



SDLSTATS

Reference



Database Search and Generation of Statistics

1. <i>Introduction</i>	2
2. <i>Parameter D00 – Date range specification</i>	3
3. <i>Parameter D01 – Time range specification</i>	4
4. <i>Parameter D02 – Datasets to be searched</i>	5
5. <i>Parameter D03 – Dataset for the selected log records</i>	6
6. <i>Parameter D04 – Definition of statistics dataset</i>	7
7. <i>Parameter D05 – Definition of Load Library</i>	8
8. <i>Parameter DEF – Definition of Filter</i>	9
9. <i>Parameter Fnn – Specification of Filter value</i>	10
10. <i>Parameter Snn – Definition of Statistic</i>	11

SDLSTATS

Reference

1. Introduction

SDLSTATS is the selection and analysis tool for the log database which is generated by SDLTRACE. The term database here refers to all datasets belonging to a particular application (as defined by the common levels of the dataset name), with exactly one dataset containing all log events for one day. Thus the complete log data of a year for a particular application is contained in up to 365 (or 366 for leap years) datasets. The last level of the dataset name specifies the day in the format: Dyyymmdd.

SDLSTATS has two distinct functions:

1. Selection of all log records satisfying specified search criteria,
2. Generation of statistics tables according to defined specifications.

The first function is used to select all those log records which match the strings defined through the **DEF** and the **Fnn** parameters. The second function, specified through **Snn** parameters, is used to compute statistics for the selected log records. The databases to be analyzed and the time frames are defined by the **Dnn** parameters.

The operation of SDLSTATS is controlled by definitions specified in member SDLSPARM (or any other name specified in the EXEC statement that calls SDLSTATS). The functions to be defined are labeled with identifiers named **Dnn**, **DEF**, **Fnn** and **Snn** located in position 1 through 3 in each parameter line, followed by additional information starting in position 5. A blank, an asterisk '*' or a slash '/' in position 1 indicates a comment, so that parameters which are temporarily not needed may easily be "commented out" by replacing D01 with *01 or /01 for example. A sample member may look like this:

```
/* -----*/
/* Parameter file for SDLSTATS - Rel 1.2                15.07.2016 */
/* -----*/
D00 D160630 D991231 /* Date range for selection */
  01 10:30 14:00 /* Time range for selection */
D02 user-id.SDLALL /* Input DSN Hilevel qualifier */
D03 user-id.SDLALL.SELDATA /* Target-DSN for selected recs*/
D04 user-id.SDLALL.STATS,20 /* Target-DSN for Statistics */
D05 user-id.SDLTRACE.DEMO.LOAD /* LoadLibraryDSN for SDLSLECT */
/*----- Definition of Filters -----*/
*EF Fnn Stats-Name (Part) (Column) (Length)
DEF F01 LogRecords 1 1 4
DEF F02 Time run 1 119 5
DEF F03 Task time 1 76 5
DEF F04 Progr Name 1 81 8
/*----- Specification of Filter Values -----*/
F01 Data
F01 More
F01 Mult
F01 Log
/* All log records that contain "Data", "More", "Mult" or "Log " */
/* in position 1 to 4 will be selected. */
/*----- Specification of Statistics -----*/
S01
S12 F01 F02
S21 F02 F01
S14 F01 F04 SortName Time
/*-----*/
```

In the following chapters the parameters will be explained in detail.

SDLSTATS

Reference

2. Parameter D00 – Date range specification

Parameter D00 is used to specify the start date and the end date of the range which is to be used for the search for matching log records. The format for D00 is as follows:

D00 Start-Date End-Date

and the format may be one of the following:

Dyymmdd yyyy-mm-dd dd.mm.yy

Examples for D00 are:

D00 D160630 D161231
D00 2016-06-30 2016-12-31
D00 30.06.16 31.12.16

A range of one single day may be specified by omitting the second date.

If Parameter D00 is not specified the default value will be the current day.

