM Integrated Synchronization is running, it constantly provides performance and health information in replication event messages. These are DSNX messages with a severity of WARNING. They are written to the SYSLOG of the accelerator database. You can also view these messages from the replication event viewer of your administration client. The message content varies with respect to the status. Hence two types of messages exist: messages that are issued in regular operation mode, and messages that are issued after a restart of IBM Integrated Synchronization.

Messages in regular operation mode

These messages contain the following information:

- Replication latency in seconds
- · Relative byte address (RBA) and log record sequence number (LRSN) of the latest commit
- · Number of open transactions
- · If open transactions exist, the earliest open RBA and LRSN
- Log parser statistics: number of INSERT, UPDATE, and DELETE operations
- Database operations: number of INSERT and DELETE operations
- If the Java Virtual Machine (JVM) is operating in a low heap condition and the risk of an out-of-memory-error is increased, the tenured heap usage is displayed in percent.

Outline:

```
Integrated Synchronization status:
Latency X seconds.
Latest commit RBA/LRSN X.
Number of open transactions X.
Earliest open RBA/LRSN X.
Log parser statistics: X insert, X update, X delete.
Database operations: X insert, X delete.
Tenured heap usage X%.
```

Messages issued after a restart

These messages contain the following information:

- · Source timestamp of the last applied transaction
- · Replication latency in seconds
- Relative byte address (RBA) and log record sequence number (LRSN) of the latest commit
- · Number of open transactions
- · If open transactions exist, the earliest open RBA and LRSN
- If the Java Virtual Machine (JVM) is operating in a low heap condition and the risk of an out-of-memory-error is increased, the tenured heap usage is displayed in percent.

Outline:

```
Integrated Synchronization status after restart at <last applied transaction source timestamp>. Latency X seconds. Latest commit RBA/LRSN X. Number of open transactions X. Earliest open RBA/LRSN X. Tenured heap usage X%.
```

Service Request Blocks

When IBM Integrated Synchronization has been configured for a Db2 subsystem and an accelerator, or for a data sharing member and an accelerator, an additional service request block (SRB) is started in the DBM1 address space. This SRB collects log records for the tables that need to be synchronized.

The SRB remains active as long as the accelerator is connected to the Db2 subsystem or data sharing member. If the SRB is no longer kept active by incoming requests, it is removed after a timeout period of 60 seconds.

Each time a new SRB is started, the message DSNI090I is issued. For example:

```
08.46.07 STC18562 DSNI090I -DB11 DSNILGRT -STARTING ASYNCHRONOUS LOG 753
753 READER TASK FOR
753 SESSION ID 43299E29761A0805
753 STARTING AT 00000000026A1D883F4A
753 WITH 1 QUALIFIERS
```

Each time an SRB is terminated, either on request by the accelerator or through a timeout, the message DSNI092I is issued. For example:

```
08.46.48 STC18562 DSNI092I -DB11 DSNILGRT -NORMAL TERMINATION OF 772
772 ASYNCHRONOUS LOG READER TASK FOR
772 SESSION ID 43299E29761A0805
772 ENDING AT 00000000026A1D89405E
772 AFTER READING 92 LOG RECORDS
```

Errors during log reading

If an error occurs during log reading, the SRB is terminated, and the message DSNI091I is issued. The z/OS system log probably contains more information about the error.

Displaying the status of log reader tasks

The status of current log reader tasks can be shown with the -DISPLAY STATS command:

```
-DISPLAY STATS(LOGREADERTASKS)
```

or, short:

```
-DISPLAY STATS(LRT)
```

These commands display all SRBs that are collecting log records in the currently active Db2 subsystem or data sharing member. For example:

The following STATUS values might be shown for log reader tasks:

RUNNING

Task is processing log records.

READING

Task is reading log records.

SUSP READ

Task suspended due to a full buffer. The task is waiting for the accelerator to fetch more records, so that processing can continue.

SUSP EOS

Task suspended because the end of the log has been reached. The task will be woken up at regular intervals to check if new log records have arrived.

CANCEL

Task is in the process of being terminated.

Changing the Db2 user ID and password

You can change the Db2 user ID for IBM Integrated Synchronization. You can also change the password of this user ID. The changes can be made without disabling incremental updates completely. However, you must stop replication on all attached accelerators before you can change the user ID or the password.

About this task

You might want to replace the Db2 user ID for integrated synchronization. Or you might have changed the password associated with this user ID, so that it no longer matches the password that was provided initially. The function described here allows you to change the credentials of the Db2 user without disabling incremental updates completely. This means that you need not reload your replicated tables. However, an interruption of the incremental update process is inevitable because you must stop replication on all accelerators connected to the Db2 subsystem, and restart replication after the change.

Procedure

 Log on to the IBM Db2 Analytics Accelerator Console by using telnet or ssh. The preferred method is ssh.

For more information:

- "Using telnet to log on to the IBM Db2 Analytics Accelerator Console" on page 199
- "Using ssh to log on to the IBM Db2 Analytics Accelerator Console" on page 200
- 2. Press the Pause key, then Enter to display the following screen:

3. Type 2 and press Enter to display the submenu:

- 4. Type 6 and press Enter.
- 5. The enablement setup process asks you for information. Type the requested information at each prompt and press the Enter.
 - a) Select the Db2 subsystem that uses the current credentials:

Type the number in front of the appropriate Db2 subsystem and press Enter. If only one such system is available, a message like the following is displayed:

Using the only Db2 subsystem available: ABCDE11

In this case, the procedure automatically moves on to the next step.

b) Press 'y' to use PassTicket authentication. PassTicket authentication is recommended, but must be configured on the z/OS side. (y/n) [y]:

If you want to use PassTicket authentication, just press Enter. If not, type n and press Enter.

c) Enter the DB2 UserID for replication:

Type the name of an existing z/OS user ID with sufficient RACF permissions. See <u>"Assigning user IDs" on page 154</u> for more information. Press Enter when ready.

d) Enter the password:

Enter the password that has been defined for the user ID in step "5.c" on page 165.

e) Enter the password again to confirm:

Re-enter the password.

f) You see the following reminder:

```
Important: Before you continue, make sure that replication has been stopped for the selected Db2 subsystem in combination with the current Db2 user.
```

Be absolutely certain that replication has been stopped on all attached accelerators before you complete the next step. Otherwise, the process will fail.

g) Press 'y' to change the Db2 user credentials for Integrated Synchronization in connection with Db2 subsystem 'ABCDE11': Type y and press Enter. The following message is displayed:

done.

You may now start replication again.

Press <return> to continue.

h) Press Enter to return to the submenu.

To exit the console, press Enter once more. Then type x and press Enter.

Related tasks

Starting or stopping incremental updates

When a Db2 subsystem has been configured for incremental updates, you can start or stop incremental updates for attached accelerators from the corresponding accelerator views in your administration client. Using the functional links in the header starts or stops incremental updates for an entire accelerator.

Changing the IP address and the DDF secure port

You can change the IP address by which a Db2 subsystem can be reached and also the associated secure DDF port for IBM Integrated Synchronization. The changes can be made without disabling incremental updates completely, and a reload of replicated tables is not necessary. However, you must stop replication on all attached accelerators before you can change the IP address or the secure port.

About this task

A working setup of IBM Integrated Synchronization requires that an accelerator can fetch the log records of the connected Db2 subsystem over the network. To this end, the correct IP address of the Db2 subsystem and the correct secure port number are required. Without this information, the Db2 subsystem cannot be reached, and consequently, the log records cannot be read.

The function described here allows planned and unplanned changes of the IP address of a Db2 subsystem and connected DDF secure port.

- A planned change might be required if the IP address or the secure port must be changed, for example if a new range of IP addresses is supposed to be used. A planned change might also be needed to connect an accelerator to a different copy of the very same Db2 subsystem for maintenance reasons.
- An unplanned change might be required if a Db2 subsystem cannot be reached anymore and if a copy of the very same Db2 subsystem has to be used instead.

The function described here allows you to change the IP address of the Db2 subsystem and the DDF secure port without disabling incremental updates completely. This means that you need not reload your replicated tables. However, an interruption of the incremental update process is inevitable because you must stop replication on all accelerators connected to the Db2 subsystem, and restart replication after the change.

Important: The function described here has been designed for Db2 subsystems. Data sharing groups have a built-in mechanism to switch to another member (with a different IP address), so this functionality is not necessary to connect or switch between members of data sharing groups.

Procedure

 Log on to the IBM Db2 Analytics Accelerator Console by using telnet or ssh. The preferred method is ssh.

For more information:

- "Using telnet to log on to the IBM Db2 Analytics Accelerator Console" on page 199
- "Using ssh to log on to the IBM Db2 Analytics Accelerator Console" on page 200
- 2. Press the Pause key, then Enter to display the following screen:

3. Type 2 and press Enter to display the submenu:

- 4. Type 7 and press Enter.
- 5. The enablement setup process asks you for information. Type the requested information at each prompt and press the Enter.
 - a) Select the Db2 subsystem that uses the current IP and secure DDF port: Type the number in front of the appropriate Db2 subsystem and press Enter. If only one such system is available, a message like the following is displayed:

Using the only applicable database system ABCDE11

In this case, the procedure automatically moves on to the next step.

b) Enter Db2 z/OS group ip ('' or '0' to exit):

Type the new IP address (address you want to change to) and press Enter.

c) Enter Db2 z/OS secure DDF port: (Default 448) >

Type the secure port that is associated with the IP address you entered in the previous step. Then press Enter.

d) You see the following reminder:

```
Important: Before you continue, make sure that replication has been stopped for the selected Db2 subsystem.
```

Be absolutely certain that replication has been stopped on all attached accelerators before you complete the next step. Otherwise, the process will fail.

e) Press 'y' to change the IP address and secure DDF port for Integrated Synchronization in connection with Db2 subsystem 'ABCDE11':

Type y and press Enter.

The following message is displayed if the operation was successful:

```
done.
You may now start replication again.
Press <return> to continue.
```

If you forgot to stop replication for the selected Db2 subsystem, you see the following error message:

```
{\tt AQT20007} The "XYZ" accelerator is in the 'STARTED' state, which does not allow a "ChangeNetworkEndpoint" operation.
```

f) Press Enter to return to the submenu.

To exit the console, press Enter once more. Then type x and press Enter.

Managing certificates

The certificate management menu on the IBM Db2 Analytics Accelerator Console consists of several functions to manage SSL certificates for IBM Integrated Synchronization. You can list all installed certificates, remove obsolete certificates, create, download, and activate new certificates, as well as activate certificates that were already transferred to the accelerator from a secure DDF port.

About this task

IBM Integrated Synchronization uses secure-socket-layer (SSL) encryption. With this technology, a Db2 subsystem uses a certificate to authenticate itself to an accelerator. The certificate is issued by RACF, then transferred to the accelerator, from where you finally import it using the console. If you want to change the certificate, you must stop incremental updates before the change, and re-enable these after. In the past, this required a reload of all replicated tables on the accelerator. The certificate management functions on the IBM Db2 Analytics Accelerator Console make the reload step superfluous, which saves a lot of time.

Displaying activated certificates

Follow the steps here to display the activated SSL certificates for a Db2 subsystem.

Procedure

 Log on to the IBM Db2 Analytics Accelerator Console by using telnet or ssh. The preferred method is ssh.

For more information:

- "Using telnet to log on to the IBM Db2 Analytics Accelerator Console" on page 199
- "Using ssh to log on to the IBM Db2 Analytics Accelerator Console" on page 200
- 2. Press the Pause key, then Enter to display the following screen:

3. Type 2 and press Enter to display the submenu:

```
main -> 2
You have the following options:

(0) - Go back one level
(1) - Enable incremental updates
(2) - Disable incremental updates
(3) - Restart incremental update processes
(4) - Include or exclude tables not enabled for incremental updates
    in WAITFORDATA queries
(5) - Allow or disallow queries to run if WAITFORDATA timeout is reached (6) - Change Db2 for
z/OS credentials for integrated synchronization
(6) - Change Db2 for z/OS credentials for Integrated Synchronization
(7) - Change IP and secure DDF port of Db2 for z/OS for Integrated Synchronization
(8) - (Menu) Manage certificates for replication
```

4. Type 8 and press Enter.

A submenu is displayed:

```
You have the following options:

(0) - Go back one level
(1) - Show activated certificates for IBM Integrated Synchronization
(2) - Download and activate certificates for IBM Integrated Synchronization
(3) - Remove a certificate used for IBM Integrated Synchronization
(4) - Activate transferred certificates for IBM Integrated Synchronization
```

5. Select a Db2 subsystem to display the certificates that it uses:

Type the appropriate number and press Enter to select a Db2 subsystem. The output screen displays the certificates that are used to authenticate this subsystem. If just a single Db2 subsystem is configured for incremental updates, this subsystem is selected automatically, which is indicated by the following message:

Using the only applicable database system subsystem-name

The installed certificates are displayed.

Example: The screen looks like this if the selected Db2 subsystem uses just one certificate:

```
The following certificates are currently used to authenticate Db2 on z/OS for this Db2 subsystem:
Alias: 16300195287980970922
Owner: OU=IDAA DB2 SERVER CA, O=IBM, L=SVL, ST=CA, C=USA
Issuer: OU=IDAA DB2 SERVER CA, O=IBM, L=SVL, ST=CA, C=USA
Valid from: 7/23/19 7:00 AM
Valid until: 1/1/31 6:59 AM
Fingerprints:
MD5: 66:74:C8:56:BB:F4:1C:08:73:77:85:D7:59:66:58:6F
SHA1: 62:23:55:90:31:16:9B:11:7D:61:3C:18:5B:4E:FF:86:7C:74:7B:9E
SHA256:
16:53:2F:D5:7C:E2:9E:B7:BB:F3:9E:EB:D7:B3:EB:AD:2F:E8:91:58:AD:12:DD:26:DE:04:11:F4:30:B0:DD:50
Signature algorithm name: SHA1withRSA
Version: 3
```

Each certificate begins with the word Alias. If there were more than one certificate, you would see extra blocks starting with the word Alias.

- 6. Press Enter to return to the submenu.
- 7. Press Enter again to return to the incremental update menu of the console.
- 8. To exit the console, press Enter repeatedly until you reach the main menu. Then type x and press Enter.

Installing new certificates

Create, download, and install a new certificate to authenticate a Db2 subsystem by following the steps here.

Procedure

1. From your administration client, stop replication between the affected Db2 subsystem and the accelerator.

That is, select **Replication** > **Stop** from the **Accelerator** view in IBM Db2 Analytics Accelerator Studio or from the list of accelerators in IBM Data Server Manager.

2. Log on to the IBM Db2 Analytics Accelerator Console by using **telnet** or **ssh**. The preferred method is **ssh**.

For more information:

- "Using telnet to log on to the IBM Db2 Analytics Accelerator Console" on page 199
- "Using ssh to log on to the IBM Db2 Analytics Accelerator Console" on page 200
- 3. Press the Pause key, then Enter to display the following screen:

4. Type 2 and press Enter to display the submenu:

5. Type 8 and press Enter.

A submenu is displayed:

```
You have the following options:

(0) - Go back one level
(1) - Show activated certificates for IBM Integrated Synchronization
(2) - Download and activate certificates for IBM Integrated Synchronization
(3) - Remove a certificate used for IBM Integrated Synchronization
(4) - Activate transferred certificates for IBM Integrated Synchronization
```

- 6. Type 2 and press Enter.
- 7. Select the Db2 subsystem that you want to install a new certificate for:

Type the appropriate number and press Enter to select a Db2 subsystem. A new certificate will be created, downloaded, and installed to authenticate this Db2 subsystem. If just a single Db2 subsystem is configured for incremental updates, this subsystem is selected automatically, which is indicated by the following message:

Using the only applicable database system subsystem-name

8. Enter the details of the SSL endpoint from which you want to download the certificates ('' or '0' to exit):

You are asked for the IP address of the Db2 subsystem and the secure DDF port that it uses.

a) Enter IP:

Enter the IP address by which the Db2 subsystem can be reached over the network.

b) Enter SSL port: (Default 448)

Enter the number of the secure DDF port that is used by the Db2 subsystem. To select the default port 448, just press Enter.

9. Enter 'y' to download the certificates from ip-address:SSL port. (y/n) (Default y):

Enter y.

Two certificates are displayed, for example:

```
1. Certificate Name: ABCXYZ10-certificate-1
   issuer= OU=IDAA DB2 SERVER CA,O=IBM,L=SVL,ST=CA,C=USA
   subject= CN=dwb1.svl.ibm.com,OU=DwB1,O=SVLDC11,C=USA
   notBefore=Aug 14 07:00:00 2019 GMT
   notAfter=Jan 1 06:59:59 2031 GMT
   <No Alias>
   MD5 Fingerprint=61:78:68:1D:DB:1D:2B:F6:6C:4C:E9:48:2B:83:8D:83
2. Certificate Name: ABCXYZ10-certificate-2
   issuer= OU=IDAA DB2 SERVER CA,O=IBM,L=SVL,ST=CA,C=USA
   subject= OU=IDAA DB2 SERVER CA,O=IBM,L=SVL,ST=CA,C=USA
   notBefore=Jul 23 07:00:00 2019 GMT
   notAfter=Jan 1 06:59:59 2031 GMT
   <No Alias>
   MD5 Fingerprint=66:74:C8:43:BB:F4:1C:08:73:77:85:D7:59:66:58:6F
```

10. Select the certificate to be installed ('' or '0' to exit):

Enter the number of the certificate you want to download and install.

For example, to download and install ABCXYZ10-certificate-2 from the previous example, type 2 and press Enter.

11. Wait until you see the confirmation message:

The certificate has been added successfully.

```
To let the changes take effect:
```

- 1. Stop replication
- 2. Restart incremental update processes
- 3. Start replication

The listed steps in the confirmation message are just a reminder. There is nothing to select at this point. You have already completed step "1" on page 170. Steps 2 and 3 follow as part of this procedure.

- 12. Press Enter to return to the submenu.
- 13. Press Enter again to return to the incremental update menu of the console.
- 14. Type 3 and press Enter to restart the incremental update process (subscription). Follow the instructions on the screen. For more information, see "Restarting incremental update processes" on page 192.
- 15. You might want to exit the console now. Press Enter as many times as needed to return to the main menu. Then type x and press Enter to exit.
- 16. From your administration client, re-enable replication between the Db2 subsystem and the accelerator. That is, select **Replication** > **Start** from the **Accelerator** view in IBM Db2 Analytics Accelerator Studio or from the list of accelerators in IBM Data Server Manager.

Removing certificates

Using the certificate management function on the IBM Db2 Analytics Accelerator Console, you can remove SSL certificates for IBM Integrated Synchronization that have expired or that you no longer want to use.

Procedure

1. From your administration client, stop replication between the affected Db2 subsystem and the accelerator.

That is, select **Replication** > **Stop** from the **Accelerator** view in IBM Db2 Analytics Accelerator Studio or from the list of accelerators in IBM Data Server Manager.

2. Log on to the IBM Db2 Analytics Accelerator Console by using **telnet** or **ssh**. The preferred method is **ssh**.

For more information:

- "Using telnet to log on to the IBM Db2 Analytics Accelerator Console" on page 199
- "Using ssh to log on to the IBM Db2 Analytics Accelerator Console" on page 200
- 3. Press the Pause key, then Enter to display the following screen:

4. Type 2 and press Enter to display the submenu:

```
main -> 2
    You have the following options:

(0) - Go back one level
(1) - Enable incremental updates
(2) - Disable incremental updates
(3) - Restart incremental update processes
(4) - Include or exclude tables not enabled for incremental updates
    in WAITFORDATA queries
(5) - Allow or disallow queries to run if WAITFORDATA timeout is reached (6) - Change Db2 for z/OS credentials for integrated synchronization
(6) - Change Db2 for z/OS credentials for Integrated Synchronization
(7) - Change IP and secure DDF port of Db2 for z/OS for Integrated Synchronization
(8) - (Menu) Manage certificates for replication
(Default 0) >
```

5. Type 8 and press Enter.

A submenu is displayed:

```
You have the following options:

(0) - Go back one level
(1) - Show activated certificates for IBM Integrated Synchronization
(2) - Download and activate certificates for IBM Integrated Synchronization
(3) - Remove a certificate used for IBM Integrated Synchronization
(4) - Activate transferred certificates for IBM Integrated Synchronization
```

6. Type 3 and press Enter.

7. Select a Db2 subsystem to display the certificates that it uses:

Type the appropriate number and press Enter to select a Db2 subsystem. The output screen displays the certificates used to authenticate this subsystem. If just a single Db2 subsystem is configured for incremental updates, this subsystem is selected automatically, which is indicated by the following message:

Using the only applicable database system subsystem-name

The currently installed certificates are displayed. For example:

```
    Alias: 16300145287980970922
Owner: OU=IDAA DB2 SERVER CA, O=IBM, L=SVL, ST=CA, C=USA

   Issuer: OU=IDAA DB2 SERVER CA, O=IBM, L=SVL, ST=CA, C=USA Valid from: 7/23/19 7:00 AM
   Valid until: 1/1/31 6:59 AM
   Fingerprints:
   MD5: 66:74:C8:56:BB:F4:1C:08:73:77:85:D7:59:66:58:6F
   SHA1: 62:23:55:90:31:16:9B:11:7D:61:3C:18:5B:4E:FF:86:7C:74:7B:9E
    SHA256:
16:53:2F:D5:7C:E2:9E:B7:BB:F3:9E:EB:D7:B3:EB:AD:2F:E8:91:58:AD:12:DD:26:DE:04:11:F4:30:B0:DD:5
    Signature algorithm name: SHA1withRSA
   Version: 3
2. Alias: 11782909041287341148
    Owner: CN=dwb1.svl.ibm.com, OU=DWB1, O=SVLDC11, C=USA
   Issuer: OU=IDAA DB2 SERVER CA, O=IBM, L=SVL, ST=CA, C=USA Valid from: 8/14/19 7:00 AM Valid until: 1/1/31 6:59 AM
   Fingerprints
   MD5: 61:78:98:1D:DB:1D:2B:F6:6C:4C:E9:48:2B:83:8D:83
SHA1: 19:03:5D:00:9F:07:F0:B3:62:D6:51:92:46:F2:76:0B:AF:BD:85:C9
    SHA256:
E1:23:EC:26:15:7B:12:F9:66:79:90:D8:8D:8E:A9:92:7A:61:3D:19:F3:52:E5:6F:91:C4:A5:62:CC:5A:40:C
    Signature algorithm name: SHA1withRSA
   Version: 3
```

You see two certificates, and the number that marks the beginning is before the word Alias.

