IBM

Program Directory for

IBM Open Data Analytics for z/OS V1R1

Program Number 5655-OD1

FMIDs HANA110, HMDS120, HSPK120

for Use with
z/OS V2R1 or higher

Document Date: September 2017

GI13-4348-00

Note

Before using this information and the product it supports, be sure to read the general information under 7.0, "Notices" on page 29.

Contents

Figures

1.0 Introduction

This program directory is intended for system programmers who are responsible for program installation and maintenance. It contains information about the material and procedures associated with the installation of IBM Open Data Analytics for z/OS.

The Program Directory contains the following sections:

- 2.0, "Program Materials" on page 3 identifies the basic program materials and documentation for Open Data Analytics for z/OS.

- 3.0, "Program Support" on page 5 describes the IBM support available for Open Data Analytics for z/OS.

- 4.0, "Program and Service Level Information" on page 7 lists the APARs (program level) and PTFs (service level) that have been incorporated into Open Data Analytics for z/OS.

- 5.0, "Installation Requirements and Considerations" on page 9 identifies the resources and considerations that are required for installing and using Open Data Analytics for z/OS.

- 6.0, "Installation Instructions" on page 19 provides detailed installation instructions for Open Data Analytics for z/OS. It also describes the procedures for activating the functions of Open Data Analytics for z/OS, or refers to appropriate publications.

Before installing Open Data Analytics for z/OS, 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 5 tells you how to find any updates to the information and procedures in this program directory.

Open Data Analytics for z/OS 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 Open Data Analytics for z/OS are included on the CBPDO tape.

Do not use this program directory if you install Open Data Analytics for z/OS 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 Open Data Analytics for z/OS Description

IBM Open Data Analytics for z/OS V1R1 offers an integrated, optimized runtime foundation of industry-leading open source technologies that can be used to gain analytic insights at the source of data origin, by leveraging Apache Spark as well as Anaconda and Python environments.

1

1.2 Open Data Analytics for z/OS FMIDs

Open Data Analytics for z/OS consists of the following FMIDs:

    HANA110
    HMDS120
    HSPK120

2.0 Program Materials

An IBM program is identified by a program number. The program number for Open Data Analytics for z/OS is 5655-OD1.

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 Open Data Analytics for z/OS. 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 Open Data Analytics for z/OS in the *CBPDO Memo To Users Extension*.

2.2 Optional Machine-Readable Material

No optional machine-readable materials are provided for Open Data Analytics for z/OS.

2.3 Program Publications

The following sections identify the basic publications for Open Data Analytics for z/OS.

Figure 1 identifies the basic unlicensed publications for Open Data Analytics for z/OS. The documentation is available in HTML and PDF formats in the IBM Knowledge Center and in PDF format in the IBM Publications Center.

- The IBM Knowledge Center at https://www.ibm.com/support/knowledgecenter provides documentation in HTML format for Open Data Analytics for z/OS and provides links to the documentation PDF files. Type the product name into the search box to locate the documentation.

- The IBM Publications Center at http://www.ibm.com/shop/publications/order provides the documentation PDF files for Open Data Analytics for z/OS for downloading.

Figure 1. Basic Material: Unlicensed		
Publication Title	**Form Number**	**Media Format**
IBM Open Data Analytics for z/OS Installation and Customization Guide	SC27-9033-00	PDF and KC plug-in
IBM Open Data Analytics for z/OS User's Guide	SC27-9034-00	PDF and KC plug-in
IBM Open Data Analytics for z/OS Administrator's Guide	SC27-9035-00	PDF and KC plug-in
IBM Open Data Analytics for z/OS Solutions Guide	SC27-9036-00	PDF and KC plug-in
Program Directory for IBM Open Data Analytics for z/OS	GI13-4348-00	PDF
License Information for IBM Open Data Analytics for z/OS	GI13-4349-00	PDF

2.4 Program Source Materials

No program source materials or viewable program listings are provided for Open Data Analytics for z/OS.

