Program Directory for

IBM Db2 Analytics Accelerator for z/OS

V8.1.0

Program Number 5698-DA8

FMIDs HAQTX10 and JAQTXID

for use with
z/OS

Document Date: June 2025

GI13-6604-00

Note

Before using this information and the product it supports, be sure to read the general information under 7.0, “Notices” on page 25.

Contents

Figures

1.0 Introduction

This program directory is intended for system programmers who are responsible for program installation and maintenance. It contains information about the material and procedures associated with the installation of the 'Shared Stored Procedures for Db2 Analytics Accelerator and Data Gate products' operated in the context of IBM Db2 Analytics Accelerator for z/OS V8.1. This publication refers to this product as Db2 Analytics Accelerator.

The Program Directory contains the following sections related to Stored Procedures:

- 2.0, “Program Materials” on page 4 identifies the basic program materials and documentation.
- 3.0, “Program Support” on page 7 describes the IBM support available.
- 4.0, “Program and Service Level Information” on page 8 lists the APARs (program level) and PTFs (service level) that have been incorporated.
- 5.0, “Installation Requirements and Considerations” on page 9 identifies the resources and considerations that are required for installing and using Stored Procedures.
- 6.0, “Installation Instructions” on page 17 provides detailed installation instructions. It also describes the procedures for activating the functions of Stored Procedures, or refers to appropriate publications.

For most recent information on prerequisites and installation information, see the

- **Release Notes for the Stored Procedures used by Db2 Analytics Accelerator for z/OS" at https://www.ibm.com/support/pages/node/7160295**
- **Prerequisites & Maintenance for IBM Db2 Analytics Accelerator for z/OS, V8.1" at https://www.ibm.com/support/pages/node/7176203.**

Before installing Stored Procedures, read the *CBPDO Memo To Users* and the *CBPDO Memo To Users Extension* that are supplied with this program in softcopy format and this program directory; after which, keep the documents for your reference. Section 3.2, “Preventive Service Planning” on page 7 tells you how to find any updates to the information and procedures in this program directory.

The Stored Procedures is supplied in a Custom-Built Product Delivery Offering (CBPDO, 5751-CS3). The program directory that is provided in softcopy format on the CBPDO is identical to the hardcopy format if one was included with your order. All service and HOLDDATA for Db2 Analytics Accelerator are included on the CBPDO.

Do not use this program directory if you install Stored Procedures with a z/OSMF Portable Software Instance (z/OSMF Portable Software Instance (ServerPac)). When you use one of those offerings, use the jobs and documentation supplied with the offering. The offering will point you to specific sections of this program directory as needed.

1.1 Db2 Analytics Accelerator Description

Db2 Analytics Accelerator complements Db2 for z/OS by providing a cost-effective, high-speed query engine to run business-critical reporting and complex analytics workloads. Each query workload runs efficiently in its optimal environment, either on Db2 for z/OS or on the Accelerator. By providing high speed analysis of valuable enterprise data without expensive Extract-Transform-Load (ETL), Db2 Analytics Accelerator helps organizations processing real-time analytics on current transactional data without impacting transactional workloads. Critical data remains on IBM Z, helping to reduce batch processing cycles, decision latency, security risk, complexity, and cost.

New capabilities and enhancements include:

- With a staged delivery plan, Large Object (LOB) data will be accelerated in the same way as traditional structured data (supported with IBM Db2 13 only).
- Ability to copy tables from one accelerator to another accelerator. Supported tables are accelerator-shadow tables, accelerator-archive tables, and accelerator-only tables (AOTs).
- Built-in monitoring infrastructure provides a unified monitoring warehouse, which consolidates all monitoring data into a single, cohesive repository. The warehouse is structured in table format consisting of easy accessible AOTs.

These capabilities provide the following benefits:

- The staged support of LOB data expands the coverage of supported data types, unlocking more use cases for the Accelerator.
- The copying of data from one accelerator to another eliminates duplicate processing across multiple accelerators to populate AOTs or load tables, resulting in enhanced performance and reduced CPU consumption on IBM z/OS.
- The unified monitoring warehouse provides a consolidated monitoring view, facilitates capacity and performance analysis, and improves user access and automation, resulting in the ability to make more better informed decisions.

Db2 Analytics Accelerator consists of the following components:

- SMP/E-installable FMIDs HAQTX10 and JAQTXID containing a set of stored procedures establishing the foundation for the administration of the accelerator
- A container image serving as the Accelerator server component for installation on IBM Z hardware using IBM Z Integrated Facility for Linux (IFLs) or LinuxONE hardware
- A set of RESTful APIs, the Db2 Analytics Accelerator Administration Services, to administer the accelerator using API calls or using Db2 Administration Foundation.

The following functions are no longer supported in Db2 Analytics Accelerator V8.1:

- Changed Data Capture (CDC) replication.

Db2 Analytics Accelerator only supports Integrated Synchronization for data replication between Db2 for z/OS and the Accelerator. It provides high performance and is 97% IBM Z Integrated Information Processor (zIIP) eligible.

- Accelerator Studio plugin for Data Studio.

  Db2 Administration Foundation is the strategic user interface for Db2 Analytics Accelerator.

- Cross-drawer head node architecture for a multi-node accelerator setup.

  Db2 Analytics Accelerator only supports the confined head node architecture for a multi-node accelerator setup.

- The setting temp_working_space=unlimited in the JSON configuration file

  Db2 Analytics Accelerator supports the settings **automatic** or **none** or a Fixed Size value for temp_working_space. As an alternative to the temp_working_space setting, it is recommended to specify a setting for transient_storage or transient_devices setting.