8. Select the certificate to be removed ('' or '0' to exit):

To remove a certificate, type the appropriate number in front of the word Alias and press Enter.

For example, to remove the second certificate listed in step <u>"7" on page 173</u>, you would type 2 and press Enter.

9. Wait until you see the confirmation message, for example:

The certificate with the alias '11782909041287341148' has been removed successfully.

```
To let the changes take effect:
1. Stop replication
2. Restart incremental update processes
3. Start replication
```

The listed steps in the confirmation message are just a reminder. You do not have to select anything here. You have already completed step <u>"1"</u> on page 172. Steps 2 and 3 follow as part of this procedure.

- 10. Press Enter to return to the submenu.
- 11. Press Enter again to return to the incremental update menu of the console.
- 12. Type 3 and press Enter to restart the incremental update process (subscription). Follow the instructions on the screen. For more information, see "Restarting incremental update processes" on page 192.
- 13. You might want to exit the console now. Press Enter as many times as needed to return to the main menu. Then type x and press Enter to exit.
- 14. From your administration client, re-enable replication between the Db2 subsystem and the accelerator. That is, select **Replication** > **Start** from the **Accelerator** view in IBM Db2 Analytics Accelerator Studio or from the list of accelerators in IBM Data Server Manager.

Activating a transferred certificate

The certificate management function of the IBM Db2 Analytics Accelerator Console also offers an option to activate certificates that have already been transferred to an accelerator.

Procedure

1. From your administration client, stop replication between the affected Db2 subsystem and the accelerator.

That is, select **Replication** > **Stop** from the **Accelerator** view in IBM Db2 Analytics Accelerator Studio or from the list of accelerators in IBM Data Server Manager.

2. Log on to the IBM Db2 Analytics Accelerator Console by using **telnet** or **ssh**. The preferred method is **ssh**.

For more information:

- "Using telnet to log on to the IBM Db2 Analytics Accelerator Console" on page 199
- "Using ssh to log on to the IBM Db2 Analytics Accelerator Console" on page 200
- 3. Press the Pause key, then Enter to display the following screen:

4. Type 2 and press Enter to display the submenu:

5. Type 8 and press Enter.

A submenu is displayed:

```
You have the following options:

(0) - Go back one level
(1) - Show activated certificates for IBM Integrated Synchronization
(2) - Download and activate certificates for IBM Integrated Synchronization
(3) - Remove a certificate used for IBM Integrated Synchronization
(4) - Activate transferred certificates for IBM Integrated Synchronization
```

6. Type 4 and press Enter.

7. Select the Db2 subsystem for which additional certificates should be installed:

Type the appropriate number and press Enter to select a Db2 subsystem. The transferred certificate that you select in a later step authenticates this subsystem. If just a single Db2 subsystem is configured for incremental updates, this subsystem is selected automatically, which is indicated by the following message:

Using the only applicable database system subsystem-name

Available certificates are displayed. For example:

```
The following certificates have been transferred to the accelerator:

1. Certificate Name: idletime issuer= OU=IDAA DB2 SERVER CA,O=IBM,L=SVL,ST=CA,C=USA subject= OU=IDAA DB2 SERVER CA,O=IBM,L=SVL,ST=CA,C=USA notBefore=Jul 23 07:00:00 2019 GMT notAfter=Jan 1 06:59:59 2031 GMT <No Alias>
MD5 Fingerprint=66:74:C8:56:BB:F4:1C:08:73:77:85:D7:74:66:58:6F
```

8. Select the certificate to be installed ('' or '0' to exit):

To select a certificate, type the appropriate number in front of the term Certificate Name and press Enter. To select the certificate in step "7" on page 175, you would type 1 and press Enter.

9. Wait until you see the confirmation message:

```
The certificate has been added successfully.

To let the changes take effect:
1. Stop replication
2. Restart incremental update processes
3. Start replication
```

The listed steps in the confirmation message are just a reminder. You do not have to select anything here. You have already completed step <u>"1" on page 174</u>. Steps 2 and 3 follow as part of this procedure.

- 10. Press Enter to return to the submenu.
- 11. Press Enter again to return to the incremental update menu of the console.
- 12. Type 3 and press Enter to restart the incremental update process (subscription). Follow the instructions on the screen. For more information, see <u>"Restarting incremental update processes" on page 192.</u>
- 13. You might want to exit the console now. Press Enter as many times as needed to return to the main menu. Then type x and press Enter to exit.
- 14. From your administration client, re-enable replication between the Db2 subsystem and the accelerator. That is, select **Replication** > **Start** from the **Accelerator** view in IBM Db2 Analytics Accelerator Studio or from the list of accelerators in IBM Data Server Manager.

REORG, REORG DISCARD and LOAD REPLACE behavior

If you empty a synchronized Db2 for z/OS table or table partition completely by running REORG DISCARD or LOAD REPLACE DD DUMMY, the effect on the related accelerator-shadow table varies with respect to the type of the Db2 for z/OS table.

Mirroring

For the listed table types, the deleted rows are also deleted from the related accelerator-shadow table and IBM Integrated Synchronization remains enabled, so that the process can continue:

- Tables in classic partitioned table spaces (non-universal partitioned table spaces)
- Tables in universal table spaces partitioned by range (PBR UTS)

Table suspension

If a REORG DISCARD or LOAD REPLACE DUMMY job is run on a table partitioned by growth (PBG), the table is suspended from IBM Integrated Synchronization, query acceleration is disabled, and the table goes into *error* state.

This is a security feature. A PBG table does not have fixed partition boundaries. Hence a REORG DISCARD or LOAD REPLACE DUMMY operation on partitions of such a table does not make sense because the user does not know which rows will be affected. It is likely that not all rows of a partition or table will be deleted by this operation.

To re-enable replication for a table that was suspended as the result of a REORG DISCARD or LOAD REPLACE DUMMY operation, you must:

- 1. Remove the table from the accelerator.
- 2. Redefine (re-add) the table to the accelerator.
- 3. Re-enable replication for this table.
- 4. Load the table.

REORG jobs without DISCARD

A simple REORG operation does not lead to the suspension of a table, irrespective of its type.

In older product versions (before 7.5.5), you can witness the suspension of PBG tables after an ordinary REORG job, that is, a REORG without the DISCARD parameter. To get the current handling of reorganized PBG tables in your old product version (where a simple REORG does not lead to a suspension of the table), contact IBM support.

IBM InfoSphere Change Data Capture for z/OS

As the title suggests, this feature uses an external product, IBM InfoSphere Change Data Capture for z/OS or CDC, for change capturing and workload propagation on the z/OS side. CDC can be customized extensively, which makes the product adaptable to very specialized use cases.

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 support" on page 25. 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.

Prerequisites for tables and hints

- Tables to be included must have a unique constraint (primary key or primary index). If such a key does not exist in the accelerator-shadow table or cannot be determined, you must redefine the table and specify such a key. You cannot do this from IBM Db2 Analytics Accelerator Studio; you must run the corresponding stored procedure (SYSPROC.ACCEL_ADD_TABLES) with a special option. The columns that you choose for the key must contain unique values or form such values when they are combined. The values of the selected key columns will then be written to a hidden column in the accelerator-shadow table. This column, called informational unique constraint, then serves as the primary key or primary index. For more information, see the entry about the table_specifications parameter of the SYSPROC.ACCEL_ADD_TABLES stored procedure in the IBM Db2 Analytics Accelerator for z/OS: Stored Procedures Reference.
- 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-shadow tables.

- The Db2 source tables of your accelerator-shadow tables have an attribute named DATA CAPTURE. 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.
 - The DATA CAPTURE attribute is set by an ALTER TABLE statement, which is run as part of the SYSPROC.ACCEL_SET_TABLES_REPLICATION stored procedure. However, the attribute can only be set successfully if the ID of the user who runs the stored procedure has ALTER TABLE authorization. If this is not the case, a database administrator must set the attribute for all tables in Db2 for z/OS.
- You might have to reload or even remove tables from an accelerator after the execution of an ALTER TABLE or ALTER TABLESPACE statement in Db2 for z/OS. For more information, see What to do after changing a table or table space in Db2 for z/OSin the IBM Db2 Analytics Accelerator for z/OS: User's Guide.

Restrictions

- 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.
- It is not possible to enable tables if the table name or the schema name contains GB18030 characters (Simplified Chinese) of Unicode plane 2 (U+20000-U+2FFFF: Supplementary Ideographic Plane). Trying to do so results in an error.

Differences to product version 5

Some differences exist if the current implementation of incremental updates with CDC is compared with the implementation you find in IBM Db2 Analytics Accelerator for z/OS Version 5:

- You can no longer exclude fault tables from query acceleration.
- You can no longer use the console to "announce" the change of a group IP address or DRDA port on a Db2 data server. If such a change is required, you must disable and then re-enable replication and specify the new values during the enabling step.
- You can no longer change the IP address of the CDC capture agent from the console. If this is required, disable and re-enable replication and specify the new IP address during the enabling step. Alternatively, contact IBM support.
- A function that cleans up unused CDC metadata does not exist anymore. If you want to remove unused CDC metadata, contact IBM support.

IBM Integrated Synchronization offers all of these features. If they are of critical importance to you, consider switching to IBM Integrated Synchronization.

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 your administration client 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

Your administration client communicates with the IBM Db2 Analytics Accelerator for z/OS server process 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 an accelerator-shadow 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 CDC Agent on the IBM Integrated Analytics System, without IBM Db2 Analytics Accelerator for z/OS components in between (see the blue arrows in Figure 54 on page 178).

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:

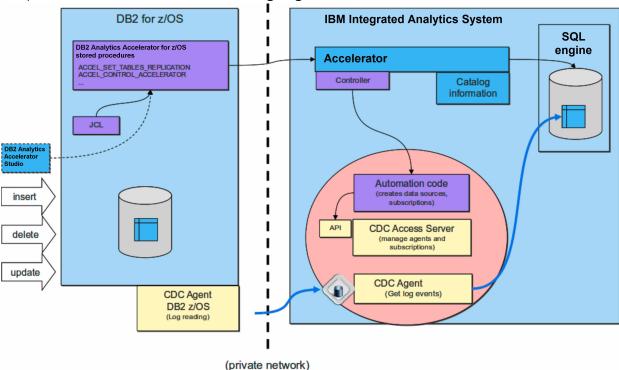


Figure 54. Interaction diagram of the components that are involved in the incremental update process.

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 an accelerator-shadow table while an incremental update process is running.

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 your administration client.

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.

Continuous incremental updates

To avoid that manual table load processes interrupt incremental update processes and that pending incremental updates are lost for that reason, a special mechanism, called *continuous incremental updates*, has been implemented.

The continuous incremental update mechanism works as follows: A marker is set when a load process sets in and also when it ends. All incremental updates that fall into the interval thus defined are written to a temporary file (called spill queue). When the load process has finished, the suspended updates are read from the spill queue and are applied to the loaded table. This is called draining. Moreover, incremental updates are not stopped for other tables in the Db2 subsystem.

This feature also allows INSERT, UPDATE, and DELETE operations against the original Db2 table during a load. Because the spill queue ensures that all updates are preserved, it is not necessary to set a lock on the table. This gives you an additional performance benefit. However, to avoid conflicts, the lock mode **Row** is recommended. Such a conflict could arise when, for example, a row is being inserted in the original Db2 table, and that very same row is captured by an ongoing load process although the insertion has not yet been completed.

Lock mode **Table**, on the contrary, blocks access to the entire Db2 table until the load is finished. Hence no incremental updates can be captured and written to the spill queue during the load. The spill queue will always be empty. The pending changes to the Db2 tables are deferred until the lock is released.

Restrictions: Draining processes have no effect on the execution time of the SYSPROC.ACCEL_CONTROL_ACCELERATOR stored procedure. That is, a <waitForReplication> element as part of the **command** parameter does not postpone the execution of the stored procedure until draining has finished.

Parallel propagation of changes

To speed up table synchronization, especially when the latency is high, incremental update changes are applied in parallel. The workload is distributed to four threads that are running at the same time.

The benefits of parallel processing come at a price: The referential integrity of the database cannot be guaranteed in all cases because the changes that are effected by the different threads must be committed, and that cannot be done at the same time for all threads. Serialization is the underlying principle of committing changes. It can therefore happen that a query "hits" an accelerator-shadow table at a stage where the changes brought about by one or more threads have been committed, but not yet all of these. This might lead to situations in which, for example:

- One and the same query, when run twice within a short time span, gives you different results because not all threads have committed their changes when the query hits the accelerator-shadow table for the first time.
- Different queries that read the same accelerator-shadow table "find" different values in the same row depending on the time when the queries hit the table.

The database of the accelerator uses multi-version concurrency control. That is, a table row is not locked for reads when it is being modified by another process (thread), so the older, unchanged values are read by a query until a change has been committed. In practice, the problem will rarely concern you because the database gives each change thread a unique transaction ID. In most cases, these IDs contain

consecutive numbers. Threads with consecutive numbers are processed in one block, and nothing can happen in between that would lead to one of the problems indicated earlier.

Faulty tables processing

A table that causes errors is suspended from the incremental update process so that processing for the other tables can continue.

To get a suspended accelerator-shadow table back into the incremental update process (after eliminating the cause of the error), it is sufficient to reload the data in the table. It is not necessary to re-enable incremental updates. However, make sure to find and eliminate the cause of the error so that the table does not interrupt the process again.

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

- 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.

 Because it makes the maintenance process easier, it is recommended that you use the same ID as the security identifier <CHCMetaID>, which is the owner of the CDC metadata tables (see steps "4.b" on page 181 and "4.c" on page 181). However, if your organization's security policy demands that you use different IDs, you can do so.

Important:

- If the ID for running the CDC started task and the security identifier *<CHCMetaID>* are the same, make sure that the password of this user ID is at least 4 characters long. It is possible to define a shorter password, but that will cause problems when you try to update IBM Db2 Analytics Accelerator in the future. If the IDs are different, the one for the CDC started task does not require a password, but the security identifier does.
- Set the password of this user ID so that it never expires. This prevents unwanted interruptions of the incremental update process.
- b) Defining an OMVS segment for this user ID because the ID needs to access TCP/IP services. This user ID 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 <u>"1" on page 25</u>). The required PTFs are listed in:
 - https://www.ibm.com/support/docview.wss?uid=swg27039487Prerequisites and Maintenance for IBM Db2 Analytics Accelerator for z/OS Version 7.1.
- 4. Complete the steps in section <u>Completing the installation using the distributed sample jobs</u> 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 (*<CHCMetaID>*) as the owner of these tables. The security identifier becomes the schema name of the metadata 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:

```
TCP/IP SERVICENAME=<port_no>
AUTODISCEXCLUDE=*
```

where <port_no> matches the port number or assigned service name that you have defined in step "2" on page 180. 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,
REPSTATSINTERVAL=5,
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,
ONTABLEERROR=IDLE,
ADDCOLUMNISSCHEMACHANGE=NO,
LOGPOLLINTERVAL=(3,ALWAYS),
LOGCACHEDELAY=5,
REPLTEMPORALTABLES=DECOUPLE,
LOGREADCOMMITINTERVAL=15,
CACHEBLOCKSIZE=2,
CACHELEVEL1RESERVED=160,
CACHELEVEL1SIZE=200
```

Important:

- Note the change from ONSCHEMACHANGE=STOP to ONSCHEMACHANGE=IDLE. This is the recommended setting if your CDC installation is at fix level APAR PI65350/UI39653 or an even newer fix level. STOP was the recommended setting for earlier product releases or fix levels. Continue to use STOP if you are at a fix level earlier than APAR PI65350/UI39653. The setting STOP causes the incremental update process to stop in case of the schema of the base tables changes, so that you can redefine and reload the corresponding accelerator-shadow table before you restart the process. With APAR PI65350/UI39653 or later, this is no longer necessary.
- The setting ONUTILIYACTION=IDLE leads to a conflict if you intend to use the High Performance Storage Saver (HPSS) on the same tables because IDLE suspends the propagation of data changes for an entire table, and not just the partitions that have been moved by the HPSS. Hence the parts

that have not been archived will not be synchronized anymore. To avoid this conflict, temporarily change the setting to ONUTILITYACTION=IGNORE, that is, before you archive or restore partitions with the HPSS. Note that you must restart the CDC capture agent for the change to take effect.

Alternatively, you can also use ONUTILITYACTION=IGNORE as a permanent setting and handle the data changes that escaped the capturing process with the help of Db2 utilities.

• In addition, the setting ONUTILIYACTION=IDLE might lead to wrong or unintended results if it is used in combination with a setup for Hybrid Transactional Analytical Processing (HTAP). Generally, if a Db2 utility action is detected for an incrementally updated table in an HTAP environment (for example, LOAD or RECOVER), the table is suspended from the ongoing incremental update process. However, if more than one table must be suspended, and queries running at the same time or shortly after the suspension reference these tables, incorrect results might be computed until the suspension has been fully completed.

If this is unacceptable, you can use the setting ONUTILITYACTION=STOP instead. Mind, however, that this setting stops the entire incremental update process when a utility action is discovered. If that happens, table changes are not propagated to the accelerator anymore.

A less drastic counter-measure is the suspend-faulty-tables function. This function does not eliminate the risk of incorrect results, but minimizes it considerably.

For more information about HTAP, see Making queries wait for incremental updates in the IBM Db2 Analytics Accelerator for z/OS: User's Guide.

- The example above uses log caching. In the past, log caching was recommended for setups in which data changes were propagated to multiple connected accelerators. However, experience has shown that log caching eliminates many issues connected with the incremental update function. Therefore, always enable it.
- ONTABLEERROR=IDLE was introduced with IBM Db2 Analytics Accelerator for z/OS Version 4.1.0, maintenance level PTF-5 (it requires the installation of the fix for CDC APAR PI26807, UI27018 or higher). Do not use this setting with earlier maintenance levels or versions.
- REPLTEMPORALTABLES = DECOUPLE was introduced with IBM Db2 Analytics Accelerator for z/OS Version 5.1.0 (it requires the installation of the fix for CDC APAR PI49337, UI31507 or higher). Do not use this setting with earlier versions. However, the setting is mandatory in Version 5.1.0.
- With this configuration, the CDC STC might allocate up to 2.2 GB (2GB staging space + 200 MB level 1 cache) of z/OS storage resources (above-the-bar storage), which might be paged to auxiliary storage. Make sure that there are sufficient system resources for this configuration. The paging might reduce the incremental update throughput significantly.

Description

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.

REPSTATSINTERVAL

Specifies the period of time, in minutes, between consecutive messages that provide status information about Db2 log scraping activities on the source server.

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 that a Db2 utility (such as LOAD or RECOVER) has run on a table space containing source tables that are being mirrored. Make sure to read the related bullets in the previous Important note.

ONDECOMPRESSIONERROR

Specifies how source data compression errors are to be handled. Decompression errors can be either soft (temporary and potentially recoverable) or hard (permanent and unrecoverable). This keyword has two positional sub-parameters that provide a recovery retry time limit for soft decompression errors and a resolution action for mitigating hard decompression errors.

Soft errors will be retried a number of times in an increasing series of intervening time intervals, with the sum of these interval not to exceed the specified retry time limit. When retrying to recover from a soft error starts, a message will be issued noting the location in the log and type of error. If the retrying attempts exceed the retry time limit, then the error will be deemed to be a hard error.

The recommended setting of ONDECOMPRESSIONERROR=(300, IDLE) means that the recovery interval for soft errors is set to 300 seconds or five minutes. If a recovery does not occur within this interval, the error is escalated to a hard error. The second sub-parameter specifies how to handle hard errors. The value IDLE suspends the incremental update process, sets the status of the affected accelerator-shadow table to *inactive*, meaning that no more updates are applied to that table, and then resumes the process for the remaining tables.

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.

ONTABLEERROR

Specifies which action to take on tables that have caused an error. The setting ONTABLEERROR=IDLE effects the stopping of incremental update activities on those tables, so that the process can continue on unaffected 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

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.

REPLTEMPORALTABLES

Specifies which action to take on system-temporal and bi-temporal tables. Setting REPLTEMPORALTABLES=DECOUPLE makes CDC work on the regular tables as well as on the history tables.

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).

Important: The example above uses log caching. In the past, log caching was recommended for setups in which data changes were propagated to multiple connected accelerators. However, experience has shown that log caching eliminates many issues connected with the incremental update function. Therefore, always enable it.

LOGREADCOMMITINTERVAL

Specifies how frequently the InfoSphere CDC for z/OS® threads that read the Db2 log will issue a Db2 COMMIT. In this case: every ten seconds. See the previous Important note under LOGCACHEDELAY.

CACHEBLOCKSIZE

Specifies the maximum size, in MB, of each buffer in the level 1 and level 2 Db2 log caches. See the previous Important note under LOGCACHEDELAY.

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. See the previous <u>Important</u> note under LOGCACHEDELAY.

CACHELEVEL1SIZE

Specifies the size, in MB, of the level 1 Db2 log cache. See the previous <u>Important</u> note under LOGCACHEDELAY.

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**.

Related information

CONFIG statement
DB2 statement in CHCDBMxx

Configuring staging space resources on z/OS

The incremental update function tracks and stages all transactions in the system memory until they are committed (and subsequently transferred to the accelerator) or rolled back. The component responsible for this task is the CDC Capture Agent on z/OS. Sufficient memory must be allocated to this component. If the memory does not suffice, a staged transaction might stop incremental updates for all connected accelerators. The error message CHC0330E is issued in such cases.

Configuration settings relevant to the staging space

The CDC configuration parameter MAXSUBSCRSTAGESIZE defines the amount of memory for the staging of incremental update transactions to be processed by an accelerator. So the total amount of memory that must be available to the CDC Capture Agent is roughly this amount multiplied by the number of accelerators. The total amount is configured by the configuration parameter **STG64LIMIT**. However, the required memory also depends on the actual workload. The default values are therefore only starting points for an evaluation. The website *Calculating staging resources* (see link at end of this topic) provides heuristics, which might serve as a computational basis.

Testing your settings

Because the staging space size depends on the workload, it is crucial that you test the configuration settings with the anticipated workload. To make sure that sufficient staging space is allocated, monitor the following events in the log of the CDC Capture Agent:

Table 3. Staging-space-related events in the log of the CDC Capture Agent			
Event	Description		
CHC0051W	The oldest open commit group in the staging space has exceeded a certain threshold. Long running transactions can occupy a significant amount of the staging space for a longer period of time and thus may lead to the space reaching its overall limits.		
CHC9618W	Indicates that the overall memory consumption has reached a defined limit.		
CHC1532W	A subscription has occupied more than 90% of the available staging space.		
CHC1524E	A subscription has used up all of the available staging space (100%).		
CHC0330E	Indicates that the staging space is full and that the system has to abort.		

For high workloads, it might prove useful to create staging space reports. These reports show the content of the staging space and other details, which might help you understand where the problems are.

Example

F <cdcjob>, REPORT, TYPE=STGSP, NAME=<subname>

As always when monitoring the memory consumption, it is helpful to monitor the storage manager.

Example

F<cdcjob>,STGMGR,STATUS

Impact of configuration settings on the CDC Capture Agent

The CDC Capture Agent uses a configurable amount of above-the-bar storage (STG64LIMIT) for the staging of transactions and for the so-called retry cache. If this additional space is not configured, the agent uses its regular amount of storage, which is no more than 2GB because the product is a 31-bit application. If the above-the-bar storage is exhausted, the component will try to keep going with the regular storage. If the regular storage is also exhausted, memory allocation (malloc) failures occur. These memory allocation failures have a varying impact depending on the part of the product in which they occur.

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 log 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. DATA CAPTURE CHANGES are mandatory for the incremental update function in IBM Db2 Analytics Accelerator for z/OS Version 7.x and later versions.

Procedure

From your interactive SQL tool of choice (SPUFI, QMF, or whatever), submit the following SQL statements:

- ALTER TABLE SYSIBM.SYSTABLES DATA CAPTURE CHANGES
- ALTER TABLE SYSIBM.SYSCOLUMNS DATA CAPTURE CHANGES
- ALTER TABLE SYSIBM.SYSTABLEPART DATA CAPTURE CHANGES

Follow the link to the CDC documentation for more information.

Enabling log caching

Experience has shown that it is always advisable to enable log caching for the capturing agent because it helps to overcome many issues related to the incremental update function.

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 reload of all the tables that take part in the process.

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.

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.

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:

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.

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:

- a) Stop the CDC started task.
- b) 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.

c) Restart the CDC started task.

Enabling incremental updates for a Db2 subsystem

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.

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

If you use IBM Db2 Analytics Accelerator on IBM Integrated Analytics System, you can enable incremental updates for up to 10 Db2 subsystems or data sharing groups that are connected to one IBM Integrated Analytics System.

If you use Db2 Analytics Accelerator on Z, the number of Db2 subsystems or data sharing groups that can be enabled for incremental updates varies with the number of processors (Integrated Facilities for Linux or IFLs) and the amount of memory configured for the accelerator LPAR. For example, if 10 IFLs and 320 GB of memory have been allocated to your accelerator LPAR, you can enable up to 5 Db2 subsystems or data sharing groups for incremental updates.

Procedure

1. Ask the network administrator or the person who did the TCP/IP setup for the IP address of the accelerator. Make a note of this information. You need to enter it as you complete the steps that follow.

For IBM Db2 Analytics Accelerator on an IBM Integrated Analytics System, this is the virtual IP or wall IP address.

For Db2 Analytics Accelerator on Z, this is the IP address of the network that you labeled DB2 in the **Appliance Installer**.

2. Log on to the IBM Db2 Analytics Accelerator Console by using **telnet** or **ssh**. The preferred method is **ssh**.

For more information:

- "Using telnet to log on to the IBM Db2 Analytics Accelerator Console" on page 199
- "Using ssh to log on to the IBM Db2 Analytics Accelerator Console" on page 200
- 3. Press the Pause key, then Enter to display the following screen:

4. Type 2 and press Enter to display the submenu:

- 5. Type 1 and press Enter.
- 6. 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) Available incremental update technologies:

```
    Integrated Synchronization
    IBM InfoSphere Change Data Capture for z/OS
    Select the incremental update technology ('' or '0' to exit):
```

Type 2 (IBM InfoSphere Change Data Capture for z/OS), and press Enter.

- c) 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.
- d) Enter the Capture Agent TCP port ('' or '0' to exit):

 Enter the port number of the CDC Capture Agent that you specified in step <u>"2" on page 180</u>. 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.

e) Enter the Db2 UserID for replication:

Incremental update processes are run under a certain user ID. This user ID must exist in z/OS and requires access to the CDC metadata tables and to all the tables that you want to enable. It is recommended that you use the same user ID as for the started task that runs the CDC Capture Agent because this ID already has the required authorizations on the metadata tables if the CDC installation process was followed properly. If you decided to use a different user ID, you must grant this user ID the following authorizations on the *CHCMetaID>*. DMMD_SIGNALS table:

- SELECT
- INSERT
- UPDATE
- DELETE

where *<CHCMetaID>* stands for the ID of the owner or the schema name of the CDC metadata tables (security identifier). You provided this ID when you ran the CDCMDMUT customization job.

In addition, mind that this ID must also have SELECT authorization on all tables that you want to enable.

Important:

- Make sure that the password of this user ID is at least 4 characters long. It is possible to define a shorter password, but that will cause problems when you try to update IBM Db2 Analytics Accelerator in the future.
- 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.

f) Enter the password:

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

g) Enter the password again to confirm:

Repeat the password and press Enter to confirm it.

h) Enter the Db2 group IP for client connections:

Type the *private network IP address* of the selected Db2 subsystem. If the target system is a Db2 data sharing group, type the IP address that is used to reach the data sharing group over the private network from the accelerator (usually, this is a shared DVIPA address). Then press Enter.

i) Enter the Db2 DRDA port for client connections: (Default 446) >

This is the listening port belonging to the IP address. Press Enter if you use the default port. If you use a different port, type the port number and press Enter.

j) Enter the schema name of the CDC capture task metadata tables:

Type the schema name and press Enter. The CDC metadata tables were created by the CHCMDMUT sample job at the time when IBM InfoSphere Change Data Capture for z/OS was installed and prepared for use with Db2.

The schema name is the ID of the owner of the CDC metadata tables (security identifier). Make sure that the user ID under which your incremental update processes are run has the required authorizations on the CDC metadata tables. See sub-step "6.e" on page 189.

k) 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.

l) Press the Enter key repeatedly until you reach the main menu of the IBM Db2 Analytics Accelerator Console.

7. 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 your administration client.

Related information

InfoSphere CDC for z/OS authorization requirements

Starting or stopping incremental updates

When a Db2 subsystem has been configured for incremental updates, you can start or stop incremental updates for attached accelerators from the corresponding accelerator views in your administration client. Using the functional links in the header starts or stops incremental updates for an entire accelerator.

About this task

When a subsystem has been configured for incremental updates, the header of the accelerator view shows an additional entry, which is labeled **Replication**. Next to this label, you find a status indicator and a functional link. Status indicator and link vary according to the current status of the accelerator. For example, if incremental updates have been stopped, the status indicator shows *Stopped* and a functional link labeled **Start**. This is reversed to *Started* and **Stop** when incremental updates have been started. The following statuses are possible:

Disabled

The incremental update function has not been configured for this accelerator and is therefore not available

Error

The incremental update function has been configured, but an error occurred. In most cases, this means that a component cannot be reached. Therefore, it cannot be started or controlled. The event log might provide useful information about the root cause of the error.

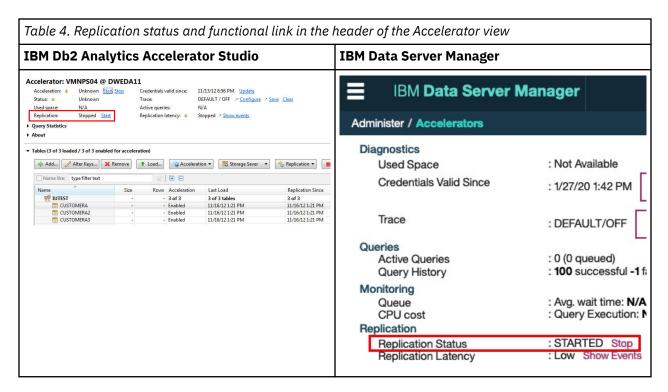
Started

The incremental update function is running and updates are applied to accelerator-shadow tables that have been included in the process.

Stopped

The incremental update function is available and accelerator-shadow tables have been included in the process. However, the function is not running right now.

See Table 4 on page 191.



Important: Using the functional links starts or stops the incremental update process for all tables that have been enabled for incremental updates on the selected accelerator.

Procedure

- 1. Start your administration client.
- 2. Connect to the Db2 subsystem that is associated with the proper accelerator.

IBM Db2 Analytics Accelerator Studio	Double-click the database-connection-profile icon for that subsystem.
IBM Data Server Manager	From the drop-down list on the top banner, select the corresponding database connection.

3. Select the accelerator for which you want to start or stop incremental updates:

IBM Db2 Analytics Accelerator Studio	a. In the Object List Editor , double-click the accelerator name.
IBM Data Server Manager	Select the accelerator (row) in the list.

- 4. Start or stop incremental updates by taking one of the following actions:
 - To start incremental updates, click Start in the header of the accelerator view. The status, which is labeled Replication or Replication Status in the header of the Accelerator view, changes to Started.
 - To stop incremental updates, click **Stop** in the header of the accelerator view. This opens the **Stop Replication** window, in which you can select between the following options for stopping the incremental update process:

Controlled

Completes all work in progress before stopping the process.

Immediate

Stops the process without completing work in progress. Compared with the **Controlled** option, using this option causes a restart of the process to take longer.

Restriction: If you stop replication inadvertently while a regular reload of a replication-enabled accelerator-shadow table is taking place (started by the load function or the SYSPROC.ACCEL_LOAD_TABLES stored procedure), the table state changes to *Suspended* after the load. That is, the table is treated like a faulty table. To re-enable incremental updates for this table, restart incremental updates on the accelerator and reload the table again.

Click **OK** after selecting an option. You return to the accelerator view, which now shows a replication status of *Stopped*.

Restarting incremental update processes

You can take various actions in case of errors or failures during incremental update processing. If restarting replication for an accelerator or disabling and re-enabling incremental updates for selected tables does not help, then the next step is to restart the associated incremental update processes.

About this task

When **Replication** > **Stop** is selected, updates are no longer applied to the accelerator-shadow tables on a particular accelerator. **Replication** > **Start** reverses this, and the application of updates sets in again.

In contrast to this, a restart of a process affects the propagation of incremental updates from the paired Db2 subsystems to the accelerator. You restart the propagation of updates either because the process has stopped for some reason, or because the current process has led to problems that cannot be solved otherwise.

Restarting the process thus means to interfere from a higher level and with broader impact. It should not be the first reaction to an incremental update problem, but neither should it be the last. As long as a problem cannot be clearly identified, the available actions should be considered in the following order:

- Disabling and re-enabling incremental updates for particular accelerator-shadow tables. See Including or excluding tables from incremental updates in the IBM Db2 Analytics Accelerator for z/OS: User's Guide.
- 2. **Replication** > **Stop** and **Replication** > **Start** on particular or on all accelerators attached to one or more Db2 subsystems. See "Starting or stopping incremental updates" on page 190.
- 3. Restarting the incremental update process.
- 4. Disabling and re-enabling incremental updates for the affected Db2 subsystems. See:
 - a. "Disabling incremental updates for a Db2 subsystem" on page 194
 - b. Depending on the chosen replication technology:

IBM Integrated Synchronization	"Enabling IBM Integrated Synchronization for a Db2 subsystem" on page 156
IBM InfoSphere Change Data Capture for z/OS (CDC)	"Enabling incremental updates for a Db2 subsystem" on page 187

Depending on the chosen replication technology, the scope of a process restart differs:

IBM Integrated Synchronization

One paired Db2 subsystem can be selected. A restart affects only the subscription to this subsystem.

IBM InfoSphere Change Data Capture for z/OS

A restart invariably affects all subscriptions based on IBM InfoSphere Change Data Capture for z/OS (CDC).

Procedure

1. Log on to the IBM Db2 Analytics Accelerator Console by using **telnet** or **ssh**. The preferred method is **ssh**.

For more information:

- "Using telnet to log on to the IBM Db2 Analytics Accelerator Console" on page 199
- "Using ssh to log on to the IBM Db2 Analytics Accelerator Console" on page 200
- 2. Press the Pause key, then Enter to display the following screen:

3. Type 2 and press Enter to display the submenu:

- 4. Type 3 and press Enter.
- 5. Provide the requested information:
 - a) Incremental update technologies:
 - Enter 1 for IBM Integrated Synchronization.
 - Enter 2 for IBM InfoSphere Change Data Capture for z/OS.
 - b) If 1. IBM Integrated Synchronization was selected in step "5.a" on page 193, select a Db2 subsystem. Otherwise go to step "5.c" on page 193

See the following example:

Select the Db2 subsystem for which incremental update processes based on integrated synchronization should be restarted:

1 : DB2WA1 2 : DB2WA2

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

Enter the number in front of a Db2 subsystem name if you want to restart the incremental update process for that subsystem.

c) Depending on your selection in <u>"5.a" on page 193</u>, you see one of the following messages on the screen:

1. IBM Integrated Synchronization

Based on the example in step "5.b" on page 193:

This will restart the integrated synchronization process on Db2 subsystem 'DB2WA1', are you sure (y/n)?:

2. IBM InfoSphere Change Data Capture for z/OS

Because you cannot select a Db2 subsystem:

```
This will restart all incremental update processes based on IBM InfoSphere Change Data Capture for z/0S, are you sure (y/n)?:
```

d) Type y and press Enter.

When this step has been completed successfully, the following message is displayed on the screen:

```
Restarting processes ... done.
Press <return> to continue
```

- e) If you selected 1. IBM Integrated Synchronization in step <u>"5.a" on page 193</u> and you want to restart the incremental update process for an additional Db2 subsystem, press Enter and repeat steps "4" on page 193 through "5.c" on page 193.
- f) Press the Enter key repeatedly until you reach the main menu of the IBM Db2 Analytics Accelerator Console.
- 6. Type x and press Enter to exit the console.

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.

About this task



Attention: If you disable incremental updates for a Db2 subsystem, you cancel all subscriptions at the same time. That is, tables on the corresponding accelerator that are currently enabled for incremental updates will be disabled. Therefore, after re-enabling the Db2 subsystem, you must also re-enable the tables by using the **Enable replication** function in IBM Db2 Analytics Accelerator Studio.

Procedure

1. Log on to the IBM Db2 Analytics Accelerator Console by using **telnet** or **ssh**. The preferred method is **ssh**.

For more information:

- "Using telnet to log on to the IBM Db2 Analytics Accelerator Console" on page 199
- "Using ssh to log on to the IBM Db2 Analytics Accelerator Console" on page 200
- 2. Press the Pause key, then Enter to display the following screen:

3. Type 2 and press Enter to display the submenu:

- 4. Type 2 and press Enter.
- 5. 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.
- 6. Type x and press Enter to exit the console.

Including or excluding loaded tables from WAITFORDATA queries

The WAITFORDATA option has been designed to add Hybrid Transactional Analytical Processing (HTAP) capabilities to the product. It minimizes the problem of replication latency for incrementally updated tables. However, a query might reference incrementally updated tables as well as tables that have not been enabled for incremental updates (loaded tables that can be referenced by accelerated queries).

About this task

If you set the WAITFORDATA option to a value greater than 0.0, loaded tables referenced by a query would, without further intervention, be waiting for incremental updates that never happen. The expiration period would always expire, and the query would fail because this is what happens when updates cannot be applied within the expiration period. It is therefore important that loaded tables can be included in queries with a WAITFORDATA setting. That is, after the passing of the specified WAITFORDATA delay, queries can go against a set of incrementally updated tables, a mixed set of tables, and even against a set of loaded tables only.

By default, loaded tables are excluded because there is a greater probability of data inconsistency. You can change this behavior and include loaded tables. This is recommended if you know that queries go against a mixed set of tables.

For a more detailed description of the WAITFORDATA option, see *Making queries wait for incremental updates* in the *IBM Db2 Analytics Accelerator for z/OS: User's Guide*.

Procedure

- 1. Log on to the IBM Db2 Analytics Accelerator Console.
- 2. Press the Pause key, then Enter to display the following screen:

3. Type 2 and press Enter to display the submenu:

```
main -> 2
You have the following options:

(0) - Go back one level
(1) - Enable incremental updates
(2) - Disable incremental updates
(3) - Restart incremental update processes
(4) - Include or exclude tables not enabled for incremental updates
    in WAITFORDATA queries
(5) - Allow or disallow queries to run if WAITFORDATA timeout is reached (6) - Change Db2 for
z/OS credentials for integrated synchronization
(6) - Change Db2 for z/OS credentials for Integrated Synchronization
(7) - Change IP and secure DDF port of Db2 for z/OS for Integrated Synchronization
(8) - (Menu) Manage certificates for replication
(Default 0) >
```

4. Type 4 and press Enter.

The following screen is displayed:

```
Specify if tables that were loaded, but not enabled for incremental updates, are to be included in delayed queries. (queries running with a setting of WAITFORDATA > 0.0)

It cannot be guaranteed that the data in these accelerator-shadow tables is up-to-date. If changes have been made to a Db2 source table, a new load process must be started to propagate these changes to the accelerator. Only then will these changes be reflected in accelerated queries.

The current setting is: 'Included'

Do you want to include loaded tables in WAITFORDATA queries? ('yes', 'no', or <return> to cancel): <cursor>
```

5. To change the value from no to yes or vice versa, type n or y at the prompt (position of <cursor> in the screen extract). Then press Enter.

To exit the dialog without changes, just press Enter.

After a change or confirmation of the present value, you see the following message:

6. Press Enter again to return to the submenu.

What to do next

To change your settings, repeat steps <u>"4" on page 196</u> through <u>"6" on page 196</u>. Specify a different value in step "5" on page 196.

Configuring the expiration behavior of WAITFORDATA queries

In the context of Hybrid Transactional Analytical Processing (HTAP), you can determine whether WAITFORDATA queries are aborted or executed despite the fact that not all incremental updates could be applied within the specified period.

About this task

WAITFORDATA is an option (specified as a special register, ZPARM, or bind option) that can be used for queries against incrementally updated tables. It effects a hold on the query execution until the most recent updates (changes committed in Db2 for z/OS) have been applied to the tables referenced by the query. For a more detailed description, see Making queries wait for incremental updates in the IBM Db2 Analytics Accelerator for z/OS: User's Guide.

Query execution will be stopped for the period specified by the WAITFORDATA option. During that time, it is expected that all pending incremental updates can be applied. You can specify how to proceed with a query in case the updates cannot be applied within the time allotted for this task:

FAIL

This setting leads to the failure and the subsequent abortion of the query. It is the default.

CONTINUE

This setting effects the execution of a query despite the fact that the updates could not be applied within the specified time. It bears the risk that the query is run on obsolete data and hence delivers results that are not up-to-date.

Procedure

- 1. Log on to the IBM Db2 Analytics Accelerator Console.
- 2. Press the Pause key, then Enter to display the following screen:

3. Type 2 and press Enter to display the submenu:

4. Type 5 and press Enter.

The following screen is displayed:

Specify which action to take after the passing of the delay for queries that were run with the WAITFORDATA option.

FAIL
means that the query will abort with SQLCODE -904 if the delay expired before the table data on the accelerator could be updated.

CONTINUE
means that the query will run on the accelerator after the passing of the delay despite the absence of current data. The query might thus run on obsolete data. Passing of the delay is indicated by SQLCODE +904. This SQL warning will only show if your application program fetches the entire result set and propagates warnings to end users.

The current setting is: FAIL

Enter your choice ('fail', 'continue', or <return> to cancel): <cursor>

5. To change the value from FAIL to CONTINUE or vice versa, type c or f at the prompt (position of <cursor> in the screen extract). Then press Enter.

To exit the dialog without changes, just press Enter.

After a change or confirmation of the present value, you see the following message:

Done. The new setting is now active and will be used with the next query that is offloaded to the accelerator.

6. Press Enter again to return to the submenu.

What to do next

To change your settings, repeat steps <u>"4" on page 197</u> through <u>"6" on page 198</u>. Specify a different value in step "5" on page 198.

Chapter 12. 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.

Logging on to the IBM Db2 Analytics Accelerator Console

All functions or features discussed in this section are accessed from the IBM Db2 Analytics Accelerator Console, so that a console log-on is required.

About this task

You can use **telnet** or **ssh** to log on to the IBM Db2 Analytics Accelerator Console. The recommended method is **ssh** because this means the connection is encrypted.

Using telnet to log on to the IBM Db2 Analytics Accelerator Console

A telnet logon to the IBM Db2 Analytics Accelerator Console is still possible for compatibility reasons, but not recommended. Consider using the more secure ssh method instead.

Procedure

- 1. 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 203.0.113.8 1600
```

4. When prompted, enter acceladm as the user ID, followed by the console password.

The initial password is dwa-1234. You must change this password at the first logon.

