



Program Directory for IBM Db2 AI for z/OS

V01.04.00

Program Number 5698-CGN

FMID HCOY120

for Use with
z/OS V2.3 or higher

Document Date: July 2021

GI13-5007-06

Note

Before using this information and the product it supports, be sure to read the general information under 7.0, "Notices" on page 27.

A form for reader's comments appears at the back of this publication. 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.

© **Copyright International Business Machines Corporation 2018, 2021.**

Note to U.S. Government Users — Documentation related to restricted rights — Use, duplication or disclosure is subject to restrictions set forth in GSA ADP Schedule Contract with IBM Corp.

Contents

1.0 Introduction	1
1.1 IBM Db2 AI for z/OS V1.4 Description	1
1.2 IBM Db2 AI for z/OS V1.4 FMIDs	3
2.0 Program Materials	4
2.1 Basic Machine-Readable Material	4
2.2 Program Publications	5
2.2.1 Optional Program Publications	5
2.3 Program Source Materials	5
2.4 Publications Useful During Installation	5
3.0 Program Support	7
3.1 Program Services	7
3.2 Preventive Service Planning	7
3.3 Statement of Support Procedures	8
4.0 Program and Service Level Information	9
4.1 Program Level Information	9
4.2 Service Level Information	9
5.0 Installation Requirements and Considerations	10
5.1 Driving System Requirements	10
5.1.1 Machine Requirements	10
5.1.2 Programming Requirements	10
5.2 Target System Requirements	11
5.2.1 Machine Requirements	12
5.2.2 Programming Requirements	12
5.2.2.1 Installation Requisites	12
5.2.2.2 Operational Requisites	13
5.2.2.3 Toleration/Coexistence Requisites	14
5.2.2.4 Incompatibility (Negative) Requisites	14
5.2.3 DASD Storage Requirements	14
5.3 FMIDs Deleted	17
5.4 Special Considerations	17
5.4.1 Dynamic Enablement - Update PARMLIB	17
5.4.2 Ordering details	18
6.0 Installation Instructions	19
6.1 Installing Db2ZAI	19
6.1.1 SMP/E Considerations for Installing Db2ZAI	19
6.1.2 SMP/E Options Subentry Values	19
6.1.3 Sample Jobs	19

6.1.4	Perform SMP/E RECEIVE	21
6.1.5	Allocate SMP/E Target and Distribution Libraries	21
6.1.6	Allocate, create and mount ZFS Files (Optional)	21
6.1.7	Allocate File System Paths	22
6.1.8	Create DDDEF Entries	22
6.1.9	Perform SMP/E APPLY	22
6.1.10	Perform SMP/E ACCEPT	24
6.1.11	Run REPORT CROSSZONE	25
6.1.12	Dynamic Enablement - Update PARMLIB - reminder	25
6.2	Activating Db2ZAI	26
6.2.1	File System Execution	26
6.2.2	Dynamic Enablement - Update PARMLIB - reminder	26
6.3	Product Customization	26
7.0	Notices	27
7.1	Trademarks	27
	Reader's Comments	28

Figures

1.	Program File Content	4
2.	Basic Material: Unlicensed Publications	5
3.	Publications Useful During Installation	6
4.	PSP Upgrade and Subset ID	7
5.	Component IDs	8
6.	Driving System Software Requirements	11
7.	Target System Mandatory Operational Requisites	13
8.	Total DASD Space Required by Db2ZAI	14
9.	Storage Requirements for SMP/E Data Sets	16
10.	Storage Requirements for Db2ZAI Target Libraries	16
11.	Db2ZAI File System Paths	16
12.	Storage Requirements for Db2ZAI Distribution Libraries	17
13.	SMP/E Options Subentry Values	19
14.	Sample Installation Jobs	20

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 AI for z/OS. This publication refers to IBM Db2 AI for z/OS as Db2ZAI.