- Deployment on IBM Integrated Analytics System (IIAS)

  The Accelerator server component of Db2 Analytics Accelerator can only be deployed on IBM Z or IBM LinuxONE servers.

Note:
Program information, including system requirements, a complete list of changed or removed features and functions in the Migration chapter, tech notes, training, support, and more, will be available on or before planned availability in IBM Documentation.

1.2 Db2 Analytics Accelerator FMIDs

Db2 Analytics Accelerator consists of an image to be deployed on zLinux and of a z/OS part covered by this Program Directory and comprising the following FMIDs:

HAQTX10 - Stored Procedures for Db2 Analytics Accelerator and Data Gate
JAQTXID - Analytics Accelerator on IBM Z (License Key)

1.3 Migration & Coexistence

When migrating from Db2 Analytics Accelerator V7.5, you will probably be running both Accelerator version 7.5 and 8.1 in parallel for some time.
Please note: The Stored Procedures delivered for Accelerator version 8.1 (FMID HAQTX10) also support the V7.5 Accelerator.
This means: After installing FMID HAQTX10, it is absolutely sufficient to only use the load library from Accelerator V8.1 in your WLM environment hosting the Stored Procedures.
If, however, you want to use both the V7.5 and V8.1 load libraries in your library concatenation of your WLM environment, make sure the V8.1 load library precedes the one of V7.5.

2.0 Program Materials

An IBM program is identified by a program number. The program number for Db2 Analytics Accelerator is 5698-DA8.

Basic Machine-Readable Materials are materials that are supplied under the base license and are required for the use of the product.

The program announcement material describes the features supported by Db2 Analytics Accelerator. Ask your IBM representative for this information if you have not already received a copy.

2.1 Basic Machine-Readable Material

The distribution medium for this program is physical media or downloadable files. This program is in SMP/E RELFILE format and is installed by using SMP/E. See 6.0, “Installation Instructions” on page 17 for more information about how to install the program.

You can find information about the physical media for the basic machine-readable materials for the Stored Procedures of Db2 Analytics Accelerator in the *CBPDO Memo To Users Extension*.

Figure 1 describes the program file content for the Stored Procedures.
You can refer to the *CBPDO Memo To Users Extension* to see where the files reside on the image.

Notes:

1. The data set attributes in this table must be used in the JCL of jobs that read the data sets. However, because the data sets are in IEBCOPY unloaded format, their actual attributes might be different.
2. If any RELFILEs are identified as PDSEs, ensure that SMPTLIB data sets are allocated as PDSEs.

Figure 1 (Page 1 of 2). Program File Content for the Stored Procedures

Name	ORG	RECFM	LRECL	BLK SIZE
SMPMCS	SEQ	FB	80	6400
IBM.HAQTX10.F1	PDS	FB	80	8800
IBM.HAQTX10.F2	PDS	FB	80	8800
IBM.HAQTX10.F3	PDSE	VB	1028	6144
IBM.HAQTX10.F4	PDSE	U	0	6144
IBM.HAQTX10.F5	PDSE	U	0	6144

Figure 1 (Page 2 of 2). Program File Content for the Stored Procedures

Name	ORG	RECFM	LRECL	BLK SIZE
IBM.HAQTX10.F6	PDS	FB	80	8800
IBM.HAQTX10.F7	PDS	FB	80	8800

Figure 2. Program File Content for Analytics Accelerator on IBM Z (License Key)

Name	ORG	RECFM	LRECL	BLK SIZE
SMPMCS	SEQ	FB	80	6400
IBM.JAQTXID.F1	PDSE	U	0	6144
IBM.JAQTXID.F2	PDS	FB	80	8800

2.2 Program Publications

The following sections identify the basic publications for Db2 Analytics Accelerator.

Figure 3 identifies the basic unlicensed publications for Db2 Analytics Accelerator.

Figure 3 (Page 1 of 2). Basic Material: Unlicensed Publications

Publication Title	Form Number	Media Format / Direct Link
IBM Db2 Analytics Accelerator for z/OS V8.1 License Information Agreement	LC28-8479	see table note
IBM Db2 Analytics Accelerator for z/OS V8.1 User's Guide	SC28-8481	PDF (*)
IBM Db2 Analytics Accelerator for z/OS V8.1 Installation Guide	SC28-8499	PDF (*)
IBM Db2 Analytics Accelerator for z/OS V8.1 Stored Procedures Reference	SC28-8484	PDF (*)

Figure 3 (Page 2 of 2). Basic Material: Unlicensed Publications		
Publication Title	**Form Number**	**Media Format / Direct Link**
IBM Db2 Analytics Accelerator for z/OS V8.1 Encryption of Data in Motion	SC28-8500	PDF (*)
Note: https://www.ibm.com/support/customer/csol/terms/		

(*) The PDF files can be found at **IBM Db2 Analytics Accelerator for z/OS V8.1: Guides and Manuals** (https://www.ibm.com/support/pages/node/7230341).

No optional publications are provided for Db2 Analytics Accelerator.

2.3 Program Source Materials

No program source materials or viewable program listings are provided for Db2 Analytics Accelerator.

2.4 Publications Useful During Installation

For use of the publications listed in Figure 4 during the installation of Db2 Analytics Accelerator, see the corrsponding web sites:
z/OS 2.5.0 IBM Documentation (https://www.ibm.com/docs/en/zos/2.5.0?topic=zos-smpe)
or at
z/OS 3.1.0 IBM Documentation (https://www.ibm.com/docs/en/zos/3.1.0?topic=zos-smpe).

