Device Support Facilities:

Capabilities and Usage

By Srinu Sista
GPD Santa Teresa

P. R. Dorn
Washington Systems Center

GG22-9216-00
December 1980

Washington Systems Center
Gaithersburg, Maryland
Technical Bulletin

Device Support Facilities
Capabilities and Usage

Srinu Sista, GPD Santa Teresa
Paul Dorn, Washington Systems Center

GG22-9216-0
December 1980

PREFACE

Device Support Facilities is a program used with IBM System/370 operating systems to perform various operations on direct-access storage devices. It can initialize a direct-access storage volume so that it can be used on an IBM operating system, inspect a volume for defective tracks, reformat a volume label, and examine a device with a non-removable storage mechanism to determine if there are problems with the drive or with reading or writing data stored on the volume.

This document describes many of the functions of Device Support Facilities. It also discusses the differences between Device Support Facilities and IEHDASDR.

Although this document discusses all the operating environments supported by Device Support Facilities, particular emphasis is placed on the OS/VS2 (MVS) operating system.

For more information about the Device Support Facilities product, see <u>Device Support Facilities</u>, GC35-0033.

THIS PAGE INTENTIONALLY LEFT BLANK

CONTENTS

Device Support Facilities is a DASD utility package released in December 1978. Device Support Facilities replaces the initialization functions of IEHDASDR. Device Support Facilities also replaces IBCDASDI (the stand-alone version of IEHDASDR).

One of the objectives of Device Support Facilities was to provide a common syntactical language that would span all operating systems, including the stand-alone environment. Release 1 operated in the VS1, MVS and stand-alone environments. Release 2 and 3 operate in the following environments:

- VS1 (Rel 7.0)
- MVS (Rel 3.8)
- DOS/VSE AF (Rel 2)
- Stand-alone

Device Support Facilities was developed to improve the Reliability and Maintainability of the supported devices. Initialization type functions were designed to minimize pack returns to the factory. This was achieved by implementing a recovery scheme for the Home-address and/or the record zero fields on the DASD device. This recovery strategy is described in "Appendix B. Track Descriptor Recovery Strategy" on page 27.

Another objective of Device Support Facilities was to to provide a utility package that was easily extendible to new devices (such as the IBM 3380 disk drive).

THIS PAGE INTENTIONALLY LEFT BLANK

There are five commands that can be invoked by the user of Device Support Facilities:

INIT — Initialize a direct access volume. Supported in all SCP environments, including stand-alone.

INSPECT — Inspect a direct access volume. Supported in all SCP environments, including stand-alone.

REFORMAT — Reformat first track of the volume. Supported in all SCP environments, including stand-alone.

ANALYZE — Analyze DASD with non-removable media (e.g., 3350) for defects. Supported in all SCP environments, including stand-alone.

BUILDIX — Build an index to an existing VTOC (OS -> IX). This command can also reset a volume with an Indexed VTOC to a VTOC with no INDEX. BUILDIX is only supported in the VS2 and VS1 environments, and requires the Data Facilities/Device Support (DF/DS) product 5740-AM6 (VS1) or 5740-AM7 (MVS).

The code required to effect the execution of these commands exist in modules called "Command Controllers." These command controllers are very naive as to the environment in which they are executing, unless they are to perform a task specific to an operating environment.

These Command Controllers therefore have to interact with the specific operating system via an "operating system independent" interface. This interface consists of a set of modules that provide various services to the Command Controllers and the only reason for their existence is to serve the command controllers. These service modules are referred to as "Adapters" because they adapt the command controller to a specific operating environment.

INIT COMMAND

Capabilites

There are three levels of initialization available with the INIT command. These three levels of initialization are available via keywords that are specified with the command. In most cases the user may want to make the pack ready for his specific operating system (minimal). If it is necessary to validate the Home-addresses or 'wipe' the pack of any existing bit patterns, it would be necessary to request a medial initialization. Finally the user may want to check track surfaces and cause alternate track assignment or track reclamation based on the results (maximal).