The console also offers an option to create additional console users with different passwords.

5. Press the Pause key, then Enter to display the following screen:

What to do next

Type x and press Enter to exit the **console** and close the telnet session.

Using ssh to log on to the IBM Db2 Analytics Accelerator Console

A secure shell (ssh) connection ensures that the information exchanged between your IBM Z system and the accelerator console is encrypted.

Before you begin

• An ssh client must be installed on the mainframe computer (IBM Z). Starting with z/OS V2R2, the OpenSSH package, which contains an ssh client, is part of the operating system.

If you still use z/OS V2R1, you can obtain an ssh client by installing the respective component from BM Ported Tools for z/OS. See:

https://www-03.ibm.com/systems/z/os/zos/features/unix/port_tools.html

Mind that IBM does not offer this package anymore. However, the functionality lives on in more specific packages provided by the Rocket software company (https://www.rocketsoftware.com/zos-open-source). These packages usually include the OpenSSH component that contains the ssh client, so that you can connect from the z/FS file system to the accelerator. Note that you only need the client component of OpenSSH for maintenance updates (called ssh).

 An ssh client must also be installed on the workstation used to access the IBM Db2 Analytics Accelerator for z/OS Console. You can use a tool like PuTTY, which comes free of charge.

Procedure

1. On your client workstation, open an ssh shell and connect to the mainframe computer by using the following command:

```
ssh <mainframe-ip>
```

where <mainframe-ip> is the IP address or host name of the IBM Z system that is connected to the accelerator.

- 2. Enter the logon user ID and password for the IBM Z system.
- 3. When connected, enter the following commands from the shell to connect to the IBM Db2 Analytics Accelerator Console:

```
ssh -p 2222 configuration-console@<accelerator-ip>
```

where

-p 2222

Specifies a non-default connection port.

<accelerator-ip>

Is the IP address of the accelerator.

For IBM Db2 Analytics Accelerator on an IBM Integrated Analytics System, this is the virtual IP or wall IP address.

For Db2 Analytics Accelerator on Z, this is the IP address of the network that you labeled DB2 in the **Appliance Installer**.

4. When prompted, enter acceladm as the user ID, followed by the console password.

The initial password is dwa-1234. You must change this password at the first logon.

The console also offers an option to create additional console users with different passwords.

5. Press the Pause key, then Enter to display the following screen:

What to do next

Type x and press Enter to exit the **console** and close the **ssh** session.

Setting resource limits for Db2 subsystems

If an accelerator is shared by multiple Db2 subsystems, you can limit the processing resources that can be claimed by individual Db2 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 rebalance.

About this task

By default, all of an accelerator's processing resources are shared by the connected Db2 subsystems. If only one subsystem passes a workload to the accelerator, this subsystem can use all of the resources. If more subsystems generate workloads, the resources are distributed dynamically among these subsystems. 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.

Note: In the following description, the word *shares* is used as an abbreviation for the term *soft CPU shares* as it is defined in the Db2 (for Linux, UNIX and Windows) documentation.

You can reserve a certain portion of the available resources for administrative tasks and incremental updates. These limits are percentages set with the help of IBM support. Once set, they remain fixed. In addition, a major share of the resources will always be used by the system for processing. The rest can be freely distributed among the connected Db2 subsystems. If you set a minimum for a Db2 subsystem, that limit is valid for the resource group associated with that Db2 subsystem on the accelerator.

The limit becomes effective only if a workload is assigned to the resource group. Otherwise, the limit is ignored, and all of the available resources can be used by other processes. If a workload is assigned to the resource group, the configured minimum share is reserved for the Db2 subsystem to which it has been assigned.

The value that you set as the minimum cannot be higher than 65535 for each connected Db2 subsystem. If you try to set a higher value, the following message is displayed:

Value too high! Maximum is 65535

If you give more shares to one of the connected Db2 subsystems, such that the total number of shares increases, the quotas are changed proportionally. See the following example:

Internal System Resource Share Entitlements:

Workload Name	Current Entitlement	Target Entitlement
ADMIN	222 (4.99%)	277 (4.99%)
INCREMENTAL UPDATE	222 (4.99%)	277 (4.99%)
SYSTEM	2002 (45.03%)	2002 (36.03%)

Tenant Resource Share Entitlements:

ID	Location Name	Current Entitlement	Target Entitlement
54	DWXDC1L	1000 (22.49%)	1000 (18.00%)
57	DWXDB1Q	1000 (22.49%)	2000 (36.00%)

You can see that the number of shares reserved for Db2 subsystem DWXDB1Q has increased by 1000. This affects the quotas in the two Target Entitlement columns as shown: In the lower table, the number of shares reserved for the DWXDC1L subsystem remains the same (1000); in the upper table, the number of shares used by the system also remains the same (2002). The two Db2 subsystems and the system claim 5002 shares in sum. These 5002 shares represent 90 percent of the total because the percentages for administrative tasks and incremental updates remain fixed, each at 5 percent of the total (the differences in the table were caused by rounding). So the number of shares for administrative tasks and incremental updates increases to 277 for each process, and the total number of shares increases by 1110, from 4446 to 5556.

You assign accelerator processing resources on the IBM Db2 Analytics Accelerator Configuration Console. The values that you enter are then mapped to guaranteed resource allocation (GRA) definitions for the IBM Integrated Analytics System.

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. Log on to the IBM Db2 Analytics Accelerator Console.
- 2. Press the Pause key, then Enter to display the following screen:

3. Type 1 and press Enter to display the submenu:

4. Type 11 and press Enter.

You see a screen similar to the following:

(Default 0) > 11main -> accel -> managewlm -> wlmconfig ACCELERATOR WORKLOAD MANAGEMENT CONFIGURATION

Internal System Resource Share Entitlements:

Workload Name	Current Entitlement	Target Entitlement
ADMIN	222 (4.99%)	222 (4.99%)
INCREMENTAL UPDATE	222 (4.99%)	222 (4.99%)
SYSTEM	2002 (45.03%)	2002 (45.03%)

Tenant Resource Share Entitlements:

Ŀ	ID	Location Name	Current Entitlement	Target Entitlement
	54	DWXDC1L	1000 (22.49%)	1000 (22.49%)
	57	DWXDB1Q	1000 (22.49%)	1000 (22.49%)

- (0) back Go back one level(1) changetenantentitlement Modify the resource share entitlement of a single accelerator tenant
- (2) applywlmconfig - Apply the current workload management configuration to the accelerator
- 5. Enter one of the options displayed on the screen in step "4" on page 202, that is:
 - To exit the screen and go back one level
 - To change the resource allocation
 - (2)
- To save and apply the values currently displayed
- 6. To change the current resource allocation, enter 1.

You are asked to select a Db2 subsystem:

```
Which tenant entitlement should be updated?
(0) - abort
(1) - DWXDC1L
(2) - DWXDB1Q
```

7. Select the subsystem by entering the correct number.

Using the previous example in step "4" on page 202, this would be number 2.

You are asked for the new number of shares to be allocated to the selected Db2 subsystem (tenant). Valid are integers between 1 and 65535.

8. Type an appropriate number and press Enter.

For example: 2000.

The screen in step "4" on page 202 is updated and shows the new values. For example, if you entered 2000, you would see the following values in the Target Entitlement column of the Tenant Resource Share Entitlements section:

. Internal System Resource Share Entitlements:

Workload Name	Current Entitlement	Target Entitlement
ADMIN	222 (4.99%)	277 (4.99%)
INCREMENTAL UPDATE	222 (4.99%)	277 (4.99%)
SYSTEM	2002 (45.03%)	2002 (36.03%)

Tenant Resource Share Entitlements:

ID	Location Name	Current Entitlement	Target Entitlement
54	DWXDC1L	1000 (22.49%)	1000 (18.00%)
57	DWXDB1Q	1000 (22.49%)	2000 (36.00%)

The values in the Target Entitlement column of the Internal System Resource Share Entitlements section also change accordingly, taking the Workload Management configuration of the system into account.

9. Enter 2 to save and apply the new values.

The screen in step <u>"4" on page 202</u> is updated once more. It now shows identical values in the Current Entitlement and Target Entitlement columns.

- 10. Repeat steps <u>"4" on page 202</u> through <u>"9" on page 204</u> for other Db2 subsystems as often as required.
- 11. When finished, enter 0 to leave the submenu.
- 12. Press the Enter key repeatedly until you reach the main menu of the console.
- 13. Type x and press Enter to exit the console.

Related information

Enabling and setting CPU shares

Resource assignment with service classes

Mapping of z/OS importance levels to priorities of the accelerator database engine

During query execution, z/OS importance levels are automatically converted. Queries thus run with the corresponding priorities of the database engine on the accelerator.

The z/OS importance levels map to accelerator database priorities as shown in Table 5 on page 204.

Table 5. Mapping of z/OS importance levels to accelerator database priorities		
z/OS importance level	Priority of accelerator database engine	
System	Critical	
Importance 1	Critical	
Importance 2	Critical	
Importance 3	High	
Importance 4	Medium	
Importance 5	Medium	
Discretionary	Low	

Behavior

Suppose you have a set of queries with identical resource requirements. There are queries at each priority level. Queries are admitted in a ratio of 8:4:2:1. This means for every query of low priority, 2 queries of medium priority, 4 queries of high priority, and 8 queries of critical priority are admitted. Queries of lower priority will not be suppressed by queries of a higher priority.

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 change of importance levels over time is not supported by the database engine of the accelerator.

Related tasks

Setting resource limits for Db2 subsystems

If an accelerator is shared by multiple Db2 subsystems, you can limit the processing resources that can be claimed by individual Db2 subsystems. This allows you to adapt the processing resources to the size of the query workloads generated by each subsystem.

Using the call-home function

A connected Db2 subsystem or data sharing group regularly polls the accelerator server for system health information. In turn, corresponding event messages (DSNX881I messages) are returned to the subsystem or group. These messages are written to the z/OS system log (SYSLOG) of the LPAR in which the subsystem or group is located. You can enable an automatic event notification service for the most critical DSNX881I events (call-home function). When one of these events occurs, one or two contact addresses are notified by email, and a support case is opened automatically.

Configuring the call-home function on the HMC

The Hardware Management Console (HMC) offers a call-home function for IBM Z. You must connect this function to your accelerator system so that it can be used to send notifications about accelerator-related events to selected administrators.

Before you begin

Be sure that you meet the prerequisites before you proceed with the configuration:

- Currently, the call-home function can be used for deployments of IBM Db2 Analytics Accelerator for z/OS on an IBM Integrated Analytics System only.
- To use the HMC call-home function with IBM Db2 Analytics Accelerator, your network must have been prepared accordingly:
 - Two cables are required between the IBM Integrated Analytics System and the HMC network switch.
 - A sufficient number of IP addresses must have been defined for the purpose. That is, each
 accelerator processing node requires a static IP address in the HMC network. In addition, a floating
 IP address must be defined, which is used for the head node.
 - The network connection between IBM Db2 Analytics Accelerator and the HMC must be a switched connection (Layer-2 network), where the IP addresses of both devices are in the same subnet. A routed network connection (Layer-3 network) won't work.
 - For the call-home function to work properly, it must be possible to reach IBM support servers from the HMC. Configure the HMC so that it can reach the following IP addresses over port 443:

Table 6. Support server IP addresses	
Host name	IP addresses
esupport.ibm.com	129.42.60.189
	129.42.54.189
	129.42.56.189

For more information, see Configuring access to the IBM Support servers.

In many cases, the support engineer makes the necessary provisions during the initial installation of IBM Db2 Analytics Accelerator on the basis of the submitted site survey. However, if you do not yet meet the network requirements, you must open a support case.

Procedure

- 1. Enter the URL of the HMC in a web browser.
- 2. Log on to the HMC with the proper user name and password.
- 3. Click **Service Management** in the navigator on the left.
- 4. On the Service Management page, select Configure IDAA Callhome.

The following page is displayed:



Figure 55. The **Configure IDAA Callhome** page on the HMC

5. To add an IP address, click **Add**, then enter the address. Repeat this procedure for all the IP addresses defined for your accelerators in the HMC network (see previous *Before you begin* section).

Hint: The radio buttons that appear in front of the IP addresses are for the removal of addresses from the list. To this end, you pick an address by selecting the radio button, and then click **Delete**.

- 6. Click **Close** when you are finished adding IP addresses.
- 7. The following step is required only if it has not been completed before: Provide a domain name and password for the HMC domain.
 - a) In the navigator on the left, click **HMC Management**.
 - b) On the **HMC Management** page, click **Domain Security**.You see the following page:

IBM Hardware Management Console		
Home Domain Security	Ľ×	
0— Domain Security		
Current domain name: Domain name:	×2,777.21	
Current password status: New password:		
Verify password:		
Select an action: Apply to the Hardware Managen Apply to defined objects and the Reset to manufacturing default of	Hardware Management Console	
	rmation about using this task correctly, of applying a customized domain name, or its defined objects.	

Figure 56. The **Domain Security** page on the HMC

- c) Enter a domain name in the **Domain name** field.
- d) Provide a password in the **New password** field.
- e) Confirm the password by re-typing it in the **Verify password** field.
- f) Click **OK** to save the settings.
- 8. You can log out of the HMC now.

Configuring the call-home function on the accelerator

Use this submenu of the IBM Db2 Analytics Accelerator Console to provide the information required to successfully send call-home notifications.

Procedure

- 1. Log on to the IBM Db2 Analytics Accelerator Console.
- 2. Press the Pause key, then Enter to display the following screen:

3. Type 4 and press Enter to display the submenu:

```
main -> 4
------
You have the following options:

(0) - Go back one level
(1) - Configure call home
(2) - Enable call home
(3) - Disable call home
(4) - Display settings
(5) - Test connectivity
```

4. Type 1 and press Enter.

Provide the following information:

- 5. Enter the company name ('' or '0' to exit):. Type the name of your organization and press Enter.
- 6. Enter the company address, line 1/3 ('' or '0' to exit):. Start to type your organization's address here (you have three lines for this information). Press Enter at the end of the first line.
 - It is mandatory to provide information on this and the following line (see next step). Line 3/3 is optional.
- 7. Enter the company address, line 2/3 ('' or '0' to exit):. Continue with your organization's address. Press enter at the end of the second line.
- 8. Enter the company address, line 3/3, optional ('0' to exit):. Complete your organization's address on this line if needed and press Enter when finished. Otherwise, just press Enter.
- 9. Enter the IBM Customer Number ('' or '0' to exit):. Type your IBM customer number and press Enter.
- 10. Do you want to define separate contact persons? [n]:. Decide whether you want to specify two different email contacts for call-home notifications, one for software and one for hardware issues.
 - If you enter y, you will have to provide information about the secondary contact in the steps that follow. That is, you will have to answer the questions <u>"11" on page 208</u> through <u>"14" on page 208</u> for the secondary contact as well.
- 11. Enter the name of the primary contact for PMRs ('' or '0' to exit):. Type the name of the person in your organization who is supposed to receive the call-home notification emails, for example Bill Murkham. Then press Enter.
- 12. Enter the phone number of the primary contact for PMRs ('' or '0' to exit):. Type the phone number of the person you entered in step <u>"11" on page 208</u>, for example 001-487-185-6123, and press Enter.
- 13. Enter the email address of the primary contact for PMRs ('' or '0' to exit): . Type the email address of the person you entered in step <u>"11" on page 208</u>, for example billym@xyz.com, and press Enter.
- 14. Enter the cell phone number of the primary contact for PMRs, optional ('0' to exit):. Type the cell-phone number of the person you entered in step <u>"11" on page 208</u>, for example 001-414-767-6741, and press Enter.

Under the heading System Information:

15. Enter the system description ('.' to exit):

Enter a description that identifies the system or accelerator for which you want to enable call home.

- 16. Enter the system location ('.' to exit):
 - Enter a description that identifies the location of the system or accelerator for which you want to enable call home. That is, something like Boston site, bldg. 2.
- 17. Enter the 2-digit ISO country code for the country where the system is located ('.' to exit):

Enter the appropriate country code, for example US for the United States.

Under the heading HMC Information:

- 18. Enter the System z HMC IP address ('.' to exit):

 Enter the IP address of the HMC that is used to monitor the accelerator LPARs. Ask an IBM Z systems programmer or network administrator for this IP address if you don't know it.
- 19. Enter the System z HMC domain name ('.' to exit):

 Enter the domain name that you specified on the **Domain Security** page of the HMC. See step <u>"7.c"</u> on page 207.
- 20. Enter the System z HMC domain password ('.' to exit):

 Enter the domain password that you specified on the **Domain Security** page of the HMC. See step "7.e" on page 207.
- 21. Press 'y' to save your configuration or any other key to cancel: Type y and press Enter to confirm and save your information.
- 22. Press Enter to return to the submenu.

Enabling the call-home function

You must enable the call-home function on the IBM Integrated Analytics System before you can use it.

Before you begin

Make sure that you have completed the previous step, *Configuring the call-home function on the accelerator* successfully.

Procedure

- 1. Log on to the IBM Db2 Analytics Accelerator Console.
- 2. Press the Pause key, then Enter to display the following screen:

3. Type 4 and press Enter to display the submenu:

- 4. Type 2 and press Enter.
- 5. At the prompt Press 'y' to enable call home or any other key to cancel [y]:, type y and press Enter.

You see the following message:

Done.
Press <return> to continue

6. Press Enter to return to the submenu.

Disabling the call-home function

To disable the call-home function on an IBM Integrated Analytics System, complete the steps in this topic.

Procedure

- 1. Log on to the IBM Db2 Analytics Accelerator Console.
- 2. Press the Pause key, then Enter to display the following screen:

3. Type 4 and press Enter to display the submenu:

- 4. Type 3 and press Enter.
- 5. At the prompt Press 'y' to disable call home or any other key to cancel [y]:, type y and press Enter.

You see the following message:

```
Done.
Press <return> to continue
```

6. Press Enter to return to the submenu.

Displaying call-home settings

To display or verify your settings for the call-home function, you can use another option on the call-home submenu of the IBM Db2 Analytics Accelerator Console.

Procedure

- 1. Log on to the IBM Db2 Analytics Accelerator Console.
- 2. Press the Pause key, then Enter to display the following screen:

3. Type 4 and press Enter to display the submenu:

```
main -> 4
------
You have the following options:

(0) - Go back one level
(1) - Configure call home
(2) - Enable call home
(3) - Disable call home
(4) - Display settings
(5) - Test connectivity
```

- 4. Type 4 and press Enter.
- 5. Check the output on the screen.

Note: For security reasons, the IP address, domain name, and domain password specified on the Hardware Management Console (HMC) and on the accelerator (compare steps <u>"18" on page 209</u> through "20" on page 209) are not displayed.

Example

The following example shows a setup with two notification addresses, one to be used in case of software events, the other in case of hardware events.

```
Company Name : XYZ GmbH
Company Address 1/3 : Schoenaicher First 220
Company Address 2/3 : 71032 Boeblingen
Company Address 3/3 : Germany
IBM Customer Number : 1717171

Contact Name : Manfred Koebel
Contact Phone Number : +49-1234-567-8990
Contact Email Address : mkoeb@xyz.de
Contact Cell Phone Number : +49-9876-543-2112
Contact Event Setting : SOFTWARE

Contact Name : Speedy Gonzales
Contact Phone Number : 321-321-4321
Contact Email Address : schnellste.maus@xyz.mx
Contact Cell Phone Number : 123-123-1234
Contact Event Setting : HARDWARE

Press <return> to continue
```

Testing connectivity

To check whether the network setup for the call-home function works correctly, you can run a connectivity test from the IBM Db2 Analytics Accelerator Console.

Procedure

1. Log on to the IBM Db2 Analytics Accelerator Console.

2. Press the Pause key, then Enter to display the following screen:

3. Type 4 and press Enter to display the submenu:

```
main -> 4
-----
You have the following options:

(0) - Go back one level
(1) - Configure call home
(2) - Enable call home
(3) - Disable call home
(4) - Display settings
(5) - Test connectivity
```

4. Type 5 and press Enter.

The following text is displayed on the screen:

```
This option verifies that the call home configuration is correct and that the call home service is able to communicate with the IBM support servers. Press 'y' to test call home connectivity or any other key to cancel [y]:
```

5. To start the test, type y and press Enter.

A message on the screen informs you whether the test succeeded. If it fails, hints at possible causes are given, for example:

```
Connectivity test failed.

Make sure that the HMC domain name and password are correct.

Done
Press <return> to continue
```

6. Press Enter to exit the screen and return to the previous submenu.

Enabling queries against system-temporal and bi-temporal tables

Temporal tables, introduced with Db2 10 for z/OS, are fully supported, which means that you can run accelerated queries against all columns of these tables, including the timestamp columns that define the time interval. System-temporal and bi-temporal tables are special because they create history tables. Accelerated queries can succeed only if both, the base and the history tables are defined and loaded on the accelerator. IBM Db2 Analytics Accelerator therefore checks for their presence. This is why queries against these two types of temporal tables must be enabled explicitly.

About this task

• To enable accelerated queries, you must define the base tables and the related history tables separately as accelerator-shadow tables and also load both types of tables.

Business-temporal tables need not be enabled.

• Keep the base table and the related history table in sync on the accelerator. You might get confusing results if the load times of these tables do not match.

• A history table can be archived with the High Performance Storage Saver (HPSS), provided that it is partitioned by range. However, the base table of a system-temporal or bi-temporal table that is partitioned in the same manner cannot be archived by the HPSS.

Procedure

- To enable accelerated queries against system-temporal or bi-temporal accelerator-shadow tables, set or include option 5 in the value of the QUERY_ACCEL_OPTIONS parameter or ZPARM.

 For more information, follow the appropriate link at the end of this topic.
- If you want to enable incremental updates for system-temporal and bi-temporal tables, you must "tell" IBM InfoSphere Change Data Capture for z/OS that it needs to work on the regular tables as well as on the history tables.