2.5 Publications Useful During Installation

You might want to use the publications listed in Figure 2 during the installation of Open Data Analytics for z/OS.

Figure 2. Publications Useful During Installation		
Publication Title	**Form Number**	**Media Format**
IBM SMP/E for z/OS Messages, Codes, and Diagnosis	GA22-7770	http://www.ibm.com/shop/ publications/order/
IBM SMP/E for z/OS Commands	SA22-7771	http://www.ibm.com/shop/ publications/order/
IBM SMP/E for z/OS Reference	SA22-7772	http://www.ibm.com/shop/ publications/order/
IBM SMP/E for z/OS User's Guide	SA22-7773	http://www.ibm.com/shop/ publications/order/

3.0 Program Support

This section describes the IBM support available for Open Data Analytics for z/OS.

3.1 Program Services

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

3.2 Preventive Service Planning

Before you install Open Data Analytics for z/OS, 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.Function.ApacheSpark fix category in SMP/E to ensure you have all the recommended service installed. Use the **FIXCAT(IBM.Function.ApacheSpark)** 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 Open Data Analytics for z/OS as part of a CBPDO, HOLDDATA is included.

If the CBPDO for Open Data Analytics for z/OS 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 Open Data Analytics for z/OS are included in Figure 3.

Figure 3. PSP Upgrade and Subset ID		
UPGRADE	**SUBSET**	**Description**
ZODAV1R1	HANA110	z/OS IzODA Anaconda
ZODAV1R1	HMDS120	z/OS IzODA MDS
ZODAV1R1	HSPK120	z/OS IzODA Spark

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 4 on page 6 identifies the component IDs (COMPID) for Open Data Analytics for z/OS.

Figure 4. Component IDs			
FMID	**COMPID**	**Component Name**	**RETAIN Release**
HANA110	5655AAB04	Bash	110
HANA110	5655AAB03	z/OS IzODA Anaconda	110
HMDS120	5655AAB02	z/OS IzODA MDS	120
HSPK120	5655AAB01	z/OS IzODA Spark	120

4.0 Program and Service Level Information

This section identifies the program and relevant service levels of Open Data Analytics for z/OS. 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 Open Data Analytics for z/OS have been incorporated into this release. They are listed by FMID.

- FMID HANA110

  No APARs have been incorporated at this time.

- FMID HMDS120

  PI59007 PI59392 PI59571 PI59573 PI59820 PI61083 PI61770
  PI62162 PI62799 PI62979 PI63427 PI63769 PI64279 PI64572
  PI64781 PI64805 PI65063 PI65542 PI65640 PI65706 PI65818
  PI66065 PI66225 PI66440 PI66740 PI67163 PI68220 PI68396
  PI68491 PI68737 PI69821 PI69822 PI70069 PI70929 PI71404
  PI71419 PI71431 PI72012 PI72124 PI72656 PI73181 PI73433
  PI73631 PI75066 PI75318 PI75841 PI75842 PI75843 PI75981
  PI79545 PI81794 PI83086

- FMID HSPK120

  PI66047 PI70302 PI72218 PI74114 PI75540 PI75937

4.2 Service Level Information

No PTFs against this release of Open Data Analytics for z/OS have been incorporated into the product package.

Frequently check the Open Data Analytics for z/OS 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.Function.ApacheSpark)** 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 Open Data Analytics for z/OS. 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 Open Data Analytics for z/OS.

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 5. Driving System Software Requirements				
Program Number	Product Name	Minimum VRM	Minimum Service Level will satisfy these APARs	Included in the shipped product?
Any **one** of the following:				
5650-ZOS	z/OS	V2R1 or higher	N/A	No
5650-ZOS	z/OS ICSF	V2R1 or higher	N/A	No
5655-G44	IBM SMP/E for z/OS	V3R6	N/A	No

Note: ICSF is a requirement for Installation. ICSF must be configured and started on the driving system.