Figure 4. Publications Useful During Installation
Publication
IBM SMP/E for z/OS User's Guide
IBM SMP/E for z/OS Reference
IBM SMP/E for z/OS Commands
IBM SMP/E for z/OS Messages, Codes, and Diagnosis

3.0 Program Support

This section describes the IBM support available for Db2 Analytics Accelerator.

3.1 Program Services

Contact your IBM representative for specific information about available program services.

3.2 Preventive Service Planning

Before you install Db2 Analytics Accelerator, make sure that you review the PSP bucket information for IBM Z products document (https://www.ibm.com/support/pages/node/7127792). It contains the latest information concerning the installation of IBM products, including the latest service recommendations and cross-product dependencies. This information was previously available in traditional PSP buckets, which are no longer published for IBM Z products.

For support, access the Software Support Website at https://www.ibm.com/mysupport/

3.3 Statement of Support Procedures

Report any problems which you feel might be an error in the product materials to your IBM Support Center. You may be asked to gather and submit additional diagnostics to assist the IBM Support Center in their analysis.

Figure 5 identifies the component IDs (COMPID) for Db2 Analytics Accelerator.

Figure 5. Component IDs

FMID	COMPID	Component Name	Release
HAQTX10	5698DA800	Stored Procedures for Db2 Analytics Accelerator and Data Gate	X10
JAQTXID	5698DA800	Analytics Accelerator on IBM Z (License Key)	XID

4.0 Program and Service Level Information

This section identifies the program and relevant service levels of Db2 Analytics Accelerator. The program level refers to the APAR fixes that have been incorporated into the program. The service level refers to the PTFs that have been incorporated into the program.

4.1 Program Level Information

The following APAR fixes against previous releases of Db2 Analytics Accelerator have been incorporated into this release. They are listed by FMID.

- FMID HAQT710

PH66192	PH53602	PH45639
PH61845	PH51109	PH42174
PH57138	PH48263	

4.2 Service Level Information

No PTFs against this release of Db2 Analytics Accelerator have been incorporated into the product package.

5.0 Installation Requirements and Considerations

The following sections identify the system requirements for installing and activating Db2 Analytics Accelerator. The following terminology is used:

- *Driving system*: the system on which SMP/E is executed to install the program.

  The program might have specific operating system or product level requirements for using processes, such as binder or assembly utilities during the installation.

- *Target system*: the system on which the program is configured and run.

  The program might have specific product level requirements, such as needing access to the library of another product for link-edits. These requirements, either mandatory or optional, might directly affect the element during the installation or in its basic or enhanced operation.

In many cases, you can use a system as both a driving system and a target system. However, you can make a separate IPL-able clone of the running system to use as a target system. The clone must include copies of all system libraries that SMP/E updates, copies of the SMP/E CSI data sets that describe the system libraries, and your PARMLIB and PROCLIB.

Use separate driving and target systems in the following situations:

- When you install a new level of a product that is already installed, the new level of the product will replace the old one. By installing the new level onto a separate target system, you can test the new level and keep the old one in production at the same time.
- When you install a product that shares libraries or load modules with other products, the installation can disrupt the other products. By installing the product onto a separate target system, you can assess these impacts without disrupting your production system.

5.1 Driving System Requirements

This section describes the environment of the driving system required to install Db2 Analytics Accelerator.

5.1.1 Machine Requirements

The driving system can run in any hardware environment that supports the required software.

5.1.2 Programming Requirements

Figure 6. Driving System Software Requirements

Program Number	Product Name	Minimum VRM	Minimum Service Level will satisfy these APARs	Included in the shipped product?
Any **one** of the following:				
5650-ZOS	z/OS	2.5.0	N/A	No
5655-ZOS	z/OS	3.1.0	N/A	No

Note: SMP/E is a requirement for Installation and is an element of z/OS.

Note: Installation might require migration to new z/OS releases to be service supported. See https://www.ibm.com/support/lifecycle/

Db2 Analytics Accelerator is installed into a file system. Before installing Db2 Analytics Accelerator, you must ensure that the target system file system data sets are available for processing on the driving system. OMVS must be active on the driving system and the target system file data sets must be mounted on the driving system.

zFS must be active on the driving system. Information on activating and using zFS can be found in z/OS Distributed File Service zSeries File System Administration, SC24-5989.

5.2 Target System Requirements

This section describes the environment of the target system required to install and use Db2 Analytics Accelerator.

Db2 Analytics Accelerator installs in the DBS (P115) SREL.

5.2.1 Machine Requirements

The target system can run in any hardware environment that supports the required software.

5.2.2 Programming Requirements

5.2.2.1 Installation Requisites

Installation requisites identify products that are required and *must* be present on the system or products that are not required but *should* be present on the system for the successful installation of this product.

Mandatory installation requisites identify products that are required on the system for the successful installation of this product.

Db2 Analytics Accelerator has no mandatory installation requisites.

Conditional installation requisites identify products that are *not* required for successful installation of this product but can resolve such things as certain warning messages at installation time.

Db2 Analytics Accelerator has no conditional installation requisites.

5.2.2.2 Operational Requisites

Operational requisites are products that are required and *must* be present on the system or products that are not required but *should* be present on the system for this product to operate all or part of its functions.

Mandatory operational requisites identify products that are required for this product to operate its basic functions.

