



SMS Workbench (SWB)

Version 1.41

General Information Manual

What SMS Workbench Is

SMS Workbench, SWB for short, is a program to aid in the development of ACS routines for IBM's z/OS SMS services. It greatly improves on the primitive facilities of ISMF for the testing of ACS routines by providing:

- A database of SMS test cases (allocation records) containing original input (IGDACERO) and output (SMS Classes) of actual SMS allocations
- Simulation of ACS routine execution
- Tracing of simulated ACS routine execution
- Testing of ACS exits
- Comparison of original SMS class allocation results, reference results and test results
- Comprehensive filter and data manipulation facilities for creating/updating test records and recognition of changes (wanted or unwanted) in assigned SMS classes.

Introduction

System managed storage

IBM's introduction of System Managed Storage as part of MVS/ESA ushered in a new era of space management capabilities for IBM installations. The later unification of DFP, SMS and other space management products such as DF/DSS and DF/HSM under the DFSMS umbrella and the introduction of SMS support for media other than DASD was a recognition of the strategic role which such products play in large scale data processing operations. These products control the availability of data by managing the allocation of datasets on various media, providing normal and disaster backup, and optimising the use of permanent storage by providing for example migration of unreferenced datasets to various levels of lower cost storage, and allocating certain datasets to devices with special properties, such as mirrored DASD or high-speed cache.

ACS routines

Central to SMS is the process of allocating datasets. SMS datasets are those stored on SMS managed volumes. The properties of such datasets are defined by a set of SMS classes, which define not only the physical volume(s) on which a particular dataset resides, but other critical properties such as the backup attributes which will subsequently be used. These classes are assigned by executing a series of special routines, called Automatic Class Selection (ACS) routines each time a dataset is created or reallocated (e.g., as a consequence of recalling an HSM-migrated dataset).

ACS routines are written in a CLIST-like language, and compiled and activated using the facilities of the Interactive Storage Management Facility (ISMF) ISPF dialog. The compiled ACS routines and associated class definitions are stored in SMS control datasets.

Testing of ACS routines

Due to the critical nature of the ACS routines (errors in such routines can effectively prevent an z/OS system from executing) and the importance of assigning correct attributes to critical datasets, it is essential that these routines be adequately tested. Unfortunately, the test facilities provided by ISMF are very limited. Using ISMF, only individual dataset test cases can be created, and individually executed. There is no means of comparing results, or running a standard set of tests after changes have been made to ensure that only the

required changes have been made. In the case of errors, the ACS routine programmer is reduced to examining his ACS code to attempt to determine why the results are not as expected. For installations which use ACS exits (assembler routines invoked during the allocation process) there is no means of testing these routines, other than executing them live (for example on a special z/OS test system).

The SMS Workbench solution

SMS Workbench, SWB for short, is a program to aid in the development of ACS routines for IBM's SMS facility. It greatly improves on the primitive facilities of ISMF for the testing of ACS routines by providing:

- A database of SMS test cases (allocation records) containing original input (IGDACERO) and output (SMS Classes) of actual SMS allocations
- Simulation of ACS routine execution
- Tracing of simulated ACS routine execution
- Testing of ACS exits
- Comparison of original SMS class allocation results, reference results and test results
- Comprehensive filter and data manipulation facilities for creating/updating test records and recognition of changes (wanted or unwanted) in assigned SMS classes.

SWB Components

The SMS Workbench product consists of the following components.

SWB database

The SWB database is a VSAM linear dataset. It contains a series of records, called test case records, each of which corresponds to an 'allocation event', i.e., the allocation or reallocation of an individual dataset. It is important to notice the distinction between a dataset and its allocation. A single dataset may be reallocated many times during its lifetime, since it may be moved, or migrated and recalled, deleted and recovered. For each such allocation, an allocation event test case record could exist in the SWB database. An SWB test case record (subsequently referred to simply as an SWB record) contains the following information:

- The original allocation information available to the ACS routines (the contents of the IGDACERO control block passed on the SMS subsystem request for the dataset allocation). This is all the information available at the time of allocation, such as the dataset name, unit, volume(s), space and dataset type parameters and so on. The actual information depends on the manner in which the dataset was allocated, the attributes specified by the user, the system environment, and various system parameters, such as whether certain RACF options are active, and can change for subsequent allocations of an individual dataset.
- The original SMS class results of the allocation: data class, storage class, management class and storage group(s). In the case of a dataset allocated to a non-SMS volume, there will be no SMS class information present.
- Reference test results. These are results provided by running the SWB reference test, which simulates execution of the ACS routines in a particular SMS CDS (control dataset) for all the records in an SWB database.
- Bulk test results. These are generated by running the SWB bulk test, which simulates ACS routine execution for a subset of the SWB test records (which

can also be the entire set). These results are typically used to determine the effect of changes in the ACS routines. They may be compared with the original or reference results values for the records.

