IBM Db2 Analytics Accelerator for z/OS
Version 7.5.1

Installation Guide

IBM

Note

Before you use this information and the product it supports, read the information in "Notices" on page 231.

Contents

Figures

x

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

Who should read this book

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:

- 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 and for an overview of the components involved in the setup. The overviews show examples; other, more sophisticated configurations are thinkable.

1

Db2 Analytics Accelerator on Z
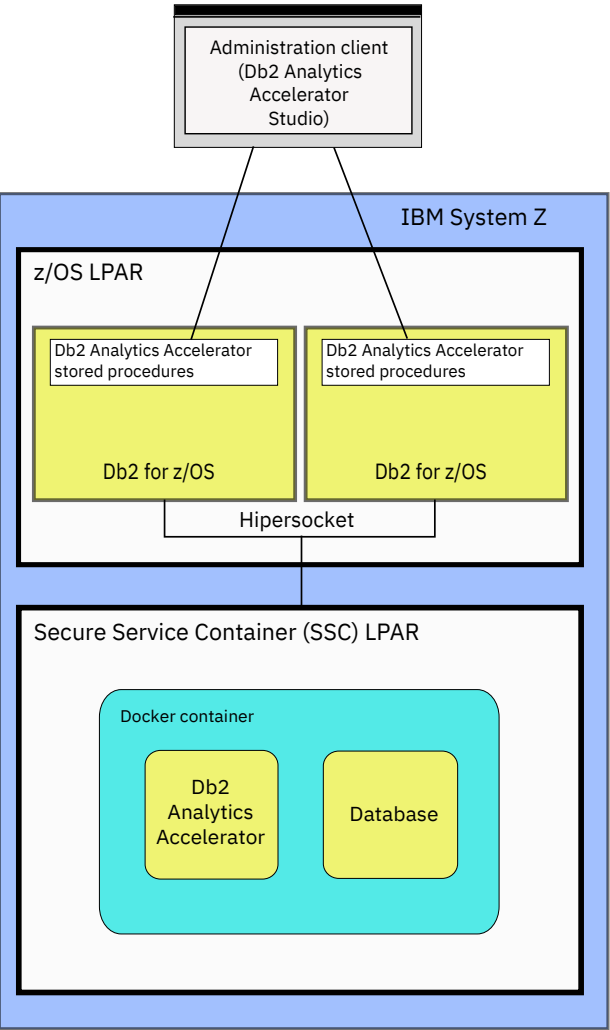
Figure 1. Solution overview of Db2 Analytics Accelerator on Z

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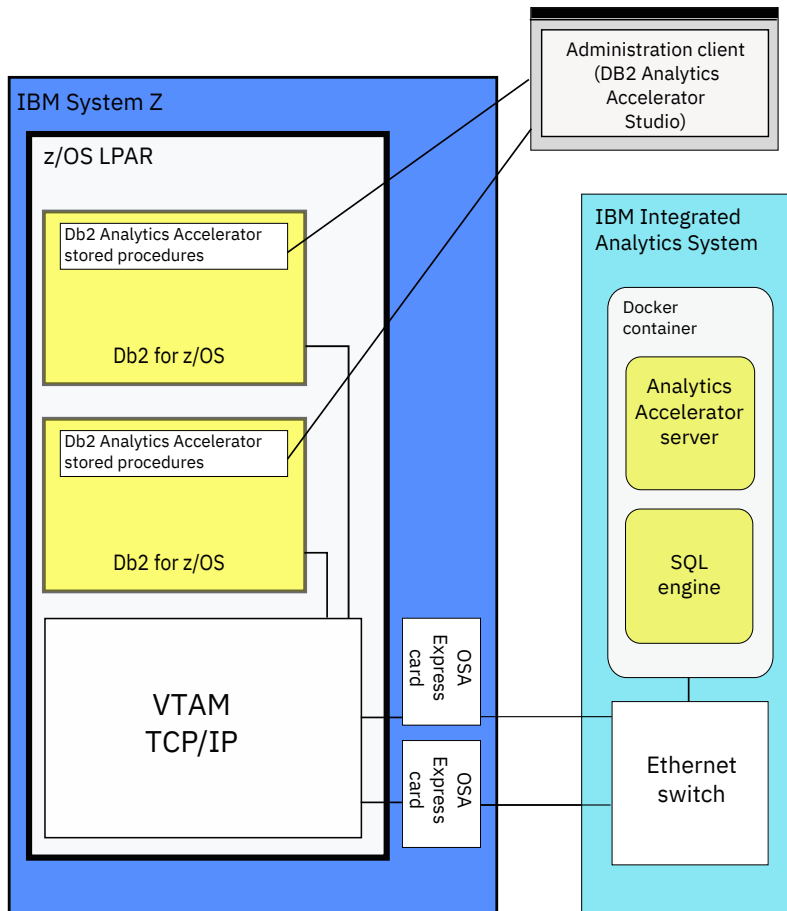
Figure 2. Solution overview of Db2 Analytics Accelerator on the IBM Integrated Analytics System

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.
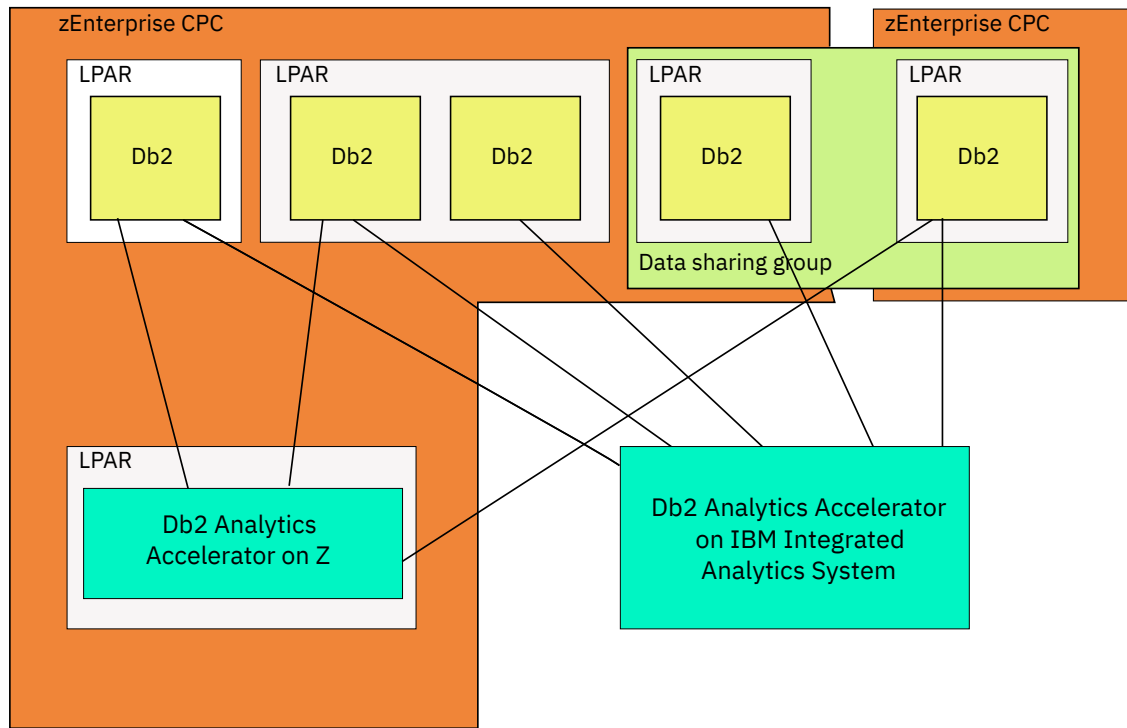
Figure 3. Possible connections

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)
- An administration client (GUI), called IBM Db2 Analytics Accelerator Studio, which runs on a workstation attached to the IBM Z®. IBM Db2 Analytics Accelerator Studio is an IBM Data Studio plugin and has to be downloaded from an IBM Web Membership page.

  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, V7.1
  - IBM DB2 Analytics Accelerator Studio (All releases, All platforms)

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

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:

http://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:

https://www.ibm.com/support/servicerequest

IBM Z documentation
Read the appropriate manual for your IBM zEnterprise® server:

- 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:

https://www.ibm.com/support/pages/prerequisites-and-maintenance-ibm-db2-analytics-accelerator-zos-v7

Installation task flow

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.

Installation focal point

A1	1/2 day
Coordinates installation and fills out site survey	

Systems programmer **Database administrator** **Network administrator**

B0	1/2 day
Capacity planning Chooses deployment options	

IBM Engineer

B1'	1/2 h
Provides LPAR for Secure Service Container on z14 Installation Guide	

B1	1/2 day
Checks and installs z/OS and DB2 for z/OS prerequisites Prerequisites website	

C1	1/2 h
Installs IBM DB2 Analytics Accelerator Studio Installation Guide	

D1	1/2 day
Optional: installs OSA cards * OSA installation manual	

* IBM Db2 Analytics Accelerator on Z can use OSA cards or Hipersockets to connect.

B2'	1/2 h
Installs Db2 Analytics Accelerator on Z Installation Guide	

B2	2 h
Installs product FMIDs using SMP/E Program directory	

C2	1 h
Evaluates WLM environment Installation Guide	

E2	1/2 day
Provides cables and connects OSA cards with IBM Integrated Analytics System. Installation Guide (this book)	

B3	2 h
Installs latest product PTFs using SMP/E Closing information	

C3	1 h
Customizes JCLs for stored procedure setup Installation Guide	

E3	2 h
Configures OSA cards and TCP/IP stack (continuous address range, min. 4 IP addresses) Installation Guide	

C4	2 h
Tests stored procedures Installation Guide	

B4	1/2 day
Assists with installation Installation Guide	

C5	1 h
Connects product to DB2 for z/OS Installation Guide	

D2	2 days
Installs product hardware and software, helps to connect product to DB2 for z/OS Installation Guide	

C6	1 day
Defines, loads, and enables tables for acceleration; runs accelerated queries User's Guide	

Legend:
Black = common tasks
Red = IBM Db2 Analytics Accelerator on IBM Integrated Analytics System
Green = IBM Db2 Analytics Accelerator on Z

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*.

Preparation tasks

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

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 IBM Db2 Analytics Accelerator Studio on a client workstation.

   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 IBM Db2 Analytics 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 *IBM Db2 Analytics Accelerator for 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 137.

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 must exist between the IBM Z server and your IBM Integrated Analytics System. This means that you cannot reuse existing subnets that are already defined on other interfaces. The number of required TCP/IP addresses depends on the configuration. For the minimum configuration, provide at least two TCP/IP addresses in the same subnet. For the recommended configuration, provide four IP addresses.

- One IP address for the IBM Integrated Analytics System.
- 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)
```

During the network configuration for Db2 Analytics Accelerator on Z, an interface needs to be tagged that accepts incoming connection requests from Db2 for z/OS subsystems. This is done by editing the network interface name and changing it to *DB2*.

Note: Only a single interface can be tagged in this way. If you plan to connect the accelerator to multiple Db2 subsystems, all these subsystems must be accessible through this interface.

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 web-browser interface. Make sure that access to this interface is limited, that the switches restrict routing to the confines of the PDN, and that the initial password has been changed. Someone with access to the interface can change the configuration of the network and might, for example, open the PDN to the intranet or internet.
- The trace function of the product allows you to select components for tracing that 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. 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 "8" on page 21, 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.

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 data server

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

The prerequisites are listed at:

http://www.ibm.com/support/docview.wss?uid=swg27050440

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 HAQT710
   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 Data Replication for z/OS components

2. If you want to use the incremental update function, follow these steps:
   a) Complete the installation of IBM InfoSphere Data Replication for z/OS.
   b) Configure IBM InfoSphere Data Replication 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

The incremental update function of IBM Db2 Analytics Accelerator for z/OS requires components of IBM InfoSphere Data Replication for z/OS (CDC) and additional components on the accelerator, called the Access Server and the replication engine. Both components are already installed and configured as part of the accelerator setup. They are updated automatically by the software update function. CDC requires additional deployment and configuration steps. If you want to use this feature, you are advised to configure at least the CDC components at this point. To read an introduction to the incremental update function or go straight away to the CDC configuration section, follow the appropriate link at the end of this topic.

Creating the IBM Db2 Analytics Accelerator database

Customize and submit the DSNTIJAS job to create the database and tables for IBM Db2 Analytics Accelerator for z/OS in Db2 for z/OS.

About this task

Complete this task before you create the IBM Db2 Analytics Accelerator stored procedures by following the steps in "Customizing and running AQTTIJSP" on page 47.

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 11 for z/OS

In Db2 11 for z/OS, you can set the ZPARMs for query acceleration on the DSNTIP82 and DSNTIP8A installation panels.

About this task

The following ZPARMs are available:

ACCEL=COMMAND | AUTO | NO

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

   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        |
```

```
|   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    NO or YES          |
|                                                                   |
|   7 YYYYMMDD   LOCAL DATE FORMAT    ===> NO    NO or YES          |
|   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. 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

If you do not use Db2 11 for z/OS or prefer a different method, you can customize and run a sample job to modify the ZPARM settings for IBM Db2 Analytics Accelerator for z/OS.

About this task

The following ZPARMs are available:

ACCEL=COMMAND | AUTO | NO

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

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        |
|  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    NO or YES        |
|                                                                     |
|  7 YYYYMMDD   LOCAL DATE FORMAT         ===> NO    NO or YES        |
|  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

The incremental update function of IBM Db2 Analytics Accelerator for z/OS requires components of IBM InfoSphere Data Replication for z/OS (CDC) and additional components on the accelerator, called the Access Server and the replication engine. Both components are already installed and configured as part of the accelerator setup. They are updated automatically by the software update function. CDC requires additional deployment and configuration steps. If you want to use this feature, you are advised to configure at least the CDC components at this point. To read an introduction to the incremental update function or go straight away to the CDC configuration section, follow the appropriate link at the end of this topic.

Related tasks

Completing the installation and configuring CDC

Read how to complete the installation and configure IBM InfoSphere Data Replication 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

<HLQSP>.SAQTLICI
    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

    AQTLICZ
        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.

<HLQSP>.SAQTMOD
    Contains shared libraries and load modules for the stored procedures.

    Important: If you want to use the IBM Db2 Analytics Accelerator *Loader* in connection with this product, this data set must be APF-authorized.

Procedure

1. 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:

```
Appl Environment Name . . DSNWLMV9
Description . . . . . . . Db2 V11 default Stored Procedures for IDAA
Subsystem type  . . . . . Db2
Procedure name  . . . . . DSNWLM
Start parameters  . . . . Db2SSN=&IWMSSNM,APPLENV=DSN
. . . . . . . . . . . . . WLMV9
```

Important:

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

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

```
//****************************************************************
//* PROCEDURE NAME = DSNWLM
//*
//*     JCL FOR RUNNING THE WLM-ESTABLISHED STORED PROCEDURES
//*     ADDRESS SPACE
//*        RGN      -- THE MVS REGION SIZE FOR THE ADDRESS SPACE.
//*        Db2SSN   -- THE Db2 SUBSYSTEM NAME.
//*        NUMTCB   -- THE NUMBER OF TCBS USED TO PROCESS
//*                    END USER REQUESTS.
//*        APPLENV -- THE MVS WLM APPLICATION ENVIRONMENT
//*                    SUPPORTED BY THIS JCL PROCEDURE.
//*
//*        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
//OUT1     DD    SYSOUT=A
//UTPRINT  DD    SYSOUT=A
//DSSPRINT DD    SYSOUT=A
//SYSPRINT DD    SYSOUT=A
//AQTENV   DD    DSN=<HLQACTIVE>.SAQTSAMP(AQTENV),DISP=SHR
//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 IBM Data Studio.

  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 statement to the procedure that starts the address space:

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

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

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

- You might want to change the default settings for IBM Db2 Analytics Accelerator 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:

<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.

<HLQSP>.SAQTMOD
    Contains shared libraries and load modules for the stored procedures.

    Important: If you want to use the IBM Db2 Analytics Accelerator *Loader* in connection with this product, this data set must be APF-authorized.

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:

```
Appl Environment Name . . DSNWLMV9
Description . . . . . . . Db2 V11 default Stored Procedures for IDAA
Subsystem type  . . . . . Db2
Procedure name  . . . . . DSNWLM
Start parameters  . . . . Db2SSN=&IWMSSNM,APPLENV=DSN
. . . . . . . . . . . . . WLMV9
```

Important:

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

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

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
//*       RGN    -- THE MVS REGION SIZE FOR THE ADDRESS SPACE.
//*       Db2SSN  -- THE Db2 SUBSYSTEM NAME.
//*       NUMTCB  -- THE NUMBER OF TCBS USED TO PROCESS
//*                  END USER REQUESTS.
//*       APPLENV -- THE MVS WLM APPLICATION ENVIRONMENT
//*                  SUPPORTED BY THIS JCL PROCEDURE.
//*
//*       <HLQXML4C1.10> is the HLQ where you have installed the
//*                  XML Toolkit for z/OS
//*
//*       Db2VERS -- Db2-VERSION (I.E. V910)
//*                  SET BY APPLICATION ENVIRONMENT
//*                      DSNWLMV9 ==> V910
//*
//*    The user ID that is used to start the task must have
//*    read access to the <HLQs> in the STEPLIB statement
//*
//****************************************************************
//DSNWLMV9  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>.SIXMLOD1
//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
//AQTDEF6 DD    DSN=<HLQACTIVEV7>.SAQTSAMP(AQTDEF6),DISP=SHR
//AQTDEFTR DD    DSN=<HLQACTIVEV5>.SAQTSAMP(AQTDEFTR),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 IBM Data Studio.

  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 Integrated Analytics System is supposed to use. To do so, add the following statement to the procedure that starts the address space:

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

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

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

- You might want to change the default settings for IBM Db2 Analytics Accelerator 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 IBM Db2 Analytics Accelerator Studio 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 IBM Db2 Analytics Accelerator Studio:

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

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

- 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

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
--------------------------------------------------------------------
---
               Modify Rules for the Subsystem Type      Row 1 to 1
of 1
 Command ===> _____    SCROLL ===>
PAGE

 Subsystem Type . : DDF        Fold qualifier names?   Y  (Y or N)
 Description  . . . Distributed requests