Figure 7. Target System Mandatory Operational Requisites

Program Number	**Product Name and Minimum VRM/Service Level**
Any **one** of the following:	
5650-DB2	Db2 12 for z/OS (V12.1.0)
5770-AF3	Db2 12 for z/OS Value Unit Edition (V12.1.0)
5698-DB2	Db2 13 for z/OS (V13.1.0)
5698-DBV	Db2 13 for z/OS Value Unit Edition (V13.1.0)

Note:

1. The Db2 12 for z/OS versions mentioned above will be withdrawn from service on December 31, 2025.
2. We strongly recommend to install RSU level 2409 or PUT level 2409 for the supported Db2 for z/OS versions.
3. When using Db2 12 for z/OS
   - The High Availability (HA) support for Accelerator-Only Tables (AOTs) requires
     - FL 509
     - The installation of the PTFs for the APARs PH30574, PH31723, and PH33015.
   - The use of Integrated Synchronization requires
     - FL 500
     - The installation of the PTFs for the following APARs:
       PI91620, PI96258, PH00574, PH00933
       PH06628, PH10672, PH13109, PH23895
       PH28849, PH37808, PH40261, PH45962
       PH53700, PH59973, PH61826
4. When using Db2 13 for z/OS

- The use of Integrated Synchronization
  - Mandatorily requires the installation of the PTFs for the APARs PH45854, PH53700, PH59973, PH61826
  - Optionally requires the installation of the PTF for APAR PH64099 and the activation of the Compression Dictionary Dataset Support (CDDS).
- The acceleration of Large Object (LOB) data requires the installation of the PTF for APAR PH65709.

5. Installation might require migration to new releases to obtain support. See https://www.ibm.com/support/lifecycle/

Conditional operational requisites identify products that are *not* required for this product to operate its basic functions but are required at run time for this product to operate specific functions. These products are specified as IF REQs.

Db2 Analytics Accelerator has no conditional operational requisites.

5.2.2.3 Toleration/Coexistence Requisites

Toleration/coexistence requisites identify products that must be present on sharing systems. These systems can be other systems in a multisystem environment (not necessarily sysplex), a shared DASD environment (such as test and production), or systems that reuse the same DASD environment at different time intervals.

Db2 Analytics Accelerator has no toleration/coexistence requisites.

5.2.2.4 Incompatibility (Negative) Requisites

Negative requisites identify products that must *not* be installed on the same system as this product.

Db2 Analytics Accelerator has no negative requisites.

5.2.3 DASD Storage Requirements

Db2 Analytics Accelerator libraries can reside on all supported DASD types.

Figure 8 lists the total space that is required for each type of library.

Figure 8 (Page 1 of 2). Total DASD Space Required by Db2 Analytics Accelerator

Library Type	Total Space Required in 3390 Trks	Description
Target	4145 tracks	Stored Procedures
Distribution	5282 tracks	Stored Procedures

Figure 8 (Page 2 of 2). Total DASD Space Required by Db2 Analytics Accelerator

Library Type	Total Space Required in 3390 Trks	Description
File System(s)	750 tracks	/usr/lpp/IBM/aqt/vxr1m0

Notes:

1. For non-RECFM U data sets, IBM recommends using system-determined block sizes for efficient DASD utilization. For RECFM U data sets, IBM recommends using a block size of 32760, which is most efficient from the performance and DASD utilization perspective.
2. Abbreviations used for data set types are shown as follows.

   U Unique data set, allocated by this product and used by only this product. This table provides all the required information to determine the correct storage for this data set. You do not need to refer to other tables or program directories for the data set size.

   S Shared data set, allocated by this product and used by this product and other products. To determine the correct storage needed for this data set, add the storage size given in this table to those given in other tables (perhaps in other program directories). If the data set already exists, it must have enough free space to accommodate the storage size given in this table.

   E Existing shared data set, used by this product and other products. This data set is *not* allocated by this product. To determine the correct storage for this data set, add the storage size given in this table to those given in other tables (perhaps in other program directories). If the data set already exists, it must have enough free space to accommodate the storage size given in this table.

   If you currently have a previous release of this product installed in these libraries, the installation of this release will delete the old release and reclaim the space that was used by the old release and any service that had been installed. You can determine whether these libraries have enough space by deleting the old release with a dummy function, compressing the libraries, and comparing the space requirements with the free space in the libraries.

   For more information about the names and sizes of the required data sets, see 6.1.6, “Allocate SMP/E Target and Distribution Libraries” on page 20.
3. Abbreviations used for the file system path type are as follows.

   N New path, created by this product.
   X Path created by this product, but might already exist from a previous release.
   P Previously existing path, created by another product.

   ligible
4. All target and distribution libraries listed have the following attributes:

   - The default name of the data set can be changed.
   - The default block size of the data set can be changed.
   - The data set can be merged with another data set that has equivalent characteristics.

- The data set can be either a PDS or a PDSE, with some exceptions.
  - **If the value in the "ORG" column specifies "PDS(E)", the data can be either a PDS or a PDSE.**
  - If the value in the "ORG" column specifies "PDS", the data set must be a PDS.
  - If the value in "DIR Blks" column specifies "N/A", the data set must be a PDSE.

5. All target libraries listed have the following attributes:

   - These data sets can be SMS-managed, but they are not required to be SMS-managed.
   - These data sets are not required to reside on the IPL volume.
   - The values in the "Member Type" column are not necessarily the actual SMP/E element types that are identified in the SMPMCS.

6. All target libraries that are listed and contain load modules have the following attributes:

   - These data sets can be in the LPA, but they are not required to be in the LPA.
   - These data sets can be in the LNKLST.
   - The SAQTMOD dataset must be APF authorized.
   - Db2 Analytics Accelerator requires that the SMPLTS data set must be a PDSE. If your existing SMPLTS is a PDS, you will need to allocate a new PDSE and copy your existing SMPLTS into it and then change the SMPLTS DDDEF entry to indicate the new PDSE data set.