ACS Gatherer

The SWB ACS Gatherer program is used to create SWB database records. It is an z/OS started task which intercepts actual SMS allocation requests, and writes these requests as SMF records to the system SMF datasets. In order to provide accurate SWB records for subsequent testing, it is necessary to use the Gatherer program to generate the records.

SMF programs

There are two SMF programs. These programs read SMF data as extracted from the system SMF datasets by the z/OS IFASMFDP program. The records are then converted to SWB database format records, and added to an SWB database. Such a database may be a newly allocated VSAM linear dataset, in which case it will be formatted as an SWB database before the records are added, or may be an SWB database containing SWB records, in which case the new records will be added to the records already in the database.

There are two versions of the SMF program:

- Program to process ACS Gatherer records. This program processes the SMF user records created by the ACS Gatherer program.
- Standard SMF record program. This program attempts to create SWB records by examining various standard SMF records. Although useful for initial testing, these records do not correspond to actual information on allocation requests. In particular, no allocation result information (original ACS classes) is present in these records.

Two additional utility programs allow filtering of ACS Gatherer SMF records (e.g., removing temporary dataset records and suppressing multiple identical allocation records) before they are passed to the ACS Gatherer SMF program.

All the SWB SMF programs are executed as simple z/OS batch jobs.

SWB dialog

The SWB Dialog is an ISPF CUA dialog which provides all the database manipulation and test functions of the SMS Workbench product. The SMS programmer works with the SWB dialog whilst developing his ACS routines. The dialog uses the SWB database created by one or other of the SWB SMF programs for the input test data records. The ACS routines to be processed are precompiled routines contained in an SMS control dataset (SCDS). In order to create and modify the ACS routines to be tested, the ISPF editor and normal ISMF functions are used. The actual execution of the compiled routines is performed by the SWB ACS Simulator, which is a part of the SWB dialog. The SWB dialog is discussed in more detail in the next chapter.

The SWB Dialog

Introduction

The SWB dialog is an ISPF dialog in CUA format. All options for controlling the dialog are accessed via drop down menu options or commands defined by ISPF keylists. The user is presented with a single consistent test record list on which all test record manipulations are performed, and to which all test requests are applied. Such test requests execute reference or bulk tests for all records, or a subset defined by use of the SWB filter functions, or can be a traced test of an individual record.

The initial panel presented to the user on startup of the SWB dialog is shown in Figure 1.

```

Database  Test  Filters  Options  Help
-----
SMS-Workbench Main Display
Command ===> _____ SCROLL ==> HALF

      SSSSSS      WWWWW      WWWWW      BBBB BBBB
    SSSSSSSS     WW      WWW      WW      BBBB BBBB
      SS         WW  WW  WW  WW      BB      BBB
    SSSSSSS     WW  WW  WW  WW      BBBB BBBB
    SSSSSSS     WW WW  WW WW      BBBB BBBB
      SS         WW WW  WW WW      BB      BBB
    SSSSSSS     WWW      WWW      BBBB BBBB
    SSSSSS      WWW      WWW      BBBB BBBB
                    P R E S E N T E D   B Y
                I N T E R C H I P   A G
  