Summarizing, the three levels of initialization are:

- Minimal Initialization
- Medial Initialization
- Maximal Initialization

MINIMAL INITIALIZATION

This is the fastest type of initialization available with the INIT command. The direct access volume is made ready for the operating system with very little checking of the volume for defective tracks, invalid track descriptors, etc. Note that such additional checking is provided by the Medial and Maximal levels of initialization. The following actions are taken on the volume:

1. Construction of the first track of the volume

   IPL bootstrap records 1 and 2 are formatted on the first track of the volume. The Volume label is formatted following these bootstrap records. This label contains the volume serial and, optionally, any owner identification that may be desired. If requested, user labels are then formatted on the volume. An IPL program record may optionally be formatted after the labels.

2. Creation of the Volume Table of Contents (VTOC)

   A VTOC is created on the volume in a user specified location or at a default location chosen by the program. This VTOC contains a Format 4 DSCB, a Format 5 DSCB (indicating all space available), and as many format 0 DSCBs as will fit on the number of tracks that the VTOC is to occupy.

   For an MSS Staging pack, the VTOC reflects that no space is available on the volume.

   In the DOS/VSE environment, one can request that a DOS VTOC be formatted on the volume. In this case, exactly the same kind of VTOC is created, but the Format 5 DSCB is marked invalid in the Format 4 DSCB.

3. Examination of alternate tracks and compression

   All the alternate tracks are examined for association with primary tracks. Complete association is assured. If there is a problem with this association, a warning message is issued but corrective action is not taken. Also, if there is an unnecessary fragmentation of the assigned alternate tracks, the alternate tracks are re-assigned to primary tracks in order. Thus the count of alternate tracks in the Volume Table of Contents and the pointer to the next available alternate track will be valid after this command is processed.

MEDIAL INITIALIZATION

This process performs more volume verification than the Minimal initialization previously described. Naturally this process requires additional CPU and I/O resources, and takes longer than a minimal initialization. The following occurs on a medial initialization:

1. All the functions of Minimal initialization

2. Validation of every track

   The home-address on every track is read and its validity checked. In the event that the home-address is invalid, it is re-written and thus the user is assured of the existence of a valid home-address. The standard-length Record zero is always written on every track; since this is a formatted write operation, the rest of the track is wiped clean of any existing bit patterns.

3. Automatic assignment of alternate tracks

   If during validation of a track, the defect-flag is set in the Home-address, then an alternate track is automatically assigned to it. Therefore, this level of initialization honors the current status of the defect-flag setting in the home-address. That is, tracks flagged as "defective" are not reclaimed.

MAXIMAL INITIALIZATION

Maximal initialization is the most time and system resource consuming level of initialization available with the INIT command. The following occurs on a maximal initialization:

1. All the functions of Medial initialization

2. Checking of every track surface for defects

3. Conditional alternate track assignment

   If a track is found defective during track surface checking, an alternate track is assigned to the primary track. Also, if reclamation is not permitted and the defect-flag is set, an automatic assignment of an alternate track will occur regardless of the results of track surface checking.

4. Conditional reclamation of tracks

   If a track was originally flagged defective and a surface check indicates a good track, this track can be reclaimed if the user specifies the RECLAIM parameter. If the NORECLAIM parameter is chosen, an alternate is assigned to the primary track regardless of the results of the track surface checking.

Usage

Following are examples of initializing a direct access volume at each of the three previously described levels. The common syntax used in the various operating systems is demonstrated.

- Minimal Initialization

  This is the quickest form of Initialization available. Note that in the following examples a volume is being re-initialized that already has a volume serial on it. For first time initialization, it would be necessary to specify the volume serial via the VOLID parameter. Also note that except for the last example, it is assumed that an Indexed VTOC is desired. In the event that a regular OS VTOC is desired, it is only necessary to remove the INDEX parameter.

  - /* VS1 and VS2 - volume is online */

    ```
    INIT DDNAME(ddname1) NVFY VTOC(0,1,15) INDEX(2,0,4)
    ```

  - /* VS1 and VS2 - volume is offline */

    ```
    INIT UNIT(2C0) NVFY VTOC(0,1,15) INDEX(2,0,4)
    ```

  - /* Stand-alone environment */

    ```
    INIT UNIT(2C0) DEVTYP(3380) NVFY VTOC(0,1,15) INDEX(2,0,4)
    ```

  - /* DOS/VSE - volume is online */

    ```
    INIT SYSNAME(SYS001) NVFY DVTOC(0,1,15)
    ```

- Medial Initialization

  To perform a Medial initialization, one need only add the "VALIDATE" parameter to the previous examples.