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 V3R6.

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.

Open Data Analytics for z/OS invokes UNIX shell scripts during installation. The userid under which the job executes must have the following:

- UID(0) or READ access or higher to the BPX.SUPERUSER facility class

- be connected to a group that has a GID

- have READ access or higher to the BPX.FILEATTR.PROGCTL and BPX.FILEATTR.APF and BPX.FILEATTR.SHARELIB facility classes

- have WRITE access to the "/usr/lpp/IBM/izoda/spark" and "/usr/lpp/IBM/izoda/anaconda" paths

Open Data Analytics for z/OS is installed into a file system, either HFS or zFS. Before installing Open Data Analytics for z/OS, 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 Open Data Analytics for z/OS 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 Open Data Analytics for z/OS.

Open Data Analytics for z/OS installs in the z/OS (Z038) 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 6. Target System Mandatory Installation Requisites				
Program Number	Product Name	Minimum VRM	Minimum Service Level will satisfy these APARs	Included in the shipped product?
5650-ZOS	z/OS	V2R1 or higher	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. These products are specified as IF REQs.

Open Data Analytics for z/OS has no conditional installation requisites.

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

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

Figure 7. Target System Mandatory Operational Requisites	
Program Number	**Product Name and Minimum VRM/Service Level**
5650-ZOS	z/OS V2R1 with PTFs UA75046, UA75273, UA75855, UA77518, and UA80506
5655-ZOS	z/OS V2R2 with PTF UA80505
5650-ZOS	z/OS ICSF V2R1 or higher
5655-DGH	IBM 64-bit SDK for z/OS, Java Technology Edition Version 8, Service Refresh 4, Fix Pack 6

Note: ICSF is a requirement for Execution. ICSF must be configured and started on the target system.

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.

Open Data Analytics for z/OS 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.

Open Data Analytics for z/OS 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.

Open Data Analytics for z/OS has no negative requisites.

5.2.3 DASD Storage Requirements

Open Data Analytics for z/OS libraries can reside on all supported DASD types.

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

Figure 8 (Page 1 of 2). Total DASD Space Required by Open Data Analytics for z/OS	
Library Type	**Total Space Required in 3390 Trks**
Target	8992
Distribution	20047

Figure 8 (Page 2 of 2). Total DASD Space Required by Open Data Analytics for z/OS	
Library Type	**Total Space Required in 3390 Trks**
File System	89860

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.

   For more information about the names and sizes of the required data sets, see 6.1.6, "Allocate SMP/E Target and Distribution Libraries" on page 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.

5. All target libraries listed have the following attributes:

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

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

- These data sets can be in the LPA, but they are not required to be in the LPA.
- These data sets can be in the LNKLST.
- These data sets are not required to be APF-authorized.

Figure 9. Storage Requirements for SMP/E Work Data Sets						
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
SMPWRK1	S	PDS	FB	80	5	2
SMPWRK2	S	PDS	FB	80	5	2
SMPWRK3	S	PDS	FB	80	5	2
SMPWRK4	S	PDS	FB	80	5	2
SMPWRK6	S	PDS	FB	80	5	2
SYSUT1	U	SEQ	--	--	2	0
SYSUT2	U	SEQ	--	--	1	0
SYSUT3	U	SEQ	--	--	1	0
SYSUT4	U	SEQ	--	--	1	0

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

Figure 10. Storage Requirements for SMP/E Data Sets						
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
SMPLTS	E	PDSE	U	0	15	--
SMPMTS	E	PDS	FB	80	15	2
SMPPTS	E	PDS	FB	80	15	2
SMPSCDS	E	PDS	FB	80	15	2
SMPSTS	E	PDS	FB	80	15	2

The following figures describe the target and distribution libraries and file system paths required to install Open Data Analytics for z/OS. The storage requirements of Open Data Analytics for z/OS 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.

