



**Program Directory for
IBM DB2 Analytics Accelerator for z/OS**

V05.01.00

Program Number 5697-DA5

FMIDs HAQT510, HCHCA21

for Use with
z/OS

HCHCA21 - Service Updated 31 July 2015

Document Date: December 2015

GI19-5018-00

Note

Before using this information and the product it supports, be sure to read the general information under 7.0, "Notices" on page 31.

Contents

1.0 Introduction	1
1.1 DB2 Analytics Accelerator Description	2
1.2 DB2 Analytics Accelerator FMIDs	4
2.0 Program Materials	5
2.1 Basic Machine-Readable Material	5
2.2 Optional Machine-Readable Material	6
2.3 Program Publications	7
2.3.1 Optional Program Publications	7
2.4 Program Source Materials	7
2.5 Publications Useful During Installation	7
3.0 Program Support	9
3.1 Program Services	9
3.2 Preventive Service Planning	9
3.3 Statement of Support Procedures	10
4.0 Program and Service Level Information	11
4.1 Program Level Information	11
4.2 Service Level Information	12
5.0 Installation Requirements and Considerations	13
5.1 Driving System Requirements	13
5.1.1 Machine Requirements	13
5.1.2 Programming Requirements	13
5.2 Target System Requirements	14
5.2.1 Machine Requirements	14
5.2.2 Programming Requirements	14
5.2.2.1 Installation Requisites	15
5.2.2.2 Operational Requisites	15
5.2.2.3 Toleration/Coexistence Requisites	16
5.2.2.4 Incompatibility (Negative) Requisites	16
5.2.3 DASD Storage Requirements	16
5.3 FMIDs Deleted	20
5.4 Special Considerations	20
6.0 Installation Instructions	21
6.1 Installing DB2 Analytics Accelerator	21
6.1.1 SMP/E Considerations for Installing DB2 Analytics Accelerator	21
6.1.2 SMP/E Options Subentry Values	21
6.1.3 SMP/E CALLLIBS Processing	22
6.1.4 Sample Jobs	22

6.1.5 Allocate SMP/E CSI (Optional)	24
6.1.6 Initialize CSI zones (Optional)	25
6.1.7 Perform SMP/E RECEIVE	25
6.1.8 Allocate SMP/E Target and Distribution Libraries	25
6.1.9 Allocate, create and mount ZFS Files (Optional)	25
6.1.10 Allocate File System Paths	26
6.1.11 Create DDDEF Entries	27
6.1.12 Perform SMP/E APPLY	27
6.1.13 Perform SMP/E ACCEPT	29
6.1.14 Run REPORT CROSSZONE	30
6.2 Activating DB2 Analytics Accelerator	30
6.2.1 File System Execution	30
6.2.2 Product Customization	30
7.0 Notices	31
7.1 Trademarks	31
Reader's Comments	32

Figures

1. Program File Content for DB2 Analytics Accelerator	5
2. Program File Content for InfoSphere Data Replication	6
3. Basic Material: Unlicensed	7
4. Publications Useful During Installation	8
5. PSP Upgrade and Subset ID	9
6. Component IDs	10
7. Driving System Software Requirements	14
8. Target System Mandatory Installation Requisites	15
9. Total DASD Space Required by DB2 Analytics Accelerator	16
10. Storage Requirements for DB2 Analytics Accelerator Target Libraries	18
11. Storage Requirements for InfoSphere Data Replication Target Libraries	18
12. DB2 Analytics Accelerator File System Paths	18
13. InfoSphere Data Replication File System Paths	19
14. Storage Requirements for DB2 Analytics Accelerator Distribution Libraries	19
15. Storage Requirements for InfoSphere Data Replication Distribution Libraries	19
16. SMP/E Options Subentry Values	21
17. Sample Installation Jobs	22

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 IBM DB2 Analytics Accelerator for z/OS. This publication refers to IBM DB2 Analytics Accelerator for z/OS as DB2 Analytics Accelerator.

The Program Directory contains the following sections:

- 2.0, "Program Materials" on page 5 identifies the basic program materials and documentation for DB2 Analytics Accelerator.
- 3.0, "Program Support" on page 9 describes the IBM support available for DB2 Analytics Accelerator.
- 4.0, "Program and Service Level Information" on page 11 lists the APARs (program level) and PTFs (service level) that have been incorporated into DB2 Analytics Accelerator.
- 5.0, "Installation Requirements and Considerations" on page 13 identifies the resources and considerations that are required for installing and using DB2 Analytics Accelerator.
- 6.0, "Installation Instructions" on page 21 provides detailed installation instructions for DB2 Analytics Accelerator. It also describes the procedures for activating the functions of DB2 Analytics Accelerator, or refers to appropriate publications.