  - /* VS1 and VS2 - volume is online */

    ```
    INIT DDNAME(ddname1) NVFY VTOC(0,1,15) INDEX(2,0,4) -
    VALIDATE
    ```

- Maximal Initialization

  This is the most time consuming level of initialization for the volume. It is requested by specifying one parameter, CHECK(n), in addition to those parameters required for a Minimal initialization. This parameter indicates the number of times (between 1 and 10) that you want each track to be checked for recording errors during initialization. Because there are two special alternating bit patterns used during checking, you may wish to specify two cycles of checking to allow both bit patterns to be used to detect surface defects. Note, however, that each additional checking cycle **significantly** increases the time required for initialization.

  If there are defective tracks on the pack and it is desirable to reclaim them, it is also necessary to specify RECLAIM since unconditional reclamation is not available.

  - /* Use two bit patterns */

    ```
    INIT DDNAME(ddname1) NVFY VTOC(0,1,15) INDEX(2,0,4) -
    CHECK(2)
    ```

  - /* Re-claim the track if no errors detected */

    ```
    INIT DDNAME(ddname1) NVFY VTOC(0,1,15) INDEX(2,0,4) -
    CHECK(2) RECLAIM
    ```

INSPECT COMMAND

Capabilties

The INSPECT command can be used to effect selective repair actions on a direct access volume. The repair actions can be done on discrete tracks. Repair actions may be needed because of a previous history of bad performance on the subject track(s). Note that it is possible to INSPECT any tracks on the volume except the tracks on the CE cylinder. Hence it is possible to make any track on the volume unavailable. It is also possible to RECLAIM any track that has been made unavailable erroneously.

The capability exists to INSPECT the whole volume using the ALLTRACKS option. However, this option should be used sparingly as it can be quite time consuming to inspect a full volume, especially if the volume has lots of data on it and the PRESERVE option is selected. Note that Release 3 significantly improves the performance of of the PRESERVE option. This still does not suggest that the installation INSPECT an entire volume using the PRESERVE option.

The ALLTRACKS option should typically be used to obtain a status of the volume without any formatting write operations occurring on the volume. It is possible to do this if the keywords NOASSIGN, NORECLAIM, NOCHECK and MAP are specified.

CONDITIONAL ALTERNATE TRACK ASSIGNMENT

It is possible to inspect a set of tracks (or all the tracks) on a volume with a view of assigning alternate tracks if the said tracks are found to be defective via a track surface check. This is called a conditional assignment of alternate tracks. This type of inspection can occur on either primary or alternate tracks. In the case of alternate tracks, this kind of inspection could result in either marking the track defective if it is not associated with a primary track, or re-assigning the associated primary track to a new alternate track and flagging the alternate track under inspection defective.

UNCONDITIONAL ALTERNATE TRACK ASSIGNMENT

By specifying NOCHECK and ASSIGN, one can effect an unconditional assignment of alternate tracks. Note that due to the very nature of this combination of keywords, one cannot do this type of inspection on all the tracks of a volume. Unconditional alternate track assignment constitutes assignment of an alternate track with no further checking of the primary track for defects:

- An unassociated primary track will be unconditionally assigned an alternate track.

- A previously assigned primary track will be assigned a new alternate track on the assumption that the user does not wish to use the old alternate for his data. The old alternate is flagged defective and made unusable.

- When an alternate track is being inspected with this set of options, it is made unavailable by assigning the primary associated with it to a new alternate and flagging the subject alternate defective. If the subject alternate is not associated to a primary it is merely flagged defective and thus made unavailable.

CONDITIONAL TRACK RECLAMATION

Using the CHECK and RECLAIM options, it is possible to inspect a set of tracks that have been previously assigned alternate tracks, so that if no defects are found during track surface checking the tracks in question are reclaimed. This process is called "conditional track reclamation." No capability exists for an unconditional reclamation of a track. Inspection of an alternate track that is flagged defective with these options effects a reclamation of the alternate track and hence makes the alternate track available again. If a primary track is reclaimed, its associated alternate is made available for reassignment.

DATA PRESERVATION

It is possible to specify the PRESERVE option for each of the functions provided with the INSPECT command. This will cause the data that exists on a track to be read and written back on the track or its assigned alternate after any processing on the track is completed. (If a primary track is reclaimed, the data is written to the primary). Note that if any errors occur during the reading of the data on this track, processing of this specific track is bypassed. No attempt is made to correct the problem with that track. The recovery of the data from that track is not the function of this utility.

