



**Program Directory for
IBM Data Virtualization Manager for z/OS**

V01.02.00

Program Number 5698-DVM

FMID HAVZ120

for Use with
z/OS

Document Date: March 2024

GI13-4693-01

Note

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

Contents

1.0 Introduction	1
1.1 Data Virtualization Manager Description	1
1.2 Data Virtualization Manager 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	6
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	11
5.2.2 Programming Requirements	11
5.2.2.1 Installation Requisites	11
5.2.2.2 Operational Requisites	12
5.2.2.3 Toleration/Coexistence Requisites	12
5.2.2.4 Incompatibility (Negative) Requisites	12
5.2.3 DASD Storage Requirements	12
5.3 FMIDs Deleted	16
5.4 Special Considerations	16
6.0 Installation Instructions	18
6.1 Installing Data Virtualization Manager	18
6.1.1 SMP/E Considerations for Installing Data Virtualization Manager	18
6.1.2 SMP/E Options Subentry Values	18
6.1.3 SMP/E CALLLIBS Processing	19
6.1.4 Sample Jobs	19
6.1.5 Allocate SMP/E CSI (Optional)	20

6.1.6 Initialize CSI zones (Optional)	20
6.1.7 Perform SMP/E RECEIVE	20
6.1.8 Allocate SMP/E Target and Distribution Libraries	21
6.1.9 Create DDDEF Entries	21
6.1.10 Perform SMP/E APPLY	21
6.1.11 Perform SMP/E ACCEPT	23
6.1.12 Run REPORT CROSSZONE	23
6.2 Activating Data Virtualization Manager	24
6.3 Product Customization	24
7.0 Notices	25
7.1 Trademarks	25
Reader's Comments	26

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. Total DASD Space Required by Data Virtualization Manager	13
8. Storage Requirements for SMP/E Work Data Sets	14
9. Storage Requirements for Data Virtualization Manager Target Libraries	14
10. Storage Requirements for Data Virtualization Manager Distribution Libraries	15
11. SMP/E Options Subentry Values	18
12. Sample Installation Jobs	19

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 Data Virtualization Manager for z/OS. This publication refers to IBM Data Virtualization Manager for z/OS as Data Virtualization Manager.

The Program Directory contains the following sections:

- 2.0, “Program Materials” on page 4 identifies the basic program materials and documentation for Data Virtualization Manager.
- 3.0, “Program Support” on page 7 describes the IBM support available for Data Virtualization Manager.
- 4.0, “Program and Service Level Information” on page 9 lists the APARs (program level) and PTFs (service level) that have been incorporated into Data Virtualization Manager.
- 5.0, “Installation Requirements and Considerations” on page 10 identifies the resources and considerations that are required for installing and using Data Virtualization Manager.
- 6.0, “Installation Instructions” on page 18 provides detailed installation instructions for Data Virtualization Manager. It also describes the procedures for activating the functions of Data Virtualization Manager, or refers to appropriate publications.

Before installing Data Virtualization Manager, 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.

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

Do not use this program directory if you install Data Virtualization Manager with a 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 Data Virtualization Manager Description

IBM Data Virtualization Manager for z/OS, V1.2.0 (5698-DVM) technology resides directly on the mainframe. It enables organizations to virtualize z/OS data safely and securely along with data from other sources in real time to provide comprehensive insights in a form easily leveraged by today's analytics and Artificial Intelligence for IT Operations (AIOps) applications, enabling you to unlock your data like never before to support your Data Fabric and lakehouse initiatives. The program offers:

- A data services layer that integrates data and content on demand from disparate sources in real time.
- A layer of abstraction that shields developers from unique data implementations and the complexities of the mainframe.
- A metadata catalog to keep track of data, location, availability, and state.
- Integration with the IBM watsonx.data lakehouse using PrestoDB connector.
- Support for transactions that write back to the original data sources regardless of location.
- z/OS resident optimization for improved performance and total cost of ownership (TCO) including zIIP enablement.