The Program Directory contains the following sections:

- 2.0, “Program Materials” on page 4 identifies the basic program materials and documentation for Db2ZAI.
- 3.0, “Program Support” on page 7 describes the IBM support available for Db2ZAI.
- 4.0, “Program and Service Level Information” on page 9 lists the APARs (program level) and PTFs (service level) that have been incorporated into Db2ZAI.
- 5.0, “Installation Requirements and Considerations” on page 10 identifies the resources and considerations that are required for installing and using Db2ZAI.
- 6.0, “Installation Instructions” on page 19 provides detailed installation instructions for Db2ZAI. It also describes the procedures for activating the functions of Db2ZAI, or refers to appropriate publications.

Before installing Db2ZAI, 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.

Db2ZAI is supplied in a Custom-Built Product Delivery Offering (CBPDO, 5751-CS3). The program directories that are provided in softcopy format on the CBPDO are identical to the hardcopy format if they were included with your order. All service and HOLDDATA for Db2ZAI are included on the CBPDO.

Do not use this program directory if you install Db2ZAI 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 these program directories as needed.

1.1 IBM Db2 AI for z/OS V1.4 Description

IBM continues to infuse AI into Db2 for z/OS with Db2 AI for z/OS 1.4. With this new release, Db2 for z/OS will operate smarter and more efficiently, helping to improve Db2 application performance with these additional key features:

- Enterprise security readiness:
 - Designed to help simplify user management for Db2 AI for z/OS through the use of default Db2 Resource Access Control Facility (RACF) methods, which can help clients improve their security and audit requirements for user management
 - Supports multi-factor authentication.

- Sysplex Support and simplified migration for Db2 AI for z/OS:
 - Enhanced high-availability support for business-critical and large workloads to help eliminate single point of failure.
 - Designed to simplify the migration process for clients upgrading to a newer Db2ZAI release
 - SQL Optimization for accurate and faster database queries:
 - Designed to improve the automated SQL regression detection and resolution feature, enabling regressions to be detected and resolved more quickly, and in some cases in real-time, even before the query has completed execution
 - Expands sort optimization to OLTP queries, which can result in savings of CPU resources
 - Improved sort optimization when combined with the IBM z15 hardware sort assist for SQL, whereby learning can improve exploitation of the z15 sort feature.
 - Improved UI to help simplify the value that Db2ZAI is providing for SQL optimization by highlighting progress, benefits, and actions.
 - *Progress* indicates the amount of progress that Db2ZAI has made in learning about static SQL packages and dynamic SQLs.
 - *Benefits* indicates the savings that have been provided by Db2ZAI for the SQLs through which progress has been made
 - *Actions* indicates the recommended actions for the user to take that help accelerate Db2ZAI's progress.
 - Db2ZAI SQL optimization UI supports limiting the scope of the view to the most important CPU-consuming packages or SQL, and also limited by user-specified time windows.
- Db2 System Health Assessment:
 - Correlates data from System Assessment and SQL Optimization by identifying SQL statements that might be contributing to System Assessment exceptions. New drill-down capability from System Assessment graph is designed to assist clients in determining the cause of system-level anomalies and help speed resolution.
 - This function delivers the use case for row identifier (RID) exceptions where the graphs directly highlight SQL statements involved in system-level RID exceptions.
 - The correlation to SQL execution data can help identify SQL statements that were executing during the exception.
 - Enables clients to map data collected from SQL performance monitors that is stored in Db2 tables to more accurately correlate System Assessment exceptions to SQL activity.
 - Provides enhanced visual interpretation of exception thresholds used for the assessment for each work period for user consumption.
 - Provides an option to hide and unhide the exceptions reported by System Assessment, improving clients' capability to manage exception lists.
- Distributed Connection Control:
 - Extends protection for Db2 for z/OS subsystems to control access from dynamic IP addresses in addition to the existing static IP address protection. Clients can set and monitor threshold by dynamic IP addresses.