 Action codes:   A=After     C=Copy       M=Move     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  UI         *        ___              DDF          _____
 ____  2   PR        ACCEL_L* ___             DDFLOAD_      _____
```

Figure 5. 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.

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 install new versions of the database objects, keep the value `INSTALL` for the `MODE` parameter in job step AQTTRIN.

   • 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.

   • 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 AQTSJI00 JCL by running SAQTSAMP(AQTSJI00).

Results

The job returns the following information:

- The contents of the SYSACCEL.SYSACCELERATORS table.

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

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

- Does the Db2 Communication Database (CDB) exist?
- Do the following tables exist?
  - ACCEL_NAMES
  - ACCEL_QUERY_INFO
  - ACCEL_TRACE_ACCELERATOR
- Have all stored procedures been defined that are used or provided by IBM Db2 Analytics Accelerator for z/OS, and if so, what are their run options (RUNOPTS) and Workload Manager (WLM) settings?

What to do next

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

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

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.

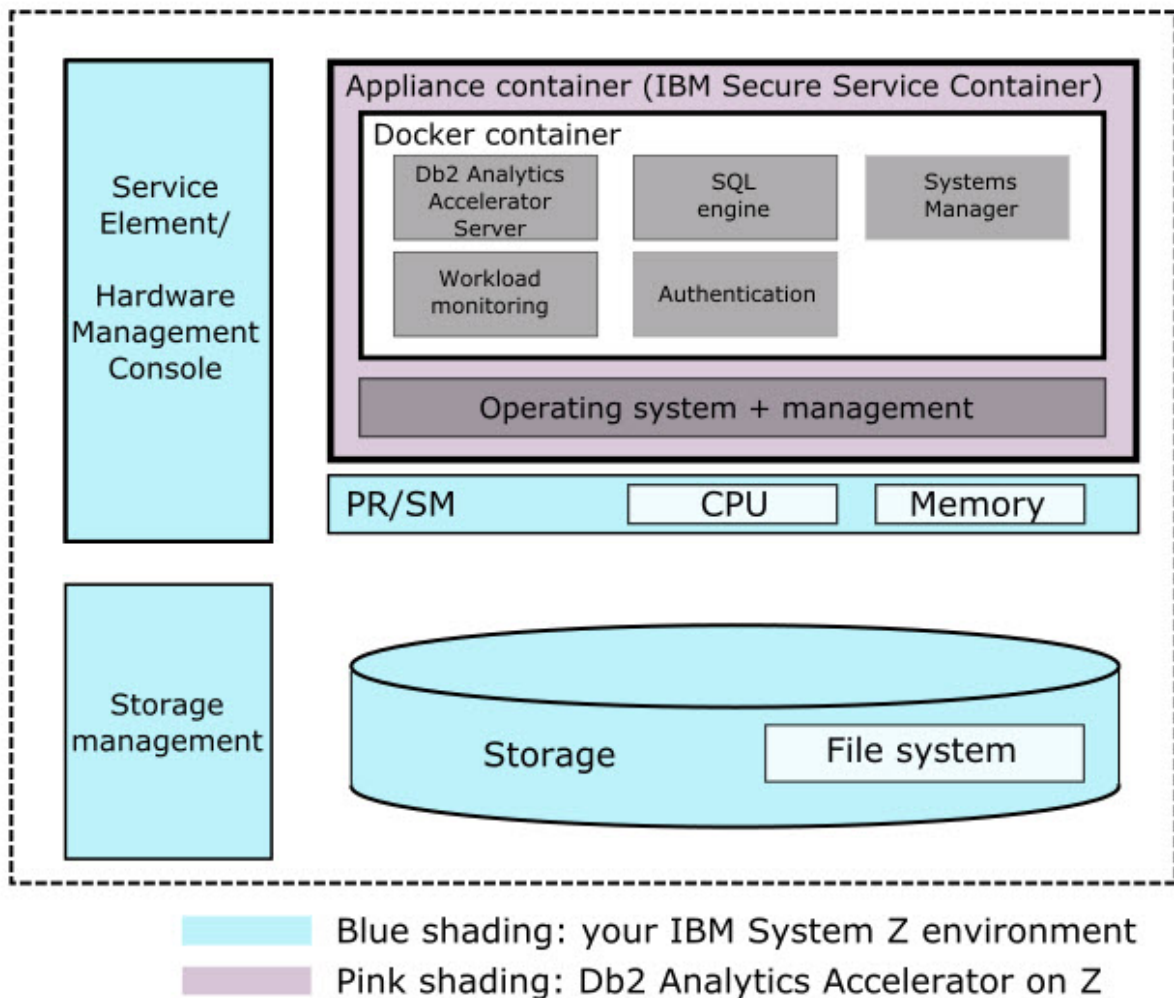*Figure 6. Architecture*

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 only.

Figure 7 on page 52 shows two sites, (CEC A and CEC B), with an active and a passive accelerator (IDAA-on-Z LPAR A and LPAR B).
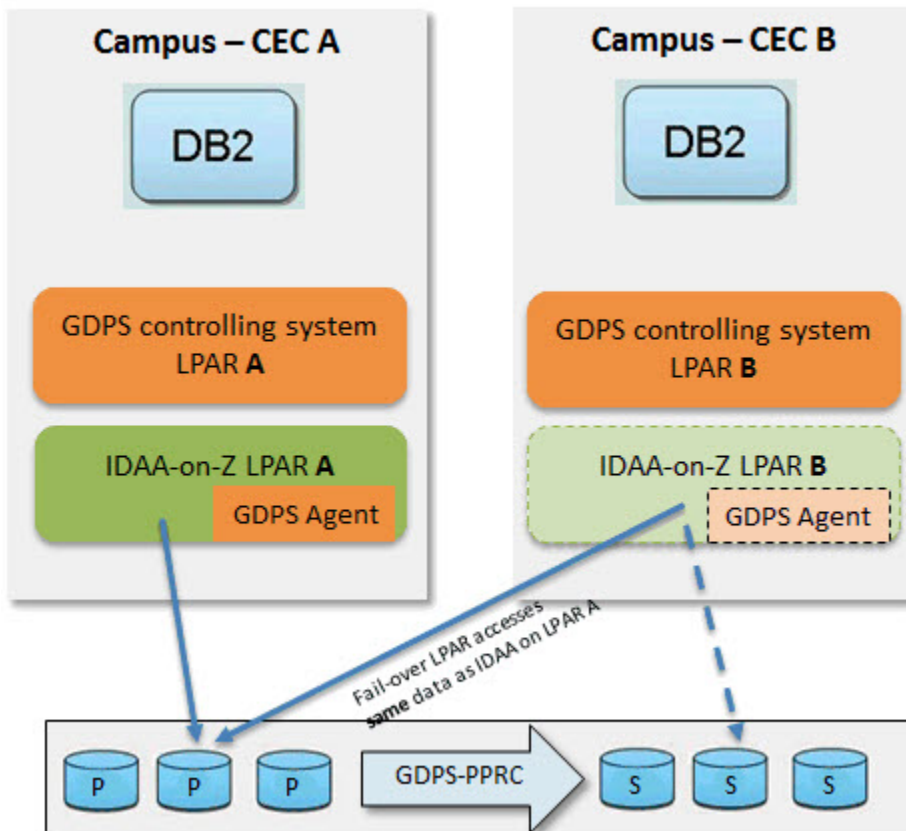
Figure 7. Failover support for a single-nodeDb2 Analytics Accelerator on Z with GDPS: schematic overview

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**.
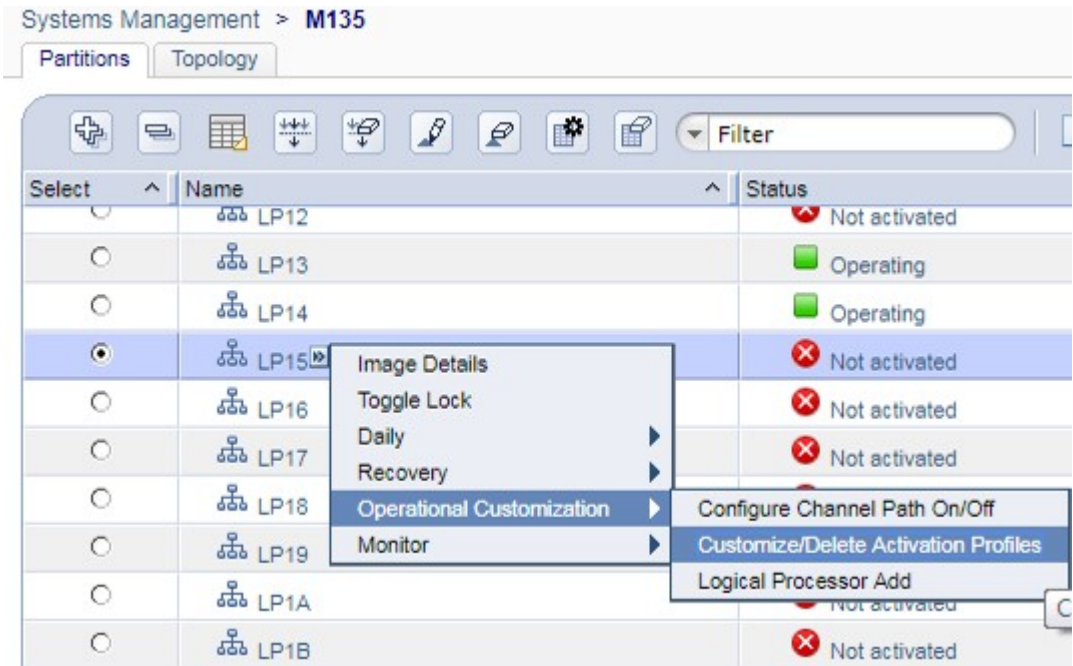
Figure 8. The Partitions window of the HMC

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 9. 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 10. SSC network adapter details on the HMC

⚠️ **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 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**:
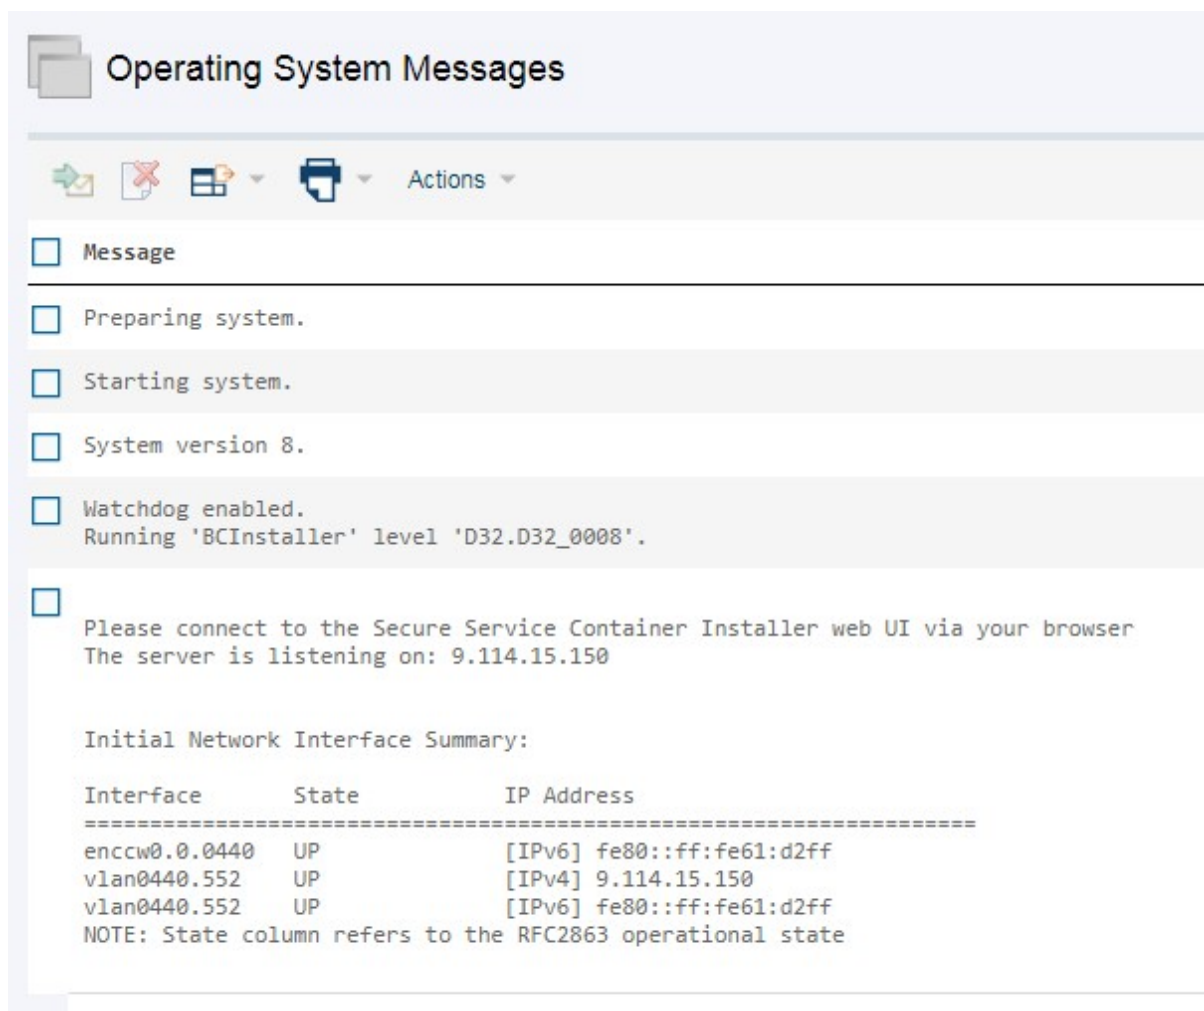
Figure 11. Confirmation: Operating System Messages

Logging on to the Appliance Installer

To complete your Db2 Analytics Accelerator on Z configuration, you must log on to the Appliance Installer. 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. Open the Appliance Installer in a web browser. Enter the IPv4 address of the Secure Service Container (SSC) LPAR (see).
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)

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

Figure 12. 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 Appliance Installer 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 Appliance Installer 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 60.

   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:

Figure 13. The Install Software Appliance window of the Appliance Installer

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.

   The following page is displayed:
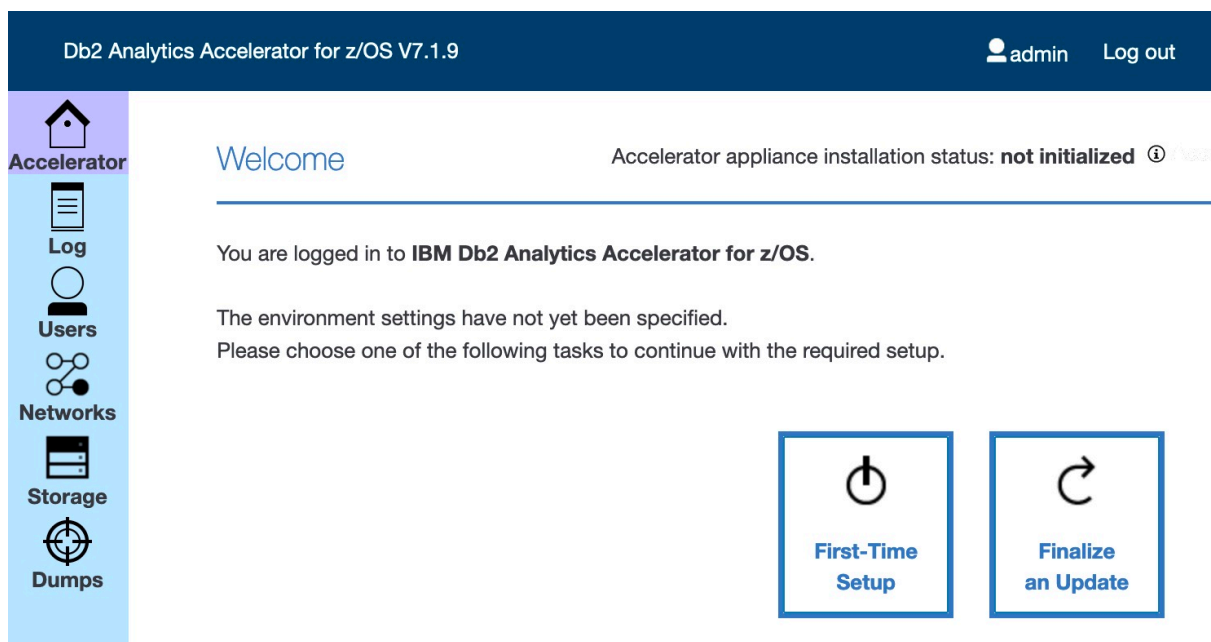
*Figure 14. The **Welcome** page of the Appliance Installer*

10. On the **Welcome** page, click **First-Time Setup**.

You see the following page:
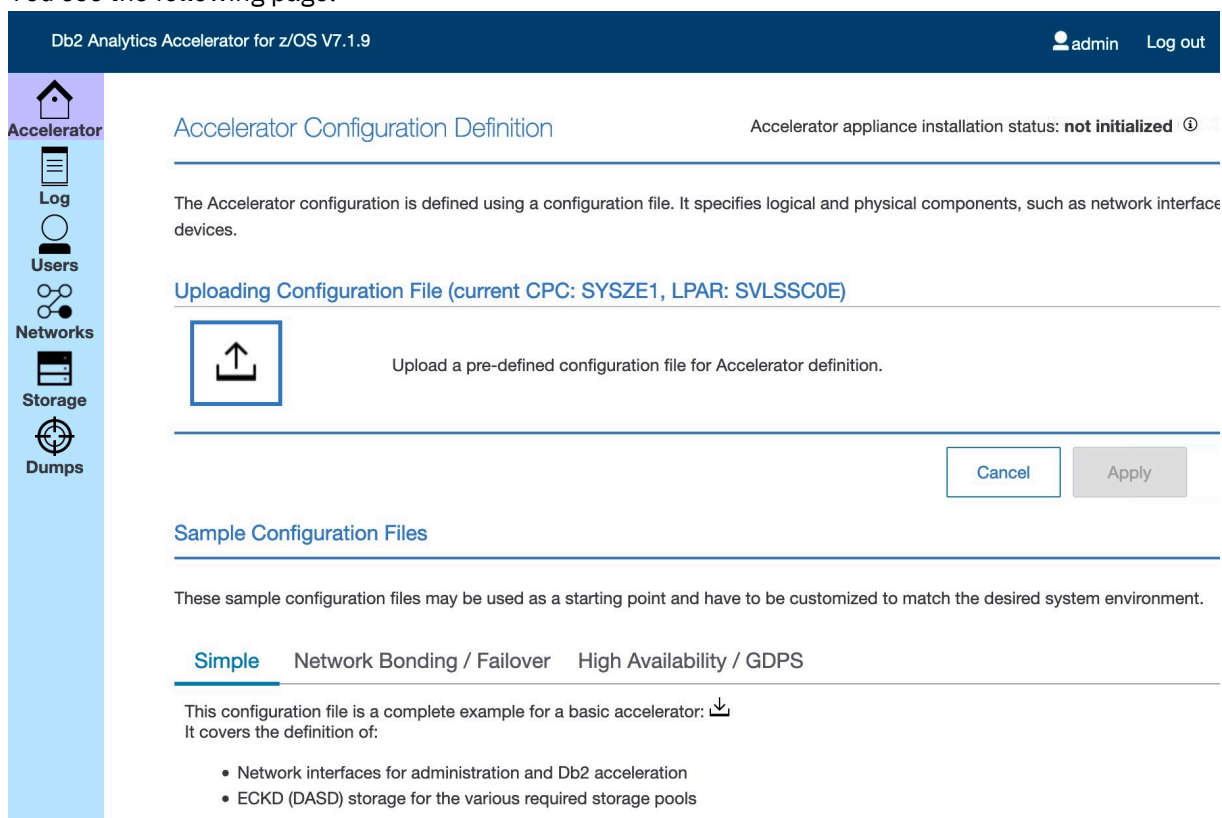
*Figure 15. **Accelerator Configuration Definition** page*

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. Table 1 on page 60 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.
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 53.
2. Log on to the Appliance Installer and proceed to the **Welcome** page.
   For a description, see "Logging on to the Appliance Installer" on page 56.
3. On the **Welcome** page, click **First-Time Setup**.
   You see the following page:
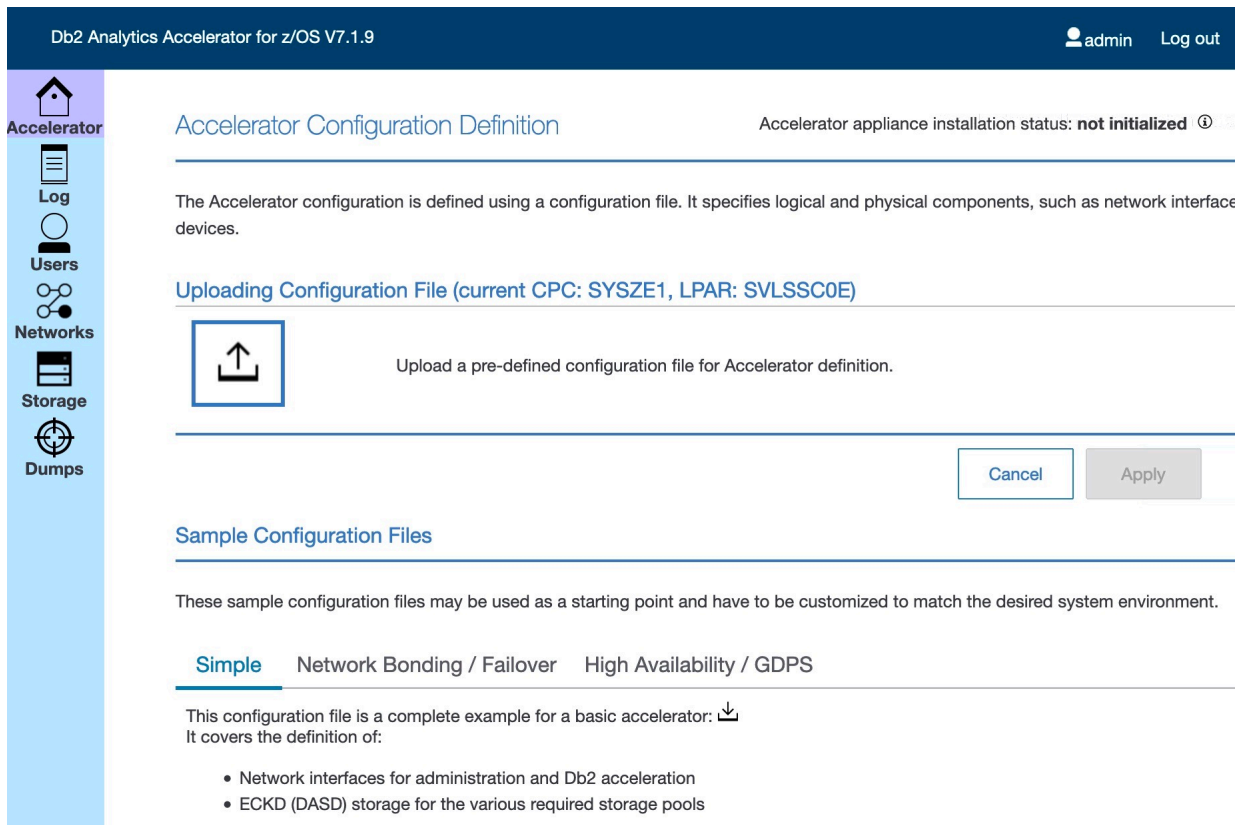
*Figure 16. **Accelerator Configuration Definition** page*

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.

   b. Click the download button ( ).

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

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

Figure 17. 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:

```
{
  "accelerator_name": "igor01",
  "accelerator_description": "SAMPLE-single-node",
  "accelerator_type": "single-node",
  "db2_pairing_ipv4": "10.1.1.101/24",
  "network_interface_bindings": {
    "mgmt_nw": "activation-profile",
    "db2_nw": "osa-0A-p0"
  },
  "runtime_environments": [
    {
      "cpc_name": "CPC001",
      "lpar_name": "IGOR001",
      "network_interfaces": [
        {
          "name": "osa-0A-p0",
          "device": "0.0.0a00",
          "port": "0"
        }
      ]
    }
  ],
  "primary_storage": {
    "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"
            ]
          ]
        }
      }
    }
  }
```

Most attributes in the configuration file are required.

"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. 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": "10.101.31.172/24"
```

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.