Data Virtualization Manager for z/OS 1.2 delivers enhanced capabilities to the IBM premier z/OS data integration solution. New capabilities enable improved access to the most critical z/OS mainframe data sources, increased performance, and integration with other IBM offerings such as IBM watsonx.data lakehouse and Data Fabric solutions. Data Virtualization Manager for z/OS unlocks the value of mainframe data for real-time client and business insights across the enterprise, without the cost and complexity associated with moving, replicating, or transforming data. Data remains securely in place, ready to be leveraged when needed, in the right format wherever it resides.

This new release includes:

- Use IMS Catalog Node to extract DBDs and PSBs from the IBM Catalog database in the DVM for z/OS studio.
- Created and manage multiple schemas.
- Access to DVM data sources on IBM Cloud Pak for Data.
- Create RESTful services to access all Db2 objects such as Db2 tables, views, stored procedures, and user-defined table functions (Studs) with DVM studio. For more information, see PH16070 and Creating RESTful services.
- Extend IBM z/OS Connect EE (zCEE) RESTful services to Db2 objects.
- Software configuration management capabilities for virtual views:
 - Allow for the modification of a virtual directory's map and microflow data set names during batch migration.
 - Allow for exporting a virtual table and virtual views to an SCM provider.
- Collect statistics about the data that is stored in the target database table.
- View online documentation for SMF virtual tables for the following without having to visit the IBM website:
 - SMF virtual tables
 - Columns in the SMF virtual tables
 - SMF sub-tables
 - Columns in SMF sub-tables

- Installation using IBM z/OS Management Facility (z/OSMF) workflow.
- Allow the DVM SQL engine to be skipped for execution of the queries and to process the queries at the connected DRDA datasource.
- Create views in Db2 subsystems on Linux environments and to create Db2 federation nicknames with DVM Studio. For more information, see the Db2 federation nicknames for distributed environment web page.
- Schema maps and all the maps related to the schema are now automatically added to a 64-bit storage.

1.2 Data Virtualization Manager FMIDs

Data Virtualization Manager consists of the following FMID:

HAVZ120

2.0 Program Materials

An IBM program is identified by a program number. The program number for Data Virtualization Manager is 5698-DVM.

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 Data Virtualization Manager. 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 18 for more information about how to install the program.

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

Figure 1 (Page 1 of 2). Program File Content

Name	O R G	R E C F M	L R E C L	BLK SIZE
SMPMCS	SEQ	FB	80	6400
IBM.HAVZ120.F1	PDS	FB	80	8800
IBM.HAVZ120.F2	PDSE	U	0	6144
IBM.HAVZ120.F3	PDS	FB	80	8800
IBM.HAVZ120.F4	PDS	FB	80	8800
IBM.HAVZ120.F5	PDS	FB	80	8800
IBM.HAVZ120.F6	PDS	VB	388	0
IBM.HAVZ120.F7	PDS	FB	2048	18432
IBM.HAVZ120.F8	PDS	FB	80	8800
IBM.HAVZ120.F9	PDS	FB	80	8800
IBM.HAVZ120.F10	PDS	FB	80	8800
IBM.HAVZ120.F11	PDS	FB	80	8800
IBM.HAVZ120.F12	PDS	FB	80	8800

Figure 1 (Page 2 of 2). Program File Content

Name	O R G	R E C F M	L R E C L	BLK SIZE
IBM.HAVZ120.F13	PDS	FB	2048	18432
IBM.HAVZ120.F14	PDS	FB	80	8800
IBM.HAVZ120.F15	PDS	FB	80	8800
IBM.HAVZ120.F16	PDS	FB	80	8800
IBM.HAVZ120.F17	PDS	FB	80	8800
IBM.HAVZ120.F18	PDS	FB	80	8800
IBM.HAVZ120.F19	PDS	FB	80	8800
IBM.HAVZ120.F20	PDS	FB	80	8800
IBM.HAVZ120.F21	PDS	FB	80	8800
IBM.HAVZ120.F22	PDS	FB	80	8800

2.2 Program Publications

The following sections identify the basic publications for Data Virtualization Manager.