Usage

Following are examples of using the INSPECT command to assign alternate tracks and to reclaim erroneously rejected tracks.

- Conditional alternate track assignment

  CHECK(1) and ASSIGN are the defaults, and these are the parameters that are to be specified if conditional assignment of alternate tracks is desired. The following two examples show such usage:

  - INSP DDNAME(ddname1) NVFY TRKS(3:4)

  - INSP DDNAME(ddname1) NVFY TRKS(4) NOPRESERVE

- Unconditional alternate track assignment

  There may be occasions when intermittent errors occur on a track. In this case it may be desirable to unconditionally assign an alternate track to the suspect track. In another instance an alternate may be assigned to a primary and intermittent I/O errors still occur on this alternate track. In this situation, one may want to assign a new alternate to the primary. Finally there may be an alternate track with some known defects and it is desirable never to use this track again. In this case we can cause this alternate track to be marked defective. The following example shows such a use:

  - INSP DDNAME(ddname1) NVFY -
    TRKS((3,4) (12,10)) NOCHECK

- Conditional reclamation

  It may be desirable to reclaim certain tracks that may have been assigned unnecessarily. This can be achieved only in a conditional manner after the track surface is checked. The following examples show such a use of the command:

  - INSP DDNAME(ddname1) NVFY ALLTRKS RECLAIM

  - INSP DDNAME(ddname1) NVFY TRK(4,10) RECLAIM

- Data preservation capability

  The data on a track can be preserved before Device Support Facilities processes on the track for the INSPECT command. This option is the default and, in order to preserve data integrity, the user is encouraged not to change this default. However, the following example shows how to bypass data preservation:

  - INSP DDNAME(ddname1) NVFY TRKS(3,4) -
    NOCHECK NOPRESERVE

- Pack map generation

  The INSPECT command can be used to generate a pack map that indicates all the tracks that are flagged defective. Note that if there are no flagged tracks, this map will be an empty table and "emptiness signifies goodness."

  - INSP DDNAME(ddname1) NVFY ALLTRKS NOCHECK NOASSIGN NORECLAIM MAP

REFORMAT COMMAND

Capabilities

The REFORMAT command can be used to modify certain critical fields on the first track of the volume such as the volume serial and/or the owner identification. It can also be used to format an IPL program after the volume labels on the first track, thus converting the pack to an IPL volume. The IPL bootstrap records on the volume (records 1 and 2) are appropriately modified.

This command is used by AMDSADMP to place the stand-alone dump program on the specified volume.

Usage

The following are examples of using the REFORMAT command:

- Changing the Volume Serial

  The following example changes the serial number of the volume pointed to by the DD statement "ddname1" to "newvol." The VFY parameter is used to insure the volume being changed is "oldvol."

  - RFMT DDNAME(ddname1) VFY(oldvol) VOLID(newvol)

- Changing the Owner Identification

  The following example changes the owner information of the volume pointed to by the DD statement "ddname1" to "newowner." The VFY parameter is used to insure the volume being changed is "oldvol," and that its previous owner information was "oldowner."

  - RFMT DDNAME(ddname1)
    VFY(oldvol,oldowner) -
    OWNERID(newowner)

- Conversion of a Non-IPL volume to an IPL volume

  The following example adds IPL text to the volume pointed to by the DD statement "ddname1."

  - RFMT DDNAME(ddname1) IPLDD(ipldnam) NVFY

ANALYZE COMMAND

Capabilities

The ANALYZE command is for use with IBM direct-access storage devices that have non-removeable storage media (3310, 3344, 3350, 3370, 3375, 3380). ANALYZE examines the drive and the user's data to determine if there are errors. The operator is notified of the errors and diagnostic information is also written to the system printer. Two basic types of tests occur with the ANALYZE command. They are:

- Drive test

  ANALYZE first determines if the drive is fully operational by performing I/O operations that test the drive's functional capabilities. Tests are executed in order of increasing complexity. Problems that can be detected during ANALYZE's initial tests are access arm positioning errors and the ability to read and write on the CE cylinder. For a 3344, 3350, 3375, or 3380, the drive is also checked to ensure it can detect an address mark, switch read/write heads as a result of multi-track commands, sense the disk's rotational position, and detect and skip over defective areas.

