IBM /Technical Newsletter

MVS/ESA JCL User's Guide

©Copyright IBM Corp. 1988

MVS/System Product - JES2, Program No. 5685-001
MVS/System Product - JES3, Program No. 5685-002

This newsletter contains replacement pages for *MVS/ESA JCL User's Guide* in support of MVS/SP Version 3 Release 1.0.

Before inserting any of the attached pages into the *JCL User's Guide*, read *carefully* the instructions on this cover. They indicate when and how you should insert pages.

Pages to be Removed	Attached Pages to be Inserted*
Cover - Edition Notice	Cover - Edition Notice
iii - iv	iii - iv
12-3 - 12-4	12-3 - 12-4
15-5 - 15-8	15-5 - 15-8
15-17 - 15-18	15-17 - 15-18

*If you are inserting pages from different Newsletters/Supplements and *identical* page numbers are involved. always use the page with the latest date (shown in the slug at the top of the page). The page with the latest date contains the most complete information.

A change to the text or to an illustration is indicated by a vertical line to the left of the change.

Summary of Amendments

This Technical Newsletter contains maintenance updates.

Note: Please file this cover letter at the back of the publication to provide a record of changes.

IBM MVS/ESA
JCL User's Guide

GC28-1830-0

MVS/System Product:
JES2 Version 3
JES3 Version 3

Summary of Changes

This Technical Newsletter contains maintenance updates.

This book contains information previously presented in *MVS/XA*™ *JCL User's Guide*, GC28-1351-3, which supports MVS/System Product Version 2 Release 2.0. The following summarizes the changes to that information.

New Information: MVS/SP™ Version 3 Release 1.0 supports MVS/Data Facility Product (MVS/DFP™) Version 3 Release 1.0, which introduces the Storage Management Subsystem (SMS). SMS provides new functions for data and storage management.

In this book, "with SMS" indicates information that applies when SMS is installed and active; "without SMS" indicates SMS is not installed or is not active.

For JCL, the following DD statement parameters are added for defining new data sets with SMS. Note that the system ignores these DD parameters when SMS is not installed or is not active.

The new DD parameters are:

- AVGREC - average record
- DATACLAS - data class
- KEYOFF - key offset
- LIKE - like dsname
- MGMTCLAS - management class
- RECORG - record organization
- REFDD - reference ddname
- SECMODEL - security model
- STORCLAS - storage class *

  * When a storage class is assigned to a data set, the data set is referred to as an "SMS-managed data set".

There are also changes to:

- The SPACE and VOLUME DD parameters for defining data sets with SMS.
- The defining and disposition of VSAM data sets with SMS.

There is a new appendix (Appendix D, "Data Sets with SMS") that contains a summary of SMS information.

Generations of a Generation Data Group

A generation data group is a collection of chronologically related data sets that have the same data set name. To add a generation to a generation data group or retrieve a generation, specify the generation data group name followed by the generation number. A zero is the current generation of the group, a negative number (for example, -1) is an older generation, a positive number (for example, +1) is a new generation that does not exist yet.

Examples

```
//NEWGDS  DD  DSNAME=GDS(0),DISP=(NEW,CATLG,DELETE),
//            UNIT=3380,VOLUME=SER=334455,SPACE=(CYL,20)

//OLDGDS  DD  DSNAME=GDS(-1),DISP=OLD

//NEWER   DD  DSNAME=GDS(+1),DISP=(NEW,CATLG,DELETE),
//            UNIT=3350,VOLUME=SER=222333,SPACE=(TRK,15)

//ALLG    DD  DSNAME=GDS,DISP=OLD

//SMSGDG  DD  DSNAME=A.B.C(+1),DATACLAS=DGDG1,DISP=(NEW,KEEP)
```

Areas of an Indexed Sequential Data Set

An indexed sequential data set consists of three areas: index, prime, and overflow. To create the data set, define each area by identifying the data set name followed by the area name. The area name is INDEX, PRIME, or OVFLOW. To define the data set on one DD statement, code DSNAME = dsname or DSNAME = dsname(PRIME). To retrieve the data set, code only the data set name.