```

Figure 1. SWB initial panel

From the main panel, the database menu options can be accessed by placing the cursor on the menu options 'Database' (the first field on the screen), and pressing ENTER. The options list is then displayed, as shown in Figure 2. From this list, initially only the 'Open SWB database', 'Run SMF update program' and 'Terminate SWB application' options are available. The unavailable options are prefixed with a star to indicate that they cannot currently be selected.

```

Database  Test  Filters  Options  Help
-----
| 1 Open SWB database... |
| 2 *Display status of SWB database... | ROLL ==> HALF
| 3 *Write and close SWB database... |
| 4 *Close SWB database (no write)... |
| 5 *Update selected database records |
| 6 *Delete selected database records... |
| 7 *Copy/merge selected database records... | BBBB
| 8 *Insert database record... | BBBB
| 9 Run SMF update program for SWB database... | BBB
|10 Terminate SWB application... | BBBB
-----BBBB
      SS      WW WW  WW WW      BB      BBB
    SSSSSSSS     WWW      WWW      BBBB BBBB
    SSSSSS      WWW      WWW      BBBB BBBB
                    P R E S E N T E D   B Y
                I N T E R C H I P   A G
  
```

Figure 2. Database menu options

The 'Run SMF update program' option enables the user to generate a batch job to run the ACS Gatherer or standard SMF update program, or executes these programs directly online. Selecting the 'Open SWB database' option prompts the user to enter the name of an existing SWB database. The records in the database are then displayed on the main panel, as shown in Figure 3.

Using the test case list

The test case list can be scrolled vertically (to view different test case records) and horizontally (to view different fields in the records on the current display) by use of the standard scroll function keys. Due to the large number of fields in an SWB record, the horizontal scrolling functions in segments. A segment is a set

of related fields which fit on a normal (80 column) ISPF screen. Horizontal scrolling is by segment, rather than column. Each segment has a number, allowing the segment to be jumped to directly, by entering the number in the command field, and pressing one of the horizontal scroll function keys.

The list may be sorted on most of the fields by entering a SORT command in the command line, for example "SORT DSN" sorts the list by dataset name. The column headings are used as the sort field name. After sorting, a particular record may be located with the LOCATE command. For example, after sorting on dataset name, the command "LOCATE E" would locate the first dataset the name of which begins with the letter "E", or the immediately previous record if no such record exists.

Operations are applied to the list entries by selecting appropriate menu options (such as the test options) or by entering line commands in the select field of a record (the first field on each record line), or by selecting the function from the popup menu (if the popup line selection option is set). For example, selecting "U" in the popup menu allows the corresponding record to be modified.

```

Database Test Filters Options Help
-----
                                SMS-Workbench Main Display
Command ==>                                SCROLL ==> PAGE
Record number 1 of 228, segment 1: DATASET
  DSN                                UNIT    NV    VOLSER
_ E006480.ISPF.PROFILE                SYSALLDA  1
_ FTRANS.RXDTMUC.BUV54301.TST        SYSALLDA  1
_ G5BS040.SAS.CONFIG                  3390     1    BVT1E3
_ P505544.HCD.MSGLOG                  SYSALLDA  1
_ P505544.HCD.TERM                    SYSALLDA  1
_ P505544.HCD.TRACE                   SYSALLDA  1
_ P505544.ISP0001.BACKUP              SYSALLDA  1
_ P505544.SDSF.OUT                    3390     1    BVT000
_ P505544.SDSF.OUT                    3390     1    BVT000
_ P505544.SPFLOG5.LIST                SYSALLDA  1
_ P505544.SPFLOG5.LIST                3390     1    BVT001
_ P505544.SPFLOG5.LIST                3390     1    BVT005
_ P510023.SPFLOG1.LIST                SYSALLDA  1
_ P517263.SPFLOG1.LIST                3390     1    BVT005
_ P517866.SPFLOG1.LIST                SYSALLDA  1
  
```

Figure 3: Test case list display

Data viewing and manipulation functions

As well as viewing records in the test case and filter lists, a fullscreen browse option is available for displaying specific records. It is invoked by entering the "B" menu or select code in the popup menu or select field for a record. In fullscreen mode, a single record is displayed as multiple lines on the screen (or subscreen).

Data manipulation functions allow the test case list records to be modified. This can be particularly valuable when using records created by the standard SMF program, it may be necessary to modify some of the fields in the records. Updates may be accomplished by the update line command for individual records, or by the bulk update command. The line update command is invoked by entering "U" in the popup menu or the select field for a record. The attributes of the line are changed to allow the fields on the display to be directly overtyped with the required new values. In addition, a session options allows the updates to be performed in fullscreen mode, similar to the fullscreen browse function.