The following figures describe the target and distribution libraries and file system paths required to install Db2 Analytics Accelerator. The storage requirements of Db2 Analytics Accelerator must be added to the storage required by other programs that have data in the same library or path.

Note: Use the data in these tables to determine which libraries can be merged into common data sets. In addition, since some ALIAS names may not be unique, ensure that no naming conflicts will be introduced before merging libraries. **If the value in the "ORG" column specifies "PDS(E)", the data can be either a PDS or a PDSE.**

Figure 9. Storage Requirements for Db2 Analytics Accelerator Target Libraries

Library DDNAME	Member Type	Target Volume	TYPE	ORG	RECFM	LRECL	No. of 3390 Trks	No. of DIR Blks
SAQTBASE	SAMP	Any	U	PDS(E)	FB	80	10	5
SAQTDBRM	MAC	Any	U	PDS(E)	FB	80	12	5
SAQTICU	PROGRAM	Any	U	PDSE	U	0	10	N/A
SAQTLICI	SAMP	Any	U	PDS(E)	FB	80	5	5
SAQTMOD	PROGRAM	Any	U	PDSE	U	0	4000	N/A
SAQTSAMP	SAMP	Any	U	PDS(E)	FB	80	120	5

Figure 10. Stored Procedures File System Paths

DDNAME	TYPE	Path Name
SAQTHFS	X	/usr/lpp/IBM/aqt/vxr1m0/IBM

Figure 11. Storage Requirements for Db2 Analytics Accelerator Distribution Libraries

Library DDNAME	TYPE	ORG	RECFM	LRECL	No. of 3390 Trks	No. of DIR Blks
AAQTBASE	U	PDS(E)	FB	80	10	5
AAQTDBRM	U	PDS(E)	FB	80	12	5
AAQTHFS	U	PDSE	VB	1028	1125	N/A
AAQTICU	U	PDSE	U	0	10	N/A
AAQTLICI	U	PDS(E)	FB	80	5	5
AAQTMOD	U	PDSE	U	0	4000	N/A
AAQTSAMP	U	PDS(E)	FB	80	120	5

5.3 FMIDs Deleted

Installing Stored Procedures might result in the deletion of other FMIDs. To see which FMIDs will be deleted, examine the ++VER statement in the SMPMCS of the product.

If you do not want to delete these FMIDs at this time, install Stored Procedures into separate SMP/E target and distribution zones.

Note: These FMIDs are not automatically deleted from the Global Zone. If you want to delete these FMIDs from the Global Zone, use the SMP/E REJECT NOFMID DELETEFMID command. See the SMP/E Commands book for details.

5.4 Special Considerations

PDSE Considerations:

Db2 Analytics Accelerator uses the "partitioned data set extended" or PDSE format for the SAQTMOD and SAQTICU target libraries. There are some operational differences between PDS and PDSE data sets. The PDS format may be shared by more than one z/OS system and no special precautions are

necessary. However, the PDSE format may only be shared by z/OS systems which are part of a sysplex or which are connected using Global Resource Serialization (are in a GRS complex). If z/OS systems share use of a PDSE data set outside of a sysplex or GRS environment, you may experience severe problems when the data set is updated. This is due to the fact that PDSE directory information is cached in storage, and when the data set is updated from one system the other system(s) have no knowledge of the update, and their cached directory information will be incorrect.

You must take care not to share the SAQTMOD and SAQTICU data sets between z/OS systems unless they are in a sysplex or are connected in a GRS complex. If you need to share the content of the SAQTMOD and SAQTICU data sets, a separate copy must be created for each z/OS system.

6.0 Installation Instructions

This chapter describes the installation method and the step-by-step procedures to install and to activate the functions of the Stored Procedures of Db2 Analytics Accelerator.

Please note the following points:

- If you want to install the Stored Procedures into its own SMP/E environment, consult the SMP/E manuals for instructions on creating and initializing the SMPCSI and the SMP/E control data sets.
- You can use the sample jobs that are provided to perform part or all of the installation tasks. The SMP/E jobs assume that all DDDEF entries that are required for SMP/E execution have been defined in appropriate zones.
- You can use the SMP/E dialogs instead of the sample jobs to accomplish the SMP/E installation steps.

6.1 Installing the Stored Procedures

6.1.1 SMP/E Considerations for Installing Db2 Analytics Accelerator

Use the SMP/E RECEIVE, APPLY, and ACCEPT commands to install this release of Db2 Analytics Accelerator.

6.1.2 SMP/E Options Subentry Values

The recommended values for certain SMP/E CSI subentries are shown in Figure 12. Using values lower than the recommended values can result in failures in the installation. DSSPACE is a subentry in the GLOBAL options entry. PEMAX is a subentry of the GENERAL entry in the GLOBAL options entry. See the SMP/E manuals for instructions on updating the global zone.

Figure 12. SMP/E Options Subentry Values

Subentry	Value	Comment
DSSPACE	(5000,2500,500)	3390 DASD tracks
PEMAX	SMP/E Default	IBM recommends using the SMP/E default for PEMAX.

6.1.3 SMP/E CALLLIBS Processing

Db2 Analytics Accelerator uses the CALLLIBS function provided in SMP/E to resolve external references during installation. When Db2 Analytics Accelerator is installed, ensure that DDDEFs exist for the following libraries:

- CSSLIB
- SCEECPP
- SCEELKED
- SCEELKEX
- SEZACMTX
- SEZARNT1

Note: CALLLIBS uses the previous DDDEFs only to resolve the link-edit for Db2 Analytics Accelerator. These data sets are not updated during the installation of Db2 Analytics Accelerator.