- Helps improve visualization for statistics graphs that yields data that is more meaningful and easier to use. This enables DBAs to gain insight on how remote applications are impacting Db2, and to identify which applications are running less optimally, to prevent potential outages.

1.2 IBM Db2 AI for z/OS V1.4 FMIDs

Db2ZAI delivers 5 FMIDs, 3 of which are required by Db2ZAI 1.4, and 2 which are not required by Db2ZAI 1.4, as described below.

- Db2 AI for z/OS, V1.4 FMIDs:
 - HCOY120 - Db2 AI for z/OS: Db2ZAI base (*required*)
 - HAQN230 - Watson Machine Learning: WMLz base (*required*)
 - HSPK120 - Open Data Analytics: Spark (*required*)
 - HANA110 - Open Data Analytics: Anaconda (*not required*)
 - HMDS120 - Open Data Analytics: MDS (*not required*)

Note

FMID HCOY120 was used to deliver Db2ZAI V1.2, V1.3, and now V1.4.

To upgrade to Db2ZAI V1.4 the following SMP/E work is required:

- SMP/E install FMID HCOY120 (if you have not already done so) and apply the HCOY120 PTF maintenance specified in the Mandatory Operational Requisites section of this Program Directory, as well as any subsequent PTF maintenance specified in IBM Documentation for Db2 AI for z/OS 1.4.
- SMP/E install FMID HAQN230 (the new WMLz V2.3 release) and apply any HAQN230 PTF maintenance specified in the Mandatory Operational Requisites section of the WMLz 2.3 Program Directory (included in this Product Package) as well as any subsequent PTF maintenance specified in IBM Documentation for Db2 AI for z/OS V1.4.
- SMP/E install FMID HSPK120 (if you have not already done so) and apply any HSPK120 PTF maintenance specified in the Mandatory Operational Requisites section of the Open Data Analytics Program Directory (included in this Product Package) as well as any subsequent PTF maintenance specified in IBM Documentation for Db2 AI for z/OS, V1.4.

Ensure you also have installed any prerequisite software at the maintenance level identified in the Mandatory Operational Requirements section for any of the above 3 FMIDs, or as specified in IBM Documentation for Db2 AI for z/OS, V1.4.

2.0 Program Materials

An IBM program is identified by a program number. The program number for Db2ZAI is 5698-CGN.

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 Db2ZAI. 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 19 for more information about how to install the program.

You can find information about the physical media for the basic machine-readable materials for Db2ZAI in the *CBPDO Memo To Users Extension*.

Figure 1 describes the program file content for Db2ZAI. 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. Program File Content

Name	ORG	RECFM	RECL	BLK SIZE
SMPMCS	SEQ	FB	80	6400
IBM.HCOY120.F1	PS	VS	80	8820
IBM.HCOY120.F2	PS	VS	32650	32654

2.2 Program Publications

The following sections identify the basic publications for Db2ZAI.

Figure 2 on page 5 identifies the basic unlicensed publications for Db2ZAI. 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 2. Basic Material: Unlicensed Publications</i>		
Publication Title	Form Number	Media Format
<i>IBM Db2 AI for z/OS, V1.4 License Information</i>	GC27-9135	Electronic
<i>IBM Open Data Analytics for z/OS License Information</i>	G113-4349	Electronic
<i>IBM Watson Machine Learning for z/OS License Information</i>	GC34-7802	Electronic
<i>IBM Db2 AI for z/OS V1.4 Program Directory</i>	G113-5007	Electronic
<i>IBM Open Data Analytics for z/OS Program Directory</i>	G113-4348	Electronic
<i>IBM Watson Machine Learning for z/OS Program Directory</i>	G113-5011	Electronic

To view the Unlicensed Publications content of the products that constitute Db2ZAI, refer to the following Program Directories included in this Product Package:

- G113-5011 - IBM Watson Machine Learning for z/OS, V2.3
- G113-4348 - IBM Open Data Analytics for z/OS, V1.1

You can also access additional product documentation for IBM Db2 AI for z/OS, IBM Watson Machine Learning for z/OS and IBM Open Data Analytics for z/OS in IBM Documentation.