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 administration user interface 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. 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": {
    "mgmt_nw": "activation-profile",
    "db2_nw": "my_db2_network",
    "gdps_nw": "my_gdps_network"
}
```

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 for the accelerator (SSC) LPARs. Each of these LPARs is identified by the CPC name and the LPAR name, and 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. 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 each network:

"name" (required)
> The name of the network interface. You can choose it freely.

"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".

"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:

```
"runtime_environments": [
    {
        "cpc_name": "SYSZE1",
        "lpar_name": "SVLSSC0C",
        "network_interfaces": [
            {
                "name": "my_db2_network",
                "device": "0.0.4b00",
                "port": "1",
                "vlan": "600"
```

```
        }
    ]
  }
]
```

In this example, you find a network interface definition for an accelerator (SSC) LPAR named SVLSSC0C. 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.

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"
  },
  "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"}]
        }
      ]
    }
  ],
```

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:

```
"network_interfaces": [
    {
    "name": "db2_nw_connection",
    "vlan": "700",
    "bond_settings": {
    "mode": "active-backup",
    "slaves": [
            {
            "device": "0.0.0a00",
```

```
                "port": "0"
            },
            {
                "device": "0.0.1b00",
                "port": "1"
            }
        ]
    }
```

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 slave 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.

"primary"
> Valid in *active-backup* mode only. The primary slave device. 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 slave device is selected after a failure. Specify "always", which means that an attempt will be made to make the primary device 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 slave is only changed if the currently active device fails.

"failover-MAC": "none"
> Valid in *active-backup* mode only. Allows you set all slaves 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 slave 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 slave 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 slave 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 slave 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 slave 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 slave 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 slave 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"
}
```

`"primary_storage"` (required)
This block lists all storage devices for the primary site, that is, the site on which accelerator is initially deployed and configured. 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. 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. The boot device must be a single device with at least 40 GB net storage capacity.

"runtime_devices" (required)

The runtime storage is used by the Docker container that runs the accelerator software. Specify a list of devices with a total net capacity of at least 80 GB.

"data_devices" (required)

The data storage is used to store the operative 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.

"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). It is not possible to mix ECKD and SCSI devices in a single device category or device pool.

Example (ECKD or "dasd"):

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

Example (SCSI or "zfcp"):

```
"type": "zfcp",
        "devices": [
            {
             "device": "7a01",
             "lun": "0x4001012384729172",
             "wwpn": "0xc00984738927u834"
            }
        ]
```

Notes:

- You can specify lists of SCSI devices in much the same way as you specify lists of ECKD devices.

- However, because the logical unit numbers (LUNs) and worldwide port names (WWPNs) are structured the way they are, you cannot specify ranges. That is, you have to reuse the keywords `"device"`, `"lun"` and `"wwpn"` for each additional device.

"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:

```
"primary_storage": {
                    "boot_device": {
                     "type": "dasd",
                     "device": "0.0.9986"
                    },
                    "runtime_devices": {
                     "type": "dasd",
                     "devices": [
                       "0.0.998c"
                       ]
                    },
                    "data_devices": {
                     "type": "dasd",
                     "devices": [
                       "0.0.9c00",
                       "0.0.9c01",
                       "0.0.9c02"
                       ]
```

```
            }
    }
```

"storage_maps"
   To replicate your primary storage to a secondary site.

   Example:

```
"storage_maps": [
    {
        "boot_device": "0.0.9c80",
        "map": [
            {
                "primary": ["0.0.9d00","0.0.9d0b"],
                "copy":    ["0.0.9c80","0.0.9c8b"]
            }
        ]
    }
]
```

Best practices for storage definitions:

For improved storage access and formatting of FICON®-attached storage devices (this does *not* apply to FCP-attached (SCSI) devices!), your storage administrator might have defined a set of HyperPAV aliases. Use one of the following procedures to include your HyperPAV aliases in the storage setup:

• Approach® 1:

   a. Add all physical storage devices to all the pools you intend to use.

   b. Add the HyperPAV aliases to the HyperPAV aliases pool.

• Approach 2:

   a. Define one physical storage device in each pool.

   b. Add the HyperPAV aliases to the HyperPAV aliases pool.

   c. Distribute the rest of the storage devices to the pools.

   Approach 2 results in a faster formatting of the disks.

Important: 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.

Automatic 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 primary storage configuration section in the JSON configuration file.

Example:

```
"primary_storage": {
    "boot_device": {
            .
            .
    },
    "runtime_devices": {
            .
            .
    },
    "data_devices": {
            .
            .
    },
    "data": {
            "device": [
                "0.0.121f"
            ],
            "type": "dasd"
    },
```

```
    "hyperpav": "auto"
}
```

Instead of "`auto`", you can use "`automatic`". Both keywords will be accepted.

Mind that in "`auto`" mode, the system uses all available HyperPAV alias devices. To use just a subset of the available devices, you can specify a list 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`.

For GDPS or high-availability setups, an explicit listing of HyperPAV aliases does not work. You must use automatic HyperPAV aliases in this case.

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

   On the , under the heading **Uploading Configuration File**, click the upload button:
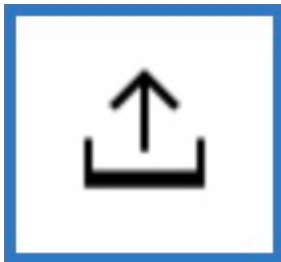
   If something is wrong with the file you uploaded, an error message is displayed on the page:
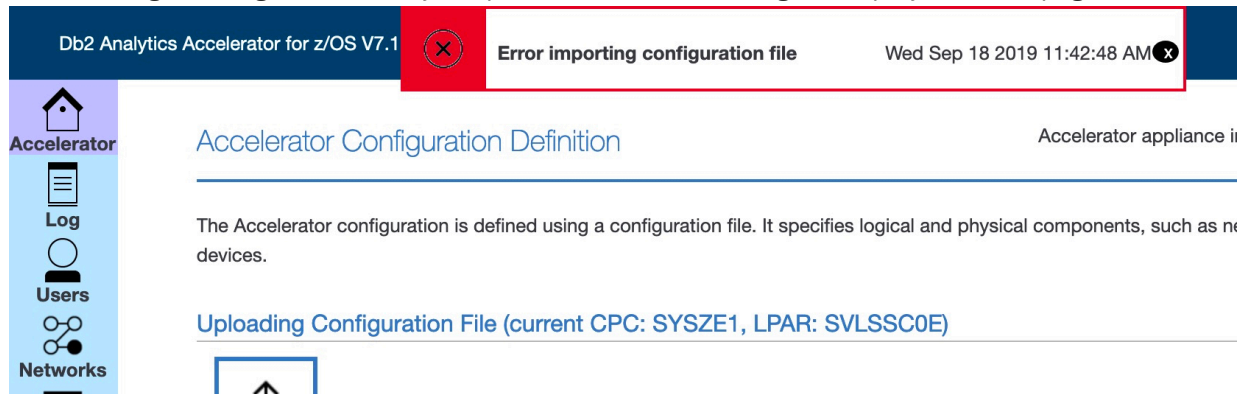
Figure 18. Error message after uploading a faulty configuration file

   You can click the error message to display further information about the error:

Figure 19. Error message expanded

7. If errors occurred, fix these and repeat the upload (steps <u>"5" on page 62</u> and <u>"6" on page 70</u>).

   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 20. Accelerator Configuration Definition after a successful configuration file upload

8. Click **Apply**.

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

Figure 21. Message window showing the progress of your Db2 Analytics Accelerator on Z configuration

Results

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

*Figure 22. The **Appliance Component Health** page is displayed after a successful configuration*

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

Related information

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 IFLs 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 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 "Defining an LPAR for Db2 Analytics Accelerator on Z" on page 53.

2. Log on to the Appliance Installer and proceed to the **Welcome** page.

   For a description, see "Logging on to the Appliance Installer" on page 56.

3. On the **Welcome** page, click **First-Time Setup**.

   You see the following page:

*Figure 23. **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.

   "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. Set this attribute to the value `"multi-node"`.

"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 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 head node is started.

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.

"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 administration user interface 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. 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": {
    "mgmt_nw": "activation-profile",
    "db2_nw": "my_db2_network",
    "cluster_nw": "my_hiper"
}
```

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 for the accelerator (SSC) 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 the Hipersocket network of your cluster. Specify the following attributes to identify an accelerator (SSC) LPAR:

"cpc_name"

The name of the CPC.

"head"

The definition of the head node. For example:

```
"head": {
        "lpar_name": "LPAR0",
        "network_interfaces": [
          {
           "name": "my_hiper",
           "device": "0.0.7f00",
           "ipv4": "172.84.0.180/23"
          },
          {
           "name": "my_db2_network",
           "device": "0.0.4b00"
          }
        ]
},
```

"lpar_name"

The name of an 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 an 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.

"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".

"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:

```
"runtime_environments": [
    {
     "cpc_name": "IBMZ1",
     "head": {
             "lpar_name": "LPAR0",
             "network_interfaces": [
               {
                "name": "my_hiper",
                "device": "0.0.7f00",
                "ipv4": "172.84.0.180/23"
               },
               {
                "name": "my_db2_network",
                "device": "0.0.4b00"
               }
```

```
            ]
      },
```

"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.

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"
  },
  "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"}]
        }
      ]
    }
  ],
```

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:

```
"network_interfaces": [
    {
    "name": "db2_nw_connection",
    "vlan": "700",
    "bond_settings": {
    "mode": "active-backup",
    "slaves": [
        {
        "device": "0.0.0a00",
        "port": "0"
        },
        {
        "device": "0.0.1b00",
        "port": "1"
        }
```

```
        ]
    }
```

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 slave 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.

"primary"

Valid in *active-backup* mode only. The primary slave device. 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 slave device is selected after a failure. Specify "always", which means that an attempt will be made to make the primary device 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 slave is only changed if the currently active device fails.

"failover-MAC": "none"

Valid in *active-backup* mode only. Allows you set all slaves 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 slave 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 slave 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 slave 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 slave 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 slave 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 slave 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 slave 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"
}
```

"data1"
The definition of the first data node. It has only one network interface: the Hipersocket network of the cluster.

Example:

```
"data1": {
          "lpar_name": "LPAR1",
          "network_interfaces": [
            {
              "name": "my_hiper",
              "ipv4": "172.84.0.181/23",
              "device": "0.0.7f00"
            }
```

```
            ]
},
```

Specify the other data nodes in the same manner, that is, create the sections `"data2"`, `"data3"`, `"data4"`, and `"data5"`

Complete example:

```
"runtime_environments": [
    {
        "cpc_name": "IBMZ1",
        "head": {
          "lpar_name": "LPAR0",
          "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",
          "network_interfaces": [
             {
                "name": "my_hiper",
                "ipv4": "172.84.0.181/23",
                "device": "0.0.7f00"
             }
          ]
        },
        "data2": {
          "lpar_name": "LPAR2",
          "network_interfaces": [
             {
                "name": "my_hiper",
                "ipv4": "172.84.0.182/23",
                "device": "0.0.7f00"
             }
          ]
        },
        "data3": {
          "lpar_name": "LPAR3",
          "network_interfaces": [
             {
                "name": "my_hiper",
                "ipv4": "172.84.0.183/23",
                "device": "0.0.7f00"
             }
          ]
        },
        "data4": {
          "lpar_name": "LPAR4",
          "network_interfaces": [
             {
                "name": "my_hiper",
                "ipv4": "172.84.0.184/23",
                "device": "0.0.7f00"
             }
          ]
        },
        "data5": {
          "lpar_name": "LPAR5",
          "network_interfaces": [
             {
                "name": "my_hiper",
                "ipv4": "172.84.0.185/23",
                "device": "0.0.7f00"
             }
          ]
        }
    }
  ],
```

"primary_storage" (required)
> This block lists all storage devices for the primary site, that is, the site on which the cluster is initially deployed and configured. During the first-time deployment, these devices are formatted, which means that the existing data on these devices is erased.

"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. 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. Specify a list of devices with a total net capacity of at least 80 GB for each node.

"data_devices" (required)
> A data device is used to store operative 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.

"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). It is not possible to mix device types, that is, use ECKD and SCSI devices, in a single device category or device pool.

> **Example (ECKD or "dasd"):**

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

> **Example (SCSI or "zfcp"):**

```
"type": "zfcp",
      "devices": [
          {
           "device": "7a01",
           "lun": "0x4001012384729172",
           "wwpn": "0xc00984738927u834"
          }
      ]
```

> **Notes:**
> - You can specify lists of SCSI devices in much the same way as you specify lists of ECKD devices.
> - However, because the logical unit numbers (LUNs) and worldwide port names (WWPNs) are structured the way they are, you cannot specify ranges. That is, you have to reuse the keywords `"device"`, `"lun"` and `"wwpn"` for each additional device.

"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:

```
"primary_storage": {
    "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"
        ]
      }
    }
  }
```

Best practices for storage definitions:

For improved storage access and formatting of FICON-attached storage devices (this does *not* apply to FCP-attached (SCSI) devices!), your storage administrator might have defined a set of HyperPAV aliases. Use one of the following procedures to include your HyperPAV aliases in the storage setup:

- Approach 1:

    a. Add all physical storage devices to all the pools you intend to use.

    b. Add the HyperPAV aliases to the HyperPAV aliases pool.

- Approach 2:

    a. Define one physical storage device in each pool.

    b. Add the HyperPAV aliases to the HyperPAV aliases pool.

    c. Distribute the rest of the storage devices to the pools.

    Approach 2 results in a faster formatting of the disks.

Important: 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.

Automatic 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 primary storage configuration section in the JSON configuration file.

Example:

```
"primary_storage": {
    "boot_device": {
          .
          .
    },
    "runtime_devices": {
          .
          .
    },
    "data_devices": {
          .
          .
    },
    "data1": {
```

```
            "device": [
                "0.0.121f"
            ],
            "type": "dasd"
        },
        "hyperpav": "auto"
    }
```

Instead of `"auto"`, you can use `"automatic"`. Both keywords will be accepted.

Mind that in `"auto"` mode, the system uses all available HyperPAV alias devices. To use just a subset of the available devices, you can specify a list 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`.

For GDPS or high-availability setups, an explicit listing of HyperPAV aliases does not work. You must use automatic HyperPAV aliases in this case.

"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 (`"primary_storage"`).

Complete storage definition (partly shortened):

```
"primary_storage": {
    "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" ]
            ]
        },
    },
    "data1": {
        "boot_device": ...,
        "data_devices": ...,
        "runtime_devices": ...
    },
    "data2": {
        "boot_device": ...,
        "data_devices": ...,
        "runtime_devices": ...
    },
    "data3": {
        "boot_device": ...,
        "data_devices": ...,
        "runtime_devices": ...
    },
    "data4": {
        "boot_device": ...,
        "data_devices": ...,
        "runtime_devices": ...
    },
    "data5": {
        "boot_device": ...,
        "data_devices": ...,
```

```
        "runtime_devices": ...
      }
    }
```

"storage_maps"
> To replicate your primary storage to a secondary site.

> **Example:**

```
"storage_maps": [
    {
      "boot_device": "0.0.9c80",
      "map": [
        {
          "primary": ["0.0.9d00","0.0.9d0b"],
          "copy":    ["0.0.9c80","0.0.9c8b"]
        }
      ]
    }
]
```

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

   On the , under the heading **Uploading Configuration File**, click the upload button:

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

Figure 24. Error message after uploading a faulty configuration file

You can click the error message to display further information about the error:
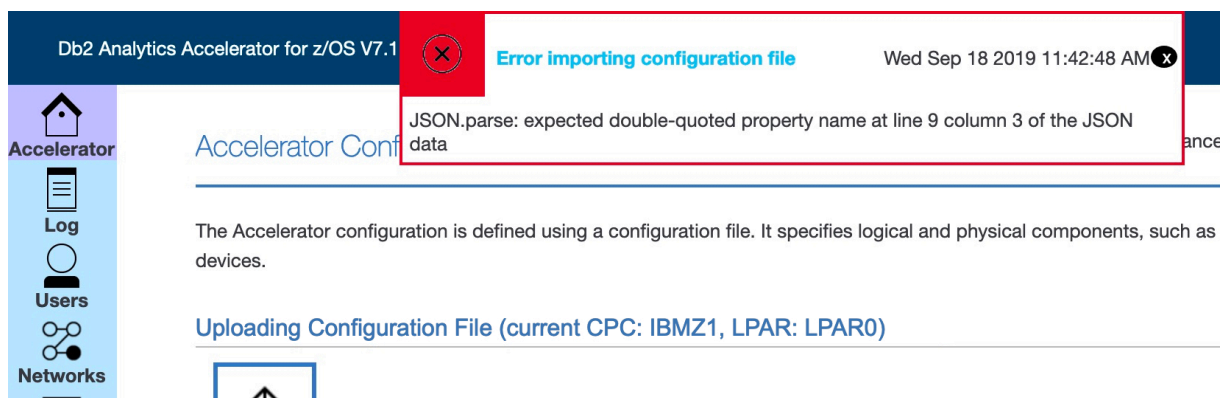
Figure 25. Error message expanded

7. If errors occurred, fix these and repeat the upload (steps "5" on page 75 and "6" on page 85).

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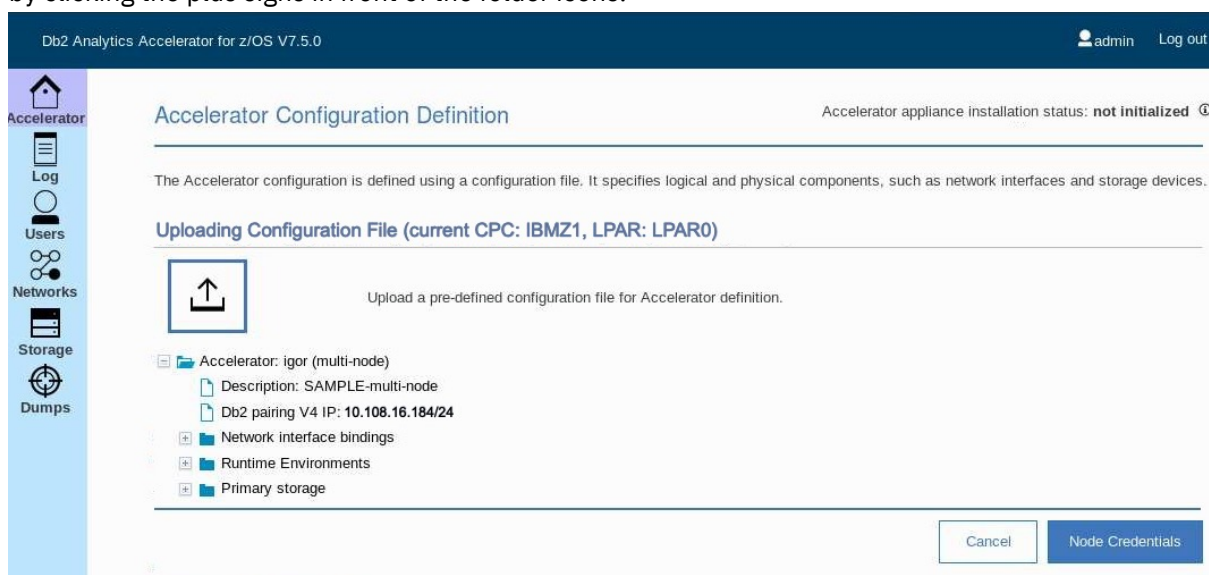
Figure 26. Accelerator Configuration Definition after a successful configuration file upload

8. Click **Node Credentials**.

9. 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 Figure 27 on page 87.

*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 **Appliance Component Health** page is displayed automatically. The page should now give you the following information:

*Figure 29. The **Appliance Component Health** page is displayed after a successful configuration*

The message `Accelerator appliance installation status: installation completed` 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 combine the performance benefit of a multi-node installation with the high availability provided by GDPS failover support. Note that this type of setup requires at least two Central Processing Complexes (CPCs) and exactly six dedicated Secure Service Container (SSC) LPARs. For a multi-node setup, your IBM Z requires 30 or more Integrated Facilities for Linux (IFLs).

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.