- Data verification test

  When the drive test ends without an error, the data verification test can optionally be performed. The user specified portion of the volume is read without data transfer to the processor. If any data checks occur during this read process, the location of the error is noted in a diagnostic message that is written on the system printer.

ANALYZE does not require exclusive control of the drive under test. Most of ANALYZE's channel programs do not hold the drive busy for more than 0.5 seconds. However, operating system error recovery as a result of I/O errors detected during ANALYZE processing may hold the drive busy for a longer period of time. ANALYZE data verification testing allows the user to control the duration of each EXCP by specifying either SPEED or NOSPEED:

SPEED	One EXCP is issued per cylinder
NOSPEED	One EXCP is issued per track

Note: The use of ANALYZE output by customer engineers or service representatives to isolate and to repair drive malfunctions should not replace the use of more detailed diagnostic tools and procedures available for that purpose. ANALYZE output in such instances should serve only as the initial reference material that shows the final problem symptom but does not specifically isolate the problem source.

Usage

The following are examples of requesting that the drive and volume be analyzed for errors:

- Test drive only
  - ANALYZE DDNAME(ddname1)
- Data verification test (full volume)
  - ANALYZE DDNAME(ddname1) SCAN
- Data verification test (partial volume)
  - ANALYZE DDNAME(ddname1) SCAN -
    LIMITS(scanlo,scanhi)
- Data verification test (performance)
  - ANALYZE DDNAME(ddname1) SCAN SPEED

BUILDIX COMMAND

Capabilities

The BUILDIX command is to be used to change an existing VTOC on an already initialized volume to an Indexed VTOC. It can also be used to change an existing Indexed VTOC into a VTOC that is non-indexed. This function is available with release 3 of Device Support Facilities, and requires the DF/DS program product. For more information on this product, see Data Facility/Device Support: General Information, GC26-3954, which will be updated when DF/DS is available.

1. OSVTOC to IXVTOC

   The VTOC index data set must have been allocated prior to issuing the BUILDIX command. The Command controller will then examine each data-set via the Data Set Control Blocks (DSCB) in the existing VTOC and construct all the appropriate index records to access the dataset via the Index to the existing VTOC.

2. IXVTOC to OSVTOC

   It may be necessary to use a volume on a system where the INDEXED VTOC is not supported. Hence the BUILDIX command can be used to to change an existing indexed VTOC to a non-indexed VTOC. The index data set can be optionally scratched via the PURGE parameter.

Usage

The following examples show to convert a volume to use an INDEXED VTOC, and how to delete the INDEXED VTOC:

- OSVTOC to IXVTOC

  - BUILDIX DDNAME(ddname1) IXVTOC

  Note: The VTOC INDEX data set must have been previously allocated.

- IXVTOC to OSVTOC (and delete the indexed VTOC dataset)

  - BUILDIX DDNAME(ddname1) OSVTOC PURGE

Release 2 and 3 of Device Support Facilities have improved the performance over Release 1 in the following areas:

1. Elapsed time to do a medial initialization.

2. Elapsed time to do a maximal initialization.

3. Elapsed time to exercise the PRESERVE option of the INSPECT command.

Performance improvements exist in Release 2 (and therefore Release 3) for the first two problems. Performance improvement for the third problem is provided in Release 3.

The following CCW mnemonics are used in this section of the document:[1]

RHA Read Home Address
This command transfers the home address (HA) area of a track to main storage.

SHA Search Home Address Equal
This command causes the storage director to compare the home address in main storage with the home address on the track. Note that this is different from the RHA command in that the RHA command uses the physical track address and the SHA uses the value specified on track's home address record.

TIC Transfer in channel
Used in our examples to loop on the previous CCW until the desired condition is found.

WRZ Write record zero (R0)
Two forms of the WRZ command are used:

- Standard length - writes a "normal" 8 byte record 0.
- Maximum length - writes a full track record 0, to insure the track can be used. This record is always replaced by a standard record 0.

RRZ Read record zero (R0)
Two forms of the RRZ command are used: one to read a "standard" R0, and one to read a full track R0.

WHA Write Home Address
This command writes a home address record on the track.

---

[1] See "Appendix C. Channel Program Description" on page 29 for a detailed explanation of the channel programs being executed by Device Support Facilities.