2.2.1 Optional Program Publications

No optional publications are provided for IBM Db2 AI for z/OS.

2.3 Program Source Materials

No program source materials or viewable program listings are provided for Db2ZAI.

2.4 Publications Useful During Installation

You might want to use the publications listed in Figure 3 during the installation of Db2ZAI.

Figure 3. 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 Db2ZAI.

3.1 Program Services

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

3.2 Preventive Service Planning

Before you install Db2ZAI, 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-REQUIRESERVICE fix category in SMP/E to ensure you have all the recommended service installed. Use the **FIXCAT(IBM.PRODUCTINSTALL-REQUIRESERVICE)** operand on the **APPLY CHECK** command. See 6.1.9, “Perform SMP/E APPLY” on page 22 for a sample APPLY command

If you obtained Db2ZAI as part of a CBPDO, HOLDDATA is included.

If the CBPDO for Db2ZAI 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.ibm.com/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 Db2ZAI are included in Figure 4.

UPGRADE	SUBSET	Description
5698CGN	HCOY120/140	IBM Db2 AI for z/OS

To view the Preventive Service Planning (PSP) of the products that constitute Db2ZAI, refer to the following Program Directories included in this Product Package:

- GI13-5011 - IBM Watson Machine Learning for z/OS, V2.3
- GI13-4348 - IBM Open Data Analytics for z/OS, V1.1

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 on page 8 identifies the component IDs (COMPID) for Db2ZAI.

<i>Figure 5. Component IDs</i>			
FMID	COMPID	Component Name	RETAIN Release
HCOY120	5698CGN00	IBM Db2 AI for z/OS	140

To view this information for the products that constitute Db2ZAI, refer to the following Program Directories included in this Product Package:

- GI13-5011 - IBM Watson Machine Learning for z/OS, V2.3
- GI13-4348 - IBM Open Data Analytics for z/OS, V1.1

4.0 Program and Service Level Information

This section identifies the program and relevant service levels of Db2ZAI. 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

No APARs have been incorporated into Db2ZAI.

4.2 Service Level Information

Frequently check the Db2ZAI 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.

To view the Service Level Information for the products that constitute Db2ZAI, refer to the following Program Directories included in this Product Package:

- GI13-5011 - IBM Watson Machine Learning for z/OS, V2.3
- GI13-4348 - IBM Open Data Analytics for z/OS, V1.1

5.0 Installation Requirements and Considerations

The following sections identify the system requirements for installing and activating Db2ZAI. 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 Db2ZAI.

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?
5650-ZOS	z/OS	V02.03.00, or higher	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 https://www-01.ibm.com/software/support/lifecycle/index_z.html.

Db2ZAI invokes UNIX shell scripts during installation. These shell scripts require the following:

- UID(0) or READ access or higher to the BPX.SUPERUSER facility class
- Be connected to a group that has a GID with READ access or higher to the following:
 - BPX.FILEATTR.APF resource in the FACILITY class
 - BPX.FILEATTR.PROGCTL resource in the FACILITY class
 - BPX.FILEATTR.SHARELIB resource in the FACILITY class
- Have WRITE access to the /usr/lpp/IBM/db2zaiv1r2 path

Note: Db2ZAI V1.4 will reuse the db2zaiv1r2 path previously used by Db2ZAI V1.2 and Db2ZAI V1.3.

Db2ZAI is installed into a file system, either HFS or zFS. Before installing Db2ZAI, 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.

If you plan to install Db2ZAI 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 Db2ZAI.

The Db2zAI FMID installs in the DBS (P115) SREL.

IBM Watson Machine Learning for z/OS and IBM Open Data Analytics for z/OS FMIDs install in the Z/OS (Z038) SREL.

Note

You can install the required FMIDs in the same SMP/E Global zone, just ensure to specify both SREL P115 and Z038. See the SMP/E reference material in IBM Documentation for more information.