For most recent information on prerequisites and installation information for DB2 Analytics Accelerator please check the document "IBM Prerequisites and Maintenance for IBM Analytics Accelerator for z/OS Version 5.1" on the IBM Support Portal at <http://www.ibm.com/support/docview.wss?uid=swg27046305>

Before installing DB2 Analytics Accelerator, 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; then keep them for future reference. Section 3.2, "Preventive Service Planning" on page 9 tells you how to find any updates to the information and procedures in this program directory.

DB2 Analytics Accelerator is supplied in a Custom-Built Product Delivery Offering (CBPDO, 5751-CS3). The program directory that is provided in softcopy format on the CBPDO tape 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 tape.

Do not use this program directory if you install DB2 Analytics Accelerator with a SystemPac or 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

IBM DB2 Analytics Accelerator for z/OS is a high-performance software appliance built to work with prerequisite hardware and which can dramatically accelerate the execution of DB2 for z/OS queries. It integrates IBM Netezza technology with z Systems technology to accelerate complex queries. This hybrid data server approach gives organizations the ability to run the entire spectrum of workloads, from transactional to complex analytics, within a single integrated environment. This provides the best of symmetric multiprocessing (SMP), through DB2 for z/OS , and the best of massively parallel processing (MPP), through the PureData, System for Analytics (powered by Netezza technology), which allows organizations to leverage their existing people, processing, and z Systems infrastructure.

DB2 Analytics Accelerator for z/OS delivers improved data insights to an organization, providing the following benefits:

- Significantly improves query performance and execution time
- Minimizes data latency, enabling near real-time applications
- Lowers the cost of storing, managing, and processing historical data
- Reduces capacity requirements for complex query workloads

DB2 Analytics Accelerator is a purpose-built, workload-optimized software appliance designed as an extension to the secure DB2 for z/OS environment. Acceleration occurs transparently to applications that submit queries to DB2 for z/OS without incurring z Systems usage charges. The DB2 for z/OS optimizer determines whether a query will run more efficiently in native DB2 or within the Accelerator and will direct the query accordingly.

DB2 Analytics Accelerator exploits two of the major trends in data warehousing and business analytics technology:

- Massively parallel processing
- Efficient data filtering by early SQL projections and restrictions through Field Programmable Gate Arrays (FPGAs)

DB2 Analytics Accelerator helps customers support their ever-changing analytic requirements, giving them the insight they need to maximize business performance. Customers using IBM DB2 Analytics Accelerator are achieving quick and dramatic improvement in the response times of eligible queries. The Accelerator offers faster business value at a lower cost by combining the simplicity and rapid deployment characteristics of an appliance with the benefits of data managed and secured by the proven database platform, DB2 for z/OS. This is opening up unprecedented opportunities to make use of the z Systems platform.

DB2 Analytics Accelerator is part of the end-to-end IBM z Systems analytic portfolio that offers data warehousing, data transformation and integration, business intelligence, and predictive analytics, including:

- DB2 for z/OS
- InfoSphere Information Server

- InfoSphere Data Replication
- Cognos BI
- DB2 Query Management Facility (QMF)
- Cognos TM1
- SPSS

What's new with DB2 Analytics Accelerator Version 5?

Version 5 of DB2 Analytics Accelerator opens up a new dimension of analytical processing by introducing **Accelerator-only tables**. Accelerator-only tables can benefit statistics and analytics tools that use temporary data for reports as the high velocity of execution allows these tools to quickly gather all required data.

Accelerator-only tables can also be used to accelerate data transformations, which are implemented via SQL statements. To take full advantage of the product's high-speed capabilities, it is now possible to store all interim results in accelerator-only tables. This allows subsequent queries or data transformations to process all relevant data on the accelerator, and as such to implement high speed in-database transformations (IDT).