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
Figure 11. Storage Requirements for Open Data Analytics for z/OS Target Libraries								
SANBINST	DATA	TVOL2	U	PDS	FB	80	3	2
SAZKBIN	DATA	TVOL2	U	PDS	VB	256	5186	2
SAZKCNTL	DATA	TVOL2	U	PDS	FB	80	18	5
SAZKDBRM	MAC	TVOL2	U	PDS	FB	80	45	4
SAZKEXEC	EXEC	TVOL2	U	PDS	FB	80	58	4
SAZKLOAD	LMOD	TVOL2	U	PDSE	U	0	2060	-
SAZKMAP	DATA	TVOL2	U	PDS	FB	2048	14	3
SAZKMENU	MSG	TVOL2	U	PDS	FB	80	7	3
SAZKOBJX	DATA	TVOL2	U	PDS	FB	80	25	2
SAZKPENU	PANL	TVOL2	U	PDS	FB	80	41	23
SAZKRPC	LMOD	TVOL2	U	PDSE	U	0	50	-
SAZKSAMP	SAMP	TVOL2	U	PDS	FB	80	35	15
SAZKSLIB	SKEL	TVOL2	U	PDS	FB	80	2	2
SAZKSMAP	DATA	TVOL2	U	PDS	FB	2048	1397	61
SAZKTENU	TBL	TVOL2	U	PDS	FB	80	2	2
SAZKXATH	DATA	TVOL2	U	PDS	FB	80	12	4
SAZKXCMD	DATA	TVOL2	U	PDS	FB	80	6	2
SAZKXEXC	DATA	TVOL2	U	PDS	FB	80	6	5
SAZKXSQL	DATA	TVOL2	U	PDS	FB	80	19	2
SAZKXTOD	DATA	TVOL2	U	PDS	FB	80	2	2
SAZKXVTB	DATA	TVOL2	U	PDS	FB	80	4	2

DDNAME	TYPE	Path Name
		Figure 12. Open Data Analytics for z/OS File System Paths
SANBROOT	N	/usr/lpp/IBM/izoda/anaconda/IBM/
SAZKS1FS	N	/usr/lpp/IBM/izoda/spark/IBM/

Figure 13 (Page 1 of 2). Storage Requirements for Open Data Analytics for z/OS Distribution Libraries

Library DDNAME	TYPE	ORG	RECFM	LRECL	No. of 3390 Trks	No. of DIR Blks
AANBBIN	U	PDS	VB	256	6670	14
AANBINST	U	PDS	FB	80	3	2
AANBTXT	U	PDS	FB	80	2	2
AAZKBIN	U	PDS	VB	256	5186	2
AAZKCNTL	U	PDS	FB	80	18	5
AAZKDBRM	U	PDS	FB	80	45	4
AAZKEXEC	U	PDS	FB	80	58	4
AAZKLOAD	U	PDSE	U	0	2020	-
AAZKMAP	U	PDS	FB	2048	14	3
AAZKMENU	U	PDS	FB	80	7	3
AAZKOBJX	U	PDS	FB	80	25	2
AAZKPENU	U	PDS	FB	80	41	23
AAZKRPC	U	PDSE	U	0	51	-
AAZKSAMP	U	PDS	FB	80	35	15
AAZKSLIB	U	PDS	FB	80	2	2
AAZKSMAP	U	PDS	FB	2048	1397	61
AAZKSPFS	U	PDS	VB	256	4242	2
AAZKTENU	U	PDS	FB	80	2	2
AAZKXATH	U	PDS	FB	80	12	4
AAZKXCMD	U	PDS	FB	80	6	2
AAZKXEXC	U	PDS	FB	80	6	5
AAZKXSQL	U	PDS	FB	80	19	2

Figure 13 (Page 2 of 2). Storage Requirements for Open Data Analytics for z/OS 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**
AAZKXTOD	U	PDS	FB	80	2	2
AAZKXVTB	U	PDS	FB	80	4	2

5.3 FMIDs Deleted

Installing Open Data Analytics for z/OS will delete the two previous FMIDs HSPK110 and HMDS110.