Figure 2 identifies the basic unlicensed publications for Data Virtualization Manager. Those that are in softcopy format can be obtained from the IBM Publications Center website at <https://www.ibm.com/resources/publications>

Figure 2. Basic Material: Unlicensed Publications

Publication Title	Form Number	Media Format
<i>IBM Data Virtualization Manager for z/OS License Information</i>	GC27-8873	http://www.ibm.com/software/sla/slabd.nsf

2.2.1 Optional Program Publications

No optional publications are provided for Data Virtualization Manager.

2.3 Program Source Materials

No program source materials or viewable program listings are provided for Data Virtualization Manager.

2.4 Publications Useful During Installation

You might want to use the publications listed in Figure 3 on page 6 during the installation of Data Virtualization Manager.

<i>Figure 3. Publications Useful During Installation</i>	
Publication Title	Form Number
<i>IBM SMP/E for z/OS User's Guide</i>	SA23-2277
<i>IBM SMP/E for z/OS Commands</i>	SA23-2275
<i>IBM SMP/E for z/OS Reference</i>	SA23-2276
<i>IBM SMP/E for z/OS Messages, Codes, and Diagnosis</i>	GA32-0883

Note: These publications can be found in IBM Documentation. Use a web browser with internet access to refer to: <https://www.ibm.com/docs/en/zos/2.5.0?topic=zos-smpe>

3.0 Program Support

This section describes the IBM support available for Data Virtualization Manager.

3.1 Program Services

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

3.2 Preventive Service Planning

Before you install Data Virtualization Manager, 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 6.1.10, “Perform SMP/E APPLY” on page 21 for a sample APPLY command

If you obtained Data Virtualization Manager as part of a CBPDO, HOLDDATA is included.

If the CBPDO for Data Virtualization Manager 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:

<https://www.ibm.com/support/pages/ibmsearch>

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 <https://www.ibm.com/mysupport/>.

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 Data Virtualization Manager are included in Figure 4.

UPGRADE	SUBSET	Description
5698DVM	HAVZ120	IBM Data Virtualization Manager for z/OS

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 ID (COMPID) for Data Virtualization Manager.

<i>Figure 5. Component IDs</i>			
FMID	COMPID	Component Name	RETAIN Release
HAVZ120	5698DVM00	IBM Data Virtualization Manager for z/OS	120

4.0 Program and Service Level Information

This section identifies the program and relevant service levels of Data Virtualization Manager. 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 Data Virtualization Manager.

4.2 Service Level Information

No PTFs against this release of Data Virtualization Manager have been incorporated into the product package.

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

5.0 Installation Requirements and Considerations

The following sections identify the system requirements for installing and activating Data Virtualization Manager. 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 Data Virtualization Manager.

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	02.04.00	N/A	No

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

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

5.2 Target System Requirements

This section describes the environment of the target system required to install and use Data Virtualization Manager.

Data Virtualization Manager installs in the DBS (P115) SREL.

5.2.1 Machine Requirements

Data Virtualization Manager V01.02.00. will run on the following IBM System z servers:

- IBM z16
- IBM z15
- IBM z14
- IBM z13
- IBM z13s
- IBM zEnterprise EC12
- IBM zEnterprise BC12

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.

Data Virtualization Manager has no mandatory installation requisites.

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

Data Virtualization Manager 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.

Data Virtualization Manager has no mandatory operational requisites.

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

Data Virtualization Manager 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.

Data Virtualization Manager 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.

Data Virtualization Manager has no negative requisites.

5.2.3 DASD Storage Requirements

Data Virtualization Manager libraries can reside on all supported DASD types.

Figure 7 lists the total space that is required for each type of library.