The introduction of Accelerator-only tables extends the value of DB2 Analytics Accelerator to these additional usage scenarios:

- Deeper insight into operational status through faster reporting:
Multi-Step Reporting Applications, QMF, and home grown applications
- Accelerated campaign tuning for IBM Campaign (Unica):
Improved performance for iterative campaign tuning
- Simplifying data-transformation processes:
Delivering in-database transformation within DB2 Analytics Accelerator
- Simpler data integration:
DB2 Analytics Accelerator Loader for z/OS - load non-DB2 for z/OS data
- Deeper insight into customers and markets:
Data scientist work area
- Data mart consolidation through flexible data infrastructure:
Host data marts on z Systems, where the data originates

In summary:

It supports the full lifecycle of a real-time analytics solution on a single, integrated system combining transactional data, historical data and predictive analytics.

Additional new functionality which comes with V5:

- Analytics Server provides significantly faster table updates for real-time analytics.

- Augmented Security and improved Performance with new PureData Technology extend advanced security with 'encryption of data at rest' while taking advantage of the legendary built-in security of System z. Performance is improved up to 2X for CPU intensive workloads and load rates.
- Enhanced Systems Management Capabilities provide faster data synchronization and data load (Incremental Update) as well as improved serviceability with automatic 'Call Home' capability.

1.2 DB2 Analytics Accelerator FMIDs

DB2 Analytics Accelerator consists of the following FMIDs:

HAQT510
HCHCA21

2.0 Program Materials

An IBM program is identified by a program number. The program number for DB2 Analytics Accelerator is 5697-DA5.

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 21 for more information about how to install the program.

You can find information about the physical media for the basic machine-readable materials for DB2 Analytics Accelerator in the *CBPDO Memo To Users Extension*.

Figure 1 describes the program file content for DB2 Analytics Accelerator.

Figure 2 on page 6 describes the program file content for InfoSphere Data Replication. You can refer to the *CBPDO Memo To Users Extension* to see where the files reside on the tape.

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 DB2 Analytics Accelerator

Name	O R G	R E C F M	L R E C L	BLK SIZE
SMPMCS	SEQ	FB	80	6400
IBM.HAQT510.F1	PDS	FB	80	8800
IBM.HAQT510.F2	PDS	FB	80	8800
IBM.HAQT510.F3	PDSE	VB	1028	6144
IBM.HAQT510.F4	PDSE	VB	1028	6144
IBM.HAQT510.F5	PDSE	U	0	6144

Figure 1 (Page 2 of 2). Program File Content for DB2 Analytics Accelerator

Name	O R G	R E C F M	L R E C L	BLK SIZE
IBM.HAQT510.F6	PDSE	U	0	6144
IBM.HAQT510.F7	PDS	FB	80	8800

Figure 2. Program File Content for InfoSphere Data Replication

Name	O R G	R E C F M	L R E C L	BLK SIZE
SMPMCS	SEQ	FB	80	6400
IBM.HCHCA21.F1	PDS	FB	80	8800
IBM.HCHCA21.F2	PDSE	U	0	6144
IBM.HCHCA21.F3	PDS	FB	80	8800
IBM.HCHCA21.F4	PDS	FB	80	8800
IBM.HCHCA21.F5	PDS	FB	80	8800
IBM.HCHCA21.F6	PDS	FB	80	8800
IBM.HCHCA21.F7	PDS	FB	80	8800
IBM.HCHCA21.F8	PDS	VB	255	8800
IBM.HCHCA21.F9	PDS	VB	255	8800
IBM.HCHCA21.F10	PDS	FB	80	8800
IBM.HCHCA21.F11	PDS	VB	255	8800
IBM.HCHCA21.F12	PDS	VB	137	8800
IBM.HCHCA21.F13	PDS	VB	255	27998

2.2 Optional Machine-Readable Material

No optional machine-readable materials are provided for DB2 Analytics Accelerator.

2.3 Program Publications

The following sections identify the basic publications for DB2 Analytics Accelerator.

Figure 3 on page 7 identifies the basic unlicensed publications for DB2 Analytics Accelerator. Those that are in softcopy format publications can be obtained from the IBM Publications Center website at: <http://www.ibm.com/shop/publications/order/>

<i>Figure 3. Basic Material: Unlicensed</i>		
Publication Title	Form Number	Media Format
IBM DB2 Analytics Accelerator for z/OS License Information Agreement CD	LCD7-2644	http://www.ibm.com/software/sla/sladb.nsf
IBM DB2 Analytics Accelerator Studio User's Guide	SH12-7061	http://publibfp.dhe.ibm.com/epubs/pdf/h1270610.pdf
IBM DB2 Analytics Accelerator for z/OS Quick Start Guide	GH12-7059	http://publibfp.dhe.ibm.com/epubs/pdf/h1270590.pdf
IBM DB2 Analytics Accelerator for z/OS Installation Guide	SH12-7062	http://publibfp.dhe.ibm.com/epubs/pdf/h1270620.pdf
IBM DB2 Analytics Accelerator for z/OS Stored Procedures Reference	SH12-7063	http://publibfp.dhe.ibm.com/epubs/pdf/h1270630.pdf
IBM DB2 Analytics Accelerator for z/OS Getting Started	GH12-7060	http://publibfp.dhe.ibm.com/epubs/pdf/h1270600.pdf

2.3.1 Optional Program Publications

No optional publications are provided for DB2 Analytics Accelerator.