5.4 Special Considerations

During the APPLY processing of Open Data Analytics for z/OS fmid HANA110, ICSF must be configured and started on the driving system and target system.

If not, you will see the following error in your APPLY output: 'Fatal Python error: failed to get random numbers to initialize Python'.

6.0 Installation Instructions

This chapter describes the installation method and the step-by-step procedures to install and to activate the functions of Open Data Analytics for z/OS.

Please note the following points:

- If you want to install Open Data Analytics for z/OS 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 Open Data Analytics for z/OS

6.1.1 SMP/E Considerations for Installing Open Data Analytics for z/OS

Use the SMP/E RECEIVE, APPLY, and ACCEPT commands to install this release of Open Data Analytics for z/OS.

6.1.2 SMP/E Options Subentry Values

The recommended values for certain SMP/E CSI subentries are shown in Figure 14. Using values lower than the recommended values can result in failures in the installation. DSSPACE is a subentry in the GLOBAL options entry. PEMAX is a subentry of the GENERAL entry in the GLOBAL options entry. See the SMP/E manuals for instructions on updating the global zone.

Figure 14. SMP/E Options Subentry Values

Subentry	Value	Comment
DSSPACE	Existing target CSI value	IBM suggests using your existing target system CSI's DSSPACE value.
PEMAX	SMP/E Default	IBM recommends using the SMP/E default for PEMAX.

6.1.3 SMP/E CALLLIBS Processing

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

- SCEELKED
- SISPLOAD

Note: CALLLIBS uses the previous DDDEFs only to resolve the link-edit for Open Data Analytics for z/OS. These data sets are not updated during the installation of Open Data Analytics for z/OS.

6.1.4 Sample Jobs

The following sample installation jobs are provided as part of the product to help you install Open Data Analytics for z/OS:

Job Name	Job Type	Description	RELFILE
ANBISALC	ALLOCATE	Sample job to allocate Anaconda target and distribution libraries	'prefix'.HANA110.F3
ANBISDDD	DDDEF	Sample job to define Anaconda SMP/E DDDEFs	'prefix'.HANA110.F3
ANBISMKD	MKDIR	Sample job to invoke the supplied AZKMKDIR EXEC to allocate Anaconda file system paths	'prefix'.HANA110.F3
ANBISZFS	ALLOMZFS	Sample Job to allocate the Anaconda ZFS	'prefix'.HANA110.F3
ANBMKDIR	REXX EXEC	Envoked by ANBISMKD to define install Anaconda directory paths	'prefix'.HANA110.F3
AZKISALC	ALLOCATE	Sample job to allocate Spark and MDS target and distribution libraries	'prefix'.HSPK120.F2
AZKISDDD	DDDEF	Sample job to define Spark and MDS SMP/E DDDEFs	'prefix'.HSPK120.F2
AZKISMKD	MKDIR	Sample job to invoke the supplied AZKMKDIR EXEC to allocate Spark and MDS file system paths	'prefix'.HSPK120.F2
AZKISZFS	ALLOMZFS	Sample Job to allocate ZFS	'prefix'.HSPK120.F2
AZKMKDIR	REXX EXEC	Envoked by AZKISMKD to define Spark and MDS install directory paths	'prefix'.HSPK120.F2

Figure 15. Sample Installation Jobs

NOTE: 'prefix' is the high-level qualifier specified as the DSPREFIX value in the SMPTLIB DDDEF or the OPTIONS entry of the global zone.

You can access the sample installation jobs by performing an SMP/E RECEIVE (refer to 6.1.5, "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 15 to find the appropriate relfile data set.