Figure 7. Total DASD Space Required by Data Virtualization Manager

Library Type	Total Space Required in 3390 Trks	Description
Target	22865 Tracks	
Distribution	22580 Tracks	

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

3. All target and distribution libraries listed have the following attributes:
 - The default name of the data set can be changed.
 - The default block size of the data set can be changed.
 - The data set can be merged with another data set that has equivalent characteristics.
 - The data set can be either a PDS or a PDSE, except for AAVZLOAD, SAVZLOAD, SAVZCLOAD, and SAVZRPC which must be PDSEs.
4. 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.

5. 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 data set should be placed in the LPA, see the Special Considerations section below.
- These data sets can be in the LNKST. If so, see the Special Considerations section below.
- These data sets are not required to be APF-authorized, with some exceptions. If the data set must be APF-authorized, see the Special Considerations section below.
- Data Virtualization Manager requires that the SMPLTS data set must be a PDSE. If your existing SMPLTS is a PDS, you will need to allocate a new PDSE and copy your existing SMPLTS into it and then change the SMPLTS DDDEF entry to indicate the new PDSE data set.

Figure 8. 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
SMPLTS	S	PDSE	U	0	850	150

The following figures describe the target and distribution libraries required to install Data Virtualization Manager. The storage requirements of Data Virtualization Manager must be added to the storage required by other programs that have data in the same library.

Note: Use the data in these tables to determine which libraries can be merged into common data sets. In addition, since some ALIAS names may not be unique, ensure that no naming conflicts will be introduced before merging libraries.

Figure 9 (Page 1 of 2). Storage Requirements for Data Virtualization Manager Target Libraries

Library DDNAME	Member Type	Target Volume	T Y P E	O R G	R E C F M	L R E C L	No. of 3390 Trks	No. of DIR Blks
SAVZBASE	SAMP	Any	U	PDS	FB	80	15	5
SAVZCLOD	MOD	Any	U	PDSE	U	0	105	n/a
SAVZCNTL	DATA	Any	U	PDS	FB	80	100	20
SAVZDBRM	MACRO	Any	U	PDS	FB	80	80	10
SAVZEXEC	EXEC	Any	U	PDS	FB	80	200	10
SAVZHDOC	TEXT	Any	U	PDS	VB	388	320	10
SAVZLOAD	MOD	Any	U	PDSE	U	0	4125	n/a

Figure 9 (Page 2 of 2). Storage Requirements for Data Virtualization Manager Target Libraries

Library DDNAME	Member Type	Target Volume	T Y P E	O R G	R E C F M	L R E C L	No. of 3390 Trks	No. of DIR Blks
SAVZMAP	DATA	Any	U	PDS	FB	2048	35	10
SAVZMENU	MSG	Any	U	PDS	FB	80	20	5
SAVZOBJX	DATA	Any	U	PDS	FB	80	45	5
SAVZPENU	PNL	Any	U	PDS	FB	80	150	35
SAVZRPC	MOD	Any	U	PDSE	U	0	180	n/a
SAVZSAMP	DATA	Any	U	PDS	FB	80	25	50
SAVZSLIB	SKEL	Any	U	PDS	FB	80	5	5
SAVZSMAP	DATA	Any	U	PDS	FB	2048	150	70
SAVZTENU	TABLE	Any	U	PDS	FB	80	15	5
SAVZXATH	DATA	Any	U	PDS	FB	80	35	15
SAVZXCMD	DATA	Any	U	PDS	FB	80	20	10
SAVZXEXC	DATA	Any	U	PDS	FB	80	20	10
SAVZXRPC	DATA	Any	U	PDS	FB	80	15	5
SAVZXSQL	DATA	Any	U	PDS	FB	80	40	10
SAVZXTOD	DATA	Any	U	PDS	FB	80	15	5
SAVZXVTB	DATA	Any	U	PDS	FB	80	20	10