SDLSTATS Reference

3. Parameter D01 – Time range specification

Parameter D01 defines a time frame which is to limit the search. Out of each day specified in parameter D00 only those log records that occurred during the time defined in D01 will be selected for further comparison against the search parameters.

Examples are:

```
D01 10:30 13:59          /* Time frame is 10:30am to 2pm */
D01 22:00 23:59        /* Time frame is 10pm to midnight*/
```

If D01 is not specified then the time will not be checked.

SDLSTATS Reference

4. Parameter D02 – Datasets to be searched

Parameter D02 is used to specify the high levels of the log datasets which are to be searched. Together with parameter D00 (the date range), yielding the last level in the format Dyymmdd, D02 determines the extent of the database to be considered.

D02 may be specified more than once so that multiple databases can be grouped together in one run.

Examples are:

```
D02 ABC.LOG.APPL01  
D02 XYZ.LOG2.ANOTHER.APPL
```

At least one Parameter D02 is required, there is no default value for D02.

SDLSTATS Reference

5. Parameter D03 – Dataset for the selected log records

The dataset which is to contain the selected log records is specified in parameter D03. If the dataset does not exist it will be allocated with an initial (and secondary) default allocation of 100 cylinders. Alternatively, the allocation (in cylinders) may be specified immediately following the dataset name, separated by a comma.

Examples are:

```
D03 ABC.SELECTED.LOG.RECORDS.APPL1
```

or with specification of an initial allocation of 20 cylinders:

```
D03 XYZ.SELECTED.LOG2.RECORDS.APPL20,20
```

The selected log records are written to the output file in their original format of 133 bytes each. Therefore the generated dataset can again be used as input to yet another SDLSTATS run. The only requirement is that the last level of the dataset name is set to the format Dymmdd with a date corresponding to the value of the D00 parameter being used. The output dataset should therefore be renamed accordingly if this feature is to be used.

Please note that original datasets only contain log records whose timestamps correspond to the date of the dataset itself. This rule does not hold anymore of course if records from several days are selected into an output dataset by an SDLSTATS run.

To determine the number of log records that can be stored in a dataset with an initial allocation of 100 cylinders the following information may be of interest:

SDLSTATS uses the multivolume feature of the operating system with MAXVOL(20). Therefore the default of 100 cylinders will allow a maximum of 32000 cylinders (100 times 16 extents times 20). On a standard disk device (3390 type) there are 60 blocks of 13300 bytes each with 100 log records per block, so that 6000 records can be stored in one cylinder. Thus with 100 cylinders initial allocation the total number of log records that may be written is 192 million.

If D03 is not specified the selected log records will not be written to any output dataset. The reason for not defining D03 may be that only a statistic is required and that the individual log records are of no interest.

6. Parameter D04 – Definition of statistics dataset

The dataset which is to contain the statistics tables is specified in parameter D04. If the dataset does not exist it will be allocated with an initial (and secondary) default allocation of 10 cylinders. Alternatively, the allocation (in cylinders) may be specified immediately following the dataset name, separated by a comma.

Examples are:

```
D04 user-id.SDLALL.STATS  
D04 user-id.SDLALL.STATS,15
```

Only one parameter D04 may be specified. If it is omitted no statistics information will be generated.

SDLSTATS Reference

7. Parameter D05 – Definition of Load Library

SDLSTATS calls the Assembler program SDLSLECT to perform the requested search operations and to generate the specified statistics. It has been written to perform even complex searches of large amounts of data in an acceptable time frame.

Example:

```
D05 user-id.SDLTRACE.DEMO.LOAD
```

The module SDLSLECT is stored during installation of SDLTRACE. It is needed for every run of SDLSTATS and D05 must therefore always be specified.

If a search of a LOG database or a statistics generation is to be performed on a system other than the one where SDLTRACE was initially installed then SDLSTATS and SDLSLECT must be copied into appropriate runtime libraries. SDLSTATS is a compiled REXX procedure and should be located in a JCL or CNTL file; SDLSLECT is a load-module and must therefore be copied to a load-library. The sample procedure SDLSTRUN can be used as a template to set up the JCL for production jobs.