6.1.4 Sample Jobs

The following sample installation jobs are provided as part of the product to help you install Db2 Analytics Accelerator:

Figure 13 (Page 1 of 2). Sample Installation Jobs

Job Name	Job Type	Description	SMPTLIB Data Set
AQTALA	SMP/E	Sample job to allocate and initialize a new SMP/E CSI data set **(Optional)**	IBM.HAQTX10.F1
AQTALB	SMP/E	Sample job to allocate SMP/E data sets **(Optional)**	IBM.HAQTX10.F1
AQTRECID	RECEIVE	Sample RECEIVE job for the Stored Procedures and the Analytics Accelerator on IBM Z (License Key)	IBM.HAQTX10.F1
AQTALLOC	ALLOCATE	Sample job to allocate target and distribution libraries for the Stored Procedures and the Analytics Accelerator on IBM Z (License Key)	IBM.HAQTX10.F1
AQTZFS	ALLOMTZFS	Sample job to allocate new ZFS data set for the Stored Procedures **(Optional)**	IBM.HAQTX10.F1
AQTISMKD	MKDIR	Sample job to invoke the supplied AQTMKDIR EXEC to allocate zFS paths for the Stored Procedures	IBM.HAQTX10.F1
AQTDDDEF	DDDEF	Sample job to define SMP/E DDDEFs for the Stored Procedures	IBM.HAQTX10.F1
AQTAPPID	APPLY	Sample APPLY job for the Stored Procedures and the Analytics Accelerator on IBM Z (License Key)	IBM.HAQTX10.F1

Figure 13 (Page 2 of 2). Sample Installation Jobs

Job Name	Job Type	Description	SMPTLIB Data Set
AQTACCID	ACCEPT	Sample ACCEPT job for the Stored Procedures and the Analytics Accelerator on IBM Z (License Key)	IBM.HAQTX10.F1

You can access the sample installation jobs by performing an SMP/E RECEIVE (refer to 6.1.5, “Perform SMP/E RECEIVE” on page 19) then copy the jobs from the SMPTLIB data sets to a work data set for editing and submission. See Figure 13 on page 18 to find the appropriate data set.

You can also copy the sample installation jobs from the product files by submitting the following job. Before you submit the job, add a job card and change the lowercase parameters to uppercase values to meet the requirements of your site.

```
//STEP1    EXEC PGM=IEBCOPY
//SYSPRINT DD SYSOUT=*
//*FILEIN   DD DSN=IBM.HAQTX10.F1,UNIT=SYSALLDA,DISP=SHR,
//*         VOL=SER=filevol
//OUT      DD DSNAME=jcl-library-name,
//         DISP=(NEW,CATLG,DELETE),
//         VOL=SER=dasdvol,UNIT=SYSALLDA,
//         SPACE=(TRK,(primary,secondary,dir))
//SYSUT3   DD UNIT=SYSALLDA,SPACE=(CYL,(1,1))
//SYSIN    DD *
   COPY INDD=FILEIN,OUTDD=OUT
/*
```

See the following information to update the statements in the previous sample:

FILEIN:
: **filevol** is the volume serial of the DASD device where the downloaded files reside.

OUT:
: **jcl-library-name** is the name of the output data set where the sample jobs are stored.
: **dasdvol** is the volume serial of the DASD device where the output data set resides.

6.1.5 Perform SMP/E RECEIVE

If you have obtained Db2 Analytics Accelerator as part of a CBPDO, use the RCVPDO job in the CBPDO RIMLIB data set to receive the Db2 Analytics Accelerator FMIDs, service, and HOLDDATA that are included on the CBPDO package. For more information, see the documentation that is included in the CBPDO.

Edit and submit the sample job AQTRECID to perform the SMP/E RECEIVE for the FMIDs HAQTX10 and JAQTXID. Consult the instructions in the sample jobs for more information.

Expected Return Codes and Messages: You will receive a return code of 0 if these job runs correctly.

6.1.6 Allocate SMP/E Target and Distribution Libraries

Edit and submit sample job AQTALLOC to allocate the SMP/E target and distribution libraries for the Stored Procedures. Consult the instructions in the sample job for more information.

Expected Return Codes and Messages: You will receive a return code of 0 if this job runs correctly.

If you plan to install Stored Procedures into a new z/OS UNIX file system, you can edit and submit the optional AQTZFS job to perform the following tasks:

- Create the z/OS UNIX file system
- Create a mount point
- Mount the z/OS UNIX file system on the mountpoint

Consult the instructions in the sample job for more information.

The recommended z/OS UNIX file system type is *zFS*. The recommended mount point is:

- */usr/lpp/IBM/aqt/vxr1m0* for Db2 Analytics Accelerator.

Before running the sample jobs to create the z/OS UNIX file systems, you must ensure that OMVS is active on the driving system. zFS must be active on the driving system if you are installing Stored Procedures into a file system that is zFS.

If you create a new file system for this product, consider updating the BPXPRMxx PARMLIB member to mount the new file system at IPL time. This action can be helpful if an IPL occurs before the installation is completed.

```
MOUNT FILESYSTEM('#dsn')
 MOUNTPOINT('#mntpnt')
 MODE(RDRW)         /* can be MODE(READ) */
 TYPE(ZFS) PARM('AGGRGROW') /* zFS, with extents */
```

See the following information to update the statements in the previous sample:

#dsn is the name of the data set holding the z/OS UNIX file system.
#mntpnt is the name of the mount point where the z/OS UNIX file system will be mounted: /usr/lpp/IBM/aqt/vxr1m0.