Figure 10 (Page 1 of 2). Storage Requirements for Data Virtualization Manager 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
AAVZBASE	U	PDS	FB	80	15	5
AAVZCNTL	U	PDS	FB	80	100	20
AAVZDBRM	U	PDS	FB	80	80	10
AAVZEXEC	U	PDS	FB	80	200	10
AAVZHDOC	U	PDS	VB	388	320	10
AAVZLOAD	U	PDSE	U	0	4125	n/a
AAVZMAP	U	PDS	FB	2048	35	10
AAVZMENU	U	PDS	FB	80	20	5

Figure 10 (Page 2 of 2). Storage Requirements for Data Virtualization Manager 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
AAVZOBJX	U	PDS	FB	80	45	5
AAVZPENU	U	PDS	FB	80	150	35
AAVZSAMP	U	PDS	FB	80	25	50
AAVZSLIB	U	PDS	FB	80	5	5
AAVZSMAP	U	PDS	FB	2048	150	70
AAVZTENU	U	PDS	FB	80	15	5
AAVZXATH	U	PDS	FB	80	35	15
AAVZXCMD	U	PDS	FB	80	20	10
AAVZXEXC	U	PDS	FB	80	20	10
AAVZXRPC	U	PDS	FB	80	15	5
AAVZXSQL	U	PDS	FB	80	40	10
AAVZXTOD	U	PDS	FB	80	15	5
AAVZXVTB	U	PDS	FB	80	20	10

5.3 FMIDs Deleted

Installing Data Virtualization Manager 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 Data Virtualization Manager 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

Data Virtualization Manager requires the following non z/OS hardware and software:

- Client drivers:
 - Java Database Connectivity (JDBC)
 - Java level 1.8 or higher

- Open Database Connectivity (ODBC)
 - Windows
 - AIX
 - HP-UX
 - Linux
 - Red Hat Enterprise Linux
 - SUSE Linux
- Studio:
 - Windows:
 - Windows 11 (32-bit and 64-bit)
 - Windows 10 (32-bit and 64-bit)
 - Windows 8 (32-bit and 64 bit)
 - Windows 7 (32-bit and 64-bit)
 - Linux:
 - Ubuntu
 - Red Hat
 - Apple macOS:
 - macOS 11 onwards
 - Hard disk space:
 - A minimum of 2 GB is recommended for a full installation.
 - System memory:
 - A minimum of 4 GB is recommended.
 - A full install includes the Studio product installer (approximately 1 GB) and bundled Eclipse 4.16 and Java level 1.8 or higher

6.0 Installation Instructions

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

Please note the following points:

- If you want to install Data Virtualization Manager 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 Data Virtualization Manager

6.1.1 SMP/E Considerations for Installing Data Virtualization Manager

Use the SMP/E RECEIVE, APPLY, and ACCEPT commands to install this release of Data Virtualization Manager.

6.1.2 SMP/E Options Subentry Values

The recommended values for certain SMP/E CSI subentries are shown in Figure 11. 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 11. SMP/E Options Subentry Values

Subentry	Value	Comment
DSSPACE	(7500,500,200)	3390 DASD tracks
PEMAX	SMP/E Default	IBM recommends using the SMP/E default for PEMAX.

6.1.3 SMP/E CALLLIBS Processing

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

- SCEELKED
- SDFHLOAD
- SISPLOAD
- SEZATCP

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

6.1.4 Sample Jobs

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

Figure 12. Sample Installation Jobs

Job Name	Job Type	Description	SMPTLIB Data Set
AVZALA	SMP/E	Sample job to allocate and initialize a new SMP/E CSI data set (Optional)	IBM.HAVZ120.F22
AVZALB	SMP/E	Sample job to allocate SMP/E data sets (Optional)	IBM.HAVZ120.F22
AVZRECEV	RECEIVE	Sample RECEIVE job	IBM.HAVZ120.F22
AVZALLOC	ALLOCATE	Sample job to allocate target and distribution libraries	IBM.HAVZ120.F22
AVZDDDEF	DDDEF	Sample job to define SMP/E DDDEFs	IBM.HAVZ120.F22
AVZAPPLY	APPLY	Sample APPLY job	IBM.HAVZ120.F22
AVZACCEP	ACCEPT	Sample ACCEPT job	IBM.HAVZ120.F22

You can access the sample installation jobs by performing an SMP/E RECEIVE (refer to 6.1.7, “Perform SMP/E RECEIVE” on page 20) then copy the jobs from the SMPTLIB data sets to a work data set for editing and submission. See Figure 12 to find the appropriate data set.