MEDIAL INITIALIZATION

The elapsed time required by Device Support Facilities for a medial initialization of a direct access volume is directly dependent on the number of revolutions taken to validate the home-address on each track and write a standard-length record zero on it, thus effecting an erasure of any data that may exist on the track. Hence the way to improve the performance is to cut down the number of revolutions required for the validation of each track.

This has been achieved with Release 2 of Device Support Facilities. The actual CCW chains that are involved are compared between release 1 and release 2 or 3 of Device Support Facilities. The processing required depends upon the validity of the home address.

Valid Home Address

Release 1	Release 2
SIO1: RHA	SIO1: SHA
SIO2: RHA	TIC
SHA	WRZ (std)
TIC	RHA
WRZ (std)	RRZ (std)
RRZ (std)	

Figure 1. Medial Initialization - Valid HA

The above CCW chains are executed when the home-address on the track is valid and readable. In release 1 it takes 2 SIOs and 4 revolutions to do the job of validating a track and in release 2 the same job is achieved in 2 revolutions and one SIO.

Invalid Home Address

Figure 2 on page 17 shows the channel programs required when an invalid home address is encountered. In this case, release 1 would take 3 SIOs and 8 revolutions to validate a track. Release 2 takes only 2 SIOs and 4 revolutions to validate the track. Hence release 2 of Device Support Facilities provides an approximate 50 percent improvement over the original release.

Release 1			Release 2		
SIO1:	RHA		SIO1:	SHA	(fails)
SIO2:	RHA			TIC	
	SHA			WRZ	(std)
	TIC			RHA	
	WRZ	(std)		RRZ	(std)
	RHA		SIO2:	RHA	
	RRZ	(std)		SHA	
SIO3:	RHA			TIC	
	SHA			WHA	
	TIC			WRZ	(std)
	WRZ	(std)		RHA	
	RRZ	(std)		RRZ	(std)

Figure 2. Medial Initialization - Invalid HA

MAXIMAL INITIALIZATION

The elapsed time taken by Device Support Facilities for a maximal initialization of a direct access volume is directly dependent on the number of revolutions taken to do the following actions on a specific track:

- Validate the existing Home-address.
- Write a maximum length Record-zero [assume CHECK(1)].
- Read the maximum length Record-zero.
- Write a standard length record zero on the track.

This action must occur on every track of the direct access volume.

Hence the way to improve the performance is to cut down the number of revolutions required for the validation and surface check of each track. This has been achieved and the actual improvement is discussed below on a comparative basis. The actual CCW chains that are involved are compared between release 1 and release 2 of Device Support Facilities.

Valid Home Address

Figure 3 on page 18 shows the CCW chains executed when the home-address on the track is valid and readable. In release 1 it takes 3 SIOs and 8 revolutions to do the job of validating a track and in release 2 the same job is achieved in 5 revolutions and 2 SIOs.

Release 1	Release 2
SIO1: RHA	SIO1: SHA
SIO2: RHA	TIC
SHA	WRZ (max)
TIC	RHA
WRZ (max)	RRZ (max)
RRZ (max)	SIO2: SHA
SIO3: RHA	TIC
SHA	WRZ (std)
TIC	RRZ (std)
WRZ (std)	
RRZ (std)	

Figure 3. Maximal Initialization - Valid HA

Invalid Home Address

Release 1	Release 2
SIO1: RHA	SIO1: SHA (fails)
SIO2: SHA	TIC
TIC	WRZ (max)
WHA	RHA
WRZ (std)	RRZ (max)
RHA	SIO2: RHA
RRZ (std)	SHA
SIO3: RHA	TIC
SHA	WHA
TIC	WRZ (max)
WRZ (max)	RHA
RRZ (max)	RRZ (max)
SIO4: RHA	SIO3: SHA
SHA	TIC
TIC	WRZ (std)
WRZ (std)	RRZ (std)
RRZ (std)	

Figure 4. Maximal Initialization - Invalid HA

When the Home address is invalid it would take release 1, 4 SIOs and 11 revolutions to check a track. Release 2 takes only 3 SIOs and 7 revolutions to validate the track.

Hence release 2 of Device Support Facilities provides approximately 33 percent improvement over release 1.