6.1.5 Perform SMP/E RECEIVE

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

6.1.6 Allocate SMP/E Target and Distribution Libraries

Edit and submit sample jobs ANBISALC and AZKISALC to allocate the SMP/E target and distribution libraries for Open Data Analytics for z/OS. 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 Allocate File System Paths

The target system zFS data set must be mounted on the driving system when running the sample ANBISMKD and AZKISMKD jobs since the job will create paths in the 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 zFS file system is mounted to the driving system. zFS must be active on the driving system if you are installing Open Data Analytics for z/OS into a file system that is zFS.

If you plan to install Open Data Analytics for z/OS into new zFS file systems, you must create the mountpoints and mount the new file systems to the driving system for Open Data Analytics for z/OS.

If your installation is using a zFS, edit and submit sample job ANBISZFS to allocate the file system for Anaconda and AZKISZFS to allocate the file system for Spark. Consult the instructions in the sample jobs for more information.

The space requirements for the file system to install the Anaconda FMID are shown in the sample job ANBISZFS. The space requirements for the file system to install the Spark FMID are shown in the sample job AZKISZFS. You may increase the space requirements for the file systems in the sample jobs to accommodate future growth due to the installation of service.

The recommended mountpoint for Spark is *1usr/lpp/IBM/izoda/spark* .

The recommended mountpoint for Anaconda is *1usr/lpp/IBM/izoda/anaconda* .

Edit and submit sample jobs ANBISMKD and AZKISMKD to allocate the HFS or zFS paths for Open Data Analytics for z/OS. 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 jobs ANBISDDD and AZKISDDD to create DDDEF entries for the SMP/E target and distribution libraries for Open Data Analytics for z/OS. Consult the instructions in the sample job for more information.

Expected Return Codes and Messages:

You will receive a return code of 4 the first time this job is run (due to the 'REP' command) with the following messages:

GIM27701W DDDEF ENTRY XXXXXXXX DOES NOT EXIST. THE REPLACE OPERATION HAS BEEN CHANGED TO AN ADD.

6.1.9 Perform SMP/E APPLY

1. Ensure that you have the latest HOLDDATA; then edit and submit a job to perform an SMP/E APPLY CHECK for Open Data Analytics for z/OS.

   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.Function.ApacheSpark)
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.Function.ApacheSpark)
GROUPEXTEND
BYPASS(HOLDCLASS(HIPER)) .
  ..any other parameters documented in the program directory
```

This method is the 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.Function.ApacheSpark 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.

Figure 16 on page 24 shows a sample job that you can use to perform an SMP/E APPLY CHECK for Open Data Analytics for z/OS.

```
//SPKAPLY JOB <job parameters>
//STEP1 EXEC PGM=GIMSMP,REGION=0M,TIME=NOLIMIT
//SMPCSI DD DSN=csiname,DISP=SHR
//SMPCNTL DD *
  SET BOUNDARY(targetzone) .
  APPLY CHECK XZREQ
  FORFMID (HANA110,HMDS120,HSPK120)
  SELECT (HANA110,HMDS120,HSPK120)
  GROUPEXTEND(NOAPARS,NOUSERMODS)
  FIXCAT(IBM.Function.ApacheSpark)
  BYPASS(HOLDSYSTEM,
  HOLDUSER,HOLDCLASS(UCLREL,ERREL,HIPER)) .
 /*
```

Figure 16. Sample SMP/E APPLY Job

> **Required Updates**
>
> 1. Update the job parameters.
>
> 2. Replace csiname on the SMPCSI DD statement with your CSI name.
>
> 3. Replace targetzone with your target zone name.

Note: ICSF must be configured and started on the driving system and target system.

If not, you will see the following error in your APPLY output of HANA110: 'Fatal Python error: failed to get random numbers to initialize Python'.

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.

The following messages can be ignored; any other messages should be investigated. The Binder produces these messages during steps that store modules into target libraries:

IEW2635I 5371 THREE BYTE ADCON IN SECTION ##IO AT OFFSET 001785F1 IN CLASS B_TEXT WITH RMODE=ANY CANNOT BE RELOCATED.

IEW2635I 5371 THREE BYTE ADCON IN SECTION ##IO AT OFFSET 0016CA11 IN CLASS B_TEXT WITH RMODE=ANY CANNOT BE RELOCATED.

IEW2635I 5371 THREE BYTE ADCON IN SECTION ##IO AT OFFSET 003A98C9 IN CLASS B_TEXT WITH RMODE=ANY CANNOT BE RELOCATED.

6.1.10 Perform SMP/E ACCEPT

Edit and submit a job to perform an SMP/E ACCEPT CHECK for Open Data Analytics for z/OS.

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 only *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.