This is done by adding the following line to the CHCDBMxx configuration member:

REPLTEMPORALTABLES = DECOUPLE

Chapter 13. 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 components. However, if you must update more than one component, follow the suggested order.

Order in which to update components

- 1. IBM Db2 Analytics Accelerator software. When in doubt: It is better if the accelerator software is at a newer level than the stored procedures. See "Updating IBM Db2 Analytics Accelerator software" on page 216.
- 2. Prerequisite PTFs for Db2 11 for z/OS or Db2 12 for z/OS. See <u>Installing prerequisite PTFs for Db2 for z/OS</u>.
- 3. Administration client. See:
 - "Updating IBM Db2 Analytics Accelerator Studio" on page 223
 - "Updating IBM Data Server Manager" on page 224
- 4. Stored procedures. See "Updating the IBM Db2 Analytics Accelerator stored procedures" on page 225
- 5. Db2 Analytics Accelerator on Z. See Updating Db2 Analytics Accelerator on Z in the Installation Guide...

Installing prerequisite PTFs for Db2 for z/OS

When you update your existing installation, first apply the prerequisite program temporary fixes (PTFs) for Db2 for z/OS.

Before you begin

- Make sure that the database management system is at least Db2 11 for z/OS. Db2 10 is not supported.
- Make sure that you meet the software requirements of Db2 for IBM Db2 Analytics Accelerator for z/OS Version 7.5.x. For more information, click the link at the end of this topic.

About this task

The product continues to work after installing Db2 for z/OS PTFs, meaning that you can update and work with your accelerators without also having to update the accelerator stored procedures.

Procedure

Install the latest prerequisite Db2 for z/OS PTFs for the product.

Related tasks

Updating the IBM Db2 Analytics Accelerator stored procedures

Read how to update the Db2 for z/OS components of IBM Db2 Analytics Accelerator for z/OS.

Updating Db2 Analytics Accelerator on Z

Follow the steps in this section to upload a new image (appliance) for Db2 Analytics Accelerator on Z and activate it.

Related information

Updating IBM Db2 Analytics Accelerator software

These are software components that run on the accelerator hardware, that is, the IBM Integrated Analytics System. This chapter does not apply to updates for Db2 Analytics Accelerator on Z. Components

on the IBM Integrated Analytics System are the IBM Db2 Analytics Accelerator, the SQL engine, and the operating system of the accelerator.

Updating administration clients

Refer to the appropriate section to update your administration client:

Prerequisites and Maintenance for IBM DB2 Analytics Accelerator for z/OS 7.1

Updating IBM Db2 Analytics Accelerator software

These are software components that run on the accelerator hardware, that is, the IBM Integrated Analytics System. This chapter does not apply to updates for Db2 Analytics Accelerator on Z. Components on the IBM Integrated Analytics System are the IBM Db2 Analytics Accelerator, the SQL engine, and the operating system of the accelerator.

Related tasks

Installing prerequisite PTFs for Db2 for z/OS

When you update your existing installation, first apply the prerequisite program temporary fixes (PTFs) for Db2 for z/OS.

Updating the IBM Db2 Analytics Accelerator stored procedures

Read how to update the Db2 for z/OS components of IBM Db2 Analytics Accelerator for z/OS.

Updating Db2 Analytics Accelerator on Z

Follow the steps in this section to upload a new image (appliance) for Db2 Analytics Accelerator on Z and activate it.

Related information

Updating administration clients

Refer to the appropriate section to update your administration client:

Installing update packages for the accelerator

IBM Db2 Analytics Accelerator for z/OS software includes specific libraries and other code to be installed on the IBM Integrated Analytics System.

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 z/OS UNIX file system (z/FS).
- The user ID is authorized to run the SYSPROC.ACCEL_UPDATE_SOFTWARE2 stored procedure.

About this task

The entire installation or update procedure consists of the following steps:

- 1. Downloading the fix pack or update package.
- 2. Transferring the update packages to the accelerator. You start this process from your administration client, that is, from IBM Db2 Analytics Accelerator Studio or IBM Data Server Manager.
- 3. Activating the newly installed software. This step is also executed from your administration client.

This topic covers step "1" on page 216 only.

Procedure

Downloading the fix pack or update package. You have two options here:

Download the fix pack or update package from IBM Fix Central to your workstation and then upload
the package from your workstation using IBM Db2 Analytics Accelerator Studio. Place the checksum
file for installation verification in the same directory on your workstation.

There is a checksum file for every supported algorithm. Supported algorithms are SHA256, SHA1 and MD5. The corresponding files on IBM Fix Central are:

- sha1sums.txt
- md5sums.txt

Any of these files contains one line for each package. Every line consists of two values. The first value is the hash used for verification; the second value is the corresponding package name.

Important: During the process, a large global temporary table is created in Db2 for z/OS. This table requires up to 25 GB of disk space. Make sure that this amount of disk space is available.

- Download the fix pack or update package from IBM Fix Central and save it directly in the directory specified by the AQT_HOST_PACKAGE_DIRECTORY environment variable. To use this second option, you must:
 - a. Log on to your z/OS data server using a remote client, such as IBM Personal Communications.
 - b. Set the AQT_HOST_PACKAGE_DIRECTORY environment variable to a directory path in the z/FS. Then save the update package in this directory. Place the checksum file for installation verification in the same directory (for details, see the previous description of the first download option).

Important:

- In the z/FS, the checksum files must be available in UTF-8 format. However, if you transfer the files from an FTP client, this format is identified as a text format and the transfer mode will probably be set to ascii. You must change the transfer mode to binary because this guarantees that the file stays as it is. If ascii is used as the transfer mode, chances are that the files are converted to EBCDIC format. This leads to an error when the checksum files are used for verification.
- An update packages requires up to 25 GB of disk space. Make sure that enough space is available in the z/FS directory that AQT_HOST_PACKAGE_DIRECTORY points to.

Important: The directory that AQT_HOST_PACKAGE_DIRECTORY points to 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 the installation packages in the file system. So make sure that the directory can accommodate all the files.

Results

Your local workstation or z/FS directory (as set by the AQT_HOST_PACKAGE_DIRECTORY environment variable) on your z/OS data server 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 on an IBM Integrated Analytics System by completing the steps in this section.

Before you begin

Make sure that the following conditions apply:

- You have downloaded suitable accelerator installation packages and their respective checksum files to a workstation that is connected to your z/OS data server, or such packages exist in a z/FS directory on that server. For the latter option, the AQT_HOST_PACKAGE_DIRECTORY environment variable must point to the z/FS directory, and the checksum files must have been transferred in *binary* mode.
- The value of the AQT_HOST_PACKAGE_DIRECTORY environment variable determines where IBM Db2 Analytics Accelerator Studio looks for software updates. If it is not set correctly, the packages cannot be located and thus cannot be transferred or installed.

• The user ID has read and write access to the directory that the AQT_HOST_PACKAGE_DIRECTORY environment variable points to and to all its sub-directories. 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. 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 7.5.

About this task

You can transfer update packages from the z/OS UNIX file system (z/FS) to the accelerator, or from a workstation on which IBM Db2 Analytics Accelerator Studio is installed, and which is connected to the z/OS data server. The first method requires that the packages have been placed in the proper z/FS directory. To get them there, you can, for example, download packages from IBM Fix Central directly to the z/FS of the mainframe. To this end, you need a client application on the mainframe that is capable of handling FTPS (FTP over TLS) traffic. The download site at IBM Fix Central offers an option called *Download using bulk FTPS*. The second method requires that the update packages have been downloaded to the client workstation before. To actually use a transferred package, you must activate it. How to do this is described in a later topic. See the hint at the end of this description.

Installing a software update invokes the SYSPROC.ACCEL_UPDATE_SOFTWARE2 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 your administration client.
- 2. Select the accelerator:

IBM Db2 Analytics Accelerator Studio	a. In the Administration Explorer , select the Accelerators folder.
	b. In the Object List Editor , double-click an accelerator name to open the corresponding Accelerator view.
IBM Data Server Manager	a. Click in the upper left.
	b. Select ADMINISTER > Accelerators.
	c. In the list of accelerators, select the name of the accelerator (row).

3. Open the **Transfer Software Version** window:

IBM Db2 Analytics Accelerator Studio	a. If necessary, expand the About section. b. In the About section, click the Transfer updates link.
IBM Data Server Manager	In the lower half of the Accelerator view, click Transfer Updates .

4. In the **Transfer Software Version** or **Transfer Updates** window, you can see all software packages that are available in the zFS of your z/OS data server. Select the packages that you want to transfer:

· ·	Select the appropriate check boxes in the first column of the package list.
IBM Data Server Manager	Select the appropriate row in the package list.



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 Updates** window might be confusing, especially if it also contains older packages. To find the correct package numbers, see the closing information or the release notes for the latest program temporary fix (PTF). You find the closing information or the release notes 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 the list contains too many packages so that selecting the proper ones becomes awkward, select unwanted package names and click **Remove**. This will just remove the packages from the list or view, but won't delete the package files from the z/FS or your client workstation.
- 5. Choose the appropriate method to transfer the selected packages:

Transfer file from client

To transfer one or more downloaded packages from the connected client workstation.

Transfer file from IBM Z

To transfer one or more packages from the z/FS of your z/OS data server.

Note: If the verification message says ... package was deployed without integrity check ..., the reason might be that the checksum files were not transferred in *binary* mode. See the second **Important** note in step "1" on page 216.

What to do next

To save space in the SMP/E target directory of the z/FS, 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.

Activating an updated accelerator software version

To use a recently transferred software package, you must first activate that package by selecting the version that the package contains.

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. You can check if there are any open console windows, and, if necessary, close these windows by following these steps:

IBM Db2 Analytics Accelerator Studio	Open the appropriate Accelerator view. Scroll down to the Tables section and click List Tasks .
	3. Cancel all active console tasks from the List Tasks window.

IBM Data Server Manager	Select the appropriate accelerator (row) in the Accelerator view.
	2. Click Manage .
	3. Click List Tasks .
	4. Cancel all active console tasks from the List Tasks window.

• If incremental updates are enabled, stop these for each Db2 subsystem connected to the accelerator. To do this from your administration client:

IBM Db2 Analytics Accelerator Studio	In the Administration Explorer, select the Accelerators folder.	
	In the Object List Editor , double-click an accelerator name to open the corresponding Accelerator view.	
	3. In the header section of the Accelerator view, look for Replication: and click Stop .	
	4. Repeat the preceding steps for each Db2 subsystem connection to the same accelerator.	
IBM Data Server Manager	1. Click in the upper left.	
	2. Select ADMINISTER > Accelerators.	
	3. In the list of accelerators, select the name of the accelerator (row).	
	4. In the lower half of the window, locate the Replication section and click the Stop button next to the Replication Status .	
	5. Repeat the preceding steps for each Db2 subsystem connection to the same accelerator.	

Note: If you have connections to Db2 data sharing groups (DSGs) and prefer to run the SYSPROC.ACCEL_SET_REPLICATION stored procedure from a client command line: It is not necessary to stop incremental updates for each individual member. A DSG is treated like a single connection and stopping incremental updates for the primary member will also stop these for the other members.

About this task

- Earlier product versions allow you to switch between older and newer versions of the IBM Db2 Analytics Accelerator software. This is not possible with IBM Db2 Analytics Accelerator for z/OS Version 7.x.x. Once you have activated a particular version, you cannot go back to an older one.
- The application of an update affects all Db2 subsystems that are connected to an accelerator.
- The activation of a different accelerator software version causes a restart of the accelerator.

Activating a software update invokes the SYSPROC.ACCEL_UPDATE_SOFTWARE2 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 "Order in which to update components" on page 215.
- When the accelerator is restarted with a new accelerator software version, accelerated queries cannot be processed on that accelerator before it reaches the *online* state. The only way to avoid a delay is to use a second accelerator in a workload balancing setup. That is, a second accelerator which is connected to the same Db2 subsystem, and which contains the same set of tables. Note that the tables must be loaded and enabled.

Procedure

1. Select the accelerator you want to apply a software update to:

IBM Db2 Analytics Accelerator Studio	a. In the Administration Explorer , select the Accelerators folder.
	b. In the Object List Editor , double-click an accelerator name to open the corresponding Accelerator view.
	c. 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.
IBM Data Server Manager	a. Click in the upper left.
	b. Select ADMINISTER > Accelerators.
	c. In the list of accelerators, select the name of the accelerator (row).

2. Open the Apply Software Version wizard:

IBM Db2 Analytics Accelerator Studio	a. If necessary, expand the About section. b. In the About section, click Apply other software version.
IBM Data Server Manager	a. In the lower half of the window, scroll down to the Server section.b. Click Apply other software version.

- 3. In the **Apply Software Version** wizard, you can see all software packages currently available on the accelerator. 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.
- 4. Change the current version by selecting a different package.

1	Activate a version by selecting the appropriate check box in the File column.
IBM Data Server Manager	Select the appropriate entry (row) in the table.

5. Optionally, select Skip firmware and update only accelerator components (faster).

This limits the activation to the most essential components, such as the accelerator software, the database engine, the web console, and the DRDA gateway. Updates of less important components, such as the lower-level firmware and operating system components are skipped. This reduces the time

needed for the activation, during which the accelerator cannot be used. On average, the activation period is reduced to a time span between one and two hours.

- 6. Click Install.
- 7. In the **Apply Software Version** window, click **OK** to confirm.

Results

In some cases, migration processes require some time to finish, especially after a restart of the database container. It is therefore quite normal that, after you click the **OK** button, the accelerator is not instantly ready to process queries. However, after clicking the **OK** button, the **Installation Log** viewer opens, so that you can check the progress of the activation.



Figure 57. The Installation Log viewer

While the activation is in progress, the log viewer shows the status of all affected components in the upper part of the window. In the lower part, you see the last 20 lines of the apupgrade_logs/apupgrade.log file on the accelerator, which provides additional information.

Statuses are flagged as follows:

not_started

The activation has been scheduled, but not yet started.

started

The activation is currently in progress.

upgrade_skip_up_to_date

The activation of the update is unnecessary because the component is already up-to-date.

install complete

The activation has been completed.

postinstall complete

The activation has been completed and checks were applied successfully.

You can refresh the **Installation Log** viewer window to display the latest status changes. It also has an **all logs** check box. When selected, additional information is pulled in from other log files and displayed in the log viewer window.

If the selected updates can be activated successfully, the accelerator returns to the *online* state, so that queries can be processed.

You can also run the SYSPROC.ACCEL_CONTROL_ACCELERATOR stored procedure with the <getActivationLog> command to obtain the same information.

What to do next

To avoid a cluttered **Apply Software Version** window, you can remove packages from the accelerator that you no longer need.

For instructions on how to complete this task, follow the appropriate link under **Related tasks**.

Updating administration clients

Refer to the appropriate section to update your administration client:

Related tasks

Installing prerequisite PTFs for Db2 for z/OS

When you update your existing installation, first apply the prerequisite program temporary fixes (PTFs) for Db2 for z/OS.

Updating the IBM Db2 Analytics Accelerator stored procedures

Read how to update the Db2 for z/OS components of IBM Db2 Analytics Accelerator for z/OS.

Updating Db2 Analytics Accelerator on Z

Follow the steps in this section to upload a new image (appliance) for Db2 Analytics Accelerator on Z and activate it.

Related information

Updating IBM Db2 Analytics Accelerator software

These are software components that run on the accelerator hardware, that is, the IBM Integrated Analytics System. This chapter does not apply to updates for Db2 Analytics Accelerator on Z. Components on the IBM Integrated Analytics System are the IBM Db2 Analytics Accelerator, the SQL engine, and the operating system of the accelerator.

Updating IBM Db2 Analytics Accelerator Studio

Follow the instructions here to download and install an upgrade of IBM Db2 Analytics Accelerator Studio.

About this task

To download a refresh pack for IBM Db2 Analytics Accelerator Studio, a workstation with an internet connection is required.

If IBM Installation Manager is installed on a workstation without an internet connection, the program cannot automatically locate the repositories for updates of already installed products or additional plugins. Therefore, an administrator has to download the update repository to a central location that users without an internet connection have access to.

Workstation users without an internet connections can then let their local IBM Installation Manager installations point to the repository in the central location or copy the entire installation package from the central location to their workstations and go on from there.

You can use the newer versions of IBM Db2 Analytics Accelerator Studio with older versions of accelerators and stored procedures. However, the user interface might not show new functions until you update the stored procedures to the next higher or newest version.

Procedure

- To download a refresh pack for yourself or for others from a computer with an internet connection:
 - a) Click the following link:

Select fixes: IBM Db2 Analytics Accelerator Studio (All releases, All platforms)

- b) Select the proper refresh pack and click the **Continue** button above the list.
- c) Sign in with your IBM ID or click Create IBM id and follow the instructions to create an IBM ID first.
- d) Select your preferred download method:
 - Download using Download Director (requires Java enabled browser)
 - Download using bulk FTP
 - Download using your browser (HTTPS)
- e) Leave the check box Include prerequisites and co-requisite fixes (you can select the ones you need later) selected.
- f) In the View and accept terms window, click I agree to agree to the terms and conditions.
- g) Click **Download now** and follow the instructions.
- To update an IBM Db2 Analytics Accelerator Studio installation:
 - a) Start IBM Installation Manager.
 - b) Update IBM Db2 Analytics Accelerator Studio by using one of the following procedures:
 - The following approach works if you can still connect to a previously accessed IBM Installation
 Manager repository, and if that repository has been updated (contains the update you want):
 - a. Click Update.
 - b. Follow the instructions in the wizard. Deselect components that you do not want to update.
 - If the first method did not work, but you can access the downloaded refresh pack (see step <u>"7"</u> on page 224) on your company's network or on your local workstation:
 - a. From the main menu of IBM Installation Manager, select File > Preferences > Repositories.
 - b. Click Add Repository.
 - c. Click **Browse** and navigate to the directory that contains the refresh pack, either in the central location, or on your local workstation.
 - d. Select the refresh pack (compressed archive) and click Open.
 - e. Click **OK** twice to return the main window of IBM Installation Manager.
 - f. Click Install.
 - g. Follow the instructions in the installation wizard.

Updating IBM Data Server Manager

Follow the instructions here to download and install an upgrade of IBM Data Server Manager.

About this task

To download a refresh pack for IBM Data Server Manager, a workstation with an internet connection is required.

Procedure

- To download a refresh pack for yourself or for others from a computer with an internet connection:
 - a) Click the following link:
 - Select fixes: Information Management, Data Server Manager (All releases, All platforms)
 - b) Select the proper refresh pack.
 - **Important:** The page also lists refresh packs for IBM Db2 Data Management Console (datamgtconsole). This is not the same as IBM Data Server Manager. Make sure to select only packages that are labeled *Data-Server-Manager*.
 - c) Click Continue.

- d) Sign in with your IBM ID or click Create IBM id and follow the instructions to create an IBM ID first.
- e) Select your preferred download method:
 - Download using Download Director (requires Java-enabled browser)
 - Download using bulk FTPS/SFTP
 - Download using your browser (HTTPS)
- f) Leave the check box Include prerequisites and co-requisite fixes (you can select the ones you need later) selected.
- g) Click Continue.
- h) On the **Download files using ...** page, scroll down to the section that shows the name of the refresh pack in bold type.
- i) Click the blue file name link with the download icon in front to start the download.
- To update a IBM Data Server Manager installation:
 - a) If necessary, transfer the downloaded package to the machine on which you want to install the update.
 - b) Depending on your operating system, take one of the following actions:
 - For AIX and Linux:
 - a. Extract the content of the downloaded archive to an empty directory of your choice.
 - b. Open a terminal window.
 - c. Change to the directory that contains the extracted package content.
 - d. Run ./setup.sh.

If the web server of your previous installation is still running, the setup.sh program will tell you so and will also give you the process identification number (PID). You can then kill the process and restart the setup.sh program to continue.

 On Windows, start the update wizard by double-clicking the downloaded file. Follow the instructions in the wizard.

What to do next

To check if the upgrade works, open IBM Data Server Manager in a web browser and log on.

Updating the IBM Db2 Analytics Accelerator stored procedures

Read how to update the Db2 for z/OS components of IBM Db2 Analytics Accelerator for z/OS.

Procedure

 If you intend to accelerate static SQL queries, verify that the table SYSACCEL.SYSACCELERATEDPACKAGES and related indexes exist in the relevant Db2 subsystems. 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.

- 2. 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:
 - i) Reuse the existing WLM JCL procedure, that is, the procedure that was used to set up the WLM for the stored procedures of the previous version.
 - ii) Edit the JCL and change the high-level qualifier (HLQ) in the following statements so that the production HLQ for the data sets of this version is used:
 - DD <HLQACTIVEV7>.SAQTMOD

To run IBM Db2 Analytics Accelerator for z/OS Version 5.1.x alongside version 7.x.x, also specify:

DD <HLQACTIVEV5>.SAQTMOD

- AQTENV DD <HLQACTIVE>.SAQTSAMP(AQTENV)
- AQTDEF6 DD <hLQACTIVEV7>.SAQTSAMP(AQTDEF6)

To run IBM Db2 Analytics Accelerator for z/OS Version 5.1.x alongside version 7.x.x, also specify:

AQTDEFTR DD <HLQACTIVEV5>.SAQTSAMP(AQTDEFTR)

- b) Copy the customized settings from your old version 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 this version. Do not refresh your WLM application environment yet.
- c) Customize the AQTTIJSP job as described in "Customizing and running AQTTIJSP" on page 49.

 Unless instructed otherwise, use the INSTALL parameter for AQTTRIN in AQTTIJSP. Do not use REINSTALL because this parameter invalidates previous installations in the same Db2 subsystem. If a previous installation, such as IBM Db2 Analytics Accelerator for z/OS Version 5.1.0 exists, and REINSTALL is used, the previous version will have to be installed again.
 - Refresh the WLM application environment for the stored procedures. This ensures that all new stored procedure calls run with the updated code. You might want to verify this by starting your administration client. If the stored procedure code was updated successfully, the graphical user interface shows the features of the new version.
- d) You might want to free the older packages unless you want to run old and new versions in parallel for some time in a data sharing environment.
 - The names of older packages start with SYSACCEL. AQTDB or SYSACCEL. AQT03 and contain a version string like IDAA4.