DATA PRESERVATION FOR INSPECT

The elapsed time taken by Device Support Facilities to read all the data from a direct access volume is directly dependent on the number of revolutions taken to read all the records from a single track. In Releases 1 and 2, the "Read Multiple Count Key Data" CCW was not used. Consequently many revolutions were required to read all the records from a track. In release 3, all the records from a track are read in a minimum of one revolution and a maximum of 2 revolutions depending on whether the "Read Multiple Count Key Data" CCW is supported for the specific device.

THIS PAGE INTENTIONALLY LEFT BLANK

Four environments exist in which dataset security and protection is required:

- Stand-alone environment
- OS/VS environment (Offline mode)
- OS/VS environment (Online mode)
- DOS/VSE environment (Not discussed further)

STAND-ALONE AND OS/VS (OFFLINE MODE) ENVIRONMENTS

The following parameters provide data set security and protection:

- VERIFY parameter of the INIT, INSP and RFMT commands.

  This parameter insures that the correct VOLSER is specified. It can also insure the the volume's "owner identification" field contains the desired value.

- PURGE/NOPURGE parameter of the INIT command is ignored.

- Permission from the operator required for purging the volume contents for INIT, INSP and RFMT commands (message ICK003D).

OS/VS (ONLINE MODE) ENVIRONMENT

The following features of Device Support Facilities and MVS provide data set security and protection:

- VERIFY parameter for commands INIT, INSP and RFMT commands

  This parameter insures that the correct VOLSER is specified. It can also insure the the volume's "owner identification" field contains the desired value.

- PURGE/NOPURGE parameter for the INIT command

  If NOPURGE is specified or defaulted, execution of the command terminates if unexpired data sets, password protected data sets, RACF protected data sets, or VSAM data sets are found on the volume. If PURGE is specified, data loss is prevented as follows:

  - Unexpired data sets

    The operator is requested (once) to confirm that the contents of the volume can be purged.

- Password protected data sets

  The operator must provide the correct password for each password protected data set. If the data set is also RACF protected, password protection is ignored.

- RACF protected data sets

  Device Support Facilities insures the job has either volume authority or is authorized to scratch every RACF protected dataset on the volume.

- VSAM data sets

  VSAM datasets cannot be scratched in the online mode. Initialization is only possible in the offline mode.

- Authorized Program Facility (MVS Only)

Device Support Facilities is an authorized program and all the associated load modules exist in the authorized library SYS1.LINKLIB, except for IGG019P2 which exists in SYS1.SVCLIB. To restrict the use of Device Support Facilities, the installation could move the modules to an authorized, RACF protected library.

PROCESSING IN A SHARED ENVIRONMENT

In Release 3 of Device Support Facilities, the RESERVE macro is issued by the INIT, INSPECT, BUILDIX and REFORMAT commands prior to accessing the device for write operations. A DEQ macro is issued at the end of the command's processing. The resource name used is "SYSVTOC/volid." The RESERVE requests exclusive use of the resource.

THIS PAGE INTENTIONALLY LEFT BLANK

APPENDIX A. COMPARISON BETWEEN ICKDSF AND IEHDASDR

Minimal Initialization or 'QUICK DASDI'

The INIT command of Device Support Facilities provides initialization of a device at the minimal level as the default. The time taken to perform this function on a device is negligible. The same is true for the level of initialization provided by the "ANALYZE PASSES=0" control statement of the IEHDASDR utility. Note that "ANALYZE PASSES=0" is the only possible type of initialization available for 3330, 3340 and 3344 devices using IEHDASDR.

Medial Initialization or FORMAT

This is the only function provided by both the utilities that is comparable in function. The time taken to perform this function by both utilities should be comparable. Note that this is the only function provided in IEHDASDR that is standardized for all supported devices.

Maximal Initialization or ANALYZE with PASSES > 0

This function is provided in Device Support Facilities by the INIT command when CHECK(n) is specified. The only devices on which IEHDASDR performs comparable function are the 2314 and the 2311. On 3330, 3340 and 3344 devices, a specification of "PASSES > 0" is meaningless as "PASSES=0" is forced and the user always gets a "QUICK DASDI" on these devices. If FLAGTEST=YES is specified for a 3350, surface analysis is performed **only** on the tracks that are flagged defective. Device Support Facilities, on the other hand, will check the media surface of **all** the tracks of all supported devices. Hence only on 2311 and 2314 device types can the maximal initialization functions of the two utilities be compared. Note that track reclamation based on surface analysis is possible only on the 3350 with IEHDASDR whereas this is possible on all devices except the 2305 using Device Support Facilities.