```
//SPKACPT JOB <job parameters>
//STEP1 EXEC PGM=GIMSMP,REGION=0M,TIME=NOLIMIT
//SMPCSI DD DSN=csiname,DISP=SHR
//SMPCNTL DD *
   SET BOUNDARY(dlibzone).
   ACCEPT CHECK XZREQ
   FORFMID (HANA110,HMDS120,HSPK120)
   SELECT (HANA110,HMDS120,HSPK120)
   GROUPEXTEND(NOAPARS,NOUSERMODS)
   FIXCAT(IBM.Function.ApacheSpark)
   BYPASS(HOLDSYSTEM,
   HOLDUSER,HOLDCLASS(UCLREL,ERREL,HIPER)) .
/*
```

Figure 17. Sample SMP/E ACCEPT Job

Required Updates

1. Update the job parameters.

2. Replace csiname on the SMPCSI DD statement with your CSI name.

3. Replace dlibzone with your distribution zone name.

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

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

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

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

6.1.11 Run REPORT MISSINGFIX

To support the operation of Open Data Analytics for z/OS, fixes may be required on other IBM products that will operate in the same environment. The SMP/E REPORT MISSINGFIX command and the latest HOLDDATA from IBM can help you determine if any fixes required on other IBM products to support Open Data Analytics for z/OS exist that have not yet been installed. This REPORT command can also create APPLY commands to help you install any identified missing fixes.

After you finish installing Open Data Analytics for z/OS, it is recommended you run the REPORT MISSINGFIX command, like this:

```
SET BDY(GLOBAL).
REPORT MISSINGFIX
      ZONES(zosTgtZoneName,
            programProductTgtZoneName,
            ...)
      FIXCAT(IBM.Function.ApacheSpark).
```

The ZONES operand should identify the target zone for your z/OS where Open Data Analytics for z/OS will be running, in addition to the target zones for your z/OS, and other program products that will operate with Open Data Analytics for z/OS. This REPORT command will analyze your software installed in the specified target zones and determine which of the PTFs identified by the specified fix category are missing from those target zones. The resulting Missing FIXCAT SYSMOD Report identifies the PTFs you need to install, if any.

Fix categories, and the information SMP/E uses for the report analysis, are defined in IBM HOLDDATA. Therefore, before you run the SMP/E REPORT MISSINGFIX command you should first acquire the most current HOLDDATA from IBM. Current HOLDDATA is provided in all IBM service deliverables, including orders produced using SMP/E RECEIVE ORDER and ShopzSeries, and a HOLDDATA file can be downloaded directly from IBM (ftp://service.boulder.ibm.com/s390/holddata/full.txt). For more information about SMP/E fix categories, see HYPERLINK
"http://www-03.ibm.com/systems/z/os/zos/features/smpe/fix-category.html"
http://www.ibm.com/systems/z/os/zos/features/smpe/fix-category.html.

6.1.12 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 Open Data Analytics for z/OS, 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.

6.1.13 Cleaning Up Obsolete Data Sets, Paths, and DDDEFs

There are no obsolete data sets, paths and DDDEFs for the product.

6.2 Activating Open Data Analytics for z/OS

The publication *IBM Open Data Analytics for z/OS Installation and Customization Guide* (SC27-9033-00) contains the necessary information to customize and use Open Data Analytics for z/OS.

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.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle, its affiliates, or both.

Rocket is a registered trademark of Rocket Software, Inc.

Reader's Comments

Program Directory for IBM Open Data Analytics for z/OS, September 2017

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.

Communicating Your Comments to IBM

IBM Open Data Analytics for z/OS
z/OS V2R1 or higher

Publication No. GI13-4348-00

If you especially like or dislike anything about this book, please use one of the methods listed below to send your comments to IBM. Whichever method you choose, make sure you send your name, address, and telephone number if you would like a reply.

Feel free to comment on specific errors or omissions, accuracy, organization, subject matter, or completeness of this book. However, the comments you send should pertain to only the information in this manual and the way in which the information is presented. To request additional publications, or to ask questions or make comments about the functions of IBM products or systems, you should talk to your IBM representative or to your IBM authorized remarketer.

When you send comments to IBM, you grant IBM a nonexclusive right to use or distribute your comments in any way it believes appropriate without incurring any obligation to you.

If you are mailing a reader's comment form (RCF) from a country other than the United States, you can give the RCF to the local IBM branch office or IBM representative for postage-paid mailing.

- If you prefer to send comments by mail, use the RCF at the back of this book.
- If you prefer to send comments by FAX, use this number:
  - FAX: (International Access Code)+1+845+432-9405
- If you prefer to send comments electronically, use the following e-mail address:
  - mhvrcfs@us.ibm.com

Make sure to include the following in your note:
- Title and publication number of this book
- Page number or topic to which your comment applies

Optionally, if you include your telephone number, we will be able to respond to your comments by phone.

Reader's Comments — We'd Like to Hear from You

IBM Open Data Analytics for z/OS
z/OS V2R1 or higher

Publication No. GI13-4348-00

You may use this form to communicate your comments about this publication, its organization, or subject matter, with the understanding that IBM may use or distribute whatever information you supply in any way it believes appropriate without incurring any obligation to you. Your comments will be sent to the author's department for whatever review and action, if any, are deemed appropriate.

Note: Copies of IBM publications are not stocked at the location to which this form is addressed. Please direct any requests for copies of publications, or for assistance in using your IBM system, to your IBM representative or to the IBM branch office serving your locality.

Today's date: _____

What is your occupation?

Newsletter number of latest Technical Newsletter (if any) concerning this publication:

How did you use this publication?

[ ]　　　As an introduction　　　　　　　　　[ ]　　　As a text (student)
[ ]　　　As a reference manual　　　　　　　　[ ]　　　As a text (instructor)
[ ]　　　For another purpose (explain)

_____

_____

Is there anything you especially like or dislike about the organization, presentation, or writing in this manual? Helpful comments include general usefulness of the book; possible additions, deletions, and clarifications; specific errors and omissions.

　　Page Number:　　　　　　　Comment:

_____
Name

_____
Company or Organization

_____
Phone No.

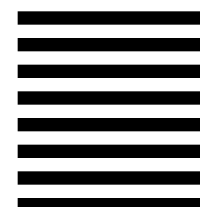Address

IBM

Fold and Tape   **Please do not staple**   Fold and Tape

NO POSTAGE
NECESSARY
IF MAILED IN THE
UNITED STATES

BUSINESS REPLY MAIL

FIRST-CLASS MAIL   PERMIT NO. 40   ARMONK, NEW YORK

POSTAGE WILL BE PAID BY ADDRESSEE

IBM Corporation
MHVRCFS, Mail Station P181
2455 South Road
Poughkeepsie, NY 12601-5400

Fold and Tape   **Please do not staple**   Fold and Tape

GI13-4348-00

IBM

Printed in USA