Examples

```
//NEWIS  DD  DSNAME=ISDS(INDEX),DISP=(NEW,CATLG,DELETE),
//           UNIT=3350,VOLUME=SER=222333,SPACE=(CYL,5)
//         DD  DSNAME=ISDS(PRIME),DISP=(NEW,CATLG,DELETE),
//           UNIT=3350,VOLUME=SER=222333,SPACE=(CYL,15)
//         DD  DSNAME=ISDS(OVFLOW),DISP=(NEW,CATLG,DELETE),
//           UNIT=3350,VOLUME=SER=222333,SPACE=(CYL,10)

//OLDIS  DD  DSNAME=ISDS,DISP=OLD
```

Temporary Data Sets

Any data set that is created and deleted in the same job is a temporary data set. Identify a temporary data set by coding:

DSNAME = &&dsname	For a temporary data set
DSNAME = &&dsname(member)	For a member of a temporary partitioned data set
DSNAME = &&dsname(area)	For an area of a temporary indexed sequential data set
No DSNAME parameter	For a temporary data set to be named by the system

SMS manages a temporary data set if (1) you specify a storage class (via the DD STORCLAS parameter) or (2) an installation-written automatic class selection (ACS) routine selects a storage class for the temporary data set.

Devices Required per DD Statement: The maximum number of tape devices or direct access devices required to satisfy any DD statement is the unit count in the UNIT parameter. However, if the UNIT parameter also specifies P, for parallel mount, the system uses the **greatest** of the following numbers to determine how many devices and volumes to allocate:

* unit-count in the UNIT parameter
* volume-count specified in the VOLUME parameter
* number of serial numbers implicitly or explicitly specified
* with SMS, volume-count in the data class

The number of devices is affected by the DD statement parameters as follows:

DD Statement Specifies	System Action
UNIT = AFF	The system obtains the device requirements from the referenced DD statement. All of the devices used for the referenced DD statement are shared with the referring statement's data set.
Generation data group (GDG)	The system determines the number of devices needed by totaling the devices needed for each generation data set. Each generation data set is handled as a single request.
VSAM data set	The system determines the number of devices needed based on the device volume configuration of the data set. If the data set is on more than one type of device, the system determines the total number of devices required and allocates them. The system may override the unit count or parallel mounting, if specified.
Unit name that includes different device types	The system allocates devices of the same type.

Devices Assigned per Step: The number of devices assigned for a job step is not necessarily the sum of the device requirements for each DD statement.

The following tend to reduce the total devices assigned for a step:

* A volume can be allocated to only one device. Therefore, when more than one DD statement asks for the same volume, the system allocates the same volume on the same device.

* Requests for direct access space on public and/or storage volumes can be allocated to the same volume. Therefore, when more than one DD statement requests such space, the system can allocate the same volume on the same device.

* Requests for the same public tape volume are allocated to that volume. Therefore, if a DD statement requests a public tape and specifies VOLUME = REF, the system can allocate the same volume on the same device.

The following tend to increase the total devices assigned for a step:

* A permanently resident or reserved volume cannot be demounted. Therefore, the system assigns a permanently resident or reserved volume to its own device, on which it is mounted. The volume is assigned to its own device even if the DD statements specify that the device was to be shared with other volumes.

* A direct access volume is requested by more than one DD statement in a step; the volume is shared by the data sets. The system assigns that volume to a device and does not assign any other volumes to that device, even if the DD statements specify that the device was to be used for other volumes.

* The system allocates additional devices for a VSAM data set, if the data set resides on more than one type of device.

* The system allocates a direct access device for a private catalog, if it is associated with and/or used to retrieve volume information about a requested data set.

Unit and Volume Affinity (Non-SMS-Managed Data Sets)

When two or more volumes are assigned the same device, the volumes are said to have **unit affinity**. Unit affinity implies deferred mounting for all except one of the volumes. When two data sets share one volume, the data sets have **volume affinity**.