2.4 Program Source Materials

No program source materials or viewable program listings are provided for DB2 Analytics Accelerator.

2.5 Publications Useful During Installation

You might want to use the publications listed in Figure 4 during the installation of DB2 Analytics Accelerator.

Figure 4. Publications Useful During Installation

Publication Title	Form Number	Media Format
<i>IBM SMP/E for z/OS User's Guide</i>	SA23-2277	http://www.ibm.com/shop/publications/order/
<i>IBM SMP/E for z/OS Commands</i>	SA23-2275	http://www.ibm.com/shop/publications/order/
<i>IBM SMP/E for z/OS Reference</i>	SA23-2276	http://www.ibm.com/shop/publications/order/
<i>IBM SMP/E for z/OS Messages, Codes, and Diagnosis</i>	GA32-0883	http://www.ibm.com/shop/publications/order/

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 have reviewed the current Preventive Service Planning (PSP) information. Review the PSP Bucket for General Information, Installation Documentation, and the Cross Product Dependencies sections. For the Recommended Service section, instead of reviewing the PSP Bucket, it is recommended you use the `IBM.ProductInstall-RequiredService` fix category in SMP/E to ensure you have all the recommended service installed. Use the **FIXCAT(IBM.ProductInstall-RequiredService)** operand on the **APPLY CHECK** command. See 6.1.12, “Perform SMP/E APPLY” on page 27 for a sample APPLY command.

If you obtained DB2 Analytics Accelerator as part of a CBPDO, HOLDDATA is included.

If the CBPDO for DB2 Analytics Accelerator is older than two weeks by the time you install the product materials, you can obtain the latest PSP Bucket information by going to the following website:

<http://www14.software.ibm.com/webapp/set2/psearch/search?domain=psp>

You can also use S/390 SoftwareXcel or contact the IBM Support Center to obtain the latest PSP Bucket information.

For program support, access the Software Support Website at <http://www-01.ibm.com/software/support/>.

PSP Buckets are identified by UPGRADEs, which specify product levels; and SUBSETs, which specify the FMIDs for a product level. The UPGRADE and SUBSET values for DB2 Analytics Accelerator are included in Figure 5.

Figure 5. PSP Upgrade and Subset ID

UPGRADE	SUBSET	Description
5697DA5	HAQT510	DB2 Analytics Accelerator
5655U96	HCHCA21/1531	InfoSphere Data Replication

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 6 on page 10 identifies the component IDs (COMPID) for DB2 Analytics Accelerator.

<i>Figure 6. Component IDs</i>			
FMID	COMPID	Component Name	RETAIN Release
HAQT510	5697AQT00	DB2 Analytics Accelerator	510
HCHCA21	5655U7600	InfoSphere Data Replication	A21

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 HAQT410

PI06638	PI19269	PI31832
PI09547	PI27503	PI31833
PI12614	PI27518	PI31834
PI12761	PI27523	PI31838
PI12771	PI27531	PI31839
PI12788	PI27537	PI35803
PI12797	PI28282	PI35806
PI12808	PI28787	PI35808
PI12810	PI29638	PI35813
PI13105	PI29639	PI35825
PI17099	PI29640	PM92398
PI19232	PI29641	PM92399
PI19238	PI29643	PM92402
PI19243	PI30282	PM92403
PI19246	PI30283	PM92404
PI19261	PI30285	PM92405
PI19267		

- FMID HCHCA20

PM84163	PM88648	PM93405
PM84370	PM88897	PM94430
PM84378	PM89846	PM95092
PM86096	PM90789	PM96087
PM86126		

4.2 Service Level Information

PTFs containing APAR fixes against this release of InfoSphere Data Replication (HCHCA21) have been incorporated into this product package. For a list of included PTFs, examine the ++VER statement in the product's SMPMCS.