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**.
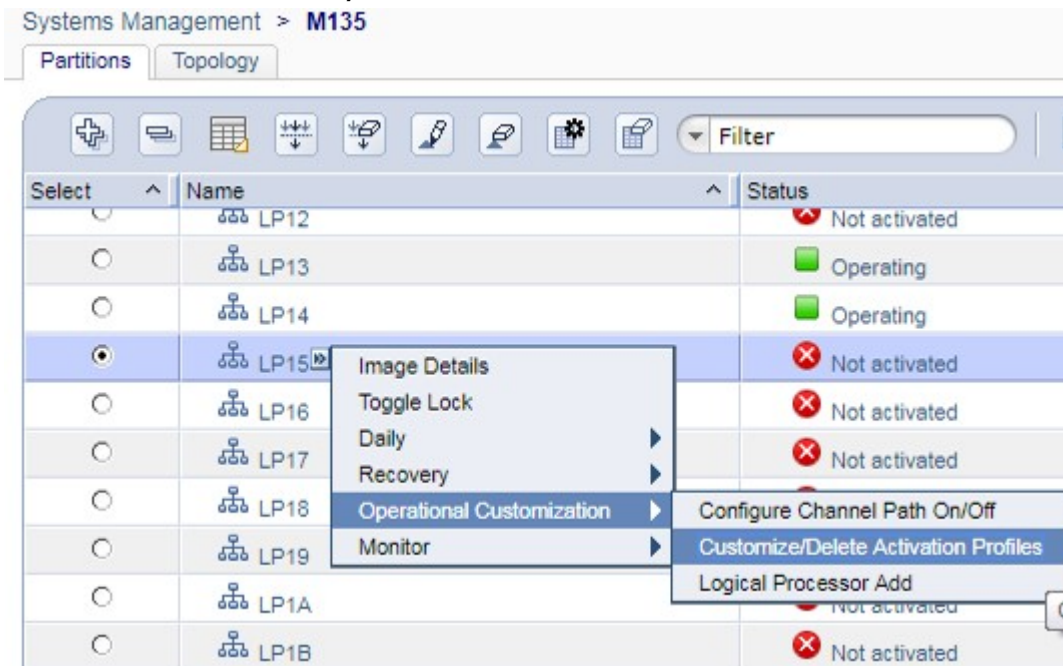
Figure 30. The Partitions window of the HMC

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.
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**:
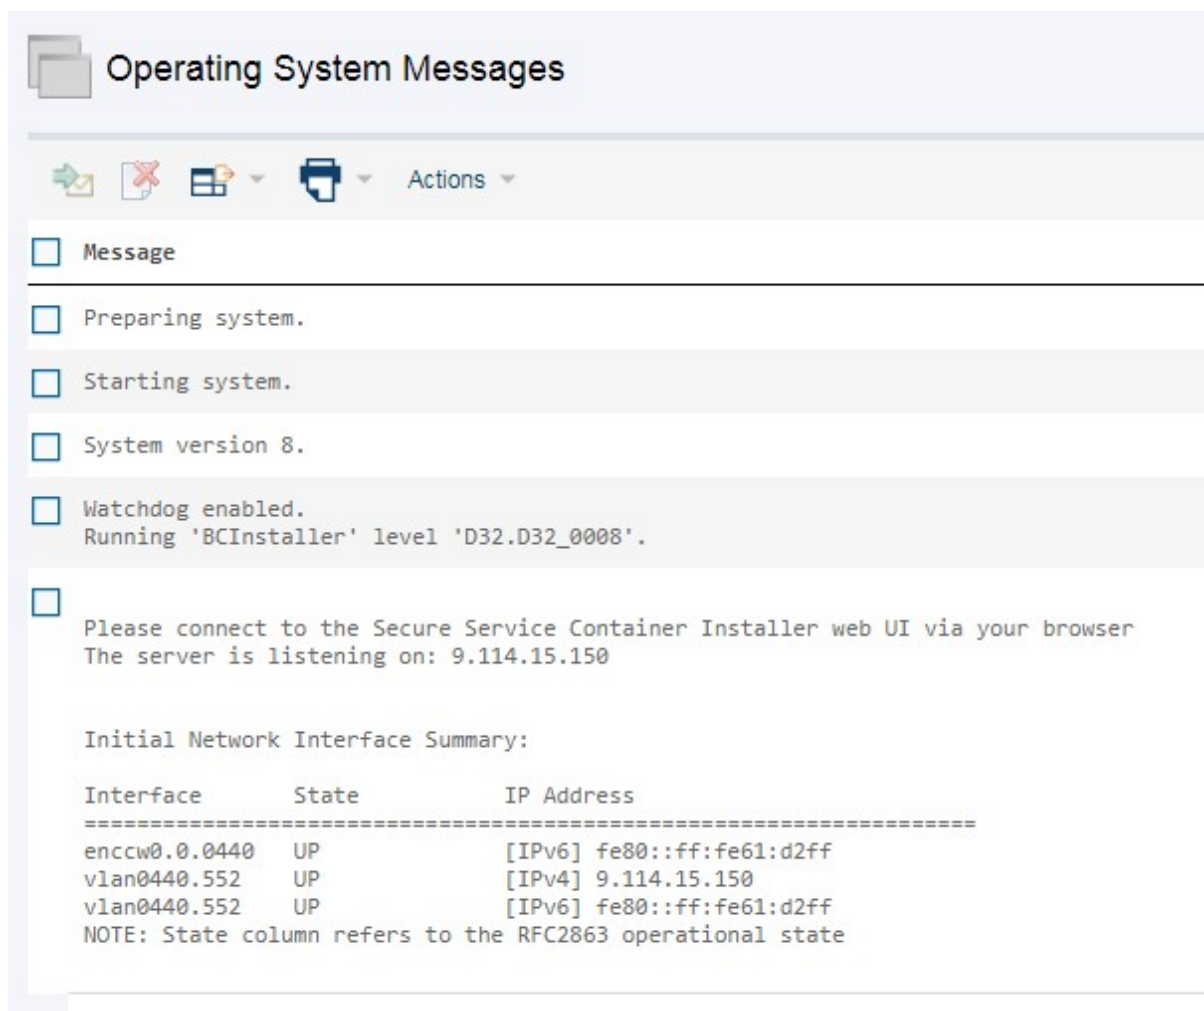
Figure 31. Confirmation: Operating System Messages

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 **Appliance Installer**.

Before you begin

Storage must be provided for various storage pools. Table 1 on page 60 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.

Table 2. 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 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 88.
2. Log on to the Appliance Installer and proceed to the **Welcome** page.

   For a description, see "Logging on to the Appliance Installer" on page 56.
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:
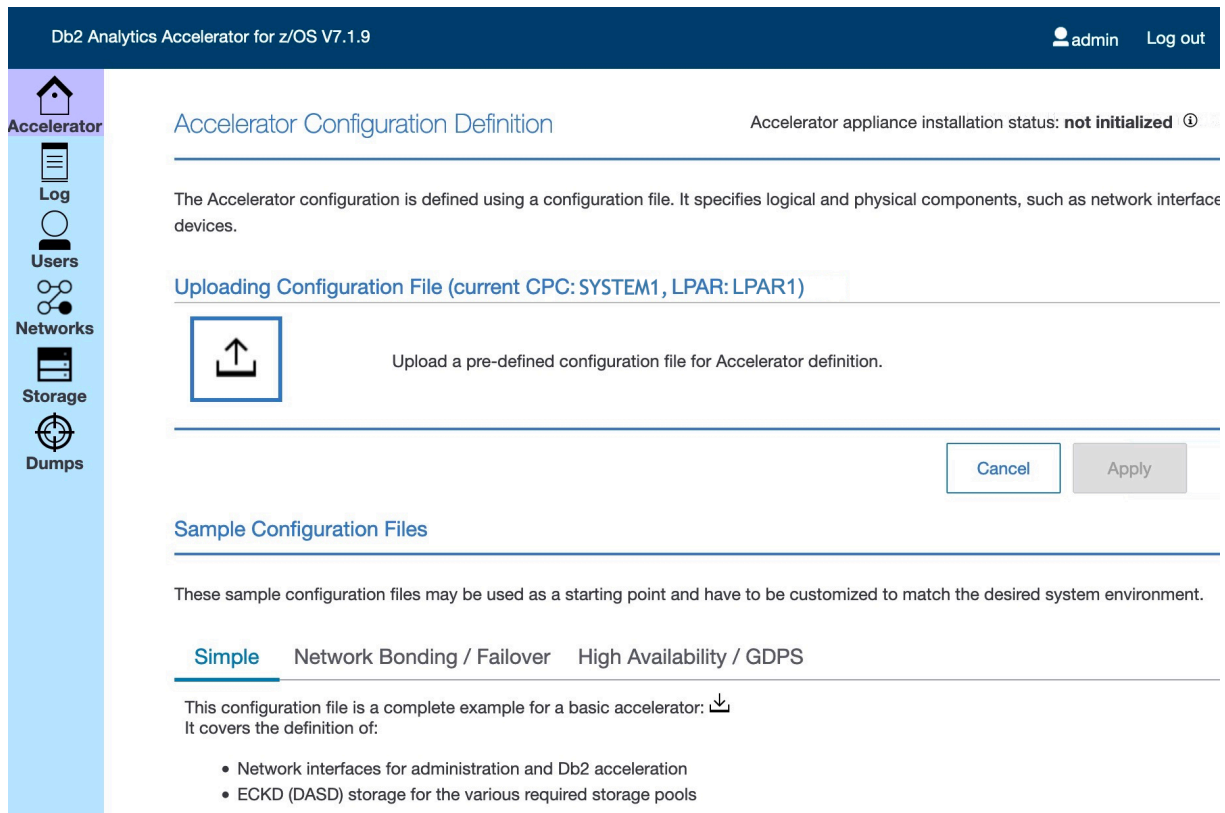
Figure 32. ***Accelerator Configuration Definition*** *page*

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.

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:

```
{
  "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": {
    "mgmt_nw": "activation-profile",
    "db2_nw": "db2_nw_connection",
    "gdps_nw": "gdps_nw_connection"
  },
  "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": "10.101.31.81/24",
              "device": "0.0.4206",
              "port": "0"
            }
          ]
        },
        {
          "cpc_name": "SYSTEM2",
          "lpar_name": "LPAR2",
          "network_interfaces": [
            {
              "name": "db2_nw_connection",
              "device": "0.0.4200",
              "port": "0"
            },
            {
              "name": "gdps_nw_connection",
              "ipv4": "10.101.31.82/24",
              "device": "0.0.4203"
            }
          ]
        }
      ],
      "primary_storage": {
        "boot_device": {
          "type": "dasd",
          "device": "0.0.9d00"
        },
        "runtime_devices": {
          "type": "dasd",
          "devices": [
            "0.0.9d01"
          ]
        },
        "data_devices": {
          "type": "dasd",
          "devices": [
            [
              "0.0.9d02",
              "0.0.9d0b"
            ]
          ]
        }
      }
    },
    "storage_maps": "GDPS",
    "gdps_servers": {
      "server1": {
        "ipv4": "9.1.2.3",
        "port": "5529"
      }
    }
  }
}
```

Most attributes in the configuration file are required.

"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 administration user interface 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. 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. The value is the interface name of the alternative data network in a failover scenario.

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 runtime environments for the accelerator (SSC) LPARs involved in the GDPS setup. 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"
          }
        ]
    },
    {
        "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"
          }
        ]
    }
  ]
```

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.

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"
  },
  "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"}]
        }
      ]
    }
  ],
```

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:

```
"network_interfaces": [
        {
          "name": "db2_nw_connection",
          "vlan": "700",
          "bond_settings": {
            "mode": "active-backup",
            "slaves": [
              {
                "device": "0.0.0a00",
                "port": "0"
              },
              {
```

```
                            "device": "0.0.1b00",
                            "port": "1"
                    }
                ]
            }
```

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 slave 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.

"primary"
Valid in *active-backup* mode only. The primary slave device. 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 slave device is selected after a failure. Specify "always", which means that an attempt will be made to make the primary device 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 slave is only changed if the currently active device fails.

"failover-MAC": "none"
Valid in *active-backup* mode only. Allows you set all slaves 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 slave 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 slave 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 slave 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 slave 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 slave 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 slave 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 slave 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"
}
```

"primary_storage" (required)
This block lists all storage devices for the primary site, that is, the site on which accelerator is initially deployed and configured. 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. 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. The boot device must be a single device with at least 40 GB net storage capacity.

"runtime_devices" (required)

The runtime storage is used by the Docker container that runs the accelerator software. Specify a list of devices with a total net capacity of at least 80 GB.

"data_devices" (required)

The data storage is used to store the operative 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.

"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). It is not possible to mix ECKD and SCSI devices in a single device category or device pool.

Example (ECKD or "dasd"):

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

Example (SCSI or "zfcp"):

```
"type": "zfcp",
        "devices": [
            {
            "device": "7a01",
            "lun": "0x4001012384729172",
            "wwpn": "0xc00984738927u834"
            }
        ]
```

Notes:

- You can specify lists of SCSI devices in much the same way as you specify lists of ECKD devices.

- However, because the logical unit numbers (LUNs) and worldwide port names (WWPNs) are structured the way they are, you cannot specify ranges. That is, you have to reuse the keywords `"device"`, `"lun"` and `"wwpn"` for each additional device.

"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:

```
"primary_storage": {
    "boot_device": {
      "type": "dasd",
      "device": "0.0.9d00"
    },
    "runtime_devices": {
      "type": "dasd",
      "devices": [
        "0.0.9d01"
      ]
    },
    "data_devices": {
      "type": "dasd",
      "devices": [
        "0.0.9d02",
        "0.0.9d0b"
      ]
    }
  }
}
```

Best practices for storage definitions:

For improved storage access and formatting of FICON-attached storage devices (this does *not* apply to FCP-attached (SCSI) devices!), your storage administrator might have defined a set of HyperPAV aliases. Use the automatic HyperPAV mode to include your HyperPAV aliases in the storage setup.

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 primary storage configuration section in the JSON configuration file.

Important:

- Make sure that only the volumes and HyperPAV aliases you want to use on a particular LPAR are visible to that LPAR. This important because in automatic mode, your accelerator (SSC) LPAR has to sift through all visible devices to determine and activate the alias devices.
- In a GDPS setup, it is not possible to list the HyperPAV aliases individually. You are limited to the automatic mode, which results in the use of all available HyperPAV devices.

Example:

```
"primary_storage": {
    "boot_device": {
            .
            .
    },
    "runtime_devices": {
            .
            .
    },
    "data_devices": {
            .
            .
    },
    "data": {
            "device": [
                "0.0.121f"
            ],
            "type": "dasd"
    },
    "hyperpav": "auto"
}
```

Instead of `"auto"`, you can use `"automatic"`. Both keywords will be accepted.

`"storage_maps"` (required)
Set this to `"GDPS"`.

`"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_servers": {
    "server1": {
        "ipv4": "9.1.2.3",
        "port": "5529"
    }
}
```

In this example, only one K-system is specified. You can specify more K-systems if needed.

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

On the , under the heading **Uploading Configuration File**, click the upload button:
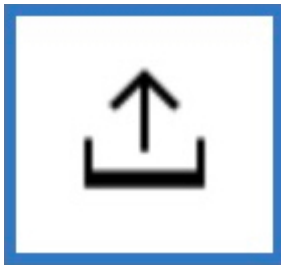
If something is wrong with the file you uploaded, an error message is displayed on the page:
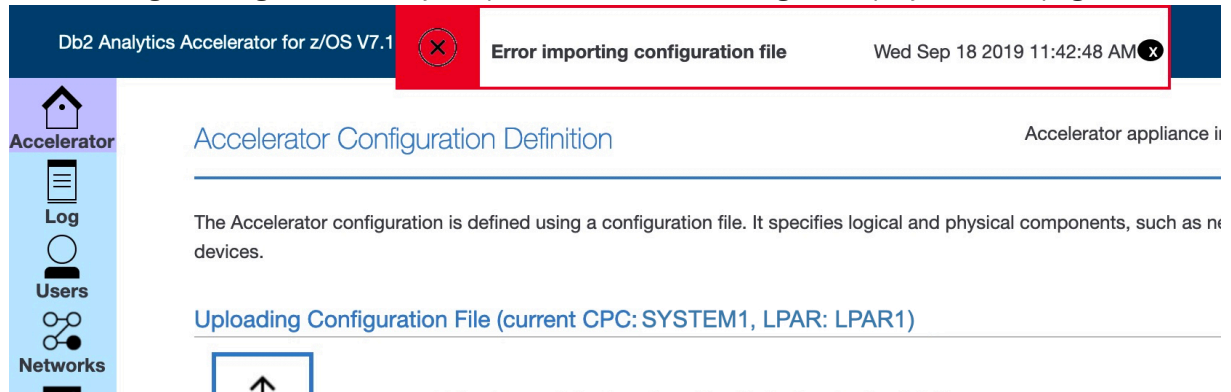
Figure 33. Error message after uploading a faulty configuration file

You can click the error message to display further information about the error:
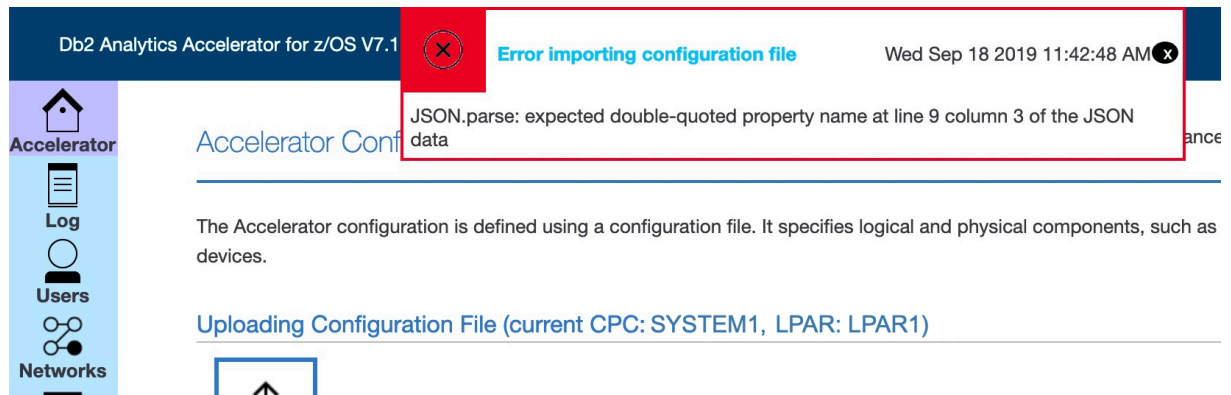
Figure 34. Error message expanded

9. If errors occurred, fix these and repeat the upload (steps <u>"7" on page 92</u> and <u>"8" on page 100</u>).

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

Results

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

*Figure 37. The **Appliance Component Health** page is displayed after a successful configuration*

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

Related information

Chapter 7, Configure Network Bonding, Red Hat Enterprise Linux 7: Networking Guide

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

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

Before you begin
The following requirements must be met for a multi-node installation with GDPS:

- 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 IFLs 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 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 "Defining an LPAR for Db2 Analytics Accelerator on Z with GDPS failover support" on page 88.
2. Log on to the Appliance Installer and proceed to the **Welcome** page.

   For a description, see "Logging on to the Appliance Installer" on page 56.
3. On the **Welcome** page, click **First-Time Setup**.

   You see the following page:

Figure 38. **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 with GDPS, 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.

   "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. Set this attribute to the value `"multi-node"`.

   "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 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 head node is started.

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.

"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 administration user interface 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. 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`.

"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`.

"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.

An example of the `"network_interface_bindings"` block:

```
"network_interface_bindings": {
    "mgmt_nw": "activation-profile",
    "db2_nw": "my_db2_network",
    "cluster_nw": "my_hiper",
    "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 for the accelerator (SSC) 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 the Hipersocket network of your cluster. Specify the following attributes to identify an accelerator (SSC) LPAR:

"cpc_name"
The name of a CPC. In a setup with GDPS, you have at least two CPCs. This means that you must define more than runtime environment.

Example:

```
{
     "cpc_name": "CPC001",
     "head": ...,
     "data1": ...,
     "data2": ...,
     "data3": ...,
     "data4": ...,
     "data5": ...
  },
  {
     "cpc_name": "CPC002",
     "head": ...,
     "data1": ...,
     "data2": ...,
     "data3": ...,
     "data4": ...,
     "data5": ...
  }
],
```

"head"
The definition of a head node. In a setup with GDPS, you have at least two head nodes because there are at least two CPCs. For example:

```
"head": {
     "lpar_name": "LPAR0",
     "network_interfaces": [
       {
         "name": "my_hiper",
         "device": "0.0.7f00",
         "ipv4": "172.84.0.180/23"
       },
       {
         "name": "my_db2_network",
         "device": "0.0.4b00"
       }
     ]
},
```

"lpar_name"
The name of an accelerator (SSC) LPAR. You have six such LPARs in the cluster and another six LPARs at the GDPS failover site.

"network_interfaces" (required)
This keyword introduces the block of network interface definitions for an 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.

"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".

"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:

```
"runtime_environments": [
    {
      "cpc_name": "IBMZ1",
      "head": {
        "lpar_name": "LPAR0",
        "network_interfaces": [
          {
            "name": "my_hiper",
            "device": "0.0.7f00",
            "ipv4": "172.84.0.180/23"
          },
          {
            "name": "my_db2_network",
            "device": "0.0.4b00"
          }
        ]
    },
```

"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.

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"
    },
    "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"}]
          }
        ]
      }
    ],
```

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:

```
"network_interfaces": [
    {
    "name": "db2_nw_connection",
    "vlan": "700",
    "bond_settings": {
    "mode": "active-backup",
    "slaves": [
            {
            "device": "0.0.0a00",
            "port": "0"
            },
            {
            "device": "0.0.1b00",
            "port": "1"
            }
        ]
    }
```

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 slave 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.

"primary"

Valid in *active-backup* mode only. The primary slave device. 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 slave device is selected after a failure. Specify `"always"`, which means that an attempt will be made to make the primary device 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 slave is only changed if the currently active device fails.