SDLSTATS

Reference

8. Parameter DEF – Definition of Filter

The Parameter DEF is used to define fields in the log record which are to be searched or for which a statistic is to be generated. The definition of a field is separate from the assignment of a value for it; values are assigned in the parameter Fnn. The definition just gives a name to the filter, specifies its description, determines its location in the log record and defines its length.

The format is as follows:

```
DEF Fnn Description Part Column Length
```

Fnn defines the name of the filter, where nn may have any value between 01 and 99.

Description starts in position 9 and is exactly 10 bytes long. It may contain any printable character including blanks. Its main use is to provide title and heading information for the statistics reports.

Part starts in position 20 and determines the part of the log record where the field is located. Since most log datasets contain only single-part records this value will in most cases be 1. Only for multi-part logs, where the length of a record is larger than the standard 80 bytes of actual log data, the value for part will be 2, 3, 4 or any other number denoting the part containing the desired field. If during processing a log record with fewer parts than defined in the filter is encountered it will simply be skipped.

Column defines the starting position of the field in the log record. Its value may be between 1 and 133.

Length defines the size of the field. Its value may be between 1 and 50.

The combined length of starting position plus size may not be larger than 134.

Some sample definitions are:

```
DEF F01 Log Types 1 1 4
DEF F02 Time run 1 119 5
DEF F03 Task time 1 76 5
DEF F04 Progr Name 1 81 8
```

Filter definitions may overlap. A definition as follows is perfectly legal:

```
DEF F23 Progrstart 1 81 4
DEF F24 Progr Name 1 81 8
```

In this example the complete program name is defined in filter F24. The first four characters of the program name are defined in filter F23. With a specification of value "CICS" for F23, all log records where the program name starts with "CICS" will be selected. To generate a statistic for these names, i.e. how often did each name occur, just specify S24.

SDLSTATS Reference

9. Parameter Fnn – Specification of Filter value

Parameter Fnn is used to assign values to the filters that were defined with parameter DEF. The format of an assignment is as follows:

```
Fnn Character string
```

or:

```
Fnn Character string Character string
```

Fnn must be one of the filters defined in parameter DEF. Therefore the definition of Fnn in a DEF parameter must precede the use of Fnn for specification of a filter value.

The length of Character string must be equal to that specified in the definition of Fnn. This is to allow blanks to appear within the string so that searches may be performed for any combination of characters.

There must be exactly one blank between Fnn and Character string.

If two strings are specified they also must be separated by exactly one blank. The two values define a range for the search with lower (first value) and upper bound (second value). All records containing a string that is higher or equal to lower bound and lower or equal to upper bound will be selected.

SDLSTATS

Reference

10. Parameter Snn – Definition of Statistic

Parameter Snn is used to generate Statistics from all the records that were selected by Fnn.

The basic form for a statistic is:

```
Snn
```

This specifies that a statistic is to be generated for Filter Fnn. It is the short form for the more general specification:

```
Sxx Fnn
```

where xx is any number from 01 to 99 and Fnn is one of the filters defined in the DEF parameter. (In the short form the statistic number corresponds to the filter number).

A statistic may also be generated for a combination of filters, specified as follows:

```
Sxx Fnn Fmm
```

In this case the two filters Fnn and Fmm are used together as one item, so that the occurrence of one string in the log record together with another string may easily be determined.

The statistic specifications must appear in the input stream after the specifications of the filters Fnn.

Any of the statistic specifications above may be followed by two optional parameters: the sort sequence and the timing information request. The desired sort sequence is specified as follows:

```
SortName or SortCount (abbreviated SN or SC).
```

SN specifies sorting by name in ascending order,
SC specifies sorting by count in descending order.

The second additional parameter is Time (abbreviated T) requesting the inclusion of CPU and elapsed times in the statistics (if the relevant fields are present in the log data).

An example of a complete statistic specification would be:

```
S12 F01 F02 SN T
```

The defaults for the sort sequence are: SC for a single filter statistic, and SN for a two filter statistic. Timing information is provided only if Time (or T) is specified.