



## SUMMARY OF AMENDMENTS - VERSION 8.0

The following items are new or revised over version 7.0 of SMFUTIL.

### 1. ASYNCHRONOUS CLEARING OF MULTIPLE MANX DATASETS

By specifying MULTICLR=YES in the SMFUFLT module or the MULTICLEAR keyword at execution time, the DUMPCLEARALL function of SMFUTIL can now clear multiple MANx dataset simultaneously.

### 2. BUFFERING ABOVE THE LINE

All BSAM I/O buffers have been moved above the line for storage addressing constraint relief. This feature allows more output datasets to be created in a single pass of the input data.

### 3. SHORTCLEAR THRESHOLD SUPPORT

A “percentage full” threshold has been added to instruct SMFUTIL to automatically switch from a requested SHORTCLEAR to the more efficient (and faster) full clearing method when the MANx dataset is full or nearly full.

### 4. SAFE CLEAR OPTIONS INCORPORATED

During CLEAR processing SMFUTIL will now write a special indicator record as the first record of the dataset prior to completing the clear. This indicator record prevents SMFUTIL from placing the MANx dataset back into service as “AVAILABLE” in the event of a system outage and subsequent IPL during the CLEAR process. The dataset will remain “DUMP REQUIRED” and an SMFUTIL restart will be performed and the CLEAR completed. In addition, if SUCLEAR=YES is specified in the defaults module or SUCLEAR specified at execution time, the same special indicator record will be used to clear the entire dataset and only the first record will be overwritten with an IBM clear indicator record. This option prevents SMF from “MOD’ing on” to a MANx dataset during the IPL process. During IPL, only datasets that are actually CLEAR will look clear to SMF initialization. All others will look full. This prevents SMF from creating out of order data by reusing a previously partially used dataset before it can be dumped and cleared.

**NOTE: If SUCLEAR is specified (or defaulted) the MAN1LEAVE option may be removed (NOMAN1LEAVE at execution time or MAN1LEAV=NO in SMFUFLT defaults module) as SMF will no longer be able to reuse the primary MANx dataset until SMFUTIL has successfully cleared it and notified SMF of it’s availability.**

### 5. LARGE BLOCK INTERFACE (LBI) SUPPORT

Support has been added for IBM’s Large Block Interface (LBI). Though the record length is still limited to 32,767, any device that supports the LBI specification can have block sizes larger than 32K. Current limits are approximately 64K for 3480 and 3490 devices and 256K for 3590 devices.

### 6. GENCENT PARAMETER REMOVED

For Year 2000 compatibility, the GENCENT defaults parameter has been removed. All suffix values containing year values will now contain both the century and the year. Sites still using GENCENT=NO in the SMFUFLT module or specifying NOGENCENT at execution time must convert their data set name structures to include the century value.

### 7. SEPARATE SYNAD SUPPRESSION CONTROL FOR INPUT AND OUTPUT DATA SETS

SMFUTIL now supports suppression of the SYNAD exits for input or output data sets for debugging purposes.

### 8. TIME AND DATE/TIME LIMITATION FILTERS FOR INPUT DATA SETS.

SMFUTIL now supports limiting the input date/time or time values accepted from specific input DD names. The ITIMERANGE allows a user to set one or more specific time value ranges that are to be kept. The IDTRANGE allows a user to select one or more specific periods of time from an input source. Each input dataset may have one or more unique limits defined.

**9. IMPROVED TYPETODD RECORD SELECTION**

The TYPETODD keyword parameter has been enhanced to allow a list of record types to be specified for a single DD name.

**10. ENHANCED SUFFIX CAPABILITY**

Several new data set suffix definitions have been added including a YEAR suffix type and an EXTENDED TIME suffix.

**11. USER CALENDAR GENERATION UTILITY INCLUDED**

A new utility, SMFUGENC, designed to automate the generation of User Calendars, is now included in the base distribution.

**12. SMF ID SYMBOLIC INCLUSION/EXCLUSION SUPPORTED**

The SYSID and XSYSID keywords now support the &DID and &SID symbolic as the first system id specified. This allows a single system to be selected (the system providing input data or the system performing the execution) without knowing the system id ahead of time.

**13. DTRANGE SYNTAX ENHANCED**

The DTRANGE keyword now accepts new parameter syntax that makes it easier to specify the desired date/time range.

