Unicenter® CA-APAS® Insight Monitor for Adabas

Writing Requests

4.1

Computer Associates®

Contents

Chapter 1: Introduction to Writing Unicenter CA-APAS Requests

Chapter 2: General Syntax Information

Chapter 3: Statement Descriptions

Chapter 4: Unicenter CA-APAS Data Fields

Appendix A: Output Record Header Formats

Appendix B: RSPCLASS Definitions

Index

Chapter 1

Introduction to Writing Unicenter CA-APAS Requests

Requests for data about Adabas command processing provide much of the power and flexibility of Unicenter CA-APAS Insight Monitor for Adabas (Unicenter CA-APAS). Unicenter CA-APAS requests are referred to in this and other documents when describing the capabilities and functioning of some Unicenter CA-APAS components. Understanding Unicenter CA-APAS requests is a basic prerequisite for understanding how to use Unicenter CA-APAS.

This section describes, in general terms, the nature and use of these requests, including examples.

The full syntax of request statements, along with descriptions of all of the data fields and summary functions that may be referenced in the statements are described in this manual.

Additional suggestions for effectively using requests are contained in the *Unicenter CA-APAS User Guide*.

About Unicenter CA-APAS Requests

Unicenter CA-APAS requests define the information to be produced by the Data Collector. Use these requests to specify:

- Which Adabas commands are reflected in outputs
- How user-defined data fields are developed
- Which data fields and summary functions are included in outputs
- How outputs are formatted
- Type of processing (detail, summary or COPY)
- Type and destination of outputs

The ways in which Unicenter CA-APAS requests accomplish these purposes are described in the remainder of this section.

Usage

Requests consist of control statements written according to the conventions stated in the chapter "General Syntax Information." Statements are free format and easy to write or modify.

Default working requests that address performance areas of general interest are provided on the distribution tape, ready for use as soon as the tape is loaded to disk.

For batch submission, request statements may be included in standard JCL input streams, or they may reside in files. For online submission through Insight, they may be stored in Natural source libraries and may be created or modified with the Natural source editor.

Unicenter CA-APAS itself imposes no limits on the number of concurrently active requests, the number of output files, line lengths for reports, etc. The only constraint is the amount of virtual storage available within the execution region. The inherent flexibility of requests and their dynamic control through Insight permit almost any kind of monitoring needed, even in situations where virtual storage is limited.

Security

Requests may be "public" or may be owned by specified individuals or groups that are defined in Natural. Non-public requests can be protected from display and/or modification by anyone except owners of the requests.

Unicenter CA-APAS Request Statements

On any given execution of the Data Collector, the following types of request statements may be used:

- A GLOBALS statement sets global parameters and options that influence other statements
- A DECLARE statement defines the user-defined data fields, if any
- An EXTRACT statement produces output that reflects selected data items from individual commands
- A SUMMARIZE statement produces output that reflects selected data items summarized across multiple commands
- A COPY statement specifies copying the entire selected Adabas Command Log records (supplemented with additional data Unicenter CA-APAS derives) to an output file
- LOG and UNLOG statements are special control request statements to dynamically control the logging of Adabas Command Log records by Adabas to the Adabas Command Log file, DDLOG/DDCLOGRn.

Statement Relationships

The figure below indicates the relationships between these statement types, input command records, and the different types of outputs.

Unicenter CA-APAS: Statements, Input Records, and Output

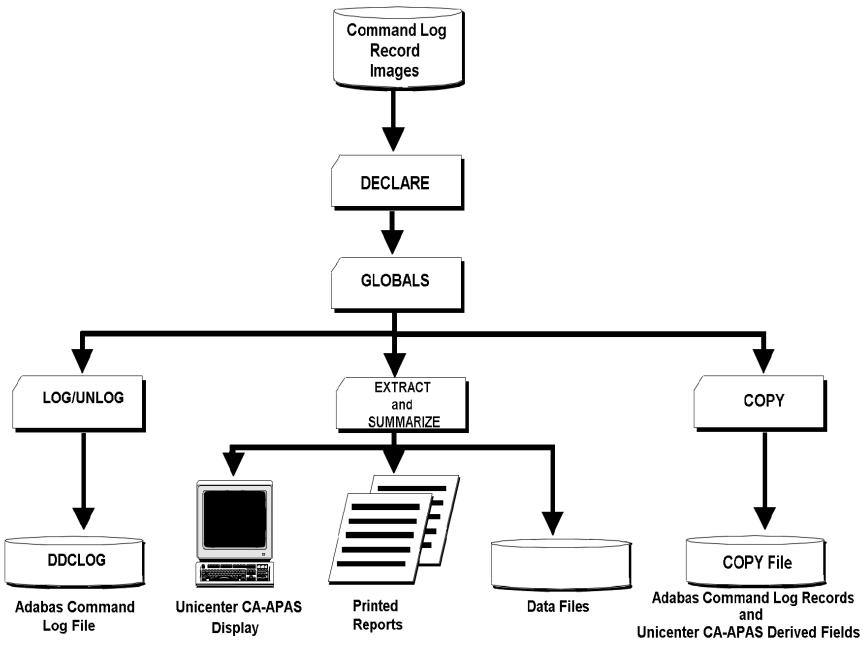

GLOBALS

A GLOBALS statement sets various options for execution. Some of the more significant options include:

- Logical criteria for excluding commands from further processing
- Factors for computing estimated command CPU time and cost
- Ranges of RABNs for developing I/O counts to specific selected database areas
- VSE system file numbers

A sample GLOBALS statement is shown below.

```
GLOBALS WHERE  FILE GE 100 AND FILE LE 199 /* limit to these files
        MAXRECS=10000                      /* limit number cmds read
        CPU-ID=SYSA DBID=PRODA             /* to identify source
        CIPS=5000000                       /* 5 MIPS CPU speed
        USER-EXIT=OLDEXIT                  /* site-specified user exit
```

DECLARE

The DECLARE statement allows redefinition of existing data fields or definition of entirely new data fields. Use of this statement type is optional. Each DECLARE assigns the name and derivation rules for one data field. User-defined fields must be defined before any other statements that reference the fields.

Sample DECLARE statements are shown below.

```
DECLARE  JOBCODE (A4 HD='JOB','CODE')
         REDEFINES JOBNAME (4,3);     /* chars 3-6 of batch jobname
*
DECLARE  DEPT (A3 HD='DEPT','NO.') VALOF
         IF      JOBCODE               <= 'C104'   RESULTIS '201'
         ELSEIF  JOBCODE > 'C104' AND   <= 'H210'   RESULTIS '460'
         ELSEIF  JOBCODE > 'H210' AND   <= 'Q360'   RESULTIS '600'
         ELSEIF  JOBCODE = 'CICS'
                 IF     TERMINAL-ID(1) = 'K'        RESULTIS '801'
                 ELSEIF TERMINAL-ID(1) = 'L'        RESULTIS '802'
                 ELSE                               RESULTIS '899'
                 IFEND
         ELSE                                       RESULTIS '999'
         IFEND;
```

EXTRACT

An EXTRACT statement specifies the output of data values from individual Adabas commands. This capability is useful for tracing categories of commands, such as those from a particular job or terminal, those to a particular file, or those within a time period. It is also a very powerful way to generate exception reports, such as commands receiving specified response codes, commands exceeding normal performance limits, or commands to sensitive files. The major elements of EXTRACT statements are:

- The data fields to be reported, optionally overriding formats or headings
- The criteria that commands must meet in order to be included
- The type of output(s) wanted (online display, printed report, and/or data file)

A sample EXTRACT statement and the report format it produces follows.

EXTRACT Request

```
EXPCMD:  EXTRACT
           DATN TIME
           NAT-LOGON
           NAT-PROGRAM
           CMD
           FILE
           DESCR-UPDATED
           TOTAL-IO DURATION
           WHERE DURATION > 1 SECS OR TOTAL-IO > 30
           PRINT TO EXPLOG
           TITLE = ('Expensive Commands',
                    'Duration Over 1 Second ' -
                    'or Total IO Over 30');
```

EXTRACT Report

```
2002-04-10 17:58:02                                                    PAGE   1
EXTRACT: EXPCMD   DATABASE NAME: TECHDB03   CPU ID: TSO44    DATABASE ID:   3
FIRST RECORD   2002-04-09 09:46:31
                          EXPENSIVE COMMANDS
                 DURATION OVER 1 SECOND OR TOTAL IO OVER 30

                                    C
                 NATURAL  NATURAL   M    FIL DES   TOT     DURAT
  DATE    TIME    LOGON   PROGRAM   D    NUM UPD    IO      SECS
-------- ------ -------- -------- -- ------ --- ----- ---------
20020409 094631  ISADMIN  ISG0010 A1     55   2    53     2.368
20020409 095654     SIRS  SIR0235 N1     29   7    53     2.552
20020409 100722                    S1     30   0   111     3.990
20020409 101507  CASHTRK  CTR0119 S1     58   0    74     3.125
20020409 102319     SIRS  SIR0235 N1     29   7    56     1.563
20020409 102412     SIRS  SIR0235 N1     29   7    54     1.374
20020409 102708    CECAS  CCS0085 N1     71  11    60     2.251
........
........
```

SUMMARIZE

A SUMMARIZE statement specifies the output of data values that have been developed from multiple Adabas commands. This capability is useful for a wide variety of measurements of Adabas workloads and performance characteristics. Major elements of SUMMARIZE statements are:

- Data fields for which summary functions, including sum, minimum, mean, maximum, percent, and rate, are to be included in outputs
- Control fields for summarization
- Criteria which commands must satisfy in order to be included in the summarization
- Interval periods for summarization
- Types and destinations of outputs

Summary reports may be requested for user-specified time intervals. Intervals begin and end at standard clock-time boundaries to allow them to be easily compared with other system performance data. The sample SUMMARIZE request and resulting report shown below illustrate the basics of this type of request.