The bulk update function is invoked via the Database menu option 5. On invoking this function, the dialog splits the screen, creating an extra line with no

field values, into which new values can be entered. The main list is also displayed, and can be scrolled, but no update or test operations may be performed on it. One or other of the subscreens may be activated by placing the cursor on the appropriate subscreen, and pressing enter. When the required values have been entered, pressing the END function key causes the entered values to be placed in all records. (Note that the records to be updated may be limited by use of the filter function described below). The bulk update function can also be used in fullscreen mode, in the same fashion as the standard update function.

Individual records may be deleted by entering the "D" select code in the select field of an individual record to be deleted. A subset of records may be deleted by selecting the records with the filter function, and invoking the Database menu option 6. This deletes all currently selected records.

The "I" Insert record function may be used on a line to insert a new record before the current record, or the Database menu option 8 may be used to insert a record before that currently displayed at the top of the screen. The function creates an empty line in update mode. Data may be entered, and the END function key pressed to finish update of the record.

Filter function

The filter function enables a subset of the records to be selected. Only the selected records will be displayed in the test case records list. In addition, certain functions affect only the currently selected records.

The filter function is invoked from option 1 of the Filter menu, or by pressing the FILTER function key. On invoking this function, the dialog splits the screen, creating four double lines, into which filter values can be entered as shown in Figure 4. The main list is also displayed, and can be scrolled, but no update or test operations may be performed on it. One or other of the subscreens may be activated by placing the cursor on the appropriate subscreen, and pressing enter. Each filter line (occupying two screen lines) enables a field value to be entered (in many cases this may be a generic mask value) and an operator such as "EQ", "GT" etc.

```

Database  Test  Filters  Options  Help
-----
SMS-Workbench Main Display
Command ==> _____ SCROLL ==> PAGE
Filter set 1, segment 1: DATASET
  OP          | OP          | OP          | OP          |
  DSN         | UNIT        | NV          | VOLSER      |
- EQ         | EQ          | EQ          | NE          |
  ** .VSAM. ** | 3390        | _____ | TEST*       |
- EQ         | EQ          | EQ          | EQ          |
  _____ | _____ | _____ | _____ |
- EQ         | EQ          | EQ          | EQ          |
  _____ | _____ | _____ | _____ |
- EQ         | EQ          | EQ          | EQ          |
  _____ | _____ | _____ | _____ |
-----
  DSN         | UNIT        | NV          | VOLSER      |
  DB2X.ARCHLOG1.A0000030 | SYSDA      | 1          | _____ |
  DB2X.ARCHLOG1.B0000030 | SYSDA      | 1          | _____ |
  DB2X.ARCHLOG2.A0000030 | SYSDA      | 1          | _____ |
  DB2X.ARCHLOG2.B0000030 | SYSDA      | 1          | _____ |
  
```

Figure 4: Filter definition display

A special set of select codes applies to filter entries. In particular, select codes allow a particular filter definition to be activated or deactivated. Only active

filters will be applied to the database records. This enables particular filters to be set up, and activated only as required. On exiting the filter function with the END function key, the active filters are applied to the records.

In addition to testing for specific field values, filters may compare two fields within a record. This enables for example the various results sets (original, reference, test) to be compared with one another, and records containing differing values to be selected.

Figure 4 shows an example of a filter definition screen. In the example, the first filter has been defined to select datasets with the qualifier "VSAM" anywhere in the dataset name, to be allocated on 3390 devices where the specified volume name (if any) does not begin with "TEST". It is important to realise that the SWB record is for an *allocation request*, and the dataset may in fact never have been allocated. (The request may have failed or been rejected by SMS). The filter does not refer to *datasets* on particular volumes and units. The final allocation could in fact be quite different.

Filters are stored in the current SWB database. Filter definitions will only be saved if the database is closed with the update option (see below).

Utility functions

The SWB dialog enables records to be copied to another SWB database, and the current contents of the database to be saved.

Records to be copied may be selected by the filter function, as described above, or the whole database may be copied. This enables special databases to be created containing test records for some particular purpose (e.g., RMM records, or records for VSAM datasets only). Such special purpose databases may be useful for special tests, and improve the speed of processing as compared to using a single database with many thousands of records.

In addition to copying database records, the filter records in the current database may also be copied to another SWB database.

The currently selected database may be closed, in order to maybe select a new database to be used. On closing, the options 'close without write' and 'close with write' may be specified. If the current status of the database is to be saved, 'close with write' should be used. All database updates by the SWB dialog are performed in dataspace storage, and no update of the actual database linear dataset occurs until the 'close with write' option is used.