Related tasks

Installing prerequisite PTFs for Db2 for z/OS

When you update your existing installation, first apply the prerequisite program temporary fixes (PTFs) for Db2 for z/OS.

Updating Db2 Analytics Accelerator on Z

Follow the steps in this section to upload a new image (appliance) for Db2 Analytics Accelerator on Z and activate it.

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.

Setting up a WLM application environment for IBM Db2 Analytics Accelerator

Follow the steps here to set up a suitable Workload Manager (WLM) application environment for the stored procedures of the product.

Setting up a WLM application environment for different product versions

To run IBM Db2 Analytics Accelerator for z/OS Version 7.x.x (Db2 Analytics Accelerator on Z or IBM Db2 Analytics Accelerator on IBM Integrated Analytics System) and an older version side-by-side, like IBM Db2 Analytics Accelerator for z/OS Version 5.1.0 PTF-2 or later, the Workload Manager (WLM) application environment must be set up to accommodate two sets of stored procedures.

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 stored procedures. Verify that this and a few other requirements are met by following the steps here.

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.

Verifying the installation of IBM Db2 Analytics Accelerator for z/OS stored procedures

An installation verification step is part of running the AQTTIJSP job. The following steps are relevant only if you encountered errors during the execution of AQTTIJSP. To verify the installation after the surfacing of AOTTIJSP errors, customize the AOTSJI00 job member (JCL) and run it.

Related reference

Setting access rights for the user who runs AQTTIJSP

The main job for the installation of the stored procedures is AQTTIJSP. You must submit this job in a later step. The ID of the user who runs this job needs SYSADM authority. Make sure that this is the case.

Related information

Updating IBM Db2 Analytics Accelerator software

These are software components that run on the accelerator hardware, that is, the IBM Integrated Analytics System. This chapter does not apply to updates for Db2 Analytics Accelerator on Z. Components on the IBM Integrated Analytics System are the IBM Db2 Analytics Accelerator, the SQL engine, and the operating system of the accelerator.

Updating administration clients

Refer to the appropriate section to update your administration client:

Updating Db2 Analytics Accelerator on Z

Follow the steps in this section to upload a new image (appliance) for Db2 Analytics Accelerator on Z and activate it.

About this task

Restriction: For the following steps, you need Mozilla Firefox or Google Chrome. Other browsers are not supported.

Procedure

- 1. Open the Appliance Installer in a web browser. Enter the IPv4 address of the Secure Service Container (SSC) LPAR (see Figure 12 on page 57).
- 2. On the **Login** page, log in with the **Master user ID** and **Master password** that you specified as you defined the SSC LPAR. (see Figure 12 on page 57)

You see the **Accelerator** page:

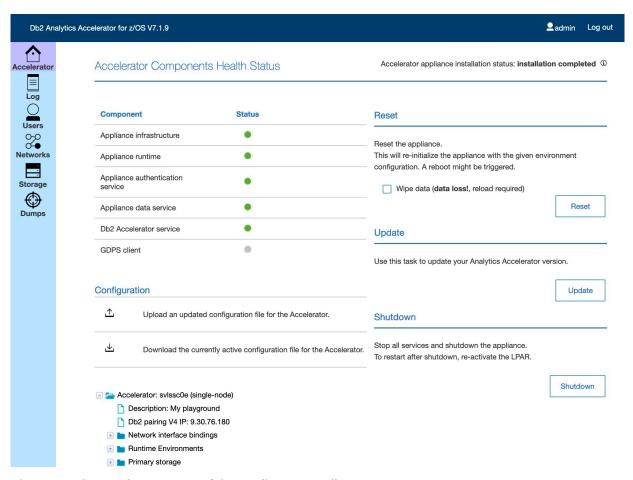


Figure 58. The Accelerator page of the Appliance Installer

3. In the lower right of the **Accelerator** page, click **Update**. You see the **Update** page:

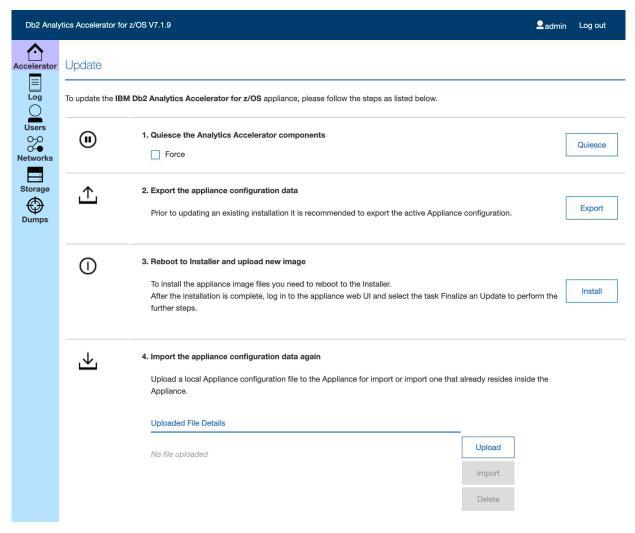


Figure 59. The Update page of the Appliance Installer

- 4. To make sure that all accelerator operations are paused during the update, click Quiesce.
- 5. By clicking **Export**, save your current configuration to a local drive, so that it can be reloaded to complete the update.
- 6. In section 3. Reboot to installer and update new image, click Install.
- 7. Select the image file and click **Install**.

You see a message saying that the appliance is rebooting.

8. Click Finalize Update

During this step, the previously exported configuration is imported and reloaded. The appliance reboots again.

Results

An information window with the heading **Invoking Secure Service Container Installer** is briefly displayed. After that, you see the following page:

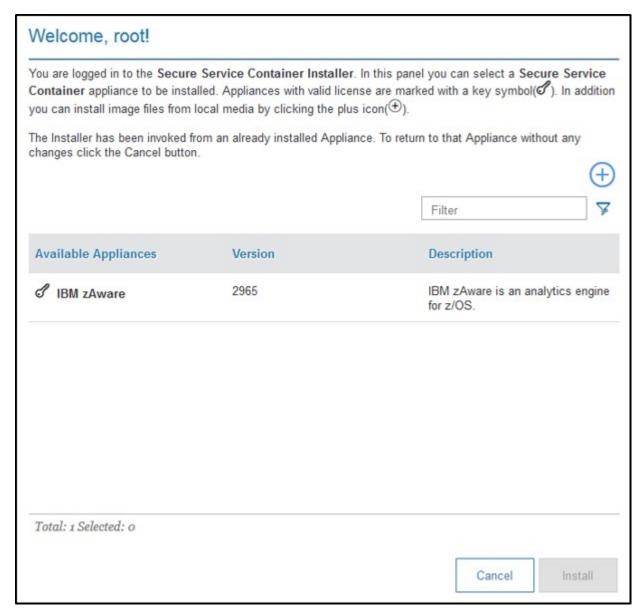


Figure 60. The Welcome page of the Secure Service Container Installer

Related tasks

Installing prerequisite PTFs for Db2 for z/OS

When you update your existing installation, first apply the prerequisite program temporary fixes (PTFs) for Db2 for z/OS.

Updating the IBM Db2 Analytics Accelerator stored procedures

Read how to update the Db2 for z/OS components of IBM Db2 Analytics Accelerator for z/OS.

Related information

Updating IBM Db2 Analytics Accelerator software

These are software components that run on the accelerator hardware, that is, the IBM Integrated Analytics System. This chapter does not apply to updates for Db2 Analytics Accelerator on Z. Components on the IBM Integrated Analytics System are the IBM Db2 Analytics Accelerator, the SQL engine, and the operating system of the accelerator.

Updating administration clients

Refer to the appropriate section to update your administration client:

Chapter 14. Removing accelerators from your Db2 for z/OS configuration

You might want to remove all accelerator-related objects and settings from a Db2 for z/OS configuration, for example after finishing a test or a proof-of-concept.

Procedure

- 1. Set DSNZPARM GET_ACCEL_ARCHIVE to NO or remove the parameter entirely.
- 2. Set DSNZPARM QUERY_ACCELERATION to NO or remove it.
- 3. Set DSNZPARM ACCEL to NO or remove it.
- 4. For each Db2 subsystem connected to an accelerator, run -STOP ACCEL
- Remove all tables from the connected accelerators.
 Use the **Remove** button in the relevant **Accelerator** views or table lists of your administration client, or run the SYSPROC.ACCEL_REMOVE_TABLES stored procedure for this purpose.
- 6. Remove all accelerators from the connected Db2 subsystems.
 To this end, use the **Remove** button in the Object List Editor of IBM Db2 Analytics Accelerator Studio or in the accelerator list of IBM Data Server Manager, or run the SYSPROC.ACCEL_REMOVE_ACCELERATOR stored procedure.
- 7. In Db2 for z/OS, drop the DSNACCEL database, including all tables and table spaces whose names start with SYSACCEL. This step requires a Db2 restart right after the drop.

Cleaning up the communications database manually

If accelerators have already been disconnected physically, tables and accelerators (steps 5 and 6 in the previous section) cannot be removed by IBM Db2 Analytics Accelerator functions anymore. In that case, clean up the Db2 communications database (CDB) manually by following the steps here.

Procedure

1. Run the following SQL query to get a list of the defined accelerators:

```
SELECT ACCELERATORNAME FROM SYSACCEL.SYSACCELERATORS
```

Delete entries in the CDB by running the following SQL statements for each defined accelerator.
 Replace <accelerator name> with the names returned by the SQL query in step "1" on page 231.

```
DELETE FROM SYSIBM.LOCATIONS WHERE LINKNAME = '<accelerator name>';
DELETE FROM SYSIBM.IPNAMES WHERE LINKNAME = '<accelerator name>';
DELETE FROM SYSIBM.USERNAMES WHERE LINKNAME = '<accelerator name>';
```

3. Drop the DSNACCEL database, including all tables and table spaces whose names start with SYSACCEL. This step requires a Db2 restart right after the drop.

Removing the remaining database objects

Remove the remaining accelerator-related database objects by running the following SQL statements:

Procedure

- 1. DROP TABLE DSNAQT.ACCEL_*;
- 2. DROP VIEW DSNAQT.ACCEL NAMES;
- 3. DROP SEQUENCE DSNAQT.UNLOADIDS;

```
4. DROP PROCEDURE SYSPROC.ACCEL_*;
5. DROP FUNCTION DSNAQT.ACCEL_*;
6. FREE PACKAGE SYSACCEL.*;
```

What to do next

Your contract with IBM might oblige you to physically erase all disks before an IBM Integrated Analytics System is moved outside of your data center. A procedure exists that guarantees and certifies the erasure of disks. If required, contact IBM support to request the execution of this procedure. If you want to reuse a machine after the execution, you must install and configure it from scratch.

Chapter 15. Sample TCP/IP configuration

In this section, you find a sample VTAM definition and a sample profile for the TCP/IP setup on your IBM Z. These samples work for both, IBM Db2 Analytics Accelerator on IBM Integrated Analytics System and Db2 Analytics Accelerator on Z.

Note: The following samples apply to a setup with OSA-Express cards.

Schematic overview

The following figures give a schematic overview of a working setup. A detailed description of this type of 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.

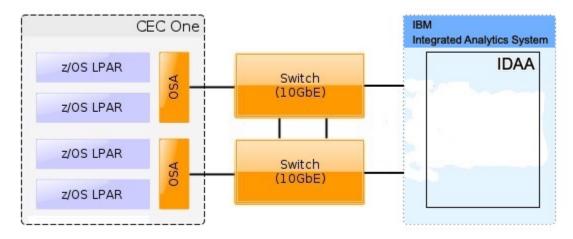


Figure 61. Recommended connections between IBM Z server and IBM Integrated Analytics System (including switches)

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 61 on page 233.

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 61 on page 233.

```
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 203.0.113.137/24
    MTU 8992
    INBPERF DYNAMIC PRIR
    VMAC ROUTEALL
    SOURCEVIPAINT VIPA1L
;Ten Gigabit Interface Definition for Portname GBE300
INTERFACE TENGBEE3 DEFINE IPAQENET
PORTNAME GBE300
    IPADDR 203.0.113.237/24
    MTU 8992
    INBPERF DYNAMIC PRIR
    VMAC ROUTEALL
    SOURCEVIPAINT VIPA1L
HOME
    203.0.113.37
                          VIPA1L
BEGINRoutes
; Direct Routes - Routes that are directly connected to my interfaces
 Destination Subnet Mask First Hop Link Name Packet Size;
ROUTE 203.0.113.0 255.255.255.0 = TENGBEE1 mtu 8992
ROUTE 203.0.113.0 255.255.255.0 = TENGBEE3 mtu 8992
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/0S.
; Otherwise, communication and tools (such as ssh) addressing the private network
; will fail because the requestor address cannot be reached.
SRCIP
  DESTINATION
                          203.0.113.0/24 203.0.113.37
ENDSRCIP
  START TENGBEE1 ; 10 GBE to Accelerator START TENGBEE3 ; 10 GBE to Accelerator
```

Related information

IBM DB2 Analytics Accelerator for z/OS: Installation Guide

Chapter 16. 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.

Important: During a customer session with IBM support, system parameters might be changed. For the parameter changes to take effect, a restart of an accelerator is often required. In such situations, the accelerator will be unavailable for operational tasks until it has come back online.

Query acceleration fails after schema change

After a schema modification of a Db2 for z/OS table, queries against this table are no longer accelerated.

Symptoms

You receive a failure message when you try to run an accelerated query against the affected table.

Causes

After the schema modification, the mechanism that handles such changes on the accelerator did not set the right timestamp in the catalog table SYSIBM.SYSACCELERATEDTABLES.

Resolving the problem

Run an UPDATE operation on the catalog table SYSIBM.SYSACCELERATEDTABLES, in which you replace the timestamp in the ALTEREDTS column with the timestamp reported in the failure message.

Firewall blocks operations because crucial ports are closed

If product functions are not executed, check your firewall settings. IBM Db2 Analytics Accelerator requires certain ports to be open. See the table in this topic.

Symptoms

- You cannot accelerate queries.
- You cannot run functions from the IBM Db2 Analytics Accelerator Console.
- Incremental updates do not work or IBM InfoSphere Change Data Capture for z/OS cannot be configured.
- You cannot open a secure-shell (ssh) connection to an accelerator.
- Services are unavailable or switches (routers) cannot be reached because ICPM (Internet Control Message Protocol) traffic is blocked.

Causes

Required networking ports are closed.

Resolving the problem

Table 7. Required networking ports for IBM Db2 Analytics Accelerator		
Function or protocol	Required port	Direction
Query execution (DRDA protocol)	1400	Both ways

Table 7. Required networking ports for IBM Db2 Analytics Accelerator (continued)			
Function or protocol	Required port	Direction	
Communication between accelerator and stored procedures	1401	Both ways	
IBM Db2 Analytics Accelerator Console	1600	Both ways	
Secure DRDA (stunnel)	11400	Both ways	
Secure communication between accelerator and stored procedures (stunnel)	11401	Both ways	
Incremental update communication	11301-11350	Outbound (from z/OS to accelerator)	
CDC configuration	5999	Inbound (generally from accelerator to z/OS; depends on the configuration of the CDC Capture Agent)	
DB2 DDF port (for continuous incremental updates)	446 (default)	Inbound (from accelerator to z/OS)	
Secure shell connection (ssh)	2222	Both ways	
ICMP traffic	N/A	Both ways	

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.

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.

Important: 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 CFEDUMP.

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, but a table space must be created before this. The following error message might be displayed:

IDC3221I 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 <u>"Installing libraries with IBM Db2 Analytics Accelerator support"</u> on page 25 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
 ** ASMAO17W Undefined keyword parameter; default to positional,
 including keyword - DSN65/ZOSMETRICS
** ASMA435I Record 3165 in Db291.ISAO.LAB.SDSNMACS(DSN6SPRM)
 on volume: TSMS08
          Avoid overflow
                                       DK153
               00000 00814 60+DSN6SPRM
 000000
 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 accelerator.

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 reveals this issue.

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 Integrated Analytics System or a Db2 Analytics Accelerator on Z was removed, but the accelerator was not removed before by running the *Remove accelerator* function from your administration

client 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 -STOP ACCEL command or the appropriate function in your administration client.
- 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

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.

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.

2. 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>

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. Your administration client returns a message window like this one:



Figure 62. SQL code -430 message window

Causes

This might be 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. If the problem persists, collect the following diagnostic information for IBM support:
 - · CEEDUMP in the WLM job
 - Stored procedure trace files in the z/FS directory /tmp. The names of these files follow this pattern:

```
aqt-trace-<YYYYMMDD-HHMM>--*
Example:
```

 $\verb|aqt-trace-20171113-1028-ACCEL_TEST_CONNECTION-ghbHxQ| \\$

Transfer the trace files in binary mode.

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_READFILE3.

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'FFFFFFB8' X'000000000' X'000000000' X'FFFFFFFF' X'000000000' X'000000000' 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_READFILE3)
- Checking available stored-procedure interface versions (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.*)
```

2. If the function is not in the state STARTED, start it with the following command:

```
-STA FUNCTION SPECIFIC (DSNAQT.*)
```

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

1. See "Adjusting WLM performance goals for SYSPROC.ACCEL_LOAD_TABLES" on page 46.

2. If you still see -471 errors, consider changing the value of the STORTIME ZPARM to NOLIMIT (the default is 180 seconds).

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.

Chapter 17. Opening a support case

You might run into a situation in which you have to contact IBM support because you cannot solve an issue by yourself. A reason might be that specific updates or patches are required to successfully complete maintenance or repair jobs on the accelerator. Such jobs are usually carried out by an IBM service engineer. Follow the steps here to open a support case 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 support case at https://www.ibm.com/support/servicerequest/Home.action. Consider that IBM support needs some time to analyze your system, coordinate maintenance actions, and update components as required.
- 4. Attach the trace file to the support case.

Transferring maintenance updates

A task that is closely related to support cases is the transfer of maintenance updates. When you have submitted the support case, 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 Integrated Analytics System.

Before you begin

Update packages must exist in the download directory for packages of this type in the z/OS UNIX file system (zFS). The target system must be able to access this directory.

Important: The AQT_HOST_PACKAGE_DIRECTORY environment variable points to this download directory. It specifies an absolute path (starting from the root directory). The AQT_HOST_PACKAGE_DIRECTORY environment is 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 updates. If it is not set correctly, the packages cannot be located and thus cannot be transferred to the IBM Integrated Analytics System. To avoid a setup with obsolete or wrong entries, use the sample AQTENV data set that comes with the latest version of IBM Db2 Analytics Accelerator for z/OS.

About this task

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

Transferring an update invokes the **SYSPROC.ACCEL_UPDATE_SOFTWARE2** 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 your administration client.

2. Select the accelerator:

IBM Db2 Analytics Accelerator Studio	a. In the Administration Explorer , select the Accelerators folder.
	b. In the Object List Editor , double-click an accelerator name to open the corresponding Accelerator view.
IBM Data Server Manager	a. Click in the upper left.
	b. Select ADMINISTER > Accelerators.
	c. In the list of accelerators, select the name of the accelerator (row).

3. Open the **Transfer Software Version** window:

IBM Db2 Analytics Accelerator Studio	a. If necessary, expand the About section. b. In the About section, click the Transfer updates link.
IBM Data Server Manager	In the lower half of the Accelerator view, click Transfer Updates .

4. In the **Transfer Software Version** or **Transfer Updates** window, you can see all software packages that are available in the zFS of your z/OS data server. Select the packages that you want to transfer:

IBM Db2 Analytics Accelerator Studio	Select the appropriate check boxes in the first column of the package list.
IBM Data Server Manager	Select the appropriate row in the package list. Restriction: In IBM Data Server Manager you can
	only select and transfer one package at a time. To transfer additional packages, you must revisit the Transfer Software Version window and repeat this and the following step.



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 Updates** window might be confusing, especially if it also contains older packages. To find the correct package numbers, see the closing information or the release notes for the latest program temporary fix (PTF). You find the closing information or the release notes 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 the list contains too many packages so that selecting the proper ones becomes awkward, select unwanted package names and click **Remove**. This will just remove the packages from the list or view, but won't delete the package files from the z/FS or your client workstation.
- 5. Choose the appropriate method to transfer the selected packages:

Transfer file from client

To transfer one or more downloaded packages from the connected client workstation.

Transfer file from IBM Z

To transfer one or more packages from the z/FS of your z/OS data server.

Note: If the verification message says ... package was deployed without integrity check ..., the reason might be that the checksum files were not transferred in *binary* mode. See the second **Important** note in step "1" on page 216.

Installing maintenance updates

While updates for IBM Db2 Analytics Accelerator for z/OS and the databaseare installed automatically at the end of the transfer process or by using the **Apply** function, other updates have to be installed manually with the help of IBM support.

Before you begin

Make sure that the following conditions apply:

- You have opened a support case as described in Chapter 17, "Opening a support case," on page 245.
- The IBM Integrated Analytics System must be fully operational. Hardware problems must not exist. Replace defective hard disks before installing the update.
- Remote access from z/OS OMVS to the accelerator via secure shell (ssh) must be enabled.
- IBM support will ask you to open or allow a screen-sharing session on your workstation. Make sure that this complies with the security guidelines of your organization or ask the responsible staff for a special permit.

About this task

An update installation requires you to stop the database container for approximately 30 minutes. The accelerator cannot be used during this time.

Procedure

1. Contact IBM support.

IBM support will ask you for remote screen-sharing access to the client machine.

2. With IBM assistance, open a login shell to the accelerator.

There are several ways to do that:

- If an ssh server and an ssh client are installed in the z/FS of your IBM Z, and if server and client have been activated, you can connect with a network client application, such as PuTTY, from your client workstation to the accelerator by opening the following connections:
 - a. From your workstation to the IBM Z server by ssh. Use your user ID and password to log in. See an example of an ssh login from a PuTTY window:

```
login as: willy willy@boedca1's password:
```

Important: If your z/OS version is V2R2 or higher, make sure that the address space for the Cryptographic Service Facility (CSF) has been started. Otherwise, the ssh connection to z/OS will fail.

b. From the IBM Z to the accelerator by ssh. See the following example:

```
ssh -p 2222 root@10.104.9.3
```

where:

2222

is the listening port on the accelerator

10.104.9.3

is the IP address of the accelerator

• If an ssh server is installed and active on your IBM Z, and if port-forwarding is enabled for that server, you can connect to the accelerator by an ssh tunnel, which looks like a single connection

from your client workstation to the accelerator. To enable port-forwarding, set AllowTcpForwarding yes in the /etc/ssh/sshd_config file of your zFS and restart the ssh server for the change to take effect.

Example: The commands for a tunnel connection would look similar to these:

a.ssh -L 10022:10.0.0.1:2222 user@systemz

where:

-L 10022

is the local port on the client workstation

10.0.0.1

is the IP address of the accelerator

2222

is the remote port or listening port on the accelerator

user@systemz

user is ID with which you log on to the UNIX System Services ssh server and systemz is the host name of that server.

b.ssh -p 10022 root@localhost

where:

-p 10022

is the local port on the client workstation

root@localhost

root is the user ID on the accelerator and localhost is the host name of the client workstation

Tip: IBM used to provide a tool package called IBM Ported Tools for z/OS. See:

https://www-03.ibm.com/systems/z/os/zos/features/unix/port_tools.html

The package does not exist anymore, but more specific packages with similar functionality are now offered by the Rocket software company (https://www.rocketsoftware.com/zos-open-source). These packages usually include an OpenSSH component, which contains an ssh client that you can use to connect from the UNIX System Services file system to the accelerator. Note that you only need the client component of OpenSSH for maintenance updates (called **ssh**).

The resulting login screen looks similar to this:

```
Using keyboard-interactive authentication.
Secure token for service password
(this token must be sent to IBM together with the serial# the first time you request a service password):