Alternate track assignment (INSPECT & GETALT)

Both the utilities will assign alternates for primary tracks. However, the manner in which this assignment occurs however varies between the utilities. Device Support Facilities will effect the assignment either conditionally or unconditionally based on a surface analysis of the track in question. Conditional assignment of a track is possible using IEHDASDR only for a 3350. Unconditional assignment is possible on all the devices. Additionally, once a track is assigned an alternate track, it is not possible to reclaim that track with IEHDASDR.

Track Reclamation (INSPECT)

Track reclamation is possible with Device Support Facilities based on the results of track surface checking. IEHDASDR has no capability to reclaim assigned tracks (without re-initializing the volume).

APPENDIX B. TRACK DESCRIPTOR RECOVERY STRATEGY

When Device Support Facilities detects a defective track, it assigns an alternate track. To accomplish this, Device Support Facilities updates the Home Address (HA) and Record Zero (R0) on the defective track to indicate the problem, and to specify the address of the alternate. However, this assumes Device Support Facilities is able to successfully process the HA and R0 of the defective track. In most cases this simply involves a re-write of the HA and the R0 on the track. However, if this fails, Device Support Facilities can, for some devices, go even further in the recovery effort by trying to restore the HA and R0 further down the track. If this recovery technique works, the track is still flagged as defective. The purpose of writing the HA and R0 is to allow the track to indicate the address of the alternate. If this technique fails, the volume is unusable.

2311 and 2314

For these two devices, it is only possible to move the HA (and R0) down a fixed number (705) of bytes. After such a move, the track has to be flagged defective and an alternate is therefore assigned to it.

3340 and 3344

These devices will tolerate a single defect on the track. This single defect on the track is reflected in the Skip displacement bytes that precede the HA. Hence on a specific track, it is possible to move HA down the track, or if the HA is readable but the R0 is not, then the R0 can be moved further down the track. Note that such a move of the track descriptor fields requires the track to be flagged defective.

3350

The 3350 device will tolerate 3 defects on a track. These 3 defects are reflected in the Skip displacement bytes that precede the HA. Hence on a specific track, it is possible to move the HA and/or R0 further down the track. Note that such a move of the track descriptor fields requires the track to be flagged defective.

3375 & 3380

Field Engineering will have a program available to provide an equivalent level of support.

THIS PAGE INTENTIONALLY LEFT BLANK

APPENDIX C. CHANNEL PROGRAM DESCRIPTION

Following is a description of the channel programs used to perform MAXIMAL initialization given a valid home address. The reader should be able to use this description to understand the other channel programs presented in this document. This table has four columns:

1. I/O - The number of the I/O request (1 or 2).
2. CCW - The CCW mnemonic.
3. X - The number of times the index marker has been passed. This is the number of revolutions required to perform the initialization of a track.
4. Description - A description of the processing performed by this CCW.

I/O	CCW	X	Description
1	SHA	1	Orients to the index marker on the specified track.
	TIC		Loop on SHA until at the index.
	WRZ (max)		Write R0, using the entire track. This is not the "real" R0.
	RHA	2	Re-orient to the HA.
	RRZ (max)		Read (and thus verify) the full-track R0.
		3	Processing the interrupt will cause a missed revolution.
2	SHA	4	Orients to the index marker on the specified track.
	TIC		Loop on SHA until at the index.
	WRZ (std)		Write the "real" record 0 on the track.
	RRZ (std)	5	Read (and thus verify) the "real" R0.
			Processing the interrupt will NOT cause a missed revolution. Have from end of RRZ to index marker to start the next I/O.

Figure 5. Valid Home Address Processing

GG22-9216-00

upport Facilities: Capabilities and Usage Printed in U.S.A.

International Business Machines Corporation
Data Processing Division
1133 Westchester Avenue, White Plains, N.Y. 10604

IBM World Trade Americas/Far East Corporation
Town of Mount Pleasant, Route 9, North Tarrytown, N.Y., U.S.A. 10591

IBM World Trade Europe/Middle East/Africa Corporation
360 Hamilton Avenue, White Plains, N.Y., U.S.A. 10601