Frequently check the DB2 Analytics Accelerator PSP Bucket for HIPER and SPECIAL attention PTFs against all FMIDs that you must install. You can also receive the latest HOLDDATA, then add the **FIXCAT(IBM.PRODUCTINSTALL-REQUIRESERVICE)** operand on your APPLY CHECK command. This will allow you to review the recommended and critical service that should be installed with your FMIDs.

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 7. 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:				
5694-A01	z/OS	V01.13.00	N/A	No
5650-ZOS	z/OS	V02.01.00	N/A	No

Note: SMP/E is a requirement for Installation and is an element of z/OS but can also be ordered as a separate product, 5655-G44, minimally V03.06.00.

Note: Installation might require migration to new z/OS releases to be service supported. See http://www-03.ibm.com/systems/z/os/zos/support/zos_eos_dates.html.

DB2 Analytics Accelerator is installed into a file system, either HFS or zFS. 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 system data sets must be mounted on the driving system.

If you plan to install DB2 Analytics Accelerator in a zFS file system, this requires that zFS 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.

Figure 8. Target System Mandatory Installation Requisites

Program Number	Product Name	Minimum VRM	Minimum Service Level will satisfy these APARs	Included in the shipped product?
Any one of the following:				
5605-DB2	DB2 for z/OS	V10.01.00	N/A	No
5697-P31	DB2 for z/OS Value Unit Edition	V10.01.00	N/A	No
5615-DB2	DB2 for z/OS	V11.01.00	N/A	No
5697-P43	DB2 for z/OS Value Unit Edition	V11.01.00	N/A	No

Note: Installation might require migration to new z/OS releases to be service supported. See http://www-03.ibm.com/systems/z/os/zos/support/zos_eos_dates.html.

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.

DB2 Analytics Accelerator has no mandatory operational requisites.

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.

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

<i>Figure 9. Total DASD Space Required by DB2 Analytics Accelerator</i>	
Library Type	Total Space Required in 3390 Trks
Target	1865 tracks for DB2 Analytics Accelerator
	854 tracks for InfoSphere Data Replication
Distribution	79865 tracks for DB2 Analytics Accelerator
	1632 tracks for InfoSphere Data Replication
File System(s)	105 000 tracks for /usr/lpp/IBM/aqt/v5r1m0
	750 tracks for /usr/lpp/chc/V10R2M1

Notes:

- 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.
- 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 6.1.8, “Allocate SMP/E Target and Distribution Libraries” on page 25.

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.

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, except for SAQTICU, SAQTMOD, SCHCLOAD, AAQTHFS, AAQTHFS1, AAQTICU, AAQTMOD, and ACHCLOAD, which must be PDSEs.

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 LNKLIST.
- 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.

Figure 10. Storage Requirements for DB2 Analytics Accelerator Target Libraries

Library DDNAME	Member Type	Target Volume	T Y P E	O R G	R E C F M	L R E C L	No. of 3390 Trks	No. of DIR Blks
SAQTBASE	SAMP	Any	U	PDS	FB	80	6	5
SAQTDBRM	MAC	Any	U	PDS	FB	80	3	5
SAQTICU	PROGRAM	Any	U	PDSE	U	0	10	N/A
SAQTMOD	PROGRAM	Any	U	PDSE	U	0	1800	N/A
SAQTSAMP	SAMP	Any	U	PDS	FB	80	46	5

Figure 11. Storage Requirements for InfoSphere Data Replication Target Libraries

Library DDNAME	Member Type	Target Volume	T Y P E	O R G	R E C F M	L R E C L	No. of 3390 Trks	No. of DIR Blks
SCHCASM	SAMP	Any	U	PDS	FB	80	3	3
SCHCC	SAMP	Any	U	PDS	VB	256	5	3
SCHCCNTL	SAMP	Any	U	PDS	FB	80	21	5
SCHCCOB	SAMP	Any	U	PDS	FB	80	6	3
SCHCDATA	DATA	Any	U	PDS	FB	80	6	3
SCHCDBRM	DATA	Any	U	PDS	FB	80	19	5
SCHCH	SAMP	Any	U	PDS	VB	256	3	3
SCHCLOAD	MOD	Any	U	PDSE	U	0	776	N/A
SCHCMAC	SAMP	Any	U	PDS	FB	80	3	3
SCHCNOTC	SAMP	Any	U	PDS	VB	137	6	1
SCHCTTL	DATA	Any	U	PDS	VB	256	6	3