MHu2cK9FLC2pIkyAEb0oyDxK1yU7UQ01VPyXVviVxwGRtAC3gBiXiio5YLXzQTCQRfL/TRyuG8zFBC2CkSe3
GXPrFjM9KgB41gEFY0oxSEhjXY0S89fKGa7E/M+/VCdnsfHX2k2GpBq0XpyVPps6o0Q==
Enter Service Password (Date: '20170506' Serial#: '7130001' Rev: '4'):
```

- 3. As you enter the command to log on to the machine, the secure token, which contains the encrypted machine password, is displayed on the screen. In the previous example, this is the cryptic string that starts with MHu2cK. For the first service session, submit this token and the machine serial number to IBM support. IBM support will then generate a service password.
 - You have to submit the secure token only for the first service session. The token is stored in an IBM service database for subsequent requests. This database cannot be accessed from the internet. If a further service session is required, you just have to give IBM support the serial number of the machine.
- 4. IBM support will enter the service password for you. The password will be valid exclusively for your accelerator machine and for the current date only.
- 5. Navigate to the directory on the accelerator in which the transferred update files are stored.

- 6. To proceed, follow the instructions of IBM support.
- 7. Having installed the updates, delete the files that are not needed anymore.
- 8. Terminate the ssh session by entering exit at the command prompt.

Appendix A. Members of SAQTSAMP

The SAQTSAMP data set contains various samples for installing, configuring or running the product's stored procedures. The following table lists all these members and provides brief descriptions of their functions.

Table 8. Members of the SAQTSAMP data set		
Member name	Function	
AQTDEF6	Default trace specification and other settings for the message input parameter. The default input is used when stored procedures are called with a NULL value or empty string in the message input parameter. The sample uses a verbosity level of INFO for all stored procedure calls and specifies that the traces are to be kept in a specific directory if the procedure ends with an error.	
	To use this member, you must add an appropriate DD statement (DD AQTDEF6) to the JCL that starts the Workload Manager (WLM) environment for IBM Db2 Analytics Accelerator stored procedures.	
AQTENV	Default environment variable settings for the stored procedures.	
AQTLCABA	Sample C program that implements the rerouting and load-callback functions for the high-availability load sample.	
AQTLMAIN	Sample C program that implements the main functionality for the high-availability load sample.	
AQTLSHMC	Sample C program that implements inter-process communication (IPC) functionality for the high-availability load sample.	
AQTLSHMH	Header files for the inter-process communication (IPC) functionality for the high-availability load sample.	
AQTSCALL	Sample application program in the C programming language with embedded SQL CALL statements for calling IBM Db2 Analytics Accelerator 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)	
	3. 3131 NOC.D3NOTILO(ONLOAD)	

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.	
	2. SYSPROC.ACCEL_TEST_CONNECTION	
	3. SYSPROC.ACCEL_ADD_TABLES	
	4. SYSPROC.ACCEL_GET_TABLES_INFO	
	5. SYSPROC.ACCEL_LOAD_TABLES	
	To enforce a reload of the accelerator-shadow tables and enable these for quer acceleration:	
	1. SYSPROC.ACCEL_LOAD_TABLES	
	2. SYSPROC.ACCEL_SET_TABLES_ACCELERATION(ON)	
	To execute the query:	
	1. START -ACCEL	
	2. Query execution	
	3. SYSPROC.ACCEL_GET_QUERIES	
	4. STOP -ACCEL	
	To check various other functions:	
	1. SYSPROC.ACCEL_REMOVE_TALBES	
	SYSPROC.ACCEL_CONTROL_ACCELERATOR (including the various subfunctions)	
	3. 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.	
AQTSFEDR	Sample JCL to create a RACF profile for federated access. The profile provides users with appropriate RACF permissions to run the following stored procedure for federated access:	
	ACCEL_GRANT_TABLES_REFERENCE ACCEL_BEVOYE_TABLES_REFERENCE	
	ACCEL_REVOKE_TABLES_REFERENCE	
AQTSJI00	JCL that collects information about required IBM Db2 Analytics Accelerator databases and tables.	
AQTSJI01	JCL that calls the Db2 command line processor for verifying IBM Db2 Analytics Accelerator stored procedures without an accelerator.	

Member name	f the SAQTSAMP data set (continued) Function	
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 stored procedures. For a successful completion of the job, a running accelerator is required.	
AQTSJI04	JCL for a manual recovery of a partition that was archived with the High Performance Storage Saver (HPSS). With the help of this JCL, you can recover an archived 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-intime, 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 archived partitions to Db2 for z/OS, and delete the corresponding accelerator-archive table when the accelerator is back online.	
AQTSJL01	JCL to compile, link, and bind the high-availability load sample programs.	
AQTSJL02	JCL to compile, link, and bind the high-availability load sample programs.	
AQTSJL03	JCL that starts the high-availability load sample program.	
AQTSSCHK	UNIX System Services shell script that checks the output of Db2 command line processor scripts containing IBM Db2 Analytics Accelerator 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 stored procedures. Divide this return code by 256 to determine the severity of an error:	
	All stored procedures that were called returned information messages Warning (return code 1024) Error (return code 2048) Severe (return code 3072)	
AQTSSCPY	UNIX System Services shell script that copies SAQTSAMP members to the /tmp/ivp directory in the z/OS UNIX file system.	
	Note: Db2 command line processor scripts cannot be run from a data set.	
AQTSSLSA	Sample JCL that creates a CA certificate for IBM Integrated Synchronization, stores the certificate in a key ring and exports it, so that it can be imported into a Java key store. In addition, appropriate RACF permissions are given to the user who runs IBM Integrated Synchronization.	
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.	

Table 8. Members of the SAQTSAMP data set (continued)		
Member name	Function	
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 stored procedure without executing its actual function. The XML code is used as the value of the MESSAGE input parameter.	
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 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 stored procedures.	

Appendix B. Environment variables

The job control language (JCL) for the configuration of the Workload Manager (WLM) environment for IBM Db2 Analytics Accelerator 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.
- 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.

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 archives partition data from Db2 for z/OS on 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

&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.

&UNIO.

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 archives partition data from Db2 for z/OS on 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 archives partition data from Db2 for z/OS on 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 archives partition data from Db2 for z/OS on an accelerator. For more details, see AQT_ARCHIVE_COPY1.

AQT_HOST_PACKAGE_DIRECTORY

Update packages can be downloaded to and transferred from a directory in the UNIX file system (z/FS) of your z/OS data server or from a workstation. If a z/FS directory is used, you must set AQT_HOST_PACKAGE_DIRECTORY to an absolute path to point to the directory containing the update packages because your administration client and the SYSPROC.ACCEL_UPDATE2 stored procedure look for update packages in this directory. Most update packages must be downloaded from IBM Shopz or from IBM Fix Central. In the following steps, they are transferred to the accelerator, from where they are deployed and activated. There might be packages that you cannot install by yourself, but require the help of IBM support. For these types of updates, you must open a support case first. For more information, follow the link to *Opening a support case* at the end of this topic.

AQT DO SYSLOGGING

Causes messages about the start and the completion of load operations to be written to the SYSLOG of the z/OS LPAR in which the SYSPROC.ACCEL_LOAD_TABLES stored procedure runs. Messages contain the name of the accelerator, the task number of the load jobs, and information about the reduction or increase of parallel processing threads (up to the limit set by AQT_MAX_UNLOAD_IN_PARALLEL). The variable is set by default (AQT_DO_SYSLOGGING=SET). To suppress these messages, comment out this line.

AQT_KEEPIDLE_INTERVAL

Specifies an interval that determines for how long a TCP/IP connection can be idle before keep-alive-packages are sent to the connected accelerator. The stored procedures build TCP/IP connections to the accelerator. To avoid connection failures caused by implicit connection closures, routers, or VPN (virtual private network) tunnels, TCP sends keep-alive-packages or probes after a period of time during which the network was idle. This is necessary to keep the connections up during long running load operations.

This setting affects all stored procedures and overwrites the KEEPALIVEOPTIONS settings for the TCP/IP stack that is specified in the Worload Manager (WLM) environment of the IBM Db2 Analytics Accelerator stored procedures. By default, the interval is set to 60 seconds. If the setting should turn out to be the source of problems, change the value or disable the mechanism entirely by commenting out the corresponding line in the AQTENV data set.

Note that socket connections between the stored procedures and the accelerator are also configured to send keep-alive-packages, and that the NO_DELAY and NON_BLOCKING options are set for the sending of these packages.

AQT_MAX_RETRIES_DSNUTILU

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.

AQT MAX UNLOAD IN PARALLEL

The maximum number of parallel DSNUTILU invocations used by the SYSPROC.ACCEL_LOAD_TABLES stored procedure when loading data from a partitioned Db2 table. Increasing the value leads to a better performance, provided that enough processors are available to handle additional parallel processes. The default value is 4.

AQT_QUOTEFIX1026

Converts Turkish characters that would otherwise cause errors during table load operations. This variable is set by default (AQT_QUOTEFIX1026=SET).

AQT_RTS_EXTERNALIZATION_AUTH_FAILURE

Determines the behavior if the Db2 command -ACCESS DATABASE (ACCESS DB) is called by a stored procedure, but cannot be invoked because the user who runs the stored procedure lacks the proper authorization. Possible values:

ERROR (default)

Processing stops and an error message is displayed.

WARNING

Processing continues, but a warning is displayed at the end of the process. There will be just one generic warning.

IGNORE

Processing continues without a warning or error message.

AQT_SECONDS_BEFORE_RETRY_DSNUTILU

Specifies the interval in seconds between DSNUTILU calls (retries). By default, this interval is set to 60 seconds.

AQT_SKIP_UNLOAD_EMPTY_PARTS

Suppresses invocations of the Db2 UNLOAD utility (DSNUTILU) when SYSPROC.ACCEL_LOAD_TABLES was called to load an accelerator-shadow table or partition, but the Db2 source table or partition is empty. This improves the performance of the load process. The detection of empty tables relies on Db2 real-time statistics.

Important: If the <externalTool> element is used as part of the *table_load_specification* parameter for SYSPROC.ACCEL_LOAD_TABLES, the Db2 UNLOAD utility is always called, irrespective of the setting of this variable.

AQT_SKIP_VALIDATE_XML

Skips the validation of XML code when it is returned in result sets or output parameter strings. The omission of this step leads to a better performance, especially when large result sets are returned. To disable this setting (turn XML validation on), comment out the line in the AQTENV data set.

AQT_SORTDEVT

Specifies the device type to use when sorting temporary data sets dynamically. This setting is used by the sort program when an IBM Db2 Analytics Accelerator stored procedure sends a request to a Db2 Utility that includes a sort job like this.

AQT_UTILITY_TMP

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.

where

&USERID.

ID of the user who runs the stored procedure

&UNIQ.

Causes the creation of a unique identifier

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.

ΤZ

Sets the correct timezone for stored procedures running in the WLM address space that has been configured for IBM Db2 Analytics Accelerator. The setting affects the formatting of timestamps in trace files.

Related information

TCP_KeepAlive socket option

IBM DB2 Analytics Accelerator for z/OS: Stored Procedures Reference
z/OS V1R12.0 XL C/C++ Programming Guide (SC09-4765-11)

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

Appendix C. 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:
 - a) Go to the Administration Explorer.
 - b) Select the Accelerators folder.
 - c) In the Object List Editor on the right, select the accelerator.
 - d) 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 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 query acceleration. Operations like a table load job can still be run. To re-enable an

accelerator, click the button on the toolbar.

Appendix D. 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:
 - a) Go to the Administration Explorer.
 - b) Select the Accelerators folder.
 - c) In the Object List Editor on the right, select the accelerator.
 - d) Click Start on top of the Object List Editor.

The status of the accelerator changes from Stopped to Online.

Notices

This information was developed for products and services offered in the US. This material might be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

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, MD-NC119 Armonk, NY 10504-1785 USA

For license inquiries regarding double-byte character set (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

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 jurisdictions 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 provide 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 Director of Licensing IBM Corporation North Castle Drive, MD-NC119 Armonk, NY 10504-1785 USA 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.

The performance data and client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions.

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.

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 is for planning purposes only. The information herein is subject to change before the products described become available.

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 actual people or business enterprises 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. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample 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_.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corp., 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 www.ibm.com/legal/copytrade.shtml.

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.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Terms and conditions for product documentation

Permissions for the use of these publications are granted subject to the following terms and conditions.

Applicability

These terms and conditions are in addition to any terms of use for the IBM website.

Personal use

You may reproduce these publications for your personal, noncommercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative work of these publications, or any portion thereof, without the express consent of IBM.

Commercial use

You may reproduce, distribute and display these publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these publications, or reproduce, distribute or display these publications or any portion thereof outside your enterprise, without the express consent of IBM.

Rights

Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the publications or any information, data, software or other intellectual property contained therein.

IBM reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the publications is detrimental to its interest or, as determined by IBM, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

IBM MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.

GDPR considerations

For PID(s): 5697-DA7, 5697-DA5, 5697-DAB

Notice: This document is intended to help you in your preparations for GDPR readiness. It provides information about features of IBM Db2 Analytics Accelerator for z/OS that you can configure, and aspects of the product's use, that you should consider to help your organization with GDPR readiness. This information is not an exhaustive list, due to the many ways that clients can choose and configure features, and the large variety of ways that the product can be used in itself and with third-party applications and systems.

Clients are responsible for ensuring their own compliance with various laws and regulations, including the European Union General Data Protection Regulation. Clients are solely responsible for obtaining advice of competent legal counsel as to the identification and interpretation of any relevant laws and regulations that may affect the clients' business and any actions the clients may need to take to comply with such laws and regulations.

The products, services, and other capabilities described herein are not suitable for all client situations and may have restricted availability. IBM does not provide legal, accounting, or auditing advice or represent or warrant that its services or products will ensure that clients are in compliance with any law or regulation.

- 1. GDPR Overview
- 2. "Product Configuration for GDPR" on page 268
- 3. "Data Life Cycle" on page 269
- 4. "Data Collection" on page 269
- 5. "Data Storage" on page 270
- 6. "Data Access" on page 270
- 7. "Data Processing" on page 270
- 8. "Data Deletion" on page 271
- 9. "Data Monitoring" on page 271
- 10. "Responding to Data Subject Rights" on page 271

GDPR

General Data Protection Regulation (GDPR) has been adopted by the European Union ("EU") and applies from May 25, 2018.

Why is GDPR important?

GDPR establishes a stronger data protection regulatory framework for processing of personal data of individuals. GDPR brings:

- New and enhanced rights for individuals
- · Widened definition of personal data
- · New obligations for processors
- Potential for significant financial penalties for non-compliance
- · Compulsory data breach notification

Read more about GDPR

- EU GDPR Information Portal (https://www.eugdpr.org/)
- ibm.com[®]/GDPR website (https://ibm.com/GDPR)

The EU General Data Protection Regulation (GDPR) regulates data privacy in the European Union (EU). For details, see: https://www.eugdpr.org

The regulations laid out in the GDPR apply if the data controller (an organization that collects data from EU residents), or the processor (an organization that processes data on behalf of a data controller, such as a cloud service provider), or the data subject (an individual) is based in the EU (source: Wikipedia. See: https://en.wikipedia.org/wiki/General_Data_Protection_Regulation). A violation of GDPR regulations might incur severe financial penalties (fines) and reputation damages. This documentation is intended for data controllers or processors who use Db2 for z/OS together with IBM Db2 Analytics Accelerator for z/OS with the aim of storing personal data of EU residents (individuals). As regards IBM Db2 Analytics Accelerator for z/OS, the IBM company is neither the data controller, nor the data processor. Nevertheless, the IBM company feels obliged to give such entities guidance on how to follow GDPR regulations.

Your responsibilities as a data controller or processor

If you act on behalf of a data controller or data processor, see the frequently asked questions (FAQ) on the official GDPR website to find out what your responsibilities are:

https://www.eugdpr.org/gdpr-faqs.html

Furthermore, refer to chapter 4 in the GDPR regulations:

https://www.eugdpr.org/article-summaries.html

Product Configuration for GDPR

The following sections provide considerations for configuring IBM Db2 Analytics Accelerator for z/OS to help your organization with GDPR readiness.

How to configure the offering so that it can be used in a GDPR environment

The following sections describe how personal data is stored and processed by IBM Db2 Analytics Accelerator for z/OS (called 'accelerator' in the text that follows).

If personal data is to be stored on an accelerator, the following configuration steps might help you meet your compliance objectives with regard to GDPR:

- Activate the encryption of data in motion (data that is transferred over a data network between Db2 for z/OS and the accelerator).
- Activate the encryption of data at rest (encryption of disks and other storage devices).
- Do not share passwords for the operation of the IBM Db2 Analytics Accelerator Console between multiple persons. Instead, create separate user IDs with different passwords for each console user.
- Give read access to personal data (the privilege to run queries) only to authorized Db2 for z/OS users. To this end, define appropriate Db2 authorizations (SELECT, and so on) in the Db2 for z/OS catalog.
- Restrict the right to unload data from and to create image copies in Db2 for z/OS. Give this right to authorized administrators only.
- Restrict access to the 'Save Trace' function in Db2 Analytics Accelerator for z/OS because the collected trace data might contain personal data. Access should be given to authorized administrators only.
- Restrict access to the table SYSIBM.USERNAMES in the catalog of the Db2 for z/OS communication database because this table contains authentication credentials. If these credentials are read or copied by unauthorized users, these users might obtain access to sensitive personal data.

Data Life Cycle

You can create copies of any Db2 for z/OS table on an accelerator (accelerator-shadow tables). As soon as you create a copy of a Db2 for z/OS table and load it with personal data, the GDPR regulations apply in the same way as they apply to personal data in the original Db2 for z/OS tables.

An accelerator also allows you to load and store external data outside of Db2 for z/OS, such as the data of VSAM files, in accelerator-shadow tables. Products like the Db2 Analytics Accelerator Loader can be used to this end. If the external data contains personal information, the GDPR regulations apply to the data that is trasferred to the accelerator.

What types of data?

The administrative environment of an accelerator (client and configuration software) does not store personal data, except for the user IDs and authentication tokens of its administrators. However, the GDPR regulations apply to all personal data that you store in database tables on the accelerator.

Where in the process?

An accelerator is loaded with data in the course of individual load operations or by the incremental update function (automated load of table updates). Adding a table to an accelerator does not yet copy any data. It merely defines the table's structure (metadata) on the accelerator. However, if you load the table or enable it for incremental updates, the data of the source table - which might contain personal data - is copied to the accelerator.

For what purpose?

An accelerator, like Db2 for z/OS, is a general database management system for SQL queries. Such queries are run with the aim of retrieving stored information and gaining further insights (analytical queries that uncover the dependencies or connections between sets of values). An accelerator can return the results of SQL queries to authorized and authenticated users of the system. The results usually consist of subsets or aggregations of numerous data records. All these records might include personal data.

Personal data used for online contact with IBM

IBM Db2 Analytics Accelerator for z/OS clients can submit online comments/feedback/requests to contact IBM about accelerator subjects in a variety of ways, primarily:

- Public comments area on pages in the IBM Integration community on IBM developerWorks[®]
- Public comments area on pages of the IBM Db2 Analytics Accelerator for z/OS product documentation in the IBM Knowledge Center
- Feedback forms in the IBM Integration community

Typically, only the client name and email address are used, to enable personal replies for the subject of the contact, and the use of personal data conforms to the *IBM Online Privacy Statement* (https://www.ibm.com/privacy/cc/).

Data Collection

IBM Db2 Analytics Accelerator for z/OS can be used to collect personal data. When assessing your use of IBM Db2 Analytics Accelerator for z/OS and the requirements of GDPR, you should consider the types of personal data which in your circumstances are stored on an accelerator.

The use of an accelerator does not impose new requirements for the collection of personal data. It does not introduce new users or new privileges, but makes use of existing users and their privileges as defined in Db2 for z/OS. The use of an accelerator does not introduce additional obligations for the safekeeping or adminstration of data if you already observe the GDPR regulations for the storage of data in Db2 for z/OS.

Make sure that the process of collecting data and storing it in Db2 for z/OS is GDPR-compliant. The use of an accelerator will then be equally compliant.

Data Storage

The use of IBM Db2 Analytics Accelerator for z/OS involves the storage of data in a persistent manner.

An accelerator is a general-purpose data processing and retrieval engine. It is not aware of the meaning of data or any meaning that can be inferred from it. It is thus the responsibility of the data controller or processor to classify the data as personal data - for example account data - and make sure that the collection, access and processing of this data is in accordance with the GDPR regulations.

Storage in backups:

An accelerator has no backup function. However, image copies are created in Db2 for z/OS when the accelerator's archiving function is used (high performance storage saver). If these image copies contain personal data, the GDPR rules apply. This means, for example, that if individuals request their data to be deleted, it must also be deleted from all the image copies that may exist.

Storage in archives:

The high performance storage saver deletes table data from Db2 for z/OS and stores it exclusively on an accelerator. If this data contains personal information, the GDPR regulations might require that you delete the personal information from a number of records. The current version of the accelerator software does not allow DELETE operations on archived partitions. So in case you must delete records, first restore the archived partition to Db2 for z/OS, then delete then delete the personal data from Db2 for z/OS. Finally, archive the partition again (which now does not include the personal data anymore).

Data Access

Access to data on an accelerator requires a Db2 for z/OS user ID and a password. The accelerator recognizes the privileges defined in Db2 for z/OS and does not allow a user to run functions that she or he is not entitled to use.

As long as the Db2 for z/OS authorizations, such as SELECT, INSERT, UPDATE, or DELETE privileges, are defined in accordance with the GDPR regulations, access to data on the accelerator will likewise be in accordance with these regulations.

So it is crucial that the definition of roles and access rights in Db2 for z/OS effectuates an adequate separation of duties and responsibilities.

One accelerator-specific function might provide access to personal data on the accelerator: This is the "Save Trace" function, which collects diagnostic information for troubleshooting. The resulting trace files might contain personal data because they sometimes include the code of SQL queries to which personal information has been passed in the form of query literals. Trace files can also include memory dumps that contain personal data.

It is therefore important that access privileges to data, such as SELECT privileges or superuser privileges like SYSADM and DBADM, are granted with great care. The permission to use the "Save Trace" function should be given to selected administrators only.

The privilege to create image copies in Db2 for z/OS must also be handled with care. It is required, for example, if a user wants to archive table partitions on an accelerator.

Data Processing

Privileges defined in the Db2 for z/OS catalog control data processing on an accelerator.

In addition to that, the encryption features (encryption of data in motion and encryption of data at rest) should be activated. This prevents an unauthorized person without proper privileges from gaining direct access to personal data by interferring with the network communication or by obtaining physical access to the storage devices (hard disks or solid state disks).

Encryption features and safeguarding precautions:

Encryption of data in motion:

This feature encrypts the network traffic between an IBM Z system and an accelerator.

Encryption of data at rest:

This feature encrypts the data on an accelerator's physical storage devices (hard disks and solid state disks). Depending on the model of the hardware, the feature might already be turned on by default.

Encryption key ownership:

Restrict access to the encryption keys because the keys enable a person to read or decrypt personal data on a network or storage device. Put the safekeeping and maintenance (rotation) of the encryption keys in the hands of a trusted administrator.

Data Deletion

When data is deleted from Db2 for z/OS, it might still exist in table copies on the accelerator. To delete the data on an accelerator, you have the following options:

The data in accelerator-shadow tables is deleted:

- If the data in the original Db2 for z/OS tables has been deleted, and if the affected accelerator-shadow tables have been reloaded (with the contents of the now empty Db2 tables).
- If the accelerator-shadow tables are removed from the accelerator.
- If the incremental update function is used and the function replicates empty table content to the accelerator because the original Db2 for z/OS table data has been deleted.

The data in an accelerator-only table (AOT) is deleted:

- If the AOT is removed from the accelerator.
- If an explicit DELETE statement is submitted on the Db2 for z/OS side, which references the AOT and contains a predicate that specifies the row to be deleted.

The data in partitions archived by the high-performance storage saver (HPSS) is deleted:

- If the table partitions are removed from the accelerator.
- If an archived partition is restored to Db2 for z/OS.
- If the data is deleted from Db2 for z/OS because this means that the archived partition is reloaded or that the partition is archived again after the deletion from Db2 for z/OS (empty table content will be reloaded or archived).

Data Monitoring

Db2 for z/OS provides powerful auditing and monitoring functions that allow you to track the processing of personal data. Because IBM Db2 Analytics Accelerator for z/OS is a logical component of Db2 for z/OS, and keeps only copies of Db2 for z/OS data, the monitoring functions for data processing in Db2 for z/OS are sufficient.

Responding to Data Subject Rights

To identify an individual's data that is stored on an accelerator, you can submit SQL queries from Db2 for z/OS. These queries ought to search for the name, the account number, or for other identifying information about the subject using SQL syntax. Since an accelerator works on copies of Db2 for z/OS data, the same queries can be run to retrieve the data either from Db2 for z/OS or from the accelerator. Specific advice or examples of such queries can not be provided here because the structure and the content of the database tables, as defined by the data controller or data processor, is not known.

Does the offering facilitate being able to meet data subject rights?

Right to Access

Can the client provide individuals access to their data?

Db2 for z/OS and accelerators can process SQL queries in which an individual's ID can be specified as a search criterion. This way, all data stored about that individual can be retrieved. In Db2 for z/OS, a

client's data access can be restricted, so that the client can only search and retrieve a subset of the available data records, that is, those records that belong to a specific group of individuals. To this end, a feature called label based access control (LBAC) can be employed. IBM Db2 Analytics Accelerator provides a similar functionality, which restricts access to a selected set of rows based on the CURRENT SQLID special register.

Can the client provide individuals information about what data the client has about the individual? Yes. To identify an individual's data stored on an accelerator, the client can submit SQL queries from Db2 for z/OS. These queries ought to search for the name, the account number, or for other identifying information about the subject using SQL syntax.

Right to Modify

Can the client allow an individual to modify or correct their data?

Db2 for z/OS allows a client to update the data that is stored about an individual. Depending on the architecture of the application, it might even be possible to authorize the individuals, so that these can update their data by themselves. After an update in Db2 for z/OS, the modified data is copied to the accelerators attached to the database during the next reload operation, which can be started manually or automatically (by the incremental update function). During that process, incorrect data will be overwritten.

Can the client correct an individual's data for them?

Yes (see the previous question Can the client allow an individual to modify or correct their data?).

Right to Restrict Processing

Can the client stop processing an individual's data?

Yes. In Db2 for z/OS, a client can submit DELETE statements to remove data records of individuals. The removal will be reflected in the data that is stored on an attached accelerator when the next reload operation (manual or automatic) is taking place. During that process, records previously deleted from Db2 for z/OS will also be deleted from the accelerator. Furthermore, an individual's data can be excluded from SQL queries through the use of filters in query predicates. For example, a query that reads SELECT CUSTOMER_DATA FROM CUSTOMER_TABLE where custid NOT IN $(a,b,\ldots c)$ returns only data records that are not related to (do not belong to) the users in the specified group $(a,b,\ldots c)$.

Right to Object

Same as "Right to Restrict Processing" on page 272.

Right to Be Forgotten

Can the client delete an individual's data?

Yes.

Right to Data Portability

Can the customer provide an individual with the information that they have about the individual in a user-friendly/machine readable format?

Yes.

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.

A

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 by which Db2 for z/OS communicates with another relational database management system.

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.

P

port-forwarding

A networking mechanism that allows Secure Shell access to a host in a private network from the outside.

S

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.

Т

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.

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).



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.

W

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	AQTTIJSP (continued) running 49
accelerator	AQTTRIN 49
adding to configuration 132	architecture 1
defining 132	assigning accelerator resources to 201
maintenance 245	auditing 13
remove from Db2 for z/OS 231	authentication 132
repair 245	authorizations 123, 124
shared 201	automatic
accelerator-only tables 31	notification of DSNX881I events 205
accelerators	support cases 205
disabling 259	automatic update of table data 143, 176
enabling 261	<u> </u>
multiple 4	_
multiple database subsystems 4	В
shared 4	hann mundunt = /00 0F
updating software 216	base product, z/OS 25
access	bi-temporal tables 212
rights	binding Db2 packages 125
AQTTIJSP 48	business-temporal tables <u>212</u>
access rights	
for Db2 for z/OS stored procedures 123, 124	C
for IBM Db2 Analytics Accelerator for z/OS stored	
procedures 123, 124	call home
minimum 124	activating <u>209</u>
power user 123	configuration
RACF 123	accelerator console 207
access, Db2 subsystem 126	configuring
activation of software updates 219	HMC 205
administration client	disabling 210
	displaying settings 210
IBM Data Server Manager 19 IBM Db3 Applytics Accelerator Studio 19	testing connectivity 211
IBM Db2 Analytics Accelerator Studio 19	call-home function 205
updating 223	CDC
allocation, processing resources 201	configuring 36, 180
appliance with GDPS support, Db2 Analytics Accelerator on	installing 25
Z	starting as z/OS started task 187
installing	user 180
single-node 95	certificate management, IBM Integrated Synchronization
appliance, Db2 Analytics Accelerator on Z	displaying certificates 168
dump <u>111</u> , <u>118</u>	installing certificates 170
GDPS setup, single-node 95	removing certificates 172
installing	transferred certificate 174
single-node <u>61</u>	certificate, Db2 Analytics Accelerator on Z 118
logon <u>58</u>	checking Db2 Analytics Accelerator on Z
multi-node setup $\frac{76}{60}$	connection to Db2 115
single-node setup 61	network connections 115
SSL certificate 118	system 115
trace file 111	classification of workloads
update configuration 109	DDF 47
updating <u>227</u>	DSNUTILU 46
application environment	SYSPROC.ACCEL_LOAD_TABLES 46
for different product versions 40	z/OS 47
IBM Db2 Analytics Accelerator 37	configuration
AQTENV 255	update, Db2 Analytics Accelerator on Z 109
AQTSJI00 50	connection
AQTTIJSP	profile 126
access rights <u>48</u>	testing 131

continuous incremental updates <u>179</u> credentials <u>132</u> , <u>199</u> , <u>200</u> customization AQTTIJSP <u>49</u> AQTTRIN <u>49</u>	DSNTIJAS <u>26</u> , <u>49</u> DSNTIJUZ <u>31</u> DSNUTILU performance goals <u>46</u> DSNX881I <u>205</u> dump, Db2 Analytics Accelerator on Z <u>111</u>
D	E
damaged disk 13 data loss 176 data privacy 267 data sharing groups 4, 7 database installing updates 216 update packages 217 database connection profile 126 Db2 Analytics Accelerator on Z add users 112 certificate, replacing 118 checking connection to Db2 115 network connections 115 system 115 checking the Db2 network 115 dump 111 GDPS setup, single-node 95 GDPS support 92 LPAR 55 multi-node setup restart 108 shutdown 108 remove users 112 reset user password 112 single-node setup 55, 61 SSL certificate, replacing 118 trace file 111 update 227 update configuration 109 Db2 Analytics Accelerator on ZLPAR, GDPS 93 Db2 authorizations for stored procedures 123, 124 Db2 for z/OS access from administration client 125	EBCDIC GRAPHIC 26, 31 MIXED 26, 31 enabling accelerators 261 existing Db2 subsystem 25 INSERT FROM SELECT 26 INSERT INTO SELECT FROM 31 mixed and double-byte EBCDIC 26, 31 Optim Query Tuner 125 enhancements in this version xiii environment variables AQT_HOST_PACKAGE_DIRECTORY 216, 217, 245 AQTENV data set 255 European Union (EU) 267 F faulty tables processing, IBM Integrated Synchronization 143 firewall 235 G GB18030 characters 176 GDPR 267 GDPS single-node 95 General Data Protection Regulation, See GDPR glossary 273 H
access from administration client 125 access from IBM Optim Query Tuner 125 connection to IBM Db2 Analytics Accelerator for z/OS 123 disabling accelerators 259 enabling accelerators 261 Db2 for z/OS, multiple subsystems 4 Db2 subsystem 201 Db2 Unload Utility 46 Db2-supplied stored procedures setup verification 45 DDF secure port 146 DDF workload classification 47 DDVIPA setup, high availability for IBM Integrated Synchronization 160 different versions WLM application environment 40 disabling accelerators in Db2 for z/OS 259 distributed data facility, See DDF DMMD_SIGNALS table 187 documentation 6 double-byte EBCDIC (graphic) 26, 31 drain 201	Hardware Management Console, See HMC high availability availability of secure port aliases, IBM Integrated Synchronization 162 IBM Integrated Synchronization 159 location alias, IBM Integrated Synchronization 161 network setup, IBM Integrated Synchronization 160 Hipersockets 11 HMC 205 HTAP, See WAITFORDATA Hybrid Transactional Analytical Processing (HTAP), See WAITFORDATA I IBM Data Server Manager installation 21 updating 223, 224 IBM Db2 Analytics Accelerator remove from Db2 for z/OS 231 updates stored procedures 225

IBM Db2 Analytics Accelerator (continued)	incremental updates (continued)
WLM application environment 37	configuring 36
IBM Db2 Analytics Accelerator Console, logging on 199	continuous 179
IBM Db2 Analytics Accelerator for z/OS	faulty tables 180
connection testing 131	high latency 179
connection to Db2 for z/OS 123	IBM InfoSphere Change Data Capture for z/OS (CDC)
Console 132, 199, 200	176
installing updates 216	IBM Integrated Synchronization 143
ssh 200	log cache 186
telnet 199	log cache <u>100</u>
update packages 217	process restart 192
updates, installing 215	start <u>190</u>
WLM performance goals for stored procedures <u>45</u>	stop <u>190</u>
IBM Db2 Analytics Accelerator Studio	suspension of tables <u>180</u>
installation <u>20</u>	WAITFORDATA <u>195</u> , <u>197</u>
updating <u>223</u>	INSERT FROM SELECT 26
IBM Db2 Analytics Accelerator StudioIBM Data Server	INSERT INTO SELECT FROM 31
Manager 19	installation
IBM InfoSphere Change Data Capture for z/OS	DSNTIJAS 26
configuring 36	from Installation Manager 20
installing 25	IBM Data Server Manager 21
starting as z/OS started task 187	IBM Db2 Analytics Accelerator database 26
See also CDC	IBM Db2 Analytics Accelerator database 20
IBM Installation Manager	preparation 7
IBM Data Studio 20	prerequisites 7
IBM Db2 Analytics Accelerator Studio 20	testing 137
IBM Integrated Synchronization	verification <u>137</u>
assigning user IDs <u>154</u>	installation planning $\underline{1}$
AT-TLS <u>147</u>	installation, updates <u>215</u>
certificate management	installing
displaying certificates 168	base product 25
installing certificates 170	CDC package 25
removing certificates 172	Db2 Analytics Accelerator on Z
transferred certificate 174	logon 58
certificates 147	multi-node setup 55
changing Db2 user ID 164	single-node setup 55
changing DDF secure port 166	Db2 Analytics Accelerator on Z (GDPS setup, single-
changing IP address for Db2 166	node 95
	Db2 Analytics Accelerator on Z (multi-node setup) 76,
creating certificates 148	•
creating key ring <u>150</u>	108
enabling data sharing member <u>156</u>	Db2 Analytics Accelerator on Z (single-node setup 61
enabling Db2 subsystem <u>156</u>	Db2 Analytics Accelerator on Z with GDPS support 92
exporting signer certificate <u>151</u>	Db2 for z/OS libraries <u>25</u>
faulty tables processing <u>143</u>	IP address <u>11</u>
high availability	
availability of secure port aliases 162	J
location alias 161	J
network setup 160	JCLs 251
import certificate 156	JCLS <u>251</u>
LOAD REPLACE DD DUMMY 175	
permissions for schema change support 155	L
prerequisites for tables 143	-
RACF definitions 148	latency, high for incremental updates 179
	libraries 25
REORG 175	limit
REORG DISCARD 175	for Db2 subsystem 201
restrictions <u>143</u>	on processing resources 201
schema change support <u>143</u>	
secure port 146	LOAD REPLACE DD DUMMY, IBM Integrated Synchronization
SSL encryption for DDF 147	175
TTLS rules, data sharing group 151, 153	location alias, high availability for IBM Integrated
TTLS rules, single Db2 subsystem 151	Synchronization 161
improvements in this version xiii	location name <u>134</u>
incremental updates	log cache size, incremental updates <u>186</u>
applying changes in parallel 179	logon console <u>199</u>
applying changes in parallel 1/2	

logon to Admin UI, Db2 Analytics Accelerator on Z <u>58</u> loss of data <u>176</u> LPAR for Db2 Analytics Accelerator on Z <u>55</u> LPAR for Db2 Analytics Accelerator on Z with GDPS <u>93</u>	query acceleration testing defining tables <u>137</u> loading and enabling tables <u>137</u> running a test query <u>137</u>
M	R
maintenance installing updates 247 transferring updates 245 mixed EBCDIC 26, 31 multi-node setup, Db2 Analytics Accelerator on Z restart 108 shutdown 108	RACF access for IBM Db2 Analytics Accelerator for z/OS stored procedures 123 remove IBM Db2 Analytics Accelerator from Db2 for z/OS 231 REORG DISCARD,IBM Integrated Synchronization 175 REORG, IBM Integrated Synchronization 175 repair
multi-version concurrency control 179 multiple accelerators 4 MVCC, See multi-version concurrency control	installing updates <u>247</u> transferring updates <u>245</u> replacing certificate, Db2 Analytics Accelerator on Z <u>118</u> requirements for Db2 for z/OS data server 25
N	resource allocation 201
network <u>11</u> new features <u>xiii</u>	limit for Db2 subsystem <u>201</u> restart incremental update process 192
0	subscriptions 192 restrictions
Optim Query Tuner 125 OSA-Express card 11 overview 1	IBM Integrated Synchronization <u>143</u> running AQTTIJSP 49
P	AQTTRIN 49
pairing code 132 parallel, incremental updates 179 performance goals DSNUTILU 46 SYSPROC.ACCEL_LOAD_TABLES 46 WLM 45 planning 1 port-forwarding 247 ports 235 post-installation 10 power user 123 prerequisites Db2 for z/OS data server 25 IBM Data Server Manager 21 IBM Db2 Analytics Accelerator Studio 20 prerequisites for tables, IBM Integrated Synchronization 143 priorities of accelerator database engine 204 priority database engine 204 of queries 204 z/OS 204 privacy, See data privacy private data network (PDN) 13 privileges for schema change support, IBM Integrated Synchronization 155 product administration 19	sample job members 251 SAQTSAMP 251 SAQTSAMP(AQTSJI00) 50 schema change 235 schema change support, IBM Integrated Synchronization 143, 155 schema filters 126 secure port aliases, high availability for IBM Integrated Synchronization 162 secure port for DDF 146 security 13 sensitive information 13 service password 13 setting
query acceleration failure 235 schema change 235	activation <u>219</u> ssh logon to console <u>200</u> SSL certificate, Db2 Analytics Accelerator on Z <u>118</u> starting CDC <u>187</u>

starting (santinued)	upon IDo 104
starting (continued)	user IDs <u>126</u>
IBM InfoSphere Change Data Capture for z/OS <u>187</u>	users
starting incremental updates 190	Db2 Analytics Accelerator on Z
stopping incremental updates 190	add 112
stored procedures	remove 112
access rights 123, 124	reset password 112
Db2-supplied 45	
installation verification	
	W
SAQTSAMP(AQTSJI00) <u>50</u>	
subnet 11	WAITFORDATA
subscription	expiration 197
restart 192	include loaded tables 195
support case 205, 245	WLM
SYSADM 48	
SYSPROC.ACCEL_LOAD_TABLES 46	performance goals
system architecture 1	for IBM Db2 Analytics Accelerator for z/OS stored
	procedures <u>45</u>
system monitoring	for SYSPROC.ACCEL_LOAD_TABLES 46
automatic notifications 205	WLM application environment
automatic support cases 205	for different product versions 40
call home 205	IBM Db2 Analytics Accelerator 37
system-temporal tables 31, 212	WLM, application environment 255
·	workload classification
T	DDF 47
	DSNUTILU <u>46</u>
TCP/IP	SYSPROC.ACCEL_LOAD_TABLES 46
configuration 11	z/OS 47
sample configuration 233	Workload Manager, See WLM
telnet logon to console 199	
temporal tables 31, 212	
	Z
testing	
query acceleration 137	z/OS
timestamp precision > 6 <u>31</u>	importance level 204
topology <u>1</u>	ZPARMs 26, 31
tracing	<u> </u>
Db2 Analytics Accelerator on Z 111	
sensitive information 13	
transferring update	
IBM Db2 Analytics Accelerator for z/OS software 217	
maintenance <u>245</u>	
repair <u>245</u>	
troubleshooting <u>235</u>	
truncation of timestamp values 31	
<u> </u>	
and the second s	
U	
10.1.	
unified store, See WAITFORDATA	
update	
administration client 223	
IBM Data Server Manager 223	
IBM Db2 Analytics Accelerator Studio 223	
required PTFs, Db2 for z/OS 215	
· · · · · · · · · · · · · · · · · · ·	
update configuration	
Db2 Analytics Accelerator on Z 109	
updateDb2 Analytics Accelerator on Z	
appliance <u>227</u>	
updating	
credentials 132, 199, 200	
database software 216	
IBM Data Server Manager 224	
IBM Db2 Analytics Accelerator software 216	
IBM Db2 Analytics Accelerator Studio from IBM	
Installation Manager 223	
updating accelerator-shadow tables continuously 143, 176	

IBW.

Product Number: 5697-DA7

SH12-7102-06