**14. PROCESS LOGGING CAPABILITY ADDED**

A new processing event recap log may be produced by the inclusion of a SYSLOG DD statement.

**15. DDSIDKEEP/DDSIDSKIP KEYWORDS ADDED**

Two new keyword, DDSIDKEEP and DDSIDSKIP, allows the user to specify specific SMF System ID's that are to acceptable or rejected for output to specific DD names.

**16. SKIPHIDT KEYWORD ADDED**

A new keyword, SKIPHIDT, allows the user to reject records with invalid future date values.

**17. FIXED FORMAT REPORT SUPPORTED**

A new keyword, REPORTFIXED, fixes all column widths in the SMFUTIL SYSPRINT report. This allows the user to develop a post processing program that can reliably read the report and summarize it.

**18. IMPROVED DATESCAN REPORTING**

The report produced by the DATESCAN function has been improved to show the record type and source DDNAME of the two records bounding any gap detected.

**19. ABILITY TO SYNCHRONIZE (SINGLE THREAD) JOBS**

A new execution parameter field, "SYNC" allows the user to single thread the execution of specific SMFUTIL functions between multiple jobs.

**20. ABILITY TO DYNAMICALLY ALLOCATE SYSPRINT OUTPUT DATASET**

If the SYSPRINT output DD statement is not available in the JCL, SMFUTIL will now dynamically allocate an output dataset for SYSPRINT. The user may direct this allocation to a spool dataset or to a sequential dataset on disk through various parameter overrides.

**21. ARCHIVE BACKUP CLEANUP CAN NOW BE AUTOMATED**

A new defaults value, ARCBNUMB, can be specified to indicate the maximum number of ARCHIVE backup datasets that SMFUTIL is to retain. SMFUTIL will automatically scratch and un-catalog older datasets above this count.

**22. FINAL RECORD CHECKPOINT DATA VALUE FOR EACH OUTPUT DDA BLOCK**

Each dynamically allocated output dataset may now have a unique checkpoint record to assist in detecting split boundaries that fall between executions. The keyword CKPTLAST can specify a unique checkpoint record name to be used for a dynamically allocated output dataset. This checkpoint record will be used in lieu of the LASTREC1 general purpose record.

**23. MQSPLIT KEYWORD IMPLEMENTED**

A new keyword, MQSPLIT, has been implemented to allow the user to split MQ Series record types 115 and 116 onto discrete output datasets based upon their subsystem name values.

**24. CHECKPOINT PDS BROWSER ADDED**

A new ISPF utility, SMFCKPUI, has been added to allow the user to examine the checkpoint records present in the Checkpoint PDS.

**25. IMPROVED ARCHIVE REPORT UTILITY**

The Archive Report Utility, SMFUARPT, has been enhanced to support a new keyword, USABLE, that causes the reports to be restricted to Archive Records that are usable to SMFUTIL for retrieving data. Defunct records will be ignored.

**26. LIST ALL RECORD TYPES REGARDLESS OF COUNTS**

A new keyword, LISTALL, causes SMFUTIL to list all record types from 0 to 255 on the "Detailed Statistics Report" even if the input count is zero for the record type.

**27. AUTOMATIC CREATION OF CHECKPOINT PDS DATASETS**

If it does not already exist and if the DACKPT parameter in SMFUDFLT is 'YES', SMFUTIL will now automatically create the CKPTDB dataset the first time it is specified. Parameter values in the SMFUDFLT module can be used to control the UNIT and VOLSER specification for this allocation.

**28. SMF ID CONSISTENCY CHECKING**

When a checkpoint dataset is present, SMFUTIL can now optionally check for an improper SMF ID in the input data. It does this by comparing the SMF ID of the input data with the ID retained in the \$PROCESS checkpoint record. If a mismatch is found, SMFUTIL will issue a warning message and set a minimum return code of 4 or optionally abort the execution.

**29. CATALOG INPUT DATA SET LIMITS EXPANDED**

The CATINPUT and GDGINPUT keywords have been enhanced to allow a specification to limit input to only a certain number of the latest datasets for an index. A 'latest' value of 1 to 255 may be specified for each data set prefix or GDG index level. When combined with the 'limit' value specification, this allows the user selective access to certain portions of a catalog list of datasets without processing the entire list.