Figure 12. DB2 Analytics Accelerator File System Paths

DDNAME	T Y P E	Path Name
SAQTHFS	X	/usr/lpp/IBM/aqt/v5r1m0/IBM

Figure 13. InfoSphere Data Replication File System Paths

DDNAME	T Y P E	Path Name
SCHCJAVA	N	/usr/lpp/chc/V10R2M1/IBM

Figure 14. Storage Requirements for DB2 Analytics Accelerator Distribution Libraries

Library DDNAME	T Y P E	O R G	R E C F M	L R E C L	No. of 3390 Trks	No. of DIR Blks
AAQTBASE	U	PDS	FB	80	6	5
AAQTDBRM	U	PDS	FB	80	3	5
AAQTHFS	U	PDSE	VB	1028	43500	N/A
AAQTHFS1	U	PDSE	VB	1028	34500	N/A
AAQTICU	U	PDSE	U	0	10	N/A
AAQTMOD	U	PDSE	U	0	1800	N/A
AAQTSAMP	U	PDS	FB	80	46	5

Figure 15 (Page 1 of 2). Storage Requirements for InfoSphere Data Replication Distribution Libraries

Library DDNAME	T Y P E	O R G	R E C F M	L R E C L	No. of 3390 Trks	No. of DIR Blks
ACHCASM	U	PDS	FB	80	3	3
ACHCC	U	PDS	VB	256	5	3
ACHCCNTL	U	PDS	FB	80	21	5
ACHCCOB	U	PDS	FB	80	6	3
ACHCDATA	U	PDS	FB	80	6	3
ACHCDBRM	U	PDS	FB	80	19	5
ACHCH	U	PDS	VB	256	3	5
ACHCJAVA	U	PDS	VB	255	541	3
ACHCLOAD	U	PDSE	U	0	1013	N/A
ACHCMAC	U	PDS	FB	80	3	3
ACHCNOTC	U	PDS	VB	137	6	1

Figure 15 (Page 2 of 2). Storage Requirements for InfoSphere Data Replication Distribution Libraries

Library DDNAME	T Y P E	O R G	R E C F M	L R E C L	No. of 3390 Trks	No. of DIR Blks
ACHCTTL	U	PDS	VB	256	6	3

5.3 FMIDs Deleted

Installing DB2 Analytics Accelerator 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 DB2 Analytics Accelerator 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, SAQTICU, and SCHCLOAD 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, SAQTICU, and SCHCLOAD 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, SAQTICU, and SCHCLOAD 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 DB2 Analytics Accelerator.

Please note the following points:

- If you want to install DB2 Analytics Accelerator 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 DB2 Analytics Accelerator

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 16. 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.

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:

<i>Figure 17 (Page 1 of 2). Sample Installation Jobs</i>			
Job Name	Job Type	Description	RELFILE
AQTALA	SMP/E	Sample job to allocate and initialize a new SMP/E CSI data set (Optional)	IBM.HAQT510.F1
AQTALB	SMP/E	Sample job to allocate SMP/E data sets (Optional)	IBM.HAQT510.F1
AQTRECEV	RECEIVE	Sample RECEIVE job for DB2 Analytics Accelerator	IBM.HAQT510.F1
CHCRECEV	RECEIVE	Sample RECEIVE job for InfoSphere Data Replication	IBM.HCHCA21.F4
AQTALLOC	ALLOCATE	Sample job to allocate target and distribution libraries for DB2 Analytics Accelerator	IBM.HAQT510.F1
CHCALLOC	ALLOCATE	Sample job to allocate target and distribution libraries for InfoSphere Data Replication	IBM.HCHCA21.F4
AQTZFS	ALLOMTZFS	Sample job to allocate new ZFS data set for DB2 Analytics Accelerator (Optional)	IBM.HAQT510.F1
CHCZFS	ALLOMTZFS	Sample job to allocate new ZFS data set for InfoSphere Data Replication (Optional)	IBM.HCHCA21.F4
AQTISMKD	MKDIR	Sample job to invoke the supplied AQTMKDIR EXEC to allocate HFS or zFS paths for DB2 Analytics Accelerator	IBM.HAQT510.F1

