

# IBM

## **XISOLATE Version 1.1.2 User's Guide**

Dr. Donald J Chesarek

IBM Corporation  
5600 Cottle Road  
San Jose, CA 95193  
Phone (408)-256-3677  
chesarek@us.ibm.com  
SJEVM5(CHESAREK)

**Note**

Before using this information and the product it supports, be sure to read the general information under Appendix B, "Notices" on page 15.

**Third Edition (February 2003)**

This edition applies to Release 1.1.2 of XISOLATE and to all subsequent releases and modifications until otherwise indicated in new editions.

**Copyright International Business Machines Corporation 1999-2003. All rights reserved.**  
US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

---

# Contents

<b>About this guide</b> . . . . .	v
Who should read this guide . . . . .	v
System requirements . . . . .	v
Where to find more information . . . . .	v
Summary of changes . . . . .	v
<b>Chapter 1. Introduction</b> . . . . .	1
Background/history . . . . .	1
Planned enhancements . . . . .	2
<b>Chapter 2. The isolation task and the XISOLATE tool set</b> . . . . .	3
The isolation task . . . . .	3
Defining data sets that need to be isolated . . . . .	3
Identifying the data sets that need to be isolated . . . . .	4
Obtaining information from the DASD subsystems . . . . .	5
The XISOLATE tool set . . . . .	5
Structure . . . . .	5
Using the tool in production . . . . .	6
Automating the XISOLATION task . . . . .	6
<b>Chapter 3. Program Directory</b> . . . . .	7
Loading the programs . . . . .	7
Customizing . . . . .	8
JOB card . . . . .	8
LINKedit jobs . . . . .	8
Data set allocation job - IEFBR14 . . . . .	8
Editing XSTAT02P . . . . .	9
Editing XISOJOB . . . . .	9
Testing . . . . .	10
Creating a sample list of cataloged data sets . . . . .	10
Additional customization of XISOJOB . . . . .	10
Executing the test . . . . .	10
Editing the data set . . . . .	10
<b>Appendix A. Data sets that should be XISOLATed</b> . . . . .	13
Sysplex couple data sets . . . . .	13
Other critical data sets . . . . .	13
General high-availability design recommendations . . . . .	14
<b>Appendix B. Notices</b> . . . . .	15
Trademarks . . . . .	15
<b>List of abbreviations</b> . . . . .	17



---

## About this guide

This guide explains the purpose and use of the XISOLATE tool.

---

## Who should read this guide

System programmers responsible for managing common tools who perform the installation and define the data set names associated with central services types of functions, as well as database administrators who define the data set names that are used with their database.

---

## System requirements

XISOLATE can be used with any DASD subsystem that works in an IBM OS/390 environment and supports the CCWs used with the IBM 3990-1. XISOLATE does not support IBM 3880-based subsystems. It can also be used with HDS 7700 and 7700E subsystems, and with IBM 3990-based DASD, IBM 9393 (RVA) and IBM 2105 (Shark) subsystems.

---

## Where to find more information

XISOLATE is available on an “as is” basis. Questions on the use of the tool should be submitted by the IBM Storage Technical Specialist to the ATS organization using ViewBlue/WWQ&A.

---

## Summary of changes

Version 1.1.2 contains the following changes:

1. Added the ability to specify UCB=(0000,0) with the XSTAT02A program which eliminates the need to specify DASD address ranges. The new option will cause the program to obtain data from all DASD devices that are included in the GEN for the LPAR that executes the program.
2. Changed to internal logic used to process IHELIST data in order to improve its use in problem determination situations.
3. Changed the file names to facilitate coexistence between this version of the code and earlier versions.

The following changes have been made to XISOLATE for Version 1.1.0 from the initial release, November 18, 1999:

1. Added checking for incorrect specification of DASD device number ranges.
2. Added ability to change data set name qualifiers from the job control language (JCL) of XISOJOB, which eliminates the need to edit Restructured Extended Executor (REXX) members. This simplifies the programmed generation of a JCL.



---

## Chapter 1. Introduction

Many software systems use, create, and manage multiple copies of control data sets (that is, software duplexing), to allow functions to continue to operate when a direct access storage device (DASD) subsystem fails and one of the copies becomes unavailable. To support this availability feature, the installation program has to ensure that each of the copies of the data sets resides on separate physical DASD subsystems. With Storage Management Subsystem (SMS) managed storage and modern migration tools, the task of ensuring that each copy of these critical data sets resides on physically isolated DASD subsystems has become increasingly complex.