5.2.1 Machine Requirements

Db2 AI for z/OS requires zEnterprise EC12, or higher, which is the minimum hardware level supported by both Db2 12 for z/OS and Watson Machine Learning for z/OS, Version 2.3.

In addition to Watson Machine Learning for z/OS, V2.3 hardware requirements, Db2 AI for z/OS requires an additional 1 GB memory per Db2 subsystem (this includes additional buffer pool usage) and may require additional zIIPs capacity.

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. These products are specified as PREs or REQs.

Db2ZAI has no mandatory installation requisites.

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.

To view the Installation Requisites of the products that constitute Db2ZAI, refer to the following Program Directories included in this Product Package:

- GI13-5011 - IBM Watson Machine Learning for z/OS, V2.3
- GI13-4348 - IBM Open Data Analytics for z/OS, V1.1

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. These products are specified as IF REQs.

Db2ZAI 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.

<i>Figure 7. Target System Mandatory Operational Requisites</i>	
Program Number	Product Name and Minimum VRM/Service Level
5650-ZOS	z/OS V02.03.00 or higher, plus Product Registration Services
5698-CGN	IBM Db2 AI for z/OS, V01.04.00 plus APAR PH35963
5698-CGN	IBM Watson Machine Learning for z/OS, V02.03.00 (contains embedded Node.js V14 and bash 4.3.48)
5698-CGN	IBM Open Data Analytics, V1.1, FMID HSPK120/Spark, plus APARs PH03469, PH04837, PH24852, PH24854, PH27181 and PH30000
5655-DGH	IBM 64-bit SDK for z/OS, Java Technology Edition, V8, SR6 or higher - Ensure that JAVA_HOME is set as your installation directory
	z/OS OpenSSh
Any one of the following:	
5650-DB2	Db2 12 for z/OS, Function Level 500 activated, plus APARs PH15742, PH16111, PH24097, PH25342, PH28506, PH29766, PH30780, PH32464, PH33408, PH34549, PH35285 and PH35964
5770-AF3	Db2 12 for z/OS Value Unit Edition, Function Level 500 activated, plus APARs PH15742, PH16111, PH24097, PH25342, PH28506, PH29766, PH30780, PH32464, PH33408, PH34549, PH35285 and PH35964

To view the Mandatory Operational Requisites of the products that constitute Db2ZAI, refer to the following Program Directories included in this Product Package:

- GI13-5011 - IBM Watson Machine Learning for z/OS, V2.3
- GI13-4348 - IBM Open Data Analytics for z/OS, V1.1

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.

Db2ZAI 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.

Db2ZAI 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.

Db2ZAI has no negative requisites.

5.2.3 DASD Storage Requirements

Db2ZAI libraries can reside on all supported DASD types.

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

<i>Figure 8. Total DASD Space Required by Db2ZAI</i>	
Library Type	Total Space Required in 3390 Trks
Target	6 3390 tracks
Distribution	25000 3390 tracks
File System(s)	52500 3390 tracks

To view the DASD Storage Requirements of the products that constitute Db2ZAI, refer to the following Program Directories included in this Product Package:

- GI13-5011 - IBM Watson Machine Learning for z/OS, V2.3
- GI13-4348 - IBM Open Data Analytics for z/OS, V1.1

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 6.1.5, "Allocate SMP/E Target and Distribution Libraries" on page 21.

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, with some exceptions. 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 not be in the LPA, with some exceptions. If the value in the "Member Type" column specifies "LPA", it is advised to place the data set in the LPA.
- These data sets can be in the LNKLIST.
- These data sets are not required to be APF-authorized, with some exceptions. If the value in the "Member Type" column specifies "APF", the data set must be APF-authorized.

The following table provides an estimate of the storage needed in the SMP/E data sets for Db2ZAI. You must add the estimates to those of any other programs and service that you install to determine the total additional storage requirements.