Figure 17 (Page 2 of 2). Sample Installation Jobs

Job Name	Job Type	Description	RELFILE
CHCISMKD	MKDIR	Sample job to invoke the supplied CHCMKDIR EXEC to allocate HFS or zFS paths for InfoSphere Data Replication	IBM.HCHCA21.F4
AQTDDDEF	DDDEF	Sample job to define SMP/E DDDEFs for DB2 Analytics Accelerator	IBM.HAQT510.F1
CHCDDDEF	DDDEF	Sample job to define SMP/E DDDEFs for InfoSphere Data Replication	IBM.HCHCA21.F4
AQTAPPLY	APPLY	Sample APPLY job	IBM.HAQT510.F1
AQTACCEP	ACCEPT	Sample ACCEPT job	IBM.HAQT510.F1

You can access the sample installation jobs by performing an SMP/E RECEIVE (refer to 6.1.7, “Perform SMP/E RECEIVE” on page 25) then copy the jobs from the RELFILES to a work data set for editing and submission. See Figure 17 on page 22 to find the appropriate relfile data set.

You can also copy the sample installation jobs from the tape or product files by submitting the following job. Depending on your distribution medium, use either the //TAPEIN or the //FILEIN DD statement and comment out or delete the other statement. 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=*
//*****
//* Make the //TAPEIN DD statement below active if you install*
//* from a CBPDO tape by uncommenting the DD statement below. *
//*****
//*TAPEIN DD DSN=IBM.HAQT510.F1,UNIT=tunit,
//* VOL=SER=volser,LABEL=(x,SL),
//* DISP=(OLD,KEEP)
//*TAPEIN2 DD DSN=IBM.HCHCA21.F4,UNIT=tunit,
//* VOL=SER=volser,LABEL=(x,SL),
//* DISP=(OLD,KEEP)
//*****
//* Make the //TAPEIN DD statement below active if you install*
//* from a product tape received outside the CBPDO process *
//* (using the optional SMP/E RECEIVE job) by uncommenting *
//* the DD statement below. *
//*****
//*TAPEIN DD DSN=IBM.HAQT510.F1,UNIT=tunit,
//* VOL=SER=AQT510,LABEL=(2,SL),
//* DISP=(OLD,KEEP)
//*TAPEIN2 DD DSN=IBM.HCHCA21.F4,UNIT=tunit,
//* VOL=SER=CHCA21,LABEL=(5,SL),
//* DISP=(OLD,KEEP)
//*****
```