Expected Return Codes and Messages: You will receive a return code of 0 if these jobs runs correctly.

6.1.7 Allocate File System Paths

The target system zFS data set must be mounted on the driving system when running the sample AQTISMKD job since the job will create paths in the file system.

Before running the sample job to create the paths in the file system, you must ensure that OMVS is active on the driving system and that the target system's zFS file system is mounted to the driving system. zFS must be active on the driving system.

If you plan to install the Stored Procedures into a new zFS file system, you must create the mountpoint and mount the new file system to the driving system for Stored Procedures.

The recommended mountpoint is */usr/lpp/IBM/aqt/vxr1m0*.

Edit and submit sample job AQTISMKD to allocate the file system paths for the Stored Procedures. Consult the instructions in the sample job for more information.

If you create a new file system for this product, consider updating the BPXPRMxx PARMLIB member to mount the new file system at IPL time. This action can be helpful if an IPL occurs before the installation is completed.

Expected Return Codes and Messages: You will receive a return code of 0 if this job runs correctly.

6.1.8 Create DDDEF Entries

Edit and submit sample job AQTDDDEF to create DDDEF entries for the SMP/E target and distribution libraries for the Stored Procedures. Consult the instructions in the sample job for more information.

Expected Return Codes and Messages: You will receive a return code of 0 if this job runs correctly.

6.1.9 Perform SMP/E APPLY

1. Ensure that you have the latest HOLDDATA; then edit and submit sample job AQTAPPID to perform an SMP/E APPLY CHECK for the Stored Procedures. Consult the instructions in the sample job for more information.

   The latest HOLDDATA is available through several different portals, including https://public.dhe.ibm.com/s390/assigns/ or https://www.ibm.com/support/pages/enhanced-holddata-zos for usage instructions. The latest HOLDDATA may identify HIPER and FIXCAT APARs for the FMIDs you will be installing. An APPLY CHECK will help you determine if any HIPER or FIXCAT APARs are applicable to the FMIDs you are installing. If there are any applicable HIPER or FIXCAT APARs, the APPLY CHECK will also identify fixing PTFs that will resolve the APARs, if a fixing PTF is available.

   You should install the FMIDs regardless of the status of unresolved HIPER or FIXCAT APARs. However, do not deploy the software until the unresolved HIPER and FIXCAT APARs have been analyzed to determine their applicability. That is, before deploying the software either ensure fixing

PTFs are applied to resolve all HIPER or FIXCAT APARs, or ensure the problems reported by all HIPER or FIXCAT APARs are not applicable to your environment.

To receive the full benefit of the SMP/E Causer SYSMOD Summary Report, do *not* bypass the PRE, ID, REQ, and IFREQ on the APPLY CHECK. The SMP/E root cause analysis identifies the cause only of *errors* and not of *warnings* (SMP/E treats bypassed PRE, ID, REQ, and IFREQ conditions as warnings, instead of errors).

Here are sample APPLY commands:

a. To ensure that all recommended and critical service is installed with the FMIDs, receive the latest HOLDDATA and use the APPLY CHECK command as follows

```
APPLY S(fmid,fmid,...) CHECK
FORFMID(fmid,fmid,...)
SOURCEID(RSU*)
FIXCAT(IBM.ProductInstall-RequiredService)
GROUPEXTEND .
```

   Some HIPER APARs might not have fixing PTFs available yet. You should analyze the symptom flags for the unresolved HIPER APARs to determine if the reported problem is applicable to your environment and if you should bypass the specific ERROR HOLDs in order to continue the installation of the FMIDs.

   This method requires more initial research, but can provide resolution for all HIPERs that have fixing PTFs available and are not in a PE chain. Unresolved PEs or HIPERs might still exist and require the use of BYPASS.

b. To install the FMIDs without regard for unresolved HIPER APARs, you can add the BYPASS(HOLDCLASS(HIPER)) operand to the APPLY CHECK command. This will allow you to install FMIDs even though one or more unresolved HIPER APARs exist. After the FMIDs are installed, use the SMP/E REPORT ERRSYSMODS command to identify unresolved HIPER APARs and any fixing PTFs.

```
APPLY S(fmid,fmid,...) CHECK
FORFMID(fmid,fmid,...)
SOURCEID(RSU*)
FIXCAT(IBM.ProductInstall-RequiredService)
GROUPEXTEND
BYPASS(HOLDCLASS(HIPER)) .
 ..any other parameters documented in the program directory
```

   This method is quicker, but requires subsequent review of the Exception SYSMOD report produced by the REPORT ERRSYSMODS command to investigate any unresolved HIPERs. If you have received the latest HOLDDATA, you can also choose to use the REPORT MISSINGFIX command and specify Fix Category IBM.PRODUCTINSTALL-REQUIREDSERVICE to investigate missing recommended service.

   If you bypass HOLDs during the installation of the FMIDs because fixing PTFs are not yet available, you can be notified when the fixing PTFs are available by using the APAR Status Tracking (AST) function of ServiceLink or the APAR Tracking function of ResourceLink.

2. After you take actions that are indicated by the APPLY CHECK, remove the CHECK operand and run the job again to perform the APPLY.

   Note: The GROUPEXTEND operand indicates that SMP/E applies all requisite SYSMODs. The requisite SYSMODS might be applicable to other functions.

   Expected Return Codes and Messages from APPLY CHECK: You will receive a return code of 0 if this job runs correctly.

   Expected Return Codes and Messages from APPLY: You will receive a return code of 0 if this job runs correctly.