<i>Figure 9. Storage Requirements for SMP/E Data Sets</i>						
Library DDNAME	T Y P E	O R G A N I Z A T I O N	R E C O R D S	L E N G T H	No. of 3390 Trks	No. of DIR Blks
SMPPTS	S	PDSE	U	0	32000	N/A

The following figures describe the target and distribution libraries and file system paths required to install Db2ZAI. The storage requirements of Db2ZAI 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.

<i>Figure 10. Storage Requirements for Db2ZAI Target Libraries</i>									
Library DDNAME	Member Type	Target Volume	T Y P E	O R G A N I Z A T I O N	R E C O R D S	L E N G T H	No. of 3390 Trks	No. of DIR Blks	
SCOYBASE	Samples	ANY	U	PDS	FB	80	8	3	

<i>Figure 11. Db2ZAI File System Paths</i>		
DDNAME	T Y P E	Path Name
SCOYHFS	N	/usr/lpp/IBM/db2zaiv1r2/IBM

Note: If you previously installed HCOY120 then you will have already defined the above zFS path and there is no need to define it again.

Figure 12. Storage Requirements for Db2ZAI 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
ACQYBASE	U	PDS	FB	80	8	3
ACQYHFS	U	PDS	VB	32630	25000	15

5.3 FMIDs Deleted

Installing Db2ZAI 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 Db2ZAI 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

5.4.1 Dynamic Enablement - Update PARMLIB

This step is mandatory to activate Db2 AI for z/OS V1.4. To ensure this step is not missed it is repeated in two more sections of this program directory but only needs to be set once.

To enable Db2ZAI to run with z/OS, the following entries must be in the IFAPRDxx PARMLIB member:

```

PRODUCT OWNER('IBM CORP')
NAME('DB2ZAI')
VERSION(1) RELEASE(*) MOD(*)
ID('5698-CGN')
STATE(ENABLED)

```

Sample registration parmlib entry COYPREG, which contains the above statements, is provided in the hlq.SCOYBASE data set for your information. Software manufacturing may include 'FEATURENAME(*)' in the above PARMLIB member, to comply with sample standards, and this addition is okay, it does not have an impact on enabling Db2ZAI.

Once the enablement records have been added to the SYS1.PARMLIB(IFAPRDxx) member, re-IPL or perform a **SET PROD=(xx)** command where xx identifies the last two characters of the IFAPRDxx member

in which the change was made. Use the z/OS console command **D PROD,STATUS** to obtain the current product enablement status.

5.4.2 Ordering details

IBM Db2 AI for z/OS FMID HCOY120 installs in SREL P115 and is orderable in feature name 'Db2 AI for z/OS **-Db2**'.

- HCOY120 (Db2 AI for z/OS, V1.4)

IBM Watson Machine Learning for z/OS and IBM Open Data Analytics FMIDs install in SREL Z038. They are orderable in feature name 'Db2 AI for z/OS - **z/OS**'

- HAQN230 (WMLz 2.3 base) *Required*
- HSPK120 (IzODA Spark) *Required*
- HANA110 (IzODA Anaconda) *Not required*
- HMDS120 (IzODA MDS) *Not required*

You are required to order both features to obtain the FMIDs required for Db2 AI for z/OS, V1.4.

Once the features have been selected and the orders processed it will be fulfilled in separate CBPDO/ServerPac orders based on SREL. Downloading the 'Db2 AI for z/OS **-z/OS**' order may take several hours due to the download of Anaconda FMID HANA110. Please note you must download Anaconda, however it is not a required FMID and you do not need to install it.

Note

You can install the required FMIDs in the same SMP/E Global zone, just ensure to specify both SREL P115 and Z038. See the SMP/E reference material in IBM Documentation for more information.