Explicit Unit Affinity: To reduce the number of devices for a step, request that an existing data set be assigned to the same device(s) assigned for an earlier DD statement in the same step. Code:

```
//ddname DD   UNIT=AFF=ddname,...
```

For concatenated data sets, code the following to assign the data sets to the same device:

```
//DD1    DD  DSNAME=dataset1,...
//       DD  DSNAME=dataset2,UNIT=AFF=DD1,...
//       DD  DSNAME=dataset3,UNIT=AFF=DD1,...
```

Implied Unit Affinity: Implied unit affinity exists among the volumes for one data set when the DD statement requests more volumes than devices.

Warning: If all of the following conditions are present, the data set on a DD statement requesting unit affinity is written over by the data set on the referenced DD statement:

- The referenced DD statement makes a nonspecific volume request.

- The data set requesting unit affinity is opened before the referenced data set.

- The tape is not unloaded before the referenced data set is opened and the LABEL parameter does not request positioning of the tape to check tape labels. A tape device allocated to more than one data set is not unloaded (1) as a result of dynamic deallocation or (2) when it is closed and FREE = CLOSE is specified.

Interaction of Unit and Volume Affinity Requests: Unit affinity, volume affinity, and/or unit and volume affinity can exist in the same step and on the same DD statement.

If both unit and volume affinity are requested in the same step, sometimes only one affinity can be honored. Figure 15-3 on page 15-8 indicates how the system honors unit and volume affinity requests for either tape or direct access devices.

For example:

```
//DD14  DD  DSNAME=DESIGNH.PGM,DATACLAS=DCLAS14,STORCLAS=STOR55,
//          DISP=(NEW,KEEP),VOLUME=SER=(223344,334455)
```

An SMS-managed data set can reside on a maximum of 59 volumes.

Nonspecific Volume Requests for SMS-Managed Data Sets

Omit the VOLUME parameter to make a nonspecific volume request for a new SMS-managed data set. SMS will select the volume to be used for the data set.

Multivolume Data Sets (SMS-Managed Data Sets)

SMS creates a preallocated multivolume data set if (1) the storage class has GUARANTEED_SPACE = YES and (2) the data class has a volume count greater than one or you specify two or more volume serial numbers or a volume count greater than one on the VOLUME parameter. If you specify specific volume serial numbers, SMS uses these volumes; otherwise, SMS selects the volumes.

Volume Allocation for Non-SMS-Managed Data Sets

Data sets on direct access and magnetic tape reside on or are written on volumes. The volumes may be permanently mounted on the device or may need to be mounted by the operator. To tell the system the volume on which an existing data set resides, make a specific volume request. To tell the system the volume on which to write a new data set, make a specific or nonspecific volume request.

With SMS, if the storage administrator has specified a system default unit, then the system uses the volumes associated with the default unit. In this case, you do not need to code the VOLUME parameter.

Volume Attributes: The system assigns volumes two attributes:

* **Use attributes,** which control how volumes are allocated, are:

  - **Private:** The volume can be allocated only when its volume serial number is explicitly or implicitly specified.

  - **Public:** The volume is eligible for allocation to temporary data sets defined with a nonspecific volume request and without a PRIVATE subparameter in the VOLUME parameter.

  - **Storage:** The volume is eligible for allocation to both temporary and permanent data sets defined with a nonspecific volume request and without a PRIVATE subparameter in the VOLUME parameter. Storage volumes usually contain permanent data sets, but can be used for temporary data sets.

* **Mount attributes,** which control how or whether volumes can be demounted after being deallocated, are:

  - **Permanently resident:** The volume, which can only be direct access, cannot be demounted. Volumes that are always permanently resident are all volumes that cannot be physically demounted, the IPL volume, and the volume containing system data sets. Permanently resident volumes have any use attribute.

  - **Reserved:** The volume remains mounted until the operator issues an UNLOAD command. Volumes that should be reserved are volumes that are used frequently by many jobs. Reserved, direct access volumes can have any use attribute; reserved, tape volumes can be only private or public.