6.1.10 Perform SMP/E ACCEPT

Edit and submit sample job AQTACCID to perform an SMP/E ACCEPT CHECK for the Stored Procedures. Consult the instructions in the sample job for more information.

To receive the full benefit of the SMP/E Causer SYSMOD Summary Report, do *not* bypass the PRE, ID, REQ, and IFREQ on the ACCEPT CHECK. The SMP/E root cause analysis identifies the cause of *errors* but not *warnings* (SMP/E treats bypassed PRE, ID, REQ, and IFREQ conditions as warnings rather than errors).

Before you use SMP/E to load new distribution libraries, it is recommended that you set the ACCJCLIN indicator in the distribution zone. In this way, you can save the entries that are produced from JCLIN in the distribution zone whenever a SYSMOD that contains inline JCLIN is accepted. For more information about the ACCJCLIN indicator, see the description of inline JCLIN in the SMP/E Commands book for details.

After you take actions that are indicated by the ACCEPT CHECK, remove the CHECK operand and run the job again to perform the ACCEPT.

Note: The GROUPEXTEND operand indicates that SMP/E accepts all requisite SYSMODs. The requisite SYSMODS might be applicable to other functions.

Expected Return Codes and Messages from ACCEPT CHECK: You will receive a return code of 0 if this job runs correctly.

If PTFs that contain replacement modules are accepted, SMP/E ACCEPT processing will link-edit or bind the modules into the distribution libraries. During this processing, the Linkage Editor or Binder might issue messages that indicate unresolved external references, which will result in a return code of 4 during the ACCEPT phase. You can ignore these messages, because the distribution libraries are not executable and the unresolved external references do not affect the executable system libraries.

Expected Return Codes and Messages from ACCEPT: You will receive a return code of 0 if this job runs correctly.

6.1.11 Run REPORT CROSSZONE

The SMP/E REPORT CROSSZONE command identifies requisites for products that are installed in separate zones. This command also creates APPLY and ACCEPT commands in the SMPPUNCH data set. You can use the APPLY and ACCEPT commands to install those cross-zone requisites that the SMP/E REPORT CROSSZONE command identifies.

After you install the Stored Procedures, it is recommended that you run REPORT CROSSZONE against the new or updated target and distribution zones. REPORT CROSSZONE requires a global zone with ZONEINDEX entries that describe all the target and distribution libraries to be reported on.

For more information about REPORT CROSSZONE, see the SMP/E manuals.

6.1.12 Cleaning Up Obsolete Data Sets, Paths, and DDDEFs

The following data sets, which were allocated and used by previous releases of this product, are no longer used in this release. You can delete these obsolete data sets after you delete the previous release from your system.

- SAQTPKGI and AAQTPKGI

The following file system paths, which were created and used by previous releases of this product, are no longer used in this release. You can delete these obsolete file system paths after you delete the previous release from your system.

- /usr/lpp/IBM/aqt/v7r1m0

The following DDDEF entries, which were created and used by previous releases of this product, are no longer used in this release. You can delete these obsolete DDDEF entries after you delete the previous release from your system.

- SAQTPKGI and AAQTPKGI

6.2 Activating the Stored Procedures of Db2 Analytics Accelerator

If you mount the file system in which you have installed Stored Procedures in read-only mode during execution, then you do not have to take further actions to activate Stored Procedures.

6.3 Product Customization

The publications *IBM Db2 Analytics Accelerator for z/OS V8.1 User's Guide* and *IBM Db2 Analytics Accelerator for z/OS V8.1 Stored Procedures Reference* contain the necessary information to customize and use Db2 Analytics Accelerator and the Stored Procedures.

7.0 Notices

This information was developed for products and services offered in the U.S.A. IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

APAR numbers are provided in this document to assist in locating PTFs that may be required. Ongoing problem reporting may result in additional APARs being created. Therefore, the APAR lists in this document may not be complete. To obtain current service recommendations and to identify current product service requirements, always refer to the instructions in the **Service Recommendation Summary and Service Recommendations** and **Cross Product Dependencies** sections of the **PSP bucket information for IBM Z products** at https://www.ibm.com/support/pages/node/7127792, to ensure you have all required service.

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to the

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, New York 10504-1785
USA

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan, Ltd.
19-21, Nihonbashi-Hakozakicho, Chuo-ku
Tokyo 103-8510, Japan

7.1 Trademarks

IBM, the IBM logo, and other IBM trademark listed on the IBM Trademarks List are trademarks or registered trademarks of International Business Machines Corporation, in the United States and/or other countries. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on ibm.com/trademark.

Reader's Comments

Program Directory for IBM Db2 Analytics Accelerator for z/OS, June 2025 We appreciate your input on this publication. Feel free to comment on the clarity, accuracy, and completeness of the information or give us any other feedback that you might have.

Send your comments by emailing us at ibmdocs@us.ibm.com, and include the following information:

```
Your name
Your Company, University, or Institution
Your email address
The URL of the topic or web page you are commenting on (when viewing
publications on the IBM Documentation website.)
The publication title and publication form number
The topic and page number related to your comment
The text of your comment
```

If you have a technical problem, do not use the feedback method listed above for providing a reader's comments. Instead, take one of the following actions:

```
Contact your IBM service representative.
Call IBM technical support.
Open a support case at https://www.ibm.com/mysupport/s
```

Note: When you send information to IBM, you grant IBM a nonexclusive right to use or distribute the information in any way it believes appropriate without incurring any obligation to you.

IBM or any other organizations will only use the personal information that you supply to contact you about the issues that you submit.

Thank you for your participation. ˚

Printed in Ireland

GI13-6604-00