The coexistence of certain data sets on the same DASD subsystem can cause performance problems. Therefore, it is important to ensure that these data sets reside on physically isolated DASD subsystems. XISOLATE is a set of programs designed to simplify this task.

The installation defines one or more groups of data sets that are to be isolated, and one of the jobs in the XISOLATE tool set is then used to perform the isolation-verification task. The batch job creates messages that go to SYSLOG to report the success or failure of the isolation analysis.

---

## Background/history

The availability requirements for the data set isolation described in the overview were violated by accident in a large installation and the DASD subsystem on which both copies had been placed failed, resulting in a major system outage. In this particular incident, the DASD subsystem involved was not an IBM product. The customer asked IBM for help in ensuring that this type of outage would be prevented in the future. DASD subsystem failures have occurred since the introduction of the IBM 305 RAMAC system, and many of the current users of XISOLATE have reported this type of system outage due to the lack of isolation on subsystems of various vendors and various vintages.

Key to ensure isolation is the ability of a program to recognize physically distinct DASD subsystems from data that is available through CCWs. The DSTAT00A tool allowed customers to obtain model and serial number information for IBM 3990, 3390 and RAMAC technology-based subsystems. It allowed the installation programs to identify physical units down to the drawer level in a RAMAC rack. The tool was developed to help installations with asset management and hardware-configuration management tasks.

Currently, there is no standard easy-to-use interface through which an installation can obtain the information required to identify physically distinct DASD subsystems. XISOLATE uses a derivative of DSTAT00A named XSTAT02A, which is designed specifically to provide the information necessary to accomplish the isolation task.

---

## Planned enhancements

For future releases of the XISOLATE tool, it is planned to give you the added ability of verifying **before** a logical volume move that the isolation will be preserved for volumes that are moved during a hardware migration or moved to performance balance across DASD subsystems.



---

## Chapter 2. The isolation task and the XISOLATE tool set

This chapter covers the steps to isolate the required data sets, and it explains the XISOLATE tool set.

---

### The isolation task

The isolation task has three basic steps:

1. Defining the data sets that need to be isolated.

This step requires user definition of the DASD address ranges that are to be included in the analysis.

2. Identifying the data sets that need to be isolated.

3. Obtaining information from the DASD subsystems that allows identification of physically distinct DASD subsystems.

The XSTAT02A program will do I/O to the defined DASD to obtain the data required for the isolation analysis.

By evaluating the information from step 1 through step 3, you can ensure that the data sets reside on physically different DASD subsystems.

### Defining data sets that need to be isolated

Data sets that need to be isolated are placed into groups. Each group contains a list of data sets that have to be placed on physically different DASD subsystems. Two data sets that are in different groups can exist on the same DASD subsystem.

Data sets that are to be isolated can be either cataloged or non-cataloged. For data sets that are non-cataloged, you must specify the volume serial (VOLSER) of the volume that contains the data set in addition to the data set name.

XISOLATE allows you to define a group that contains both cataloged and non-cataloged data sets.

Place cataloged data sets in one data set and non-cataloged data sets in a second data set.

Records in both of these data sets are similar in structure. Positions 1 and 2 in each record contain a two-character group identifier. A blank character in position 1 is used to designate a comment record. The other fields are separated from the group identifier and from each other by one or more blanks. Each record in the cataloged data set contains two fields: the group identifier and data set name. Each record in the data set that defines non-cataloged data sets contains three fields: the group identifier, data set name and the VOLSER on the volume that contains the data set.