To view the individual Special Considerations of the products that constitute Db2ZAI, refer to the following Program Directories included in this Product Package:

- GI13-5011 - IBM Watson Machine Learning for z/OS, V2.3
- GI13-4348 - IBM Open Data Analytics for z/OS, V1.1

6.0 Installation Instructions

This chapter describes the installation method and the step-by-step procedures to install and to activate the functions of Db2ZAI.

Please note the following points:

- If you want to install Db2ZAI 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 Db2ZAI

6.1.1 SMP/E Considerations for Installing Db2ZAI

Use the SMP/E RECEIVE, APPLY, and ACCEPT commands to install this release of Db2ZAI. Because this product contains FMIDs that install in more than one SREL you will need to create two Global zones. Install Db2AI in P115 and install Watson Machine Learning and IzODA in Z038.

6.1.2 SMP/E Options Subentry Values

The recommended values for certain SMP/E CSI subentries are shown in Figure 13. 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	500,500,500	3390 DASD Tracks
PEMAX	SMP/E Default	IBM recommends using the SMP/E default for PEMAX.

6.1.3 Sample Jobs

The following sample installation jobs are provided as part of the product to help you install Db2ZAI:

Figure 14. Sample Installation Jobs

Job Name	Job Type	Description	RELFILE
COYRECEV	RECEIVE	Sample RECEIVE job	IBM.HCOY120.F1
COYALLOC	ALLOCATE	Sample job to allocate target and distribution libraries	IBM.HCOY120.F1
COYZFS	zFS	Sample job to allocate, create mountpoint, & mount zFS data sets	IBM.HCOY120.F1
COYISMKD	MKDIR	Sample job to invoke the supplied COYMKDIR EXEC to allocate file system paths	IBM.HCOY120.F1
COYDDDEF	DDDEF	Sample job to define SMP/E DDDEFs	IBM.HCOY120.F1
COYAPPLY	APPLY	Sample APPLY job	IBM.HCOY120.F1
COYACCEP	ACCEPT	Sample ACCEPT job	IBM.HCOY120.F1

You can access the sample installation jobs by performing an SMP/E RECEIVE (refer to 6.1.4, "Perform SMP/E RECEIVE" on page 21) then copy the jobs from the RELFILES to a work data set for editing and submission. See Figure 14 on page 19 to find the appropriate relfile 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=*
//IN DD DSN=IBM.fmid.Fy,UNIT=SYSALLDA,DISP=SHR,
// VOL=SER=filevol
//OUT DD DSN=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=IN,OUTDD=OUT
/*
```

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

IN:

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.

6.1.4 Perform SMP/E RECEIVE

If you have obtained Db2ZAI as part of a CBPDO, use the RCPDO job in the CBPDO RIMLIB data set to receive the Db2ZAI 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 COYRECEV to perform the SMP/E RECEIVE for Db2ZAI. 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.

6.1.5 Allocate SMP/E Target and Distribution Libraries

Edit and submit sample job COYALLOC to allocate the SMP/E target and distribution libraries for Db2ZAI. Consult the instructions in the sample job for more information.

Note: You may be required to allocate an SMPPTS of 32K cylinders for each global zone.

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

6.1.6 Allocate, create and mount ZFS Files (Optional)

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

If you plan to install Db2ZAI into a new z/OS UNIX file system, you can edit and submit the optional COYZFS job to perform the following tasks:

- Create the z/OS UNIX file system
- Create a mountpoint
- 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 mountpoint is `/usr/lpp/IBM/db2zaiv1r2..`

Before running the sample job to create the z/OS UNIX file system, you must ensure that OMVS is active on the driving system. zFS must be active on the driving system if you are installing Db2ZAI 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('&mntpnt1.')
MODE(RDWR) /* 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.

&mntpnt1. is the name of the mountpoint where the z/OS UNIX file system will be mounted.

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

6.1.7 Allocate File System Paths

The target system HFS or zFS data set must be mounted on the driving system when running the sample COYISMKD job since the job will create paths in the HFS 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 Db2ZAI into a file system that is zFS.

If you plan to install Db2ZAI into a new HFS or zFS file system, you must create the mountpoint and mount the new file system to the driving system for Db2ZAI.

The recommended mountpoint is */usr/lpp/IBM/db2zaiv1r2*.

Edit and submit sample job COYISMKD to allocate the HFS or zFS paths for Db2ZAI. 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 COYDDDEF to create DDDEF entries for the SMP/E target and distribution libraries for Db2ZAI. 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 COYAPPLY to perform an SMP/E APPLY CHECK for Db2ZAI. Consult the instructions in the sample job for more information.

Perform an SMP/E APPLY CHECK for Db2ZAI.

The latest HOLDDATA is available through several different portals, including <http://service.software.ibm.com/holdata/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)) .
..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 HOLDDs 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 COYACCEP to perform an SMP/E ACCEPT CHECK for Db2ZAI. Consult the instructions in the sample job for more information.

Perform an SMP/E ACCEPT CHECK for Db2ZAI.

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.

To view the SMP/E installation information for the products that constitute Db2ZAI, refer to the following Program Directories included in this Product Package:

- GI13-5011 - IBM Watson Machine Learning for z/OS, V2.3
- GI13-4348 - IBM Open Data Analytics for z/OS, V1.1

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 Db2ZAI, 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 Dynamic Enablement - Update PARMLIB - reminder

To enable Db2ZAI to run with z/OS, the following entries must be in the IFAPRDxx PARMLIB member:

```
PRODUCT OWNER('IBM CORP')
NAME('DB2ZAI')
VERSION(1) RELEASE(*) MOD(*)
ID('5698-CGN')
STATE(ENABLED)
```

Sample registration parmlib entry COYPREG, which contains the above statements, is provided in the hlq.SCOYBASE data set for your information. Software manufacturing may include 'FEATURENAME(*)' in the above PARMLIB member, to comply with sample standards, and this addition is okay, it does not have an impact on enabling Db2ZAI.

Once the enablement records have been added to the SYS1.PARMLIB(IFAPRDxx) member, re-IPL or perform a **SET PROD=(xx)** command where xx identifies the last two characters of the IFAPRDxx member in which the change was made. Use the z/OS console command **D PROD,STATUS** to obtain the current product enablement status.

6.2 Activating Db2ZAI

6.2.1 File System Execution

If you mount the file system in which you have installed Db2ZAI in read-only mode during execution, then complete the tasks in Product Customization to activate Db2ZAI.

6.2.2 Dynamic Enablement - Update PARMLIB - reminder

To enable Db2ZAI to run with z/OS, the following entries must be in the IFAPRDxx PARMLIB member:

```
PRODUCT OWNER('IBM CORP')
NAME('DB2ZAI')
VERSION(1) RELEASE(*) MOD(*)
ID('5698-CGN')
STATE(ENABLED)
```

Sample registration parmlib entry COYPREG, which contains the above statements, is provided in the hlq.SCOYBASE data set for your information. Software manufacturing may include 'FEATURENAME(*)' in the above PARMLIB member, to comply with sample standards, and this addition is okay, it does not have an impact on enabling Db2ZAI.

Once the enablement records have been added to the SYS1.PARMLIB(IFAPRDxx) member, re-IPL or perform a **SET PROD=(xx)** command where xx identifies the last two characters of the IFAPRDxx member in which the change was made. Use the z/OS console command **D PROD,STATUS** to obtain the current product enablement status.

6.3 Product Customization

Product Customization information for Db2 AI for z/OS and the products that constitute Db2 AI for z/OS are available in the IBM Documentation Website.

7.0 Notices

The notices file for IBM Db2 AI for z/OS can be found in a file named COYNOTIC and COYNOTC2 in the zFS directory.

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 AI for z/OS, July 2021 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 ibmkc@us.ibm.com, and include the following information:

- Your name and address
- Your email address
- Your telephone or fax number
- The publication title and order number
- The topic and page number related to your comment
- The text of your comment

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

G113-5007-06