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

```

//STEP1 EXEC PGM=IEBCOPY
//SYSPRINT DD SYSOUT=*
//IN DD DSN=IBM.HAVZ120.F22,UNIT=SYSALLDA,DISP=SHR,
// VOL=SER=filevol
//OUT DD DSNAME=jcl-library-name,
// DISP=(NEW,CATLG,DELETE),
// VOL=SER=dasdvol,UNIT=SYSALLDA,
// SPACE=(TRK,(20,10,5))
//SYSUT3 DD UNIT=SYSALLDA,SPACE=(CYL,(1,1))
//SYSIN DD *
COPY INDD=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.

6.1.5 Allocate SMP/E CSI (Optional)

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

If you are allocating a new SMP/E data set for this install, edit and submit sample job AVZALA to allocate the SMP/E data set for Data Virtualization Manager. Consult the instructions in the sample job for more information.

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

6.1.6 Initialize CSI zones (Optional)

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

Edit and submit sample job AVZALB to initialize SMP/E zones for Data Virtualization Manager. Consult the instructions in the sample job for more information.

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

6.1.7 Perform SMP/E RECEIVE

If you have obtained Data Virtualization Manager as part of a CBPDO, use the RCVPDO job in the CBPDO RIMLIB data set to receive the Data Virtualization Manager FMIDs, service, and HOLDDATA that are included on the CBPDO package. For more information, see the documentation that is included in the CBPDO. You will receive a return code of 0 if this job runs correctly.

6.1.8 Allocate SMP/E Target and Distribution Libraries

Edit and submit sample job AVZALLOC to allocate the SMP/E target and distribution libraries for Data Virtualization Manager. 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 Create DDDEF Entries

Edit and submit sample job AVZDDDEF to create DDDEF entries for the SMP/E target and distribution libraries for Data Virtualization Manager. 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.

6.1.10 Perform SMP/E APPLY

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

The latest HOLDDATA is available through several different portals, including <http://service.software.ibm.com/holddata/390holddata.html>. The latest HOLDDATA may identify HIPER and FIXCAT APARs for the FMIDs you will be installing. An APPLY CHECK will help you determine if any HIPER or FIXCAT APARs are applicable to the FMIDs you are installing. If there are any applicable HIPER or FIXCAT APARs, the APPLY CHECK will also identify fixing PTFs that will resolve the APARs, if a fixing PTF is available.

You should install the FMIDs regardless of the status of unresolved HIPER or FIXCAT APARs. However, do not deploy the software until the unresolved HIPER and FIXCAT APARs have been analyzed to determine their applicability. That is, before deploying the software either ensure fixing PTFs are applied to resolve all HIPER or FIXCAT APARs, or ensure the problems reported by all HIPER or FIXCAT APARs are not applicable to your environment.

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

Here are sample APPLY commands:

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

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

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

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

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

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

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

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

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

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

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

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

6.1.11 Perform SMP/E ACCEPT

Edit and submit sample job AVZACCEP to perform an SMP/E ACCEPT CHECK for Data Virtualization Manager. Consult the instructions in the sample job for more information.

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

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

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

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

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

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

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

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

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

6.2 Activating Data Virtualization Manager

6.3 Product Customization

The publication *IBM Data Virtualization Manager for z/OS Installation and Customization Guide* (GC27-8874) contains the necessary information to customize and use Data Virtualization Manager.

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

Reader's Comments

Program Directory for IBM Data Virtualization Manager for z/OS, March 2024 We appreciate your input on this publication. Feel free to comment on the clarity, accuracy, and completeness of the information or give us any other feedback that you might have.

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

Your name 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 Ireland

GI13-4693-01