In the following example, the records for the two data sets have been concatenated into a single series of records. Comments are used to identify the data set that contains the records.

```

    This is a comment line that starts the cataloged data set list
1 SYS1.XCF24.CDS 1
1 SYS1.XCF24.CDS 2

2 SYS1.WLM28.CDS 1
2 SYS1.WLM28.CDS 2

```

```

    The following entry has as its mate a non-cataloged data set
X C816171.CAT

```

```

    This is a comment line that ends the cataloged data set list

```

```

    This is a comment line that starts the non-cataloged data set list
A SYS1.JESCKPT1     CHK3 1
A SYS1.JESCKPT2     CHK3 2

```

```

    The following entry is the mate to the X0 cataloged entry
X C816171.NOTCAT     EX14B

```

```

    This is a comment line that ends the non-cataloged data set list

```

In the example, there are four groups with identifiers: 01, 02, 0A and X0. Data sets SYS1.XCF24.CDS01 and SYS1.XCF24.CDS02 are in group 01 and they are required to be on separate physical DASD subsystems. Data sets C816171.CAT and C816171.NOTCAT are in group X0 and are required to be on separate physical DASD subsystems. C816171.CAT is cataloged and C816171.NOTCAT is non-cataloged.

## Identifying the data sets that need to be isolated

Establishing the list of names of data sets that have to be isolated is probably the most complex part of the process: There are no common standards for the data set names and they are system- or application subsystem-dependent. For example, in a large installation, the data sets could be owned by separate organizations within the enterprise.

Use the following operator commands to obtain the data set names (and VOLSERS) for some of the data sets that you have to isolate. Note that the data set names used for IMS logs, DB2 logs, CICS logs, MQSeries data sets, and so on, vary. You have to determine which ones to use.

SMS control data sets

Use the **D SMS** command. These data sets are cataloged.

Job entry subsystem (JES) control data sets

Use the **\$D CKPTDEF** command. These data sets are non-cataloged in many environments.

Sysplex couple data sets

Use the **D XCF,COUPLE** command. These data sets are cataloged.

## Obtaining information from the DASD subsystems

The XSTAT02A program is used to obtain information from the DASD subsystems, which allows the program to perform the isolation analysis.

The following code is a sample of output from the XSATA02A program. All of the logical volumes in this sample are from the same DASD subsystem:

```
XSTAT 2A DATE: 993 9 TIME: 223623
DEVN MODEL  VOLSER STORAGE  SSID MFG STG CNTL
                CONTROL          SERIAL
5   339 -3  Y2KA3  9394  3  134 IBM 13-      TC 35
5  1 339 -3  Y2KB3  9394  3  134 IBM 13-      TC 35
5  2 339 -3  Y2KSPL 9394  3  134 IBM 13-      TC 35
5  3 339 -3  Y2KMAS 9394  3  134 IBM 13-      TC 35
5  4 339 -3  Y2KM3  9394  3  134 IBM 13-      TC 35
5  5 339 -3  Y2KLIB 9394  3  134 IBM 13-      TC 35
5  6 339 -3  EI 5 6 9394  3  134 IBM 13-      TC 35
5  7 339 -3  Y2KHM2 9394  3  134 IBM 13-      TC 35
5  8 339 -3  Y2KHM3 9394  3  134 IBM 13-      TC 35
5  9 339 -3  Y2KBG2 9394  3  134 IBM 13-      TC 35
5  A 339 -3  Y2KCL1 9394  3  134 IBM 13-      TC 35
5  B 339 -3  Y2KCTL 9394  3  134 IBM 13-      TC 35
```

---

## The XISOLATE tool set

This section explains the structure and use of the XISOLATE tool set, and it covers the steps necessary to automate the XISOLATION task.

### Structure

The members of the .PDS that contain the tool can be divided by functions as follows:

Members that contain elements that are used during setup

JOBCARD A model job card  
LKEDANNX The LKED job for the SYSLOG annotation program  
LKEDXS2A The LKED job for the XSTAT02A program  
IEFBR14 A job that allocates data sets used by other jobs  
XSTAT02P A sample job to test the XSTAT02A program  
ANNOTATX A sample job to test the SYSLOG annotation program

Members that are used when the XISOLATE function is used in the batch mode

XISOJOB The model XISOLATE batch job  
XISOLATE The REXX program that performs the isolation analysis

Members that contain utility functions

ANNOTATX A sample job to test the SYSLOG annotation program  
LKEDDS2A The LKED job for the DSTAT02A program  
DSTAT02P A sample job to test the DSTAT02A program  
XISOJOBX A job that executes only the XISOLATE REXX program  
XUNIJOB A job that obtains unique serial numbers from XSTAT02A data  
XUNIQUE The REXX program that performs the "unique" analysis

## Using the tool in production

It is recommended that you perform the isolation-verification process before the start of the business day. In addition, you should run it following an IPL or a data migration. The batch program is able to complete the verification process in an environment with over 8000 volumes in about 2 minutes elapsed time. The elapsed time may be considerably longer if XSTAT02A encounters volumes that are hardware-reserved to other systems.