SUMMARIZE Request

```
SUMFILE:  SUMMARIZE  COUNT SORT DESC PERCENT
                     SUM(TOTAL-IO) PCT(TOTAL-IO)
                     MEAN(TOTAL-IO) SUM(DUR) MEAN(DUR)
                     BY FILE
          INTERVAL 4 HOURS
          PRINT TO SUMLOG
          TITLE = 'Summary of Performance by File';
```

SUMMARIZE Report

```
2002-04-10 17:58:02                                                    PAGE    2
SUMMARY: SUMFILE  DATABASE NAME: TECHDB03          CPU ID: TSO44   DATABASE  ID:    3
FIRST RECORD   2002-04-09 20:02:48                   LAST RECORD 2002-04-09 22:29:53

                          SUMMARY OF PERFORMANCE BY FILE

                  PCT             SUM   PCT      MEAN                SUM        MEAN
   FIL        CMD TOTAL             TOT   TOT       TOT              DURAT       DURAT
   NUM      COUNT  CMDS              IO    IO        IO              SECS        SECS
------ ---------- ----- -------------- ----- -------- ------------------- -----------
    30    372,291  29.7        207,002  25.3      0.6              7700.7      .02071
    31    354,054  28.3        380,317  46.5      1.1              9210.7      .02604
     0     23,354   1.9         24,719   3.1      1.1              1349.0      .05772
     7     20,200   1.7         18,765   2.3      0.9               822.4      .04079
     5      9,618   0.8            144   0.1      0.0                 8.5      .00094
    29      4,762   0.4          3,810   0.3      0.8               150.0      .03153
    27      1,821   0.2          1,985   0.3      1.1               103.6      .05698
     8      1,370   0.1          2,364   0.3      1.7                94.6      .06912
    28        372   0.1            919   0.1      2.5                31.5      .08490
    16        363   0.1            194   0.1      0.5                 4.0      .01137
```

Note that in the report the "SORT DESC" option following the data field COUNT causes the files to be reported in descending sequence of their individual total command counts rather than the default sequence, ascending file number sequence. As the number of entities (files, jobs, terminals, etc.) reported increases, SORT DESC becomes more important. It eliminates the need to visually scan long reports to find the entities with the greatest performance impact.

This feature is particularly helpful when viewing displays at terminals. By positioning the critical information at the top of the report, it eliminates the need to scroll beyond the first page of a display.

COPY

A COPY statement specifies that selected Adabas Command Log records, supplemented with additional data derived by Unicenter CA-APAS, are to be written to a sequential file for later processing.

Output records include all standard Adabas Command Log fields, many of the additional fields derived by Unicenter CA-APAS, and any specified combination of Adabas buffers.

The WHERE clause allows users to select which records are written to a COPY file. This permits the detailed analysis of particular activities without incurring the significant overhead of writing complete Adabas Command Log files, and without losing the valuable data Unicenter CA-APAS derives. This greatly simplifies the tuning of individual transactions or applications without affecting other Adabas users or losing functionality.

Below is an example of a COPY statement.

```
COPYTEST:  COPY SEARCH-BUFFER VALUE-BUFFER UEXB
           WHERE NAT-LOGON = 'SYSTEST'
           OUTPUT-FILE = (DUAL,DDCOPY1);    /* dual X/Y files
```

LOG

A LOG statement specifies that selected Adabas Command Log records and any specified combination of Adabas buffers are to be written by Adabas to the Adabas Command Log file, DDLOG/DDCLOGRn, for later processing.

The WHERE clause allows users to select which records are written to the Adabas Command Log file. This permits the detailed analysis of particular activities without incurring the significant overhead of writing complete Adabas Command Log files for every Adabas command.

Below is an example of a LOG statement.

```
LOGF412:  LOG FB SB VB
          WHERE FILE = 412;
```

UNLOG

An UNLOG statement specifies which LOG request is to be deleted from Data Collector processing.

Below is an example of a UNLOG statement.

```
LOGF412:  UNLOG;
```

Request Outputs

COPY requests produce sequential files of Adabas Command log records followed by the Unicenter CA-APAS derived data fields.

An EXTRACT or SUMMARIZE request may specify any combination of the following kinds of output:

- Print report
- Terminal display
- Sequential data file
- Automated dual files

Each of these four kinds of output is described in more detail below.

Print Reports

Reports