```

/* Make the //FILEIN DD statement below active for          *
/* downloaded DASD files.                                  *
/******
/*FILEIN DD DSN=IBM.HAQT510.F1,UNIT=SYSALLDA,DISP=SHR,
/* VOL=SER=filevol
/*FILEIN2 DD DSN=IBM.HCHCA21.F4,UNIT=SYSALLDA,DISP=SHR,
/* VOL=SER=filevol
//OUT DD DSNAME=jcl-library-name,
// DISP=(NEW,CATLG,DELETE),
// VOL=SER=dasdvol,UNIT=SYSALLDA,
// SPACE=(TRK,(20,10,5))
//SYSUT3 DD UNIT=SYSALLDA,SPACE=(CYL,(1,1))
//SYSIN DD *
COPY INDD=xxxxIN,OUTDD=OUT
COPY INDD=xxxxIN2,OUTDD=OUT
SELECT MEMBER=(CHCALLOC,CHCDDDEF,CHCRECEV)
SELECT MEMBER=(CHCISMKD,CHCMKDIR,CHCZFS)
/*

```

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

TAPEIN and TAPEIN2 :

tunit is the unit value that matches the product package.

volser is the volume serial that matches the product package.

x is the tape file number that indicates the location of the data set name on the tape.

See the documentation that is provided by CBPDO for the location of IBM.HAQT510.F1, and IBM.HCHCA21.F4 on the tape(s).

FILEIN and FILEIN2:

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.

SYSIN:

xxxxIN is either TAPEIN or FILEIN depending on your input DD statement.

xxxxIN2 is either TAPEIN2 or FILEIN2 depending on your input DD statement.

6.1.5 Allocate SMP/E CSI (Optional)

If you are using an existing CSI, do not execute this job.

If you are allocating a new SMP/E data set for this install, edit and submit sample job AQTALA to allocate the SMP/E data set for DB2 Analytics Accelerator. 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.

6.1.6 Initialize CSI zones (Optional)

If you are using an existing CSI, do not execute this job.

Edit and submit sample job AQTALB to initialize SMP/E zones for DB2 Analytics Accelerator. 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.

6.1.7 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.

You can also choose to edit and submit sample job AQTRECEV to perform the SMP/E RECEIVE for DB2 Analytics Accelerator. 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.

You can also choose to edit and submit sample job CHCRECEV to perform the SMP/E RECEIVE for InfoSphere Data Replication. 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.

6.1.8 Allocate SMP/E Target and Distribution Libraries

Edit and submit sample job AQTALLOC to allocate the SMP/E target and distribution libraries for DB2 Analytics Accelerator. Consult the instructions in the sample job for more information.

Edit and submit sample job CHCALLOC to allocate the SMP/E target and distribution libraries for InfoSphere Data Replication. Consult the instructions in the sample job for more information.

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

6.1.9 Allocate, create and mount ZFS Files (Optional)

This job allocates, creates a mountpoint, and mounts zFS data sets.

If you plan to install DB2 Analytics Accelerator 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.

If you plan to install InfoSphere Data Replication into a new z/OS UNIX file system, you can edit and submit the optional CHCZFS 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 points are:

- `/usr/lpp/IBM/aqt/v5r1m0` for DB2 Analytics Accelerator.
- `/usr/lpp/chc/V10R2M1` for InfoSphere Data Replication.

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 DB2 Analytics Accelerator, and InfoSphere Data Replication. 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 - either `/usr/lpp/IBM/aqt/v5r1m0` or `/usr/lpp/chc/V10R2M1`.

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

6.1.10 Allocate File System Paths

The target system HFS or zFS data set must be mounted on the driving system when running the sample jobs, AQTISMKD, and CHCISMKD, since these jobs will create paths in theHFS or zFS.

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 HFS or zFS file system is mounted to the driving system. zFS must be active on the driving system if you are installing DB2 Analytics Accelerator, and InfoSphere Data Replication, into a file system that is zFS.

If you plan to install DB2 Analytics Accelerator into a new HFS or zFS file system, you must create the mountpoint and mount the new file system to the driving system for DB2 Analytics Accelerator.

If you plan to install InfoSphere Data Replication into a new HFS or zFS file system, you must create the mountpoint and mount the new file system to the driving system for InfoSphere Data Replication.

The recommended mountpoint for DB2 Analytics Accelerator is `/usr/lpp/IBM/aqt/v5r1m0`.
The recommended mountpoint for InfoSphere Data Replication is `/usr/lpp/chc/V10R2M1`.

Edit and submit sample job AQTISMKD to allocate the HFS or zFS paths for DB2 Analytics Accelerator. Consult the instructions in the sample job for more information.

Edit and submit sample job CHCISMKD to allocate the HFS or zFS paths for InfoSphere Data Replication. 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 these jobs runs correctly.

6.1.11 Create DDDEF Entries

Edit and submit sample job AQTDDDEF to create DDDEF entries for the SMP/E target and distribution libraries for DB2 Analytics Accelerator. Consult the instructions in the sample job for more information.

Edit and submit sample job CHCDDDEF to create DDDEF entries for the SMP/E target and distribution libraries for InfoSphere Data Replication. Consult the instructions in the sample job for more information.

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

6.1.12 Perform SMP/E APPLY

1. Ensure that you have the latest HOLDDATA; then edit and submit sample job AQTAPPLY to perform an SMP/E APPLY CHECK for DB2 Analytics Accelerator. Consult the instructions in the sample job for more information.

The latest HOLDDATA is available through several different portals, including <http://service.software.ibm.com/holddata/390holddata.html>. 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),HOLDFIXCAT) .
..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.

6.1.13 Perform SMP/E ACCEPT

Edit and submit sample job AQTACCEP to perform an SMP/E ACCEPT CHECK for DB2 Analytics Accelerator. 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.

6.1.14 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 DB2 Analytics Accelerator, 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.

6.2 Activating DB2 Analytics Accelerator

6.2.1 File System Execution

If you mount the file system in which you have installed DB2 Analytics Accelerator in read-only mode during execution, then you do not have to take further actions to activate DB2 Analytics Accelerator.

6.2.2 Product Customization

The publication *IBM DB2 Analytics Accelerator Studio User's Guide* (SH12-7061) contains the necessary information to customize and use DB2 Analytics Accelerator.

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 contact the IBM Customer Support Center or use S/390 SoftwareXcel to obtain the current "PSP Bucket".

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 ibm.com are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at www.ibm.com/legal/copytrade.shtml.

Reader's Comments

Program Directory for IBM DB2 Analytics Accelerator for z/OS, December 2015

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.

Use one of the following methods to send us your comments:

1. Send an email to comments@us.ibm.com
2. Use the form on the Web at:

www.ibm.com/software/data/rcf

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 USA

G119-5018-00