**30. CKPTGDG AND GDGCKPTI PROCESSING IMPROVED FOR INDEX WRAP**

The CKPTGDG and GDGCKPTI keyword pair have been improved to automatically handle the situation where the GDG index wraps from 9999 to 0001. Data sets in this situation will be handled correctly and input in the correct order (i.e. order created). In addition, a BIDH sub-parameter keyword has been added for the user to be able to specify that, for ARCHIVEINPUT runs, datasets under the specified index are to have hourly BID values.

**31. FORCED CREATION OF EMPTY DATASETS NOW MAY BE SPECIFIED AT THE DDA LEVEL FOR EACH DATASET**

A keyword, DDACREATE, allow a DDA type dataset to be force allocated even if it has no output data. Alternately, a new keyword, DDANOCREATE, allows the suppression of a forced allocation for a specific dataset even though CREATEU0 or CREATES0 was specified.

**32. DEFER/ALLOCATE SPECIFICATION ELIMINATED ON THE DDAEND KEYWORD**

The positional keywords of DEFER or ALLOCATE have been removed from the syntax of the DDAEND keyword. Unused output datasets will never be allocated unless the CREATEU0, CREATES0 or DDACREATE keywords are specified as appropriate. The old syntax will still be honored but the specification will be ignored. For DDA data set names containing date sensitive suffixing or overlays, a new default parameter of CRECDATE will control what date is used for creating suffixes and overlays for empty datasets. A specification of 'YES' (the default as shipped) will cause the date from the first input record to be used. A specification of 'NO' will cause the executing systems current date to be used.

**33. DDACOPY0RC KEYWORD SPECIFIES RETURN CODE FOR EMPTY DATASET**

A new keyword, DDACOPY0RC, can now be used to indicate the return code to be issued when an output dataset receives no data. Previously the COPY0OK keyword could indicate a RC=0 for all output datasets in an entire run. Valid values are 0, 4, 8 or 12.

**34. ARCHIVE BACKUP DSNAME CAN NOW BE TIED TO THE ARCHIVE CLUSTER NAME**

If the ARCBDSN backup data set name prefix is specified as blanks, the ARCHIVE cluster name with qualifier '.BACKUP' appended to it will be used instead. It may be defined as a GDG index or it will be automatically suffixed with a DATE suffix for DAILY or specific DAY backups or with a DATE/TIME suffix if the backup frequency is ALWAYS. This has the added benefit of associating the ARCHIVE cluster with its backup datasets at the data set name prefix level. This is useful when multiple ARCHIVE data base structures are used.

**35. ARCFORCE COMMAND ALLOWS ARCHIVE UPDATES FROM INTERRUPTED RUNS**

A new keyword, ARCFORCE, allows SMFUTIL to force the update of the ARCHIVE database for ARCHIVE records that exist when a run is aborted for some reason.

**36. BID VALUES NOW BY DAY OR HOUR**

The Block ID (BID's) retained in the SMFUTIL ARCHIVE will automatically adjust to daily or hourly values depending on the output dataset split type. Daily datasets will have each BID slot represent the location of the associated hour on the tape. Weekly or Monthly datasets will have each BID slot represent the location of each day on the tape. For input data archiving or output datasets not being split, a BIDBYHOUR keyword is available within the DDA block for force hourly BID values.

**37. ALTUNITS KEYWORD ALLOWS UNIT NAME SUBSTITUTIONS**

The new ALTUNITS keyword allows a unit name to be substituted for a UCB device type to prevent emulated devices from being allocated. These devices typically cannot be used to mount real device media. An example of this is a 4mm SCSI DAT drive being used to emulate a 3490E tape drive. Even though the UCBTYPE device type strings are identical, SMFUTIL ARCHIVE SMF data on actual 3490E tapes cannot be mounted on the 4mm DAT drive. To prevent allocation from trying to use an emulated device as a real one a unit name for UCB device type substitution is required.



e3 Sciences Limited  
Kingston House, High Street, Kings Stanley, Glos, GL10 3JF  
Tel: +44 (0)207 060 6601 Fax: +44 (0)207 060 6602  
Email: [info@e3sciences.com](mailto:info@e3sciences.com) Web Page: [www.e3sciences.com](http://www.e3sciences.com)