Test functions

The SWB dialog provides three test functions:

- Reference test. This test runs the current ACS routines for every record in the database, and updates the 'reference results' ACS class fields in the records.
- Bulk test. This test runs the current ACS routines for all selected (filtered) records in the database (or for all records if no filtering is active), and updates the 'test results' ACS class fields in the records.
- Trace test. This function is applied by select line command to a single database record. It generates a trace of the execution of selected ACS routines. No results are stored in the tested record.

In order to test the ACS routines, an SMS SCDS must be selected containing the routines to be tested. The CDS should contain translated and validated ACS routines. The CDS is selected by using the Test menu option 1. On selecting this option, the user is prompted for the name of the CDS dataset to be used. This CDS is then mapped to virtual storage and freed. The mapped contents

Trace data is written to a log dataset, which can subsequently be browsed or edited. An example of a trace execution of a simple DATACLAS ACS routine is shown in Figure 6. Figure 7 shows the corresponding original ACS source.

```

**** **** S W B   T E S T   T R A C E   O U T P U T **** ****
CDS NAME:          E005240.SMSCDS.VSAM
-----
--STARTING:      DATA CLASS ACS ROUTINE
COMP-EQ          (CONINDX 1) 2
                  (OPERAND 1) &DSN(2)      = SMS
                  (OPERAND 2) SMS

---TRUE---
SET              (R/W VAR)   &DATACLAS
                  (CONSTANT) MYDC
WRITE           (CONSTANT)  PGM=
                  (VARIABLE) &PGM          = TSOSESS
                  (CONSTANT) ,JOB=
                  (VARIABLE) &JOB          = E005240
                  (CONSTANT) ,DD=
                  (VARIABLE) &DD          = ISP13153
                  (CONSTANT) *NULL*
>>PGM=TSOSESS,JOB=E005240,DD=ISP13153
STOP
--FINISHED:     DATA CLASS ACS ROUTINE
--RETCODE:      00000000
--DATACLAS:     MYDC

```

Figure 6: Sample ACS routine trace

```

PROC 1 DATACLAS
IF (&DSN(2) = 'SMS')
  THEN DO
    SET &DATACLAS = 'MYDC'
    WRITE 'PGM='&PGM',JOB='&JOB',DD='&DD'
  END
END

```

Figure 7: Sample DATACLAS routine

Testing ACS exits

SMS Workbench allows the execution of ACS exit routines when performing its simulated ACS routine processing. Such exits are defined by the Test menu option 2, which provides a panel on which the names of the data class, storage class and management class exits may be entered, and the name of the load library containing the exit modules. To avoid conflict with any system modules, the user may call his exit routines by any name during SWB testing.

Such exits are loaded before each test execution, and are passed a simulated environment which corresponds (with one or two exceptions) to that provided in 'real life'. The ability to reinvoke the ACS routine is also provided. Messages set by an exit in the IGDACERW control block are written to a trace file, if performing a trace test, and the handling of class results set by the exit or the ACS routines corresponds to that performed by actual z/OS ACS processing.

Session options

The Options menu option 2 provides a panel on which a number of session options may be specified. These include default names for datasets such as the active SWB database and the SMS CDS to be used, date format for display and entry of date values, and options for customising the appearance and function of the main test case list display. These values are all saved in the ISPF profile, and are available again the next time SWB is invoked.

Help information

SWB provides comprehensive panel and field level help. All data entry panels have associated extended help, and all display and data entry fields have associated field level help panels. Help panels are available in addition for all menus, their suboptions, and for all error messages issued by the dialog.

Summary

To sum up, SMS Workbench offers the following features:

- The ability to create comprehensive databases of accurate test records for ACS routines
- A comfortable interactive dialog for running tests of ACS routines
- The ability to run complete sets of tests for old or modified ACS routines
- Comparison of original SMS class allocation results, reference results and test results
- Tracing of simulated ACS routine execution
- Testing of ACS exits
- Comprehensive filter and data manipulation facilities for creating/updating test records and recognition of changes (wanted or unwanted) in assigned SMS classes.

This completes the introductory review of SMS Workbench. For those who need to know more, the remaining sections of this document provide comprehensive information.