"failover-MAC": "none"
Valid in *active-backup* mode only. Allows you set all slaves 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 slave 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 slave 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 slave 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 slave 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 slave 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 slave 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 slave 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"
}
```

"data1"

The definition of the first data node. It works in the same way as for a head node.

Example:

```
"data1": {
        "lpar_name": "LPAR1",
        "network_interfaces": [
          {
            "name": "my_hiper",
            "ipv4": "172.84.0.181/23",
            "device": "0.0.7f00"
          }
        ]
},
```

Specify the other data nodes in the same manner, that is, create the sections "data2", "data3", "data4", and "data5"

Example:

```
"runtime_environments": [
    {
      "cpc_name": "IBMZ1",
      "head": {
        "lpar_name": "LPAR0",
        "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",
        "network_interfaces": [
          {
            "name": "my_hiper",
            "ipv4": "172.84.0.181/23",
            "device": "0.0.7f00"
          }
        ]
      },
      "data2": {
        "lpar_name": "LPAR2",
        "network_interfaces": [
          {
            "name": "my_hiper",
            "ipv4": "172.84.0.182/23",
            "device": "0.0.7f00"
          }
        ]
      },
      "data3": {
        "lpar_name": "LPAR3",
        "network_interfaces": [
          {
            "name": "my_hiper",
            "ipv4": "172.84.0.183/23",
            "device": "0.0.7f00"
          }
        ]
```

```
      },
      "data4": {
        "lpar_name": "LPAR4",
        "network_interfaces": [
          {
            "name": "my_hiper",
            "ipv4": "172.84.0.184/23",
            "device": "0.0.7f00"
          }
        ]
      },
      "data5": {
        "lpar_name": "LPAR5",
        "network_interfaces": [
          {
            "name": "my_hiper",
            "ipv4": "172.84.0.185/23",
            "device": "0.0.7f00"
          }
        ]
      }
    }
  },
  "cpc_name": "IBMZ2",
    "head": {
      "lpar_name": "LPAR6",
      "network_interfaces": [
        .
        .
      "data1": {
        "lpar_name": "LPAR7",
        "network_interfaces": [
        .
        .
        ],
```

"primary_storage" (required)

This block lists all storage devices for the primary site, that is, the site which is usually the active site. You need not specify storage devices for the GDPS failover site. During the first-time deployment, these devices are formatted, which means that the existing data on these devices is erased.

"head" (required)

This section lists the storage devices for a 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. 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. Specify a list of devices with a total net capacity of at least 80 GB for each node.

"data_devices" (required)

A data device is used to store operative 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.

"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). It is not possible to mix device types, that is, use ECKD and SCSI devices, in a single device category or device pool.

Example (ECKD or "dasd"):

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

Example (SCSI or "zfcp"):

```
"type": "zfcp",
            "devices": [
              {
                "device": "7a01",
                "lun": "0x4001012384729172",
                "wwpn": "0xc00984738927u834"
              }
            ]
```

Notes:

- You can specify lists of SCSI devices in much the same way as you specify lists of ECKD devices.
- However, because the logical unit numbers (LUNs) and worldwide port names (WWPNs) are structured the way they are, you cannot specify ranges. That is, you have to reuse the keywords "device", "lun" and "wwpn" for each additional device.

"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) on each node.

Example:

```
"primary_storage": {
    "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"
        ]
      }
    }
}
```

Best practices for storage definitions:

For improved storage access and formatting of FICON-attached storage devices (this does *not* apply to FCP-attached (SCSI) devices!), your storage administrator might have defined a set of HyperPAV aliases. Use the automatic HyperPAV mode to include your HyperPAV aliases in the storage setup.

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 primary storage configuration section in the JSON configuration file.

Important:

- Make sure that only the volumes and HyperPAV aliases you want to use on a particular LPAR are visible to that LPAR. This important because in automatic mode, your accelerator (SSC) LPAR has to sift through all visible devices to determine and activate the alias devices.
- In a GDPS setup, it is not possible to list the HyperPAV aliases individually. You are limited to the automatic mode, which results in the use of all available HyperPAV devices.

Example:

```
"primary_storage": {
    "boot_device": {
            .
            .
    },
    "runtime_devices": {
            .
            .
    },
    "data_devices": {
            .
            .
    },
    "data": {
            "device": [
                "0.0.121f"
            ],
            "type": "dasd"
    },
    "hyperpav": "auto"
}
```

Instead of `"auto"`, you can use `"automatic"`. Both keywords will be accepted.

"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 (`"primary_storage"`).

Example of storage definition (partly shortened):

```
"primary_storage": {
    "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" ]
        ]
      },
    },
    "data1": {
      "boot_device": ...,
      "data_devices": ...,
      "runtime_devices": ...
    },
    "data2": {
      "boot_device": ...,
      "data_devices": ...,
      "runtime_devices": ...
    },
    "data3": {
      "boot_device": ...,
      "data_devices": ...,
      "runtime_devices": ...
    },
    "data4": {
      "boot_device": ...,
```

```
        "data_devices": ...,
        "runtime_devices": ...
    },
    "data5": {
        "boot_device": ...,
        "data_devices": ...,
        "runtime_devices": ...
    }
}
```

"storage_maps" (required)
> Set this to "GDPS".

"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_servers": {
    "server1": {
        "ipv4": "10.2.1.11",
        "port": "1020"
    },
    "server2": {
        "ipv4": "10.2.2.11",
        "port": "1030"
    }
}
```

> In this example, two K-systems are specified. You can specify more K-systems if needed.

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

   On the , under the heading **Uploading Configuration File**, click the upload button:

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

Figure 39. Error message after uploading a faulty configuration file

   You can click the error message to display further information about the error:
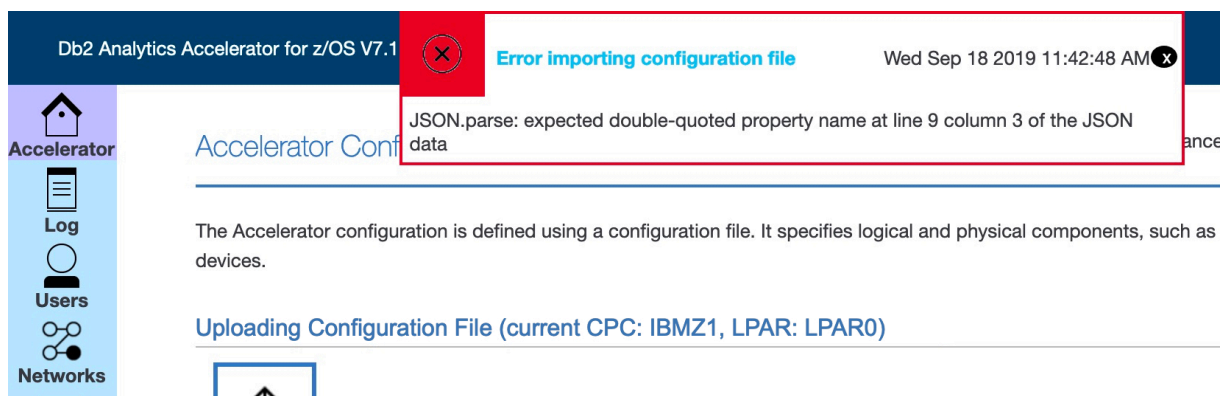
Figure 40. Error message expanded

7. If errors occurred, fix these and repeat the upload (steps "5" on page 105 and "6" on page 115).

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.
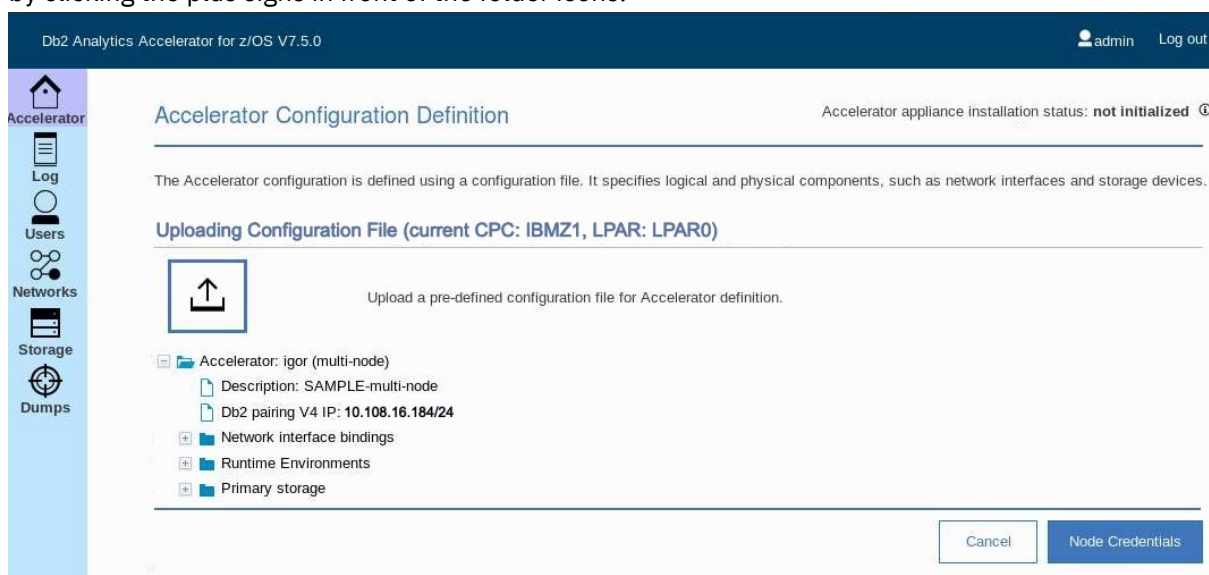
Figure 41. Accelerator Configuration Definition after a successful configuration file upload

8. Click **Node Credentials**.

9. In the **Multiple Node Deployment Credentials** window, enter the IP addresses of the data-node LPARs. Also provide the user ID and the password of each accelerator (SSC) LPAR, as defined in 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 Figure 42 on page 117.

*Figure 42. The **Multiple Node Deployment Credentials** window*

10. Click **Validate**.

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

Figure 43. Message window showing the progress of your Db2 Analytics Accelerator on Z configuration

Results

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

*Figure 44. The **Appliance Component Health** page is displayed after a successful configuration*

The message `Accelerator appliance installation status: installation completed` 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

- To shut down a cluster, click the **Shutdown** button on the **Accelerator Component Health Status** page of the Appliance Installer.
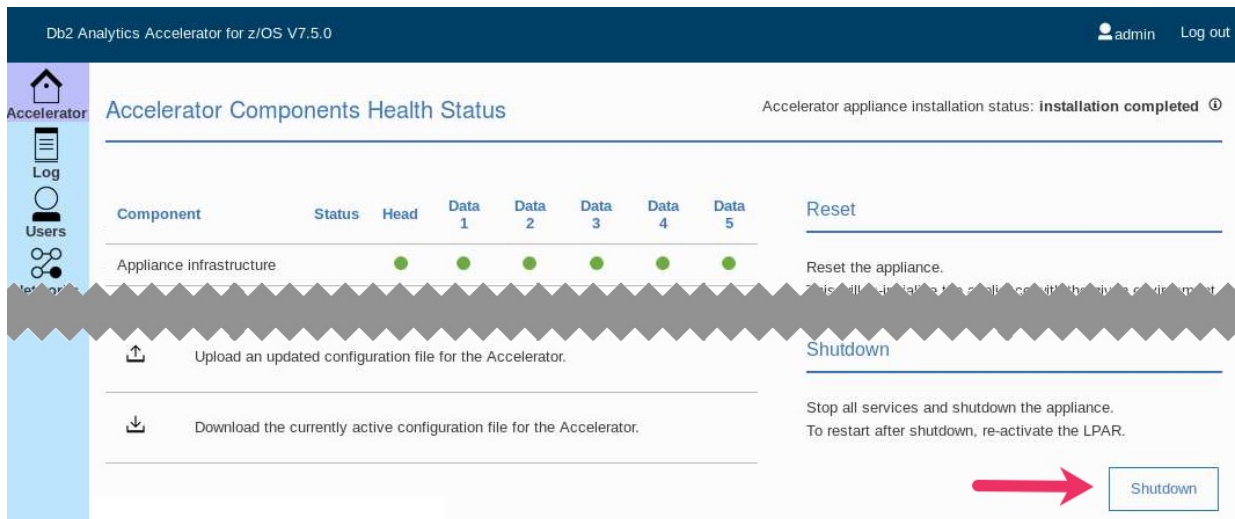
Figure 45. Shutdown from the Accelerator Component Health Status page

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

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.

To copy the data on existing storage devices to new devices, you can use the "storage_maps" parameter in the JSON configuration file. This will replicate all the data to the new device. The It is the same mechanism you would use for a failover setup, except that the shift is never reversed.

Example:

```
"storage_maps": [
    {
      "boot_device": "0.0.9c80",
      "map": [
        {
          "primary": ["0.0.9d00","0.0.9d0b"],
          "copy":    ["0.0.9c80","0.0.9c8b"]
        }
      ]
    }
  ]
```

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 10 on page 55).
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 10 on page 55)

   You see the **Appliance Component Health** page.

*Figure 46. The **Appliance Component Health** page*

3. You might want to use the existing configuration file as the basis for your modifications. If so, download it by clicking ⬇ **Download the currently active configuration file for the Accelerator**. Navigate to the file location of your choice and save the file.

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:
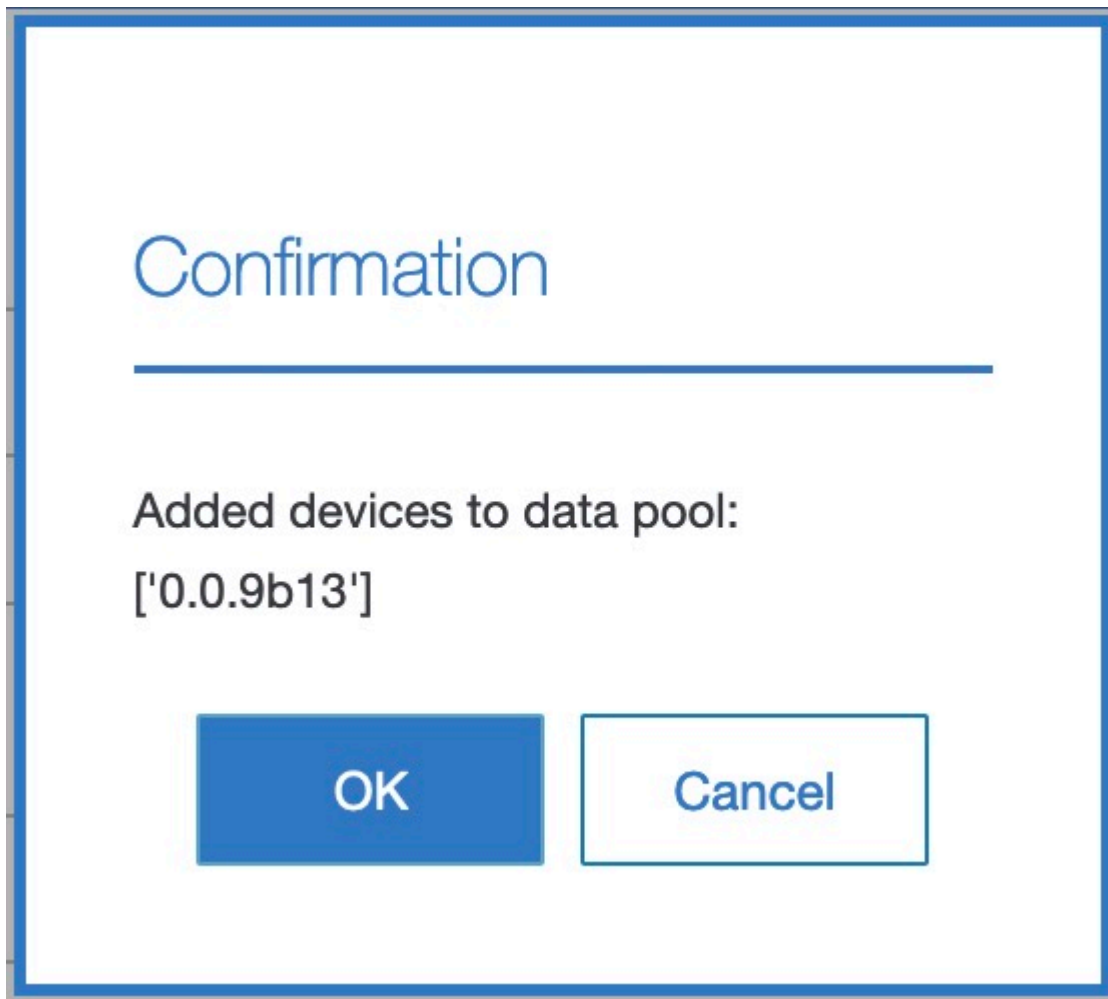
Figure 47. **Confirmation** *window*

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 Appliance Installer. The **Log** window lists recent events, as in the following example:
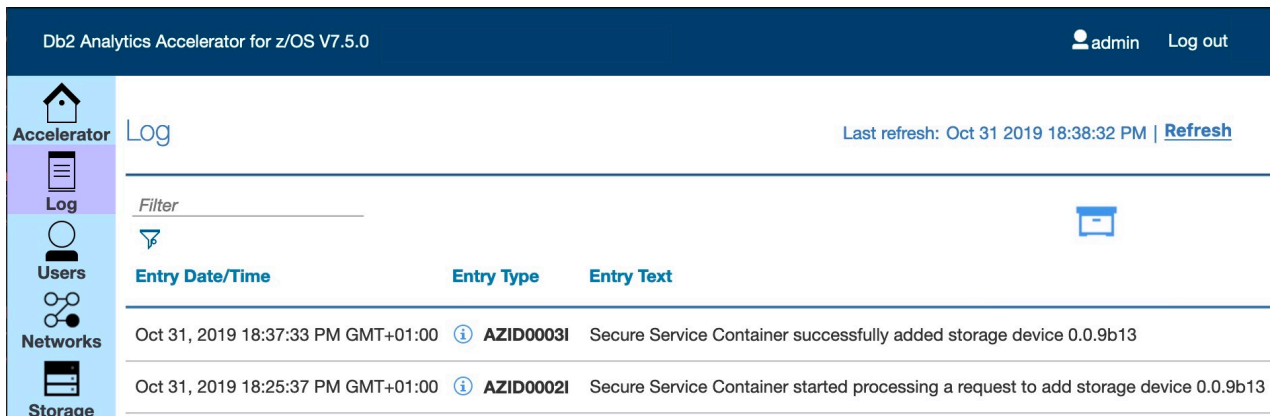
Figure 48. The Log window of the Appliance Installer

The example in Figure 48 on page 122 shows that the storage device 0.0.9b13 has been added at 18:25 on Oct. 31. The formatting was finished at 18:37. The added storage has been in use since then.

Creating a trace file (dump)

If you encounter installlation 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. In the Appliance Installer, such a trace file is called a dump. This section describes how to create a dump.

Procedure

1. Open the Appliance Installer in a web browser. Enter the IPv4 address of the Secure Service Container (SSC) LPAR (see Figure 10 on page 55)
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 10 on page 55)

   You see the **Welcome** page.
3. In the navigator on the left, click **Dumps** to show the **Dumps** page.
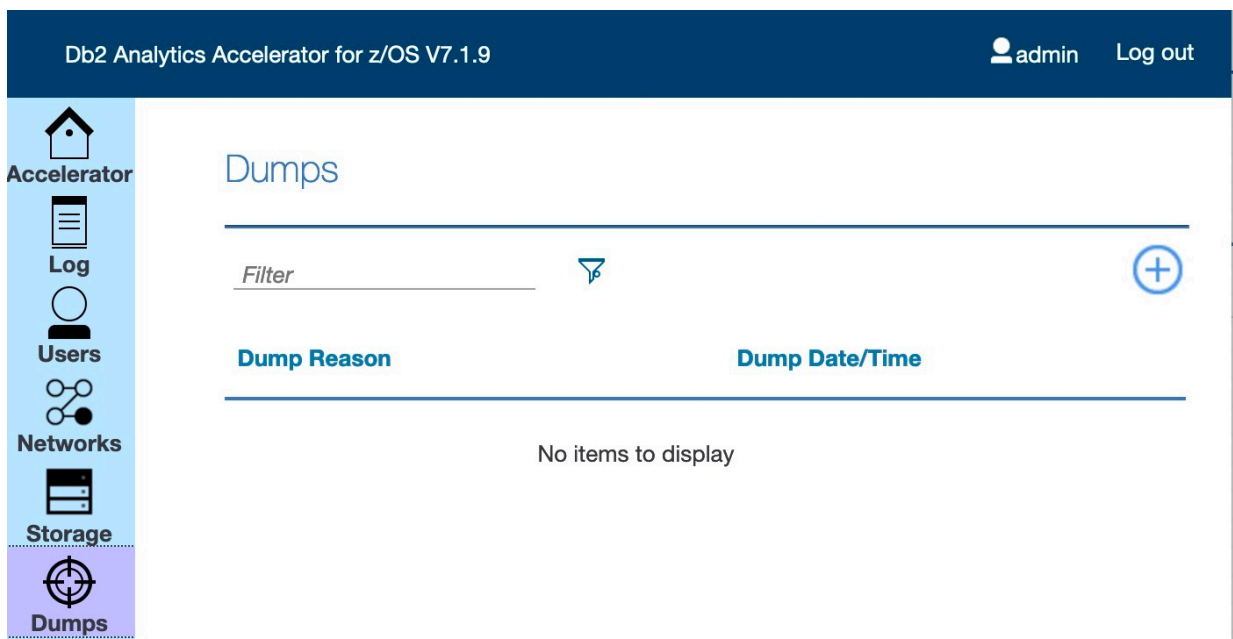
Figure 49. The Dumps page

4. Click ⊕ .

   You see the following page:

   | Db2 Analytics Accelerator for z/OS V7.1.9 | 👤admin   Log out |

   Accelerator

   Log

   Users

   Networks

   Storage

   Dumps

   # Create Dump

   You can create a dump to report detailed system information, if requested by a service representative. The dump reason is a mandatory description text.

   ⓘ Dump Type:

   ⦿ Concurrent

   ○ Disruptive

   Dump Reason*

   Characters left: 200

   Cancel   Create Dump

Figure 50. 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.

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 IBM Db2 Analytics Accelerator Studio
You might want to enable your power user to run IBM Db2 Analytics Accelerator Studio. If so, give the power user the following authorizations in Db2 for z/OS:

- SELECT on the DSNAQT.ACCEL_NAMES view. This privilege is required for the enumeration of accelerators.
- SELECT on the SYSACCEL* tables. This privilege is required to associate the tables with an accelerator pairing code.
- SELECT on the following catalog tables of the database management system:

  - SYSIBM.SYSCOLUMNS
  - SYSIBM.SYSCONTROLS
  - SYSIBM.SYSDATABASE
  - SYSIBM.SYSDUMMY1
  - SYSIBM.SYSINDEXES
  - SYSIBM.SYSRELS
  - SYSIBM.SYSTABLEPART
  - SYSIBM.SYSTABLES
  - SYSIBM.SYSVIEWS

  Read access to these tables is required for the creation of accelerator-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 IBM Db2 Analytics Accelerator Studio
The recommended minimum set of authorizations for IBM Db2 Analytics Accelerator Studio 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_3.1.1/
   com.ibm.datatools.qrytune.configothers.doc/topics/enabledb2zfromclient_ds.html

2. To grant the EXECUTE privilege to other users, proceed as follows:

   a) Select **Analyze and Tune** > **Configure for Tuning** > **Advanced Configuration and Privilege Management** as described on the following website:

   https://www.ibm.com/support/knowledgecenter/SS7L9Q_4.1.0/
   com.ibm.datatools.qrytune.workloadtune.doc/topics/verify.html

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

Creating EXPLAIN tables

To be able to display an access plan graph of your queries (accelerated or not) in IBM Db2 Analytics Accelerator Studio, you must create certain EXPLAIN tables in the Db2 subsystems that you use.