## Automating the XISOLATION task

The XISOLATION task can be easily integrated into the normal daily processing of an installation. Usually, you can do this with the following steps:

1. Create a process that runs the XISOJOB at a set time each day.
2. Create a process that traps the SYSLOG message that indicates an isolation failure and sends a notification to the appropriate people.
3. Create a process that runs the XISOJOB following each IPL.
4. Create a process that automatically develops the control cards within the XISOJOB that are used to define the DASD for the configuration.

---

## Chapter 3. Program Directory

This chapter explains the loading, customizing and testing of the programs.

---

### Loading the programs

XISOLATE is packaged as a single file. You might be able to transfer the file directly from an IBM FTP site to your MVS environment. If this is the case, go directly to step 3:

1. The programs are distributed as a single file that is designed to be easily transferred through a PC as a binary file. Because the file contains object modules, it must be handled as a binary file while on a PC. When you transfer the file from a PC to the MVS environment, the following parameters must be in effect:
  - a. The carriage return and line feed options should be set to **no**.
  - b. The ASCII-EBCDIC conversion options should be set to **no**.
2. The file should be uploaded to an MVS system as a sequential file. The file as it is received by MVS **must** have the following attributes: **RECFM=FB, LRECL=80**. The block size can be system-determined. The way in which RECFM and LRECL are set on the PC depends entirely on its file-transfer program.

**Note:** The default record format for most PC file-transfer programs for .BIN files is RECFM=V. File-transfer programs have a setup function that allows you to create a new transfer rule that gives you the required record format and record length.

Use the name DSN="userid.XISOLATE.PDS.XFRSEQ" as the target MVS data set; *userid* is your TSO user ID or a common high-level qualifier for a storage management "group." If more than one user builds XISOLATE jobs, the qualifier should be selected to give each potential user access to the common files.

3. After the file has been transferred to an MVS environment, it must be converted from the sequential format back into the .PDS format from which it was originally packaged. To perform this conversion, use the **TSO RECEIVE** command. To convert the sequential file to a .PDS file, go to the TSO command line and enter the following TSO command:

```
RECEIVE INDSN ( 'userid.XISOLATE.V112.PDS. XFRSEQ' )
```

You are prompted for the name of the .PDS file. Enter:

```
DSN( 'userid.XISOLATE.V112. PDS' )
```

The target .PDS data set does **not** have to be preallocated.

4. Take a look at DSN= 'userid.XISOLATE.V112.PDS'. The .PDS directory and members should look like they were copied using the **IEBCOPY** command from the source .PDS file.

---

## Customizing

The data set naming convention used by XISOLATE allows the installation to specify the high-level and third-level qualifiers in data set names. In the following section, the high-level qualifier are referred to as *highqual* and the third-level qualifier are referred to as *q3qual*.

On most installations, the System Management Facility (SMF) system ID is used as the third-level qualifier.

The total number of available characters in *highqual* and *q3qual* combined is 14.

## JOB card

Edit the member JOB card. Create a JOB card that can be used in customizing other XISOLATE jobs.

## LINKedit jobs

There are two LINKedit jobs that you have to run before you can use XISOLATE. The load modules **must** be placed into an APF-authorized library.

The LINKedit jobs are in the following members:

LKEDANNX console annotation module  
LKEDXS2A module used to develop physical device information

Customize the LINKedit jobs as follows:

1. Replace the job card with JOBCARD.
2. Change SYS1.DJC.LINKLIB to the name of the authorized library that you have selected to hold the load module. Record the name of this data set because you will need it in the steps that follow.
3. Submit the job. It should complete with a return code of 0.

## Data set allocation job - IEFBR14

This job allocates the following four data sets:

highqual.XISOLATE.q3qual.CONTROL.LIST.CAT  
highqual.XISOLATE.q3qual.CONTROL.LIST.NOTCAT  
highqual.XISOLATE.q3qual.CONTROL.IEHLIST.DATA  
highqual.XISOLATE.q3qual.SUMMARY.DATA

Customize the IEFBR14 job as follows:

1. Replace the job card with JOBCARD.
2. Change C816171 to the high-level qualifier you want to use.
3. Change EDS1 to the q3qual value you want to use to test the tool.
4. Submit the job. It should complete with a return code of 0.

## Editing XSTAT02P

This job develops one record that contains physical identifier information for each device in the list. It is set up as a multistep job: the first two steps are IEFBR14 steps that ensure that each time the job is executed, it starts with a new copy of the XSTAT02A.DATA data set. The third through *N*th step are steps that execute the inline PROC XSTAT02A. In the XSTAT02P model, the last six steps have been commented out.

The XSTAT02P job is customized as follows:

1. Replace the job card with JOBCARD.
2. Change the JOBPARM card as required.
3. Change SYS1.ESP.LINKLIB to the name of the authorized library that you used with the LKED jobs.
4. Change C816171 to your selected highqual value.
5. Change EDS1 to the q3qual value you are using.
6. Do a find on UNIT=0500. Change the 0500 to the MVS device number of the set of consecutive DASD addresses on your test system.
7. Submit the job. It should complete with a return code of 0.

If the program abends with either an 047 or 133 abend code, then the load library is not APF-authorized.

Browse highqual.XISOLATE.q3qual.XSTAT02A.DATA. It contains records similar to those shown earlier for this data set.

## Editing XISOJOB

Use this job to verify the isolation of data sets.

Customize the XISOJOB job as follows:

1. Replace the job card with JOBCARD.
2. Change the JOBPARM card as required.
3. Change SYS1.DJC.LINKLIB to the name of the authorized library that you have selected to hold the load module (three places).
4. Change C816171 to your selected highqual value.
5. Change EDS1 to the q3qual value you are using.
6. Do a find on UNIT=0500. Modify the lines D#0500 through D#9000 to match the DASD attached to your test system. The parameter N= defines the number of consecutive addresses defined for the DASD string.

**Note:** It takes about 2 minutes to obtain the physical identifier information for 8000 devices.

7. Submit the job. It should complete with a return code of 8. The job output should include a line like the following code:

```
JOB 1 4 XISOLATE - ISOLATION TEST FAILED
```

The last line of the job output contains a message indicating that less than two data sets were in the list. Note that you allocated the data sets that define the

data sets to be isolated, but you did not put any records with data set names into them.

---

## Testing

Test with data sets using the steps in the following sections.

### Creating a sample list of cataloged data sets

1. Edit dataset= "highqual.XISOLATE.q3qual.CONTROL.LIST.CAT"
2. Create a data set with the following four records. You have to modify the data set names to match your actual Sysplex names:

```
1 SYS1.XCF24.CDS 1
1 SYS1.XCF24.CDS 2
2 SYS1.WLM28.CDS 1
2 SYS1.WLM28.CDS 2
```

### Additional customization of XISOJOB

To implement this testing step, you need to develop a more complete list of the DASD that is on the test system. To do this, you have to edit XISOJOB and update the list of DASD addresses by adding steps as required.

### Executing the test

Submit the job in XISOJOB. If the data sets are on different physical DASD subsystems, then the job completes with a return code of 0 and the JOB log includes the following statement:

```
JOB 1 8 XISOLATE - CORRECT ISOLATION HAS BEEN VERIFIED
```

The ISOLATION TEST FAILED message indicates that there is an isolation problem that needs correction.

### Editing the data set

The data set dataset="highqual.XISOLATE.q3qual.SUMMARY.DATA" is similar to the following sample code. It contains four groups of data sets that have to be isolation-verified:

```
1 SYS1.XCF24.CDS 1           MAS  A 933B IBM 21 5    35
1 SYS1.XCF24.CDS 2           MAS  B 14DF IBM 399  A 276
2 SYS1.WLM28.CDS 1           MAS  A 933B IBM 21 5    35
2 SYS1.WLM28.CDS 2           MAS  B 14DF IBM 399  A 276
A SYS1.JESCKPT1              CHK3 1 13 1 EMC 399    781
A SYS1.JESCKPT2              CHK3 2 149A IBM 399  A 276
B SYS1.NEWCKPT1              JES9  ???? ??? ???? ?????
B SYS1.NEWCKPT2              JES9  5 8 HTC 399  35524
```

The columns contain the following values:

- Group identifier
- Data set name
- Label of volume containing the data set
- Device number of volume
- Manufacture's code
- Model of storage control



### Low-order five digits of the manufacturer's serial number

The question marks mean that the program could not obtain the required information. Question marks prevent the program from validating isolation and the result is a `FAILED` message.

Question marks usually mean that you forgot to include some DASD address ranges when you customized the `XISOJOB`.

Isolation requires that the string composed of manufacturer's code, model and serial number be different for each data set in each group. If this string is the same for any two data sets in a group, then these two data sets are in the same physical DASD subsystem.



---

## Appendix A. Data sets that should be XISOLATed

---

### Sysplex couple data sets

Sysplex	Primary/Alternate/Spare
CFRM	Primary/Alternate/Spare
SFM	Primary/Alternate/Spare
ARM	Primary/Alternate/Spare
WLM	Primary/Alternate/Spare
LOGR	Primary/Alternate/Spare

In addition to ensuring that primary, alternate, and spare couple data sets (CDSs) are allocated on a different DASD subsystems, do not allocate the primary Sysplex and the primary CFRM CDSs on the same volume (XISOLATE does not check this).

---

### Other critical data sets

RACF database (primary/backup)

JES2 Checkpoint data sets and spool volumes

If there are multiple JES2 subsystems, you have to isolate the data of each JES2 subsystem.

DFHSM control data sets

With DFSMS 1.5, multiple HSMplexes can coexist within a single Sysplex. If your system is configured for this type of environment, then you should ensure that the CDSs for each HSMplex are isolated.

NetView data sets

Each NetView subsystem has its own set of unique data sets. The data sets for each of the individual subsystems should be isolated.

SMS control data sets

ACDS data sets should be isolated.

SMF data sets

IMS

- RECON data set - all three copies (two active and one backup)
- Log data sets - WADS and OLDS
- MADS data sets
- For IMS Version 7, HALDB partitions should be isolated to minimize the impact of a single DASD failure.

DB2 BSDS and logs

- In data sharing, there are multiple DB2 members. It is recommended that the BSDSs and LOGs be isolated as much as possible at each member

level to prevent a single DASD subsystem failure from impacting all DB2 members in the same Sysplex.

- Primary and alternate BSDS should be isolated.
- There are multiple active log data sets and each log has an A and a B copy. At a minimum, LOGCOPY A/B should be isolated.
- Archive logs: normally, one copy goes to DASD, and the second one goes to tape. If both copies go to DASD, then the A copy and B copy should be isolated.

MQSeries page sets, BSDS and logs.

- BSDS and logs have the same requirements as DB2.
- Page sets for different MQ subsystems should be isolated.
- Adopt a strategy to isolate the data sets in each MQ subsystem in a manner that when a DASD subsystem failure occurs, not all subsystems are going to be impacted.

---

## General high-availability design recommendations

For high availability, it is recommended that two of the following items should be isolated:

- Master Catalog
- SYSRES volumes
- SYS1.PARMLIB
- SYS1.IPLPARM (optional)
- VTAMLST and NCPLIB
- IODF files

These data sets should be tied to an isolation strategy for OS/390 system data sets (page data sets, and so on) so that in the event that a single DASD subsystem fails, not all systems are going to be affected.

---

## Appendix B. 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.

References in this publication to IBM products, programs, or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM product, program, or service is not intended to state or imply that only IBM's product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any of IBM's intellectual property rights may be used instead of the IBM product, program, or service. Evaluation and verification of operation in conjunction with other products, except those expressly designated by IBM, are the user's responsibility.

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:

IBM Director of Licensing  
IBM Corporation  
North Castle Drive  
Armonk, NY 10504-1785  
U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

---

## Trademarks

The following are trademarks of International Business Machines Corporation in the United States, or other countries, or both:

IBM  
the IBM logo  
CICS  
DB2

IMS  
MQSeries  
MVS  
OS/390

NetView is a trademark of Tivoli Systems, Inc. in the United States or other countries or both.

Other company, product, and service names may be trademarks or service marks of others.

---

## List of abbreviations

**BSDS.** Bootstrap data set

**DASD.** Direct access storage device

**JCL.** Job control language

**JES.** Job entry subsystem

**REXX.** Restructured Extended Executor

**SMS.** Storage Management Subsystem

**SMF.** System Management Facility

**VOLSER.** Volume serial