About this task
The instructions on the following web pages were originally written for IBM Query Tuner. However, the steps can be applied to IBM Db2 Analytics Accelerator Studio likewise.

Procedure

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

   http://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:

http://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 New Accelerator** wizard. However, you can only start this wizard after connecting to a Db2 subsystem.

About this task
In IBM Db2 Analytics Accelerator Studio, the connection information is stored in profiles for reuse. Having created a profile, you can reconnect to a database by double-clicking the icon representing the profile in the **Administration Explorer**.

Procedure

1. Start IBM Db2 Analytics Accelerator Studio.
2. If the **Welcome** screen is displayed, close it.
3. On the header of the **Administration Explorer** on the left, click the downward-pointing arrow next to **New** and select **New Connection to a Database**.
4. In the **New Connection** window, decide how to name the database connection profile:
   - To use the name of the database server that you want to connect to, leave **Use default naming convention** selected.
   - To choose a different name, clear **Use default naming convention**, and type the name in the **Connection Name** field.
5. From the **Select a database manager** list, select **Db2 for z/OS**.
6. Make sure that in the **JDBC driver** drop-down list, **IBM Data Server Driver for JDBC and SQLJ (JDBC 4.0) Default** is selected.

7. In the **Location** field, type the name of the database server that you want to connect to.

   Tip: To determine the **Location**, **Host** name, and **Port number**, a Db2 for z/OS systems programmer or database administrator can issue a DIS DDF command.

8. In the **Host** field, type the host name or IP address of the data server on which the database server is located.

9. In the **Port number** field, you see that port number 446 is selected by default. Leave this setting unless the database server uses another port.

10. Select **Retrieve objects created by this user only** if you want to restrict database access to the databases, schemas, tables, and other objects that were created by the logon user.

    If you do not select this option (default), IBM Db2 Analytics Accelerator Studio will show and make selectable all databases, schemas, and tables that the logon user has access to, including those to which this user might have only read access.

11. In the **User name** field, type the user ID that you want to use to log on to the database server.

    Note that you can only use IBM Db2 Analytics Accelerator Studio successfully if this user has sufficient rights to run the stored procedures behind the IBM Db2 Analytics Accelerator Studio functions. The section *Appendix C. Required access rights* in the *IBM Db2 Analytics Accelerator for z/OS: Stored Procedures Reference* lists the privileges that are required to run a particular stored procedure. If you are uncertain, use an ID with SYSADMIN authority.

    In many organizations, it is a common practice to have personal user IDs with restricted authority and special-purpose user IDs (groups in most cases) with extensive privileges in a certain field. IBM Db2 Analytics Accelerator Studio supports this practice in that you can specify a secondary user ID, which might have the privileges that your logon user ID lacks, such as the privilege to run stored procedures. If the secondary ID is a group user ID, the logon user must of course be a member of that group. To specify a secondary user ID, follow these steps:

    a. On the **Connection Parameters** page, click the **Optional** tab.

    b. In the **Property** field, type the following statement:

    ```
    currentSQLID
    ```

    c. In the **Value** field, type the secondary user ID.

    d. Click **Add**.

    e. Click the **General** tab to return to that page and complete the logon.

12. In the **Password** field, type the password belonging to the logon user ID.

13. Leave the **Save password** check box deselected.

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

14. Leave the **Default schema** field blank.

15. Click **Test Connection** to check if you can log on to the database server.

16. Click **Finish**.

Results

After creating the profile, IBM Db2 Analytics Accelerator Studio automatically connects to the Db2 subsystem.

What to do next

If it takes too long to load all objects of the Db2 subsystem into the **Administration Explorer** (more than one minute), you can set a filter to limit the number of schemas to be loaded:

1. In the **Administration Explorer**, right-click the icon representing the Db2 subsystem (database symbol).

2. Select **Properties** from the menu.

3. In the **Properties for ...** window, select **Default Schema Filter**.

4. Clear the **Disable filter** check box. This activates the filter controls.

5. From the **Name** drop-down list, select a suitable filter mask. In the adjacent text field, type the filter string. For example, to exclude all schemas whose names starts with the characters BLU:

   a. From the **Name** drop-down list, select **Does not start with the characters**.

   b. In the text field, type BLU.

6. Click **Apply**.
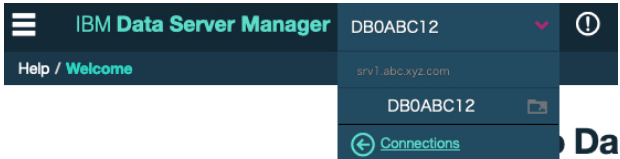
Related reference
Access rights for power users
The various IBM Db2 Analytics Accelerator for z/OS components require different authorizations. It is useful to create at least one power user with extensive authorizations, that is, a user who can run all IBM Db2 Analytics Accelerator for z/OS functions 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: STLEC1
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: Alternatively: a. Click and select **SETTINGS** > **Manage Connections**. b. Select the box in front of a listed connection.
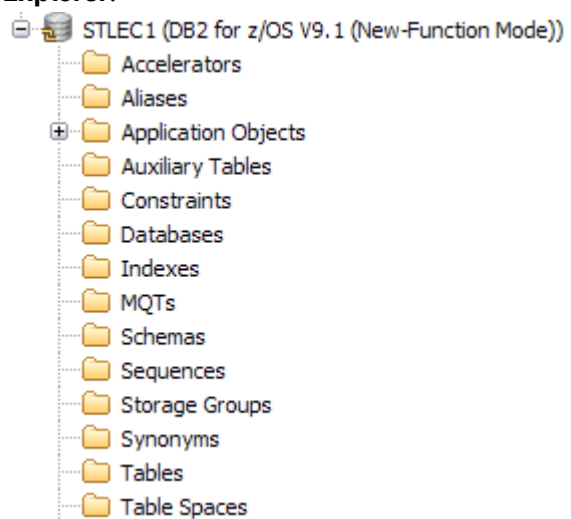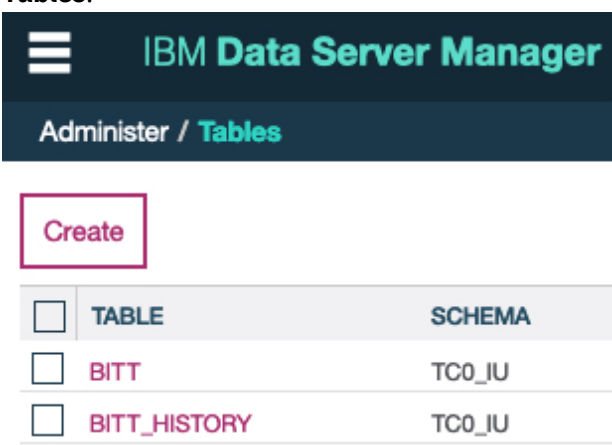
Results
You see the following results if the connection was successful:

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

IBM Db2 Analytics Accelerator Studio	After a successful connection, you can see the following folder hierarchy in the **Administration Explorer**:
IBM Data Server Manager	Tables in the connected Db2 subsystem are displayed after selecting ▤ > **ADMINISTER** > **Tables**:

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 IBM Db2 Analytics Accelerator Studio must coordinate the steps.

Procedure

1. Ask the network administrator or the person who did the TCP/IP setup for the IP address 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 173
   - "Using ssh to log on to the IBM Db2 Analytics Accelerator Console" on page 174

3. Press the Pause key, then Enter to display the following screen:

```
******************************************************************************
*          Welcome to the IBM Db2 Analytics Accelerator Console
******************************************************************************

You have the following options:
 (0) - (Menu) Manage Configuration Console
Users
 (1) - (Menu) Run Accelerator Functions
 (2) - (Menu) Manage Incremental Updates
 (3) - (Menu) Manage Encryption of Data in Motion
 (4) - (Menu) Manage Call Home


 ----------------------------------------------------------------------
 (x) - Exit the Configuration Console
```

4. Type 1 and press Enter to display the submenu:

```
main -> 1
-------------------------------------------------------
You have the following options:

(0) - Go back one level
(1) - Obtain pairing code, IP address, and port
(2) - List paired Db2 subsystems
(3) - Clear query history
(4) - Restart accelerator process
(5) - Dump extensive diagnostic information
(6) - Stop backend database
(7) - Start backend database
(8) - Transparently convert Db2/Z REAL columns into DOUBLE columns when loading a table
(9) - Set the DB2 subsystem for time synchronization
(10) - Rotate key used to generate service password
(11) - Accelerator Workload Management
```

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 IBM Db2 Analytics Accelerator Studio, start the **Add New Accelerator** wizard, and enter the information that is returned by the console. Values from 5 to 1440 are allowed. If you just press Enter, you accept the default of 30 minutes.

   ```
   Press <return> to accept the default of 30 minutes.
    Cancel the process by entering 0.


    Accelerator pairing information:
      Pairing code  : 6048
      IP address    : 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 L Pars.

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 IBM Db2 Analytics Accelerator Studio.
2. Select the **Accelerators** folder in the **Administration Explorer**.
3. On the menu bar of the **Object List Editor**, click the downward-pointing arrow next to the green plus sign.
4. From the drop-down menu, select **Add Accelerator**.
5. In the **Name** field, type a name for the accelerator.

   This name is automatically copied to the **Location** field.

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

   Restriction: An accelerator cannot be shared between two or more Db2 subsystems if the subsystems use the same location name. If you copy an entire subsystem, make sure to change the location name of the copy afterwards.
6. In the **Pairing code** field, type the pairing code.
7. In the **IP address** field, type the IP address of the accelerator.
8. In the **Port** field, type 1400. This is the fixed port for network communication between the z/OS data server and the accelerator.
9. Click **Test Connection** to check whether the accelerator with the given address can be connected to.
10. Click **OK**.

    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. In the **Administration Explorer** (upper left) of IBM Db2 Analytics Accelerator Studio, select the **Accelerators** folder.
2. In the list of the Object List Editor on the right, double-click the name of the accelerator.
   This opens the **Accelerator** view. It is empty by default, that is, it does not contain any tables.
3. Click the **Add** button.
4. In the **Add Tables** wizard, select the tables that are referenced in your test query.
5. Click **OK**.

Results

The newly created accelerator-shadow tables are shown in the **Accelerator** view.

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 IBM Db2 Analytics Accelerator Studio to hold the network connection until the process has finished.

An option that circumvents this problem is to invoke IBM Db2 Analytics Accelerator for z/OS stored procedures directly from JCL batch jobs. Such jobs can even be scheduled by using a third-party tool.

Procedure

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

   Note: The processes can take a couple of minutes to complete. If the load fails, first check whether DSNUTIL was started in Db2 for z/OS.

Running a test query

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

Before you begin

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

About this task

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

Procedure

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

   • SET CURRENT QUERY ACCELERATION = ELIGIBLE;

   Note: `<x>` stands for a counting number (integer). This means that your first SQL script is named `Script1.sql`, the second `Script2.sql`, and so on.
4. Type your query. Use a simple query, such as SELECT SUM(QUANTITY) FROM SALES.FIGURES_FACT.
5. Select **Script** > **Run SQL** from the main menu.

   The **SQL Results** pane in the lower left shows you whether the query ran successfully.
6. To check whether the accelerator was used, follow these steps:

   a) Return to the accelerator view by clicking the tab with the name of the accelerator.

   b) Click the twistie next to the heading **Query Monitoring**.

   The section that unfolds shows a table that lists the most recent queries. These are only the queries that were run on the accelerator. Inhouse Db2 queries are not listed. Your test query should appear at the top of the table. Your query might not be listed there for various reasons, for example:

   • The installation of IBM Db2 Analytics Accelerator for z/OS is incomplete.

   • The Db2 optimizer did not route the query to the accelerator because the query does not qualify.

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

   To determine the cause for the failure, use the Db2 EXPLAIN function and look up the reason code in DSN_QUERYINFO_TABLE. For more information, read the section *EXPLAIN information* in the *IBM Db2 Analytics Accelerator Studio: User's Guide*. You find a link under **Related information** at the end of this topic.

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.

If you followed the instructions in this guide, you already installed the necessary IBM InfoSphere® Data Replication for z/OS (CDC) components by completing the steps in "Installing libraries with IBM Db2 Analytics Accelerator support" on page 23. However, for updates or more detailed information, follow the link to the installation instructions in the IBM InfoSphere Data Replication 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 Data Replication 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 PTF UI63356 installed, running at function level V12R1M500
- 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, for example by using a member-specific secure port that has a local alias.
- 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.

- Tables on an accelerator have a Db2 attribute named DATA CAPTURE. The attribute can carry the value Y or N (default), for *yes* or *no*. To enable incremental updates for a table, you must set the DATA CAPTURE attribute to the value Y. Integrated synchronization differs from CDC in this respect. With CDC, the attribute is set automatically when incremental updates are enabled for a table. However, the fact that the attribute value persists, even if the table is disabled, still holds true. Bear this in mind, especially if you run applications that use the DATA CAPTURE attribute.

- 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/OS* in the *IBM Db2 Analytics Accelerator for z/OS: User's Guide*.

Restrictions

- Currently, IBM Integrated Synchronization does not support automatic failovers in high-availability environments.
- 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.

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, the connection must always go to the same member. The easiest way to achieve this is to define a unique secure port for each member by use of location aliases. With a configuration like this, all members used for connections can share the same DVIPA IP address. This makes the handling and setup of SSL certificates easier because a single, shared certificate can be used.

For that reason, you must define a secure port for every member that is supposed to serve as a connection endpoint for integrated synchronization. Again, there are different ways to achieve this:

- 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
-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.

Related information

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=*
//SYSUADSDD 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.
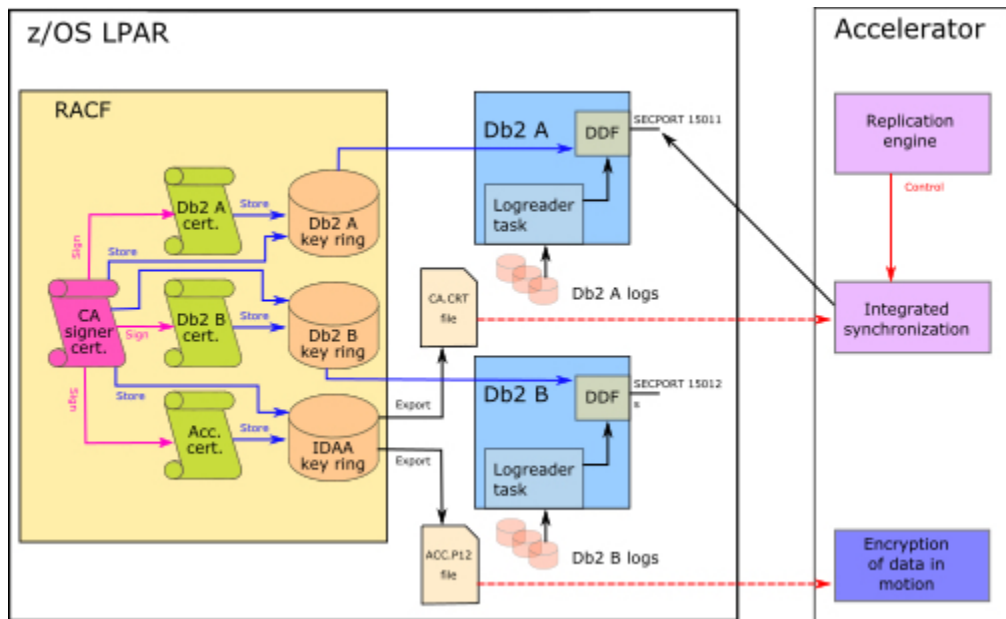
Figure 51. Your guide through the certificate jungle

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

1. 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 .

- 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:

```
RACDCERT ID(DB2USER) ADDRING(DB2AKEYRING)
RACDCERT ID(DB2USER) -
        CONNECT(CERTAUTH -
        LABEL('DB2 SERVER CA') RING(DB2AKEYRING))
RACDCERT ID(DB2USER) -
        CONNECT(ID(DB2USER) -
        LABEL('DB2ASERVER CERTIFICATE') -
        RING(DB2AKEYRING) DEFAULT)
SETR RACLIST (DIGTRING) REFRESH
SETR RACLIST (DIGTCERT) REFRESH
SETR RACLIST (FACILITY) REFRESH
```

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:

```
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:

```
#################################################
# Rules for SSL access to Db2 DDF
#################################################
TTLSRule DB11Rule15111
{
    LocalPortRange          15111
    JobName                 DB11DIST
    Direction               Inbound
    TTLSGroupActionRef      Db2SslGroup
    TTLSEnvironmentActionRef DB11SslEnv
}
```

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:

```
TTLSGroupAction Db2SslGroup
{
    TTLSEnabled       On
    CtraceClearText   On
}
```

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:

```
TTLSEnvironmentAction DB11SslEnv
{
    TTLSKeyRingParms
    {
        Keyring            DB11KEYRING
    }
    HandShakeRole          Server
    TTLSCipherParmsRef      Db2SslCipherParms
}
```

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 Db2SslCipherParms block might look like:

```
TTLSCipherParms Db2SslCipherParms
{
    V3CipherSuites      TLS_RSA_WITH_AES_128_GCM_SHA256
    V3CipherSuites      TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384
    V3CipherSuites      TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256
    V3CipherSuites      TLS_RSA_WITH_AES_128_CBC_SHA
}
```

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 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        DB2ASecurePorts
    JobName                  DB2ADIST
    Direction                Inbound
    TTLSGroupActionRef       Db2SslGroup
    TTLSEnvironmentActionRef DB2ASslEnv
}

PortGroup DB2ASecurePorts
{
  PortRange
  {
    Port 12000
  }
  PortRange
  {
    Port 15011
  }
}
```

If you compare this with the example in step , 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.

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 146 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 "1" on page 146 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 command:

   ```
   GRANT MONITOR2 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.

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 173
   - "Using ssh to log on to the IBM Db2 Analytics Accelerator Console" on page 174

3. Press the Pause key, then Enter to display the following screen:

```
**************************************************************************
*          Welcome to the IBM Db2 Analytics Accelerator Console
**************************************************************************

You have the following options:
 (0) - (Menu) Manage Configuration Console
Users
 (1) - (Menu) Run Accelerator Functions
 (2) - (Menu) Manage Incremental Updates
 (3) - (Menu) Manage Encryption of Data in Motion
 (4) - (Menu) Manage Call Home

 --------------------------------------------------------------------
 (x) - Exit the Configuration Console
```

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 138 for example.

b) `Available incremental update technologies:`

```
  1. Integrated Synchronization
  2. IBM InfoSphere Data Replication 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 138, 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 146 (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 148.

h) `Enter the password again to confirm (in TSO, use PF3 to hide input):`

Repeat the password and press Enter to confirm it.

i) `Select the certificate that identifies the DB2 subsystem:`

Import the signer certificate that you exported by following the steps in "Exporting and transferring the signer certificate" on page 143. 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.

j) `Use certificate 'CERTAUTH' (y/n) [n]:`

Confirm your certificate choice. Type y and press Enter.

k) `Press 'y' to configure 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 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 IBM Db2 Analytics Accelerator Studio.

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.

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:

```
DSNT788I  -DB11
 SESSIONID        STATUS     CURR. POSITION       NUM RECS   AGE
 ---------------- ---------- -------------------- ---------- ------
 164FA983947B0801 SUSP EOS   0000000000018898F744 4          167s
******* DISPLAY OF STATS TERMINATED ********************************
 DSN9022I  -DB11 DSNTDSTS 'DISPLAY STATS' NORMAL COMPLETION
```

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

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 173
   • "Using ssh to log on to the IBM Db2 Analytics Accelerator Console" on page 174

2. Press the Pause key, then Enter to display the following screen:

```
******************************************************************************
*          Welcome to the IBM Db2 Analytics Accelerator Console
******************************************************************************

You have the following options:
 (0) - (Menu) Manage Configuration Console
Users
 (1) - (Menu) Run Accelerator Functions
 (2) - (Menu) Manage Incremental Updates
 (3) - (Menu) Manage Encryption of Data in Motion
 (4) - (Menu) Manage Call Home

 ----------------------------------------------------------------------
 (x) - Exit the Configuration Console
```

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

(Default 0) >
```

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 "Assigning user IDs" on page 146 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 151.

   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 IBM Db2 Analytics

Accelerator Studio. Using the functional links in the header starts or stops incremental updates for an entire accelerator.

IBM InfoSphere Data Replication for z/OS

As the title suggests, this feature uses an external product, IBM InfoSphere Data Replication 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 Data Replication for z/OS (CDC) components by completing the steps in "Installing libraries with IBM Db2 Analytics Accelerator support" on page 23. However, for updates or more detailed information, follow the link to the installation instructions in the IBM InfoSphere Data Replication 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.

- Tables on an accelerator have a Db2 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.

- 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/OS* in 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.

Concepts and architecture

The incremental update function is deeply integrated into the IBM Db2 Analytics Accelerator for z/OS solution.

You administer incremental updates from IBM Db2 Analytics Accelerator Studio or call the corresponding IBM Db2 Analytics Accelerator for z/OS stored procedures directly.

Components

IBM InfoSphere Change Data Capture for Db2 for z/OS is required for an implementation of the incremental update function. For further details, refer to the list of IBM Db2 Analytics Accelerator for z/OS prerequisites on the web.

System architecture

IBM Db2 Analytics Accelerator Studio communicates with the IBM Db2 Analytics Accelerator for z/OS server process 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 52 on page 153).

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.

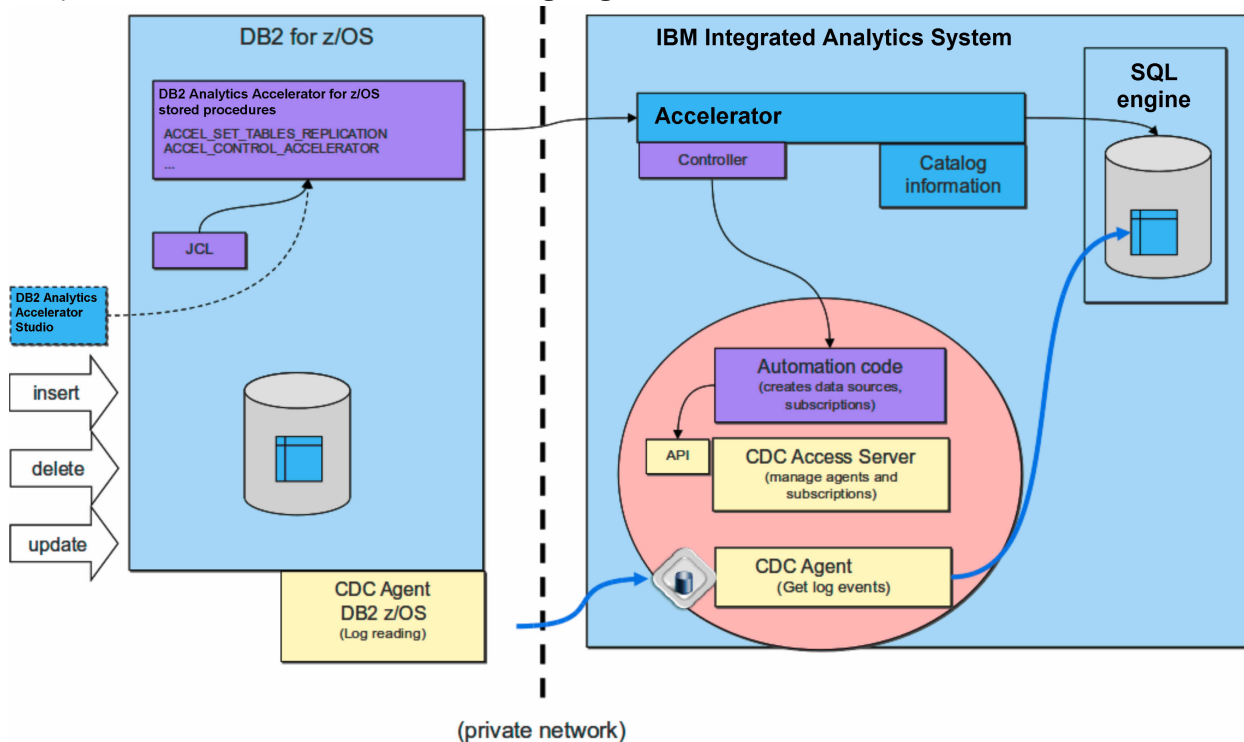The process flow is illustrated in the following diagram:

Figure 52. 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 IBM Db2 Analytics Accelerator Studio.

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

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 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 that you have found and eliminated the cause of the error, so that the table will not interrupt the process again.

Completing the installation and configuring CDC

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

Procedure

1. Check and make sure that you meet the system requirements for CDC at:

   System requirements for InfoSphere CDC for z/OS. This includes the following steps (among others):

   a) Creating a z/OS user ID for running the CDC started task.

   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 156 and "4.c" on page 156). 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 "1" on page 23). The required PTFs are listed in:

http://www.ibm.com/support/docview.wss?uid=swg27039487Prerequisites and Maintenance for IBM Db2 Analytics Accelerator for z/OS Version 7.1.

4. Complete the steps in section Completing the installation using the distributed sample jobs of the CDC information center.

   Basically, this task comprises the following steps:

   a) Creating the configuration control data set and copying members into this data set.

      The settings in this control data set are read and processed during the initialization of the CDC address space.

   b) Defining a security identifier for use of the CDC address space to ensure a controlled access to the relevant Db2 subsystems and the CDC metadata tables.

   c) Customizing and running a number of jobs to prepare your Db2 subsystems for CDC.
      For example, the CHCMDMUT job creates the metadata tables for CDC. You must customize this job and specify the security identifier (*<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 Data Replication 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 Data Replication 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 155. 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 Data Replication 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 Data Replication 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 Data Replication 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 Important note under LOGCACHEDELAY.

CACHELEVEL1SIZE

Specifies the size, in MB, of the level 1 Db2 log cache. See the previous Important note under LOGCACHEDELAY.

Important: These settings are different from those of IBM InfoSphere Data Replication for z/OS. For example, IBM InfoSphere Data Replication 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 Data Replication 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.

Installing and activating the Access Server and the replication engine

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

Enabling or disabling Db2 subsystems

After installing the capturing agent on your Db2 for z/OS data server, you must still enable incremental updates for all the Db2 subsystems that you want to use the function for. This gives you additional control and security because the enablement process requires a distinct user ID to be specified for running incremental update processes. Furthermore, if the function must be enabled, this also means that it can be disabled, which has the advantage that the capturing agent does not have to be uninstalled for maintenance, error analysis, or error recovery.

Enabling incremental updates for a Db2 subsystem

Follow the steps in this topic to enable incremental updates for a single Db2 subsystem from the IBM Db2 Analytics Accelerator Console. Repeat the process for other subsystems if required.

Before you begin

Make sure that the authentication process (pairing) between the accelerator and the Db2 subsystem that you want to enable has been completed successfully.

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 173
   - "Using ssh to log on to the IBM Db2 Analytics Accelerator Console" on page 174

3. Press the Pause key, then Enter to display the following screen:

```
*************************************************************************
*          Welcome to the IBM Db2 Analytics Accelerator Console
*************************************************************************

You have the following options:
  (0) - (Menu) Manage Configuration Console
Users
  (1) - (Menu) Run Accelerator Functions
  (2) - (Menu) Manage Incremental Updates
  (3) - (Menu) Manage Encryption of Data in Motion
  (4) - (Menu) Manage Call Home

  ----------------------------------------------------------------------
  (x) - Exit the Configuration Console
```

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

(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 of the Db2 subsystem that you want to enable and press Enter.

   b) `Available incremental update technologies:`

      ```
      1. Integrated Synchronization
      2. IBM InfoSphere Data Replication for z/OS
      Select the incremental update technology ('' or '0' to exit):
      ```

      Type 2 (IBM InfoSphere Data Replication 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 "2" on page 155. 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 Data Replication 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 163.

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 IBM Db2 Analytics Accelerator Studio.

Related information

InfoSphere CDC for z/OS authorization requirements

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 173
   - "Using ssh to log on to the IBM Db2 Analytics Accelerator Console" on page 174

2. Press the Pause key, then Enter to display the following screen:

```
**************************************************************************
*        Welcome to the IBM Db2 Analytics Accelerator Console
**************************************************************************

You have the following options:
 (0) - (Menu) Manage Configuration Console
Users
 (1) - (Menu) Run Accelerator Functions
 (2) - (Menu) Manage Incremental Updates
 (3) - (Menu) Manage Encryption of Data in Motion
 (4) - (Menu) Manage Call Home

 ----------------------------------------------------------------------
 (x) - Exit the Configuration Console
```

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

(Default 0) >
```

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.

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 IBM Db2 Analytics Accelerator Studio. 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 accelerator view for regular accelerators 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.

*Figure 53. **Replication status** and functional link in the header of the accelerator view*

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 IBM Db2 Analytics Accelerator Studio and connect to a Db2 subsystem.
2. In the **Administration Explorer**, select the **Accelerators** folder.
3. In the **Object List Editor** on the right, double-click the accelerator that you want to include in the incremental update process.
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 **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:

1. 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 all accelerators attached to one or more Db2 subsystems. See "Starting or stopping incremental updates" on page 166.

3. Restarting the incremental update process.

4. Disabling and re-enabling incremental updates for the affected Db2 subsystems. See "Enabling or disabling Db2 subsystems" on page 162.

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 Data Replication for z/OS

A restart invariably affects all subscriptions based on IBM InfoSphere Data Replication 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 173
   - "Using ssh to log on to the IBM Db2 Analytics Accelerator Console" on page 174

2. Press the Pause key, then Enter to display the following screen:

```
****************************************************************************
*          Welcome to the IBM Db2 Analytics Accelerator Console
****************************************************************************

You have the following options:
 (0) - (Menu) Manage Configuration Console
Users
 (1) - (Menu) Run Accelerator Functions
 (2) - (Menu) Manage Incremental Updates
 (3) - (Menu) Manage Encryption of Data in Motion
 (4) - (Menu) Manage Call Home

 ----------------------------------------------------------------------
 (x) - Exit the Configuration Console
```

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

 (Default 0) >
```

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 Data Replication for z/OS.

   b) If `1.IBM Integrated Synchronization` was selected in step "5.a" on page 169, select a Db2 subsystem. Otherwise go to step "5.c" on page 169

      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 "5.a" on page 169, you see one of the following messages on the screen:

      1. `IBM Integrated Synchronization`
         Based on the example in step "5.b" on page 169:

         ```
         This will restart the integrated synchronization process
         on Db2 subsystem 'DB2WA1', are you sure (y/n)?:
         ```

      2. `IBM InfoSphere Data Replication for z/OS`
         Because you cannot select a Db2 subsystem:

         ```
         This will restart all incremental update processes based on
         IBM InfoSphere Data Replication for z/OS,
         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 "5.a" on page 169 and you want to restart the incremental update process for an additional Db2 subsystem, press Enter and repeat steps "4" on page 169 through "5.c" on page 169.

   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.

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:

```
******************************************************************************
*          Welcome to the IBM Db2 Analytics Accelerator Console
******************************************************************************

You have the following options:
 (0) - (Menu) Manage Configuration Console
Users
 (1) - (Menu) Run Accelerator Functions
 (2) - (Menu) Manage Incremental Updates
 (3) - (Menu) Manage Encryption of Data in Motion
 (4) - (Menu) Manage Call Home

 --------------------------------------------------------------------
 (x) - Exit the Configuration Console
```

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

(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:

```
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 "4" on page 170 through "6" on page 171. Specify a different value in step "5" on page 171.

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:

```
*************************************************************************
*          Welcome to the IBM Db2 Analytics Accelerator Console
*************************************************************************

You have the following options:
 (0) - (Menu) Manage Configuration Console
Users
 (1) - (Menu) Run Accelerator Functions
 (2) - (Menu) Manage Incremental Updates
 (3) - (Menu) Manage Encryption of Data in Motion
 (4) - (Menu) Manage Call Home


 ------------------------------------------------------------------------
 (x) - Exit the Configuration Console
```

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

(Default 0) >
```

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 through . Specify a different value in step .

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:

   ```
   ******************************************************************************
   *           Welcome to the IBM Db2 Analytics Accelerator Console
   ******************************************************************************

   You have the following options:
    (0) - (Menu) Manage Configuration Console
   Users
    (1) - (Menu) Run Accelerator Functions
    (2) - (Menu) Manage Incremental Updates
    (3) - (Menu) Manage Encryption of Data in Motion
    (4) - (Menu) Manage Call Home

    ----------------------------------------------------------------------
    (x) - Exit the Configuration Console
   ```

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 (http://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:

```
**************************************************************************
*          Welcome to the IBM Db2 Analytics Accelerator Console
**************************************************************************

You have the following options:
 (0) - (Menu) Manage Configuration Console
Users
 (1) - (Menu) Run Accelerator Functions
 (2) - (Menu) Manage Incremental Updates
 (3) - (Menu) Manage Encryption of Data in Motion
 (4) - (Menu) Manage Call Home


 ------------------------------------------------------------------------
 (x) - Exit the Configuration Console
```

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:

```
******************************************************************************
*          Welcome to the IBM Db2 Analytics Accelerator Console
******************************************************************************

You have the following options:
 (0) - (Menu) Manage Configuration Console
Users
 (1) - (Menu) Run Accelerator Functions
 (2) - (Menu) Manage Incremental Updates
 (3) - (Menu) Manage Encryption of Data in Motion
 (4) - (Menu) Manage Call Home

 ----------------------------------------------------------------------
 (x) - Exit the Configuration Console
```

3. Type 1 and press Enter to display the submenu:

```
main -> 1
--------------------------------------------------------
You have the following options:

 (0)  - Go back one level
 (1)  - Obtain pairing code, IP address, and port
 (2)  - List paired Db2 subsystems
 (3)  - Clear query history
 (4)  - Restart accelerator process
 (5)  - Dump extensive diagnostic information
 (6)  - Stop backend database
 (7)  - Start backend database
 (8)  - Transparently convert Db2/Z REAL columns into DOUBLE columns when loading a table
 (9)  - Set the DB2 subsystem for time synchronization
 (10) - Rotate key used to generate service password
 (11) - Accelerator Workload Management
```

4. Type 11 and press Enter.

   You see a screen similar to the following:

```
(Default 0) > 11

 main -> 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:
  ┌────┬──────────────┬─────────────────────┬──────────────────────┐
  │ 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 176, that is:

   (0)
   > To exit the screen and go back one level

   (1)
   > 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 176, 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 176 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 "4" on page 176 is updated once more. It now shows identical values in the `Current Entitlement` and `Target Entitlement` columns.

10. Repeat steps "4" on page 176 through "9" on page 178 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

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 4. 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:
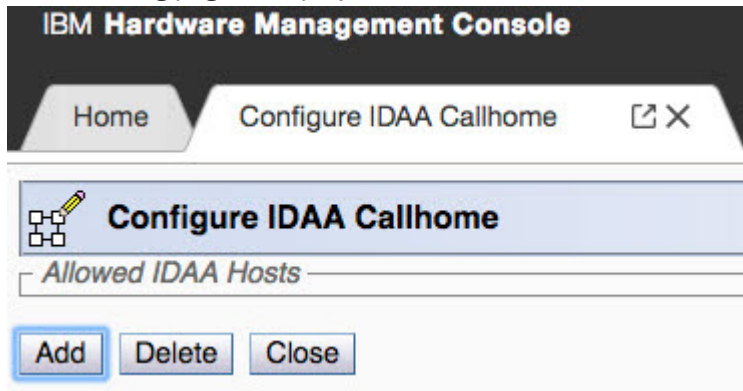
*Figure 54. The **Configure IDAA Callhome** page on the HMC*

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:

Figure 55. 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:

```
*****************************************************************************
*          Welcome to the IBM Db2 Analytics Accelerator Console
*****************************************************************************

You have the following options:
 (0) - (Menu) Manage Configuration Console
Users
 (1) - (Menu) Run Accelerator Functions
 (2) - (Menu) Manage Incremental Updates
 (3) - (Menu) Manage Encryption of Data in Motion
 (4) - (Menu) Manage Call Home

 ----------------------------------------------------------------------
 (x) - Exit the Configuration Console
```

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 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 , 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 , 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 , 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 .

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 .

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:

```
*****************************************************************************
*          Welcome to the IBM Db2 Analytics Accelerator Console
*****************************************************************************

You have the following options:
 (0) - (Menu) Manage Configuration Console
Users
 (1) - (Menu) Run Accelerator Functions
 (2) - (Menu) Manage Incremental Updates
 (3) - (Menu) Manage Encryption of Data in Motion
 (4) - (Menu) Manage Call Home

 ----------------------------------------------------------------------
 (x) - Exit the Configuration Console
```

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 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:

```
******************************************************************************
*           Welcome to the IBM Db2 Analytics Accelerator Console
******************************************************************************

You have the following options:
 (0) - (Menu) Manage Configuration Console
Users
 (1) - (Menu) Run Accelerator Functions
 (2) - (Menu) Manage Incremental Updates
 (3) - (Menu) Manage Encryption of Data in Motion
 (4) - (Menu) Manage Call Home

 ------------------------------------------------------------------------
 (x) - Exit the Configuration Console
```

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 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:

```
****************************************************************************
*          Welcome to the IBM Db2 Analytics Accelerator Console
****************************************************************************

You have the following options:
 (0) - (Menu) Manage Configuration Console
Users
 (1) - (Menu) Run Accelerator Functions
 (2) - (Menu) Manage Incremental Updates
 (3) - (Menu) Manage Encryption of Data in Motion
 (4) - (Menu) Manage Call Home


 ----------------------------------------------------------------------
 (x) - Exit the Configuration Console
```

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.

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.

```
Call Home Status         : Enabled

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:

```
*****************************************************************************
*           Welcome to the IBM Db2 Analytics Accelerator Console
*****************************************************************************

You have the following options:
 (0) - (Menu) Manage Configuration Console
Users
 (1) - (Menu) Run Accelerator Functions
 (2) - (Menu) Manage Incremental Updates
 (3) - (Menu) Manage Encryption of Data in Motion
 (4) - (Menu) Manage Call Home


  ---------------------------------------------------------------------
 (x) - Exit the Configuration Console
```

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 Data Replication 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 188.
2. Prerequisite PTFs for Db2 11 for z/OS or Db2 12 for z/OS. See Installing prerequisite PTFs for Db2 for z/OS.
3. Stored procedures. See "Updating the IBM Db2 Analytics Accelerator stored procedures" on page 193
4. Db2 Analytics Accelerator on Z. See *Updating Db2 Analytics Accelerator on Z* in the *Installation Guide*..
5. IBM Db2 Analytics Accelerator Studio. Actually, the order does not matter for IBM Db2 Analytics Accelerator Studio. You can update it any time. See "Updating IBM Db2 Analytics Accelerator Studio" on page 198

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.

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

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.

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.

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

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 IBM Db2 Analytics Accelerator Studio, the remote client application of IBM Db2 Analytics Accelerator for z/OS.
3. Activating the newly installed software. This step is also executed from IBM Db2 Analytics Accelerator Studio.

This topic covers step "1" on page 188 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 purposes.
  – 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 subdirectories. 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 IBM Db2 Analytics Accelerator Studio.
2. In the **Administration Explorer**, select the **Accelerators** folder.
3. In the Object List Editor on the right, double-click the accelerator.
4. If necessary, expand the **About** section.
5. In the **About** section, click the **Transfer updates** link.
6. In the **Transfer Updates** window, you can see all software packages that are available in the z/FS of your z/OS data server. Select the appropriate check boxes in the first column of the table to mark the packages that you want to transfer.

   ⚠️ **Attention:**

   - Make sure that you select the proper packages, that is, packages belonging to the release level that you want to upgrade to. The list in the **Transfer 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.

7. 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 <u>"1" on page 188</u>.

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. If you are not sure whether there are any open console windows, click **List Tasks** in the relevant **Accelerator** view of IBM Db2 Analytics Accelerator Studio and 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 IBM Db2 Analytics Accelerator Studio:

  1. In the **Administration Explorer** window, open the **Accelerators** folder.

  2. Select a relevant connection and open its **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.

  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 .
>
> - 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. In the **Administration Explorer**, select the **Accelerators** folder.
2. In the Object List Editor on the right, double-click the accelerator.
3. In the **Accelerator** view, from the **Refresh** drop-down list in the upper right, select **Automatic off**.

   Otherwise, you might see warnings during the activation of the new software saying that the accelerator cannot be contacted.

4. If necessary, expand the **About** section.

5. In the **About** section, click the **Apply other software version** link.

6. 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.

7. Activate a version by selecting the appropriate check box in the **File** column.

8. 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.

9. Click **Install**.

10. 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 56. 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 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

1. 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:

      1) 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.

      2) 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 47.

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 IBM Db2 Analytics Accelerator Studio. 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.

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

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 AQTTIJSP errors, customize the AQTSJI00 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 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 10 on page 55).
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 10 on page 55)

   You see the **Accelerator** page:

Figure 57. The Accelerator page of the Appliance Installer

3. In the lower right of the **Accelerator** page, click **Update**.

   You see the **Update** page:
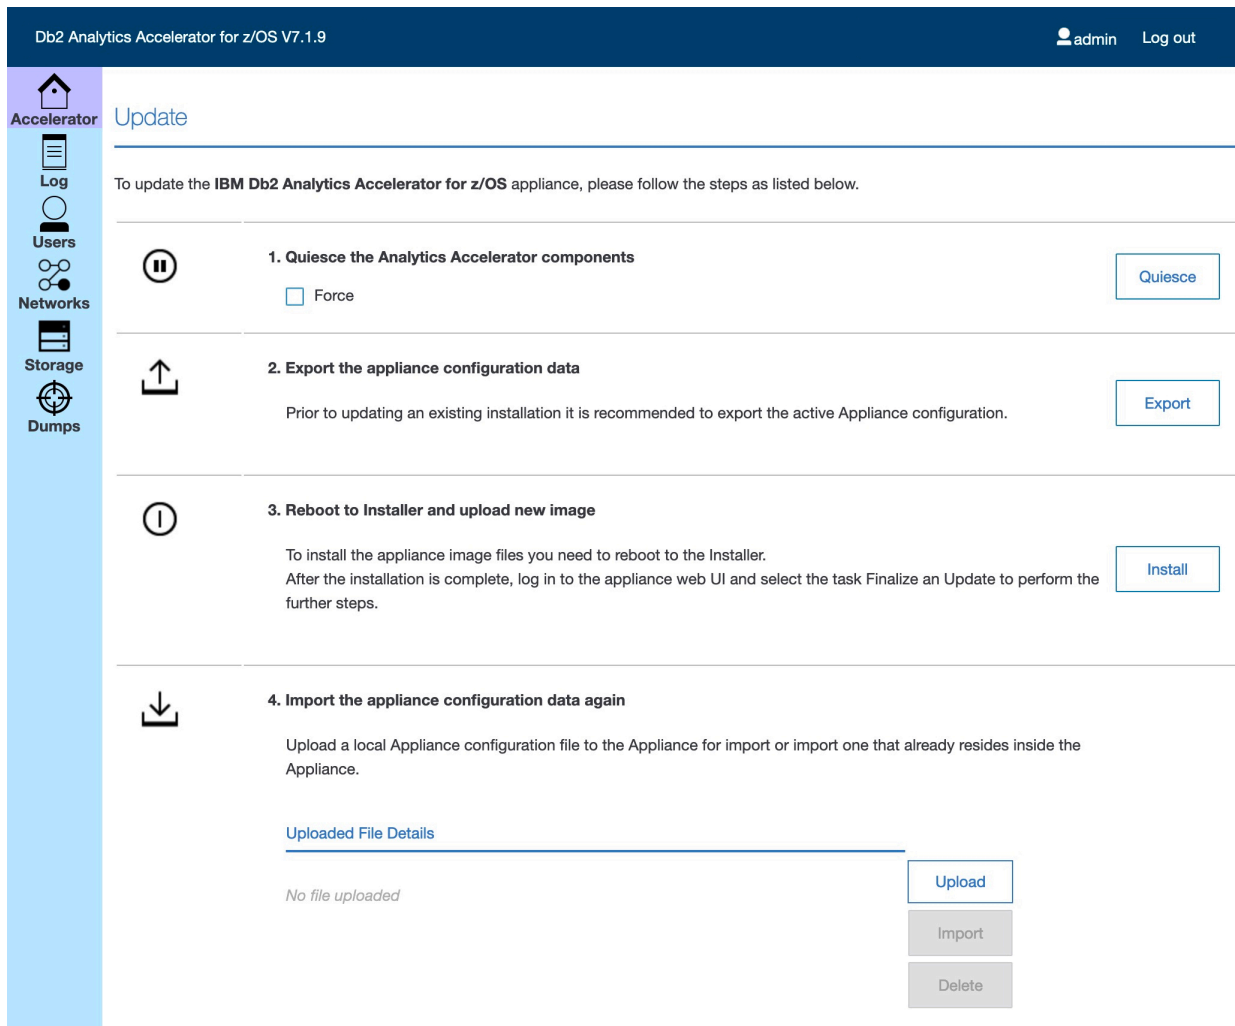
Figure 58. The Update page of the Appliance Installer

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:
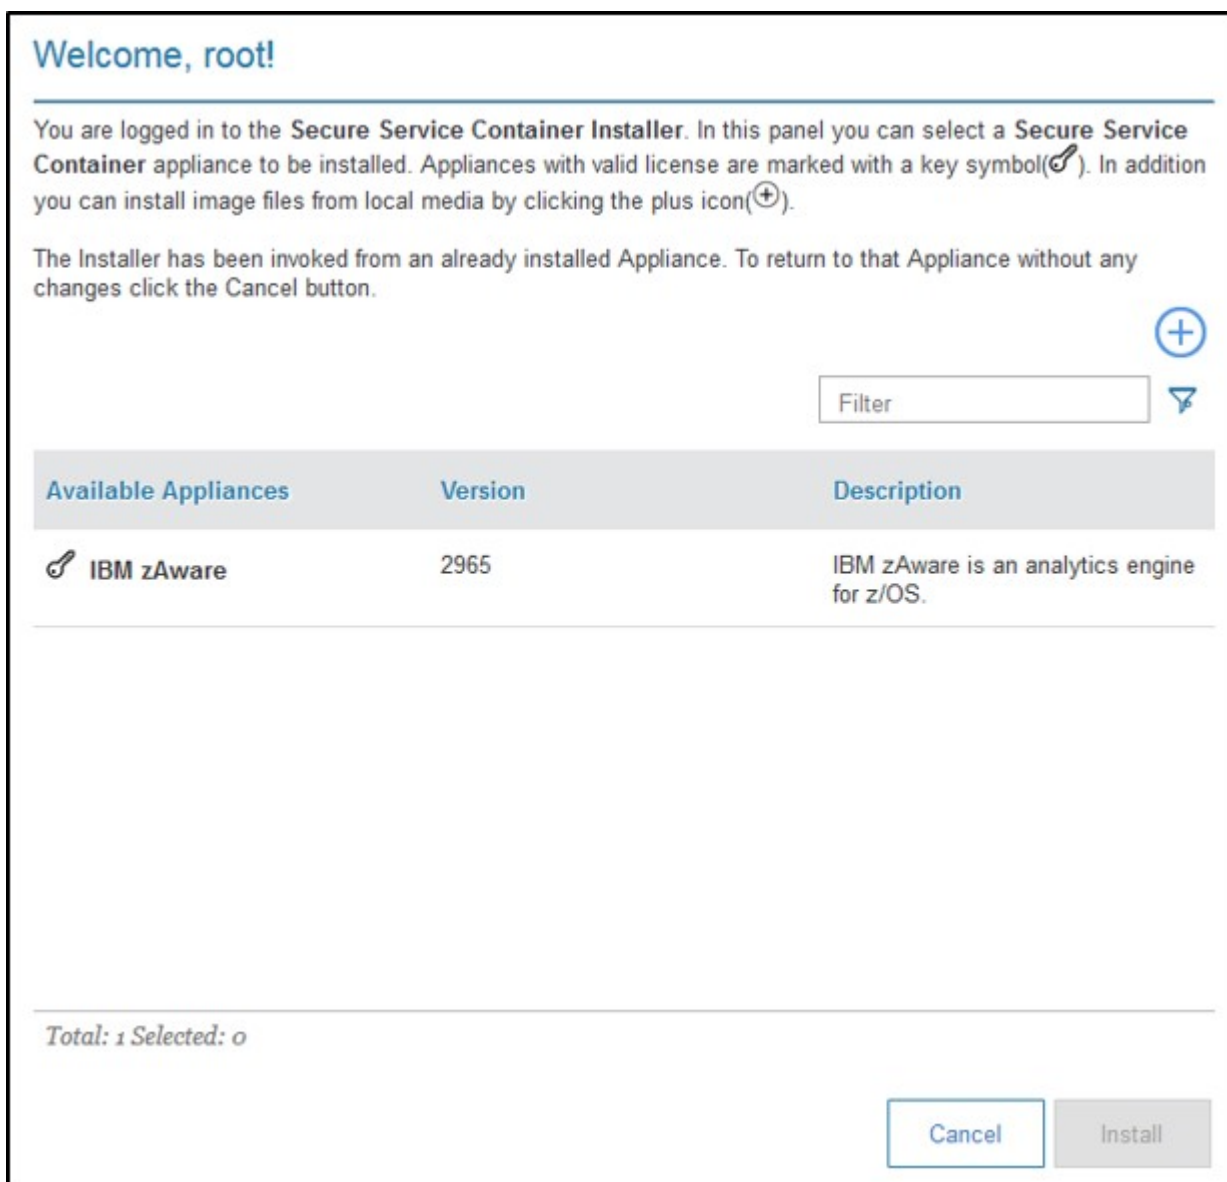
Figure 59. The Welcome page of the Secure Service Container Installer

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.

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

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 "7" on page 198) 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.

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.

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. This parameter change requires a Db2 restart, which can be postponed to the next scheduled maintenance window.
4. For each Db2 subsystem connected to an accelerator, run **-STOP ACCEL**
5. Remove all tables from the connected accelerators.

   Use the **Remove** button in the relevant **Accelerator** views of IBM Db2 Analytics Accelerator Studio or 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 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.

Cleaning up the communications database manually

If accelerators have already been disconnected physically, tables and accelerators (steps 4 and 5 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
   ```

2. 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 .

   ```
   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.

   ⚠️ **Attention:** Because this step deletes the Db2 pseudo-catalog tables, the information about other accelerators attached to the same Db2 subsystem gets lost as well. It is not just the information about the *<accelerator name>* accelerator.

Removing the remaining database objects

Remove the remaining accelerator-related database objects by running the following SQL statements:

About this task

⚠️ **Attention:** Do not run the following commands if you just want to remove a single accelerator and keep others. These commands will remove them all.

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

If you use both product features (coexistence): 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.
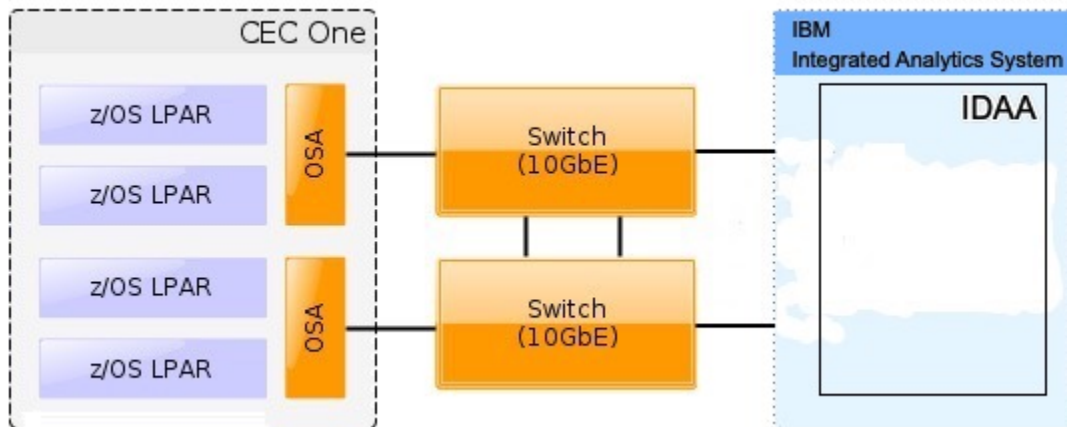
Figure 60. Recommended connections between IBM Z server and IBM Integrated Analytics System (including switches)

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 60 on page 203.

```
*---------------------------------------------*
* GIGABIT ETHERNET OSA EXPRESS                 *
* VTAM NODE : IPAMNZA                          *
*---------------------------------------------*
TRLE0   VBUILD TYPE=TRL

IPANZA0 TRLE    LNCTL=MPC,READ=E100,
                WRITE=E101,DATAPATH=E102,
                PORTNAME=GBE100,MPCLEVEL=QDIO

IPANZA1 TRLE    LNCTL=MPC,READ=E300,
                WRITE=E301,DATAPATH=E302,
                PORTNAME=GBE300,MPCLEVEL=QDIO
```

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

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

The following sample TCP/IP profile also refers to the setup in Figure 60 on page 203.

```
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/OS.
; 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.

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 Data Replication 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 5. Required networking ports for IBM Db2 Analytics Accelerator		
Function or protocol	**Required port**	**Direction**
Query execution (DRDA protocol)	1400	Both ways
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)

Table 5. Required networking ports for IBM Db2 Analytics Accelerator (continued)		
Function or protocol	Required port	Direction
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 CEEDUMP.

PRIQ value too high when creating table spaces

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

Symptoms
The problem occurs if you want to create a table, 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 "Installing libraries with IBM Db2 Analytics Accelerator support" on page 23*Installing Db2 libraries with IBM Db2 Analytics Accelerator for z/OS support* of the *IBM Db2 Analytics Accelerator for z/OS: Installation Guide*.
2. Verify that the IBM Db2 Analytics Accelerator for z/OS libraries are part of the STEPLIB statement.

Connection authorization failure (error -4214)

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

Symptoms

The full message is similar to this one:

```
[jcc][t4][2010][11246][4.7.89]
Connection authorization failure occurred.
Reason: Local security service non-retryable error.
ERRORCODE=-4214, SQLSTATE=28000
```

Causes

No access to the Db2 subsystem, although z/OS can be accessed.

Diagnosing the problem

The error is probably due to a missing RACF authorization.

Resolving the problem

Check the RACF security settings.

Errors during ZPARM compilation

You receive error messages during ZPARM compilation.

Symptoms

You receive the following error messages:

```
X00370097
VOLTDEVT=SYSDA,                      X00370098
XLKUPDLT=NO,                         X00370099
ZOSMETRICS=NO
 ** ASMA017W Undefined keyword parameter; default to positional,
 including keyword - DSN6S/OPTIOWGT
```

```
** ASMA017W Undefined keyword parameter; default to positional,
including keyword - DSN6S/OPTJBPR
** ASMA017W Undefined keyword parameter; default to positional,
including keyword - DSN6S/ZOSMETRICS
** ASMA435I Record 3165 in Db291.ISAO.LAB.SDSNMACS(DSN6SPRM)
on volume: TSMS08
57+*         Avoid overflow        DK153
000000       00000 00814    60+DSN6SPRM
CSECT        CSECT name     02-DSNDS
```

Causes

These errors are caused by the following parameters, which are unknown because they have been introduced by Program Temporary Fixes (PTFs) that are newer than those included in the base level of the libraries with IBM Db2 Analytics Accelerator for z/OS support:

- OPTIOWGT
- OPTJBPR
- ZOSMETRICS

Resolving the problem

Update the libraries with IBM Db2 Analytics Accelerator for z/OS support.

DRDA connection does not work

You can ping the accelerators, but you cannot establish a distributed relational database access (DRDA) connection between your database management system and the 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 IBM Db2 Analytics Accelerator Studio or by running the SYSPROC.ACCEL_REMOVE_ACCELERATOR stored procedure, you find invalid entries in a number of Db2 for z/OS catalog tables and system tables.

Symptoms

You find invalid entries in the following tables:

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

Resolving the problem

1. Disable the accelerator in the Db2 subsystem, by using the -STOP ACCEL command or the appropriate function in IBM Db2 Analytics Accelerator Studio.

2. To remove the invalid entries from the system tables, run the following SQL commands in the order indicated:

a.
```
DELETE FROM SYSIBM.IPNAMES
   WHERE LINKNAME=(SELECT LINKNAME FROM SYSIBM.LOCATIONS
   WHERE LOCATION=(SELECT LOCATION FROM SYSACCEL.SYSACCELERATORS
   WHERE ACCELERATORNAME=<acceleratorName>));
```

b.
```
DELETE FROM SYSIBM.USERNAMES
   WHERE LINKNAME=(SELECT LINKNAME FROM SYSIBM.LOCATIONS
   WHERE LOCATION=(SELECT LOCATION FROM SYSACCEL.SYSACCELERATORS
   WHERE ACCELERATORNAME=<acceleratorName>));
```

c.
```
DELETE FROM SYSIBM.LOCATIONS
   WHERE LOCATION=(SELECT LOCATION FROM SYSACCEL.SYSACCELERATORS
   WHERE ACCELERATORNAME=<acceleratorName>);
```

d.
```
DELETE FROM SYSACCEL.SYSACCELERATORS
   WHERE ACCELERATORNAME=<acceleratorName>;
```

e.
```
DELETE FROM SYSACCEL.SYSACCELERATEDTABLES
   WHERE ACCELERATORNAME=<acceleratorName>;
```

Related tasks
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. In IBM Db2 Analytics Accelerator Studio, the **Administration Explorer** returns a message window like this one:

Figure 61. 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>-<procedure name>-*
     ```

     Example:

     ```
     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'FFFFFFD8' X'00000000' X'00000000' X'FFFFFFFF'
X'00000000' X'00000000' SQL DIAGNOSTIC INFORMATION
```

Causes

In addition to externally published stored procedures in the SYSPROC schema, IBM Db2 Analytics
Accelerator for z/OS employs user-defined functions internally. These functions are included in the
DSNAQT schema and serve the following purposes:

• Reading temporary trace files from IBM Db2 Analytics Accelerator for z/OS stored procedures (function
  DSNAQT.ACCEL_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 IBM Db2 Analytics Accelerator Studio.

2. In the **Administration Explorer**, select the **Accelerators** folder.

3. In the Object List Editor on the right, double-click the accelerator.

4. If necessary, expand the **About** section.

5. In the **About** section, click the **Transfer updates** link.

6. In the **Transfer Updates** window, you can see all software packages that are available in the z/FS of your z/OS data server. Select the appropriate check boxes in the first column of the table to mark the packages that you want to transfer.

> ⚠️ **Attention:**
>
> - Make sure that you select the proper packages, that is, packages belonging to the release level that you want to upgrade to. The list in the **Transfer 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.

7. 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 188.

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 215.

- 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 (http://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):

MHu2cK9FLC2pIkyAEbOoyDxK1yU7UQO1VPyXVviVxwGRtAC3gBiXiio5YLXzQTCQRfL/TRyuG8zFBC2CkSe3
GXPrFjM9KgB41gEFYOoxSEhjXY0S89fKGa7E/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 6. 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)

Table 6. Members of the SAQTSAMP data set (continued)	
Member name	**Function**
AQTSCI02	Db2 command line processor script that calls the following stored procedures and Db2 commands in the order indicated:
	1. SYSPROC.ACCEL_ADD_ACCELERATOR for the initial setup
	Running this stored procedure requires a valid pairing code. You must therefore edit AQTSCI02 so that a valid pairing code is provided when SYSPROC.ACCEL_ADD_ACCELERATOR is invoked.
	Important: In the sample script, the call of this stored procedure has been commented out. To activate the call, you must uncomment the corresponding line.
	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 query 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
	2. 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.
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.
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.

Table 6. Members of the SAQTSAMP data set (continued)	
Member name	**Function**
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 *online* state at a later point-in-time, the data might not be in sync anymore with the data in Db2 for z/OS. This might result in different query results. To solve this problem, restore all 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:
	
	0
	    All stored procedures that were called returned information messages
	4
	    Warning (return code 1024)
	8
	    Error (return code 2048)
	12
	    Severe (return code 3072)

AQTSSCPY	UNIX System Services shell script that copies SAQTSAMP members to the /tmp/ivp directory in the z/OS UNIX file system. **Note:** Db2 command line processor scripts cannot be run from a data set.
AQTSXADD	XML definition of a simple table specification (`<tableSpecifications>` element). This XML code can be used as input for the SYSPROC.ACCEL_ADD_TABLES stored procedure.
AQTSXALT	XML definition of a table specification as input for the SYSPROC.ACCEL_ALTER_TABLES stored procedure.
AQTSXCN0	XML code for the deletion of collected trace data. The XML code is used as the value of the **COMMAND** parameter of the SYSPROC.ACCEL_CONTROL_ACCELERATOR stored procedure.
AQTSXCN1	XML code for the retrieval of accelerator status information. The XML code is used as the value of the **COMMAND** parameter of the SYSPROC.ACCEL_CONTROL_ACCELERATOR stored procedure.
AQTSXCN2	XML code for the configuration of accelerator tracing. The XML code is used as the value of the **COMMAND** parameter of the SYSPROC.ACCEL_CONTROL_ACCELERATOR stored procedure.

Table 6. Members of the SAQTSAMP data set (continued)	
Member name	**Function**
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.

> **&UNIQ.**
> > Causes the creation of a unique identifier

> The template specification in the example could, for instance, result in the following image copy name:

> ```
> BCKE.V4L1.BCKERTSE.CKRANGE3.P00001.D72R4KHN
> ```

- All template variables that are documented for the Db2 COPY utility can be used, with the exception of &SEQ (&SQ), &LIST (&LI), and &DSNUM.
- The chosen variables must ensure the uniqueness of image-copy data-set names. It is therefore recommended that you use at least the &PART. and &UNIQ. template variables.
- Templates must resolve to valid z/OS data set names.
- The template data-set names that you use must have been mapped to suitable data classes in the DFSMS.

Important: The AQT_ARCHIVECOPY_HLQ environment variable that was used with earlier versions of the HPSS is deprecated. Remove it from the AQTENV data set.

AQT_ARCHIVE_COPY2
Specifies the name of the second image-copy data-set to be created when the SYSPROC.ACCEL_ARCHIVE_TABLES stored procedure 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 IBM Db2 Analytics Accelerator Studio 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. Note also that increasing the value of AQT_MAX_UNLOAD_IN_PARALLEL to more than 8 might not increase the throughput any further. The default value is 4.

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 anIBM 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.

TZ
> 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 button and select **Force** from the menu to disable the selected accelerator and cancel all running queries. The status of the accelerator changes to *Stopped* immediately.

  Note: Disabling does not make an accelerator unusable or remove it from the configuration. It just deactivates it. To re-enable an accelerator, click the button on the toolbar.

Appendix 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

Trademarks

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 236
3. "Data Life Cycle" on page 237
4. "Data Collection" on page 237
5. "Data Storage" on page 238
6. "Data Access" on page 238
7. "Data Processing" on page 238
8. "Data Deletion" on page 239
9. "Data Monitoring" on page 239
10. "Responding to Data Subject Rights" on page 239

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 (http://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* (http://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,...c) returns only data records that are not related to (do not belong to) the users in the specified group (a,b,...c).

Right to Object
Same as "Right to Restrict Processing" on page 240.

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.

L

lock
  A means of preventing uncommitted changes made by one application process from being perceived by another application process and for preventing one application process from updating data that is being accessed by another process. A lock ensures the integrity of data by preventing concurrent users from accessing inconsistent data.

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.

T

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).

V

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

Workload Manager, *See* WLM

Z

ZPARMs 24, 29

IBM. ®

Product Number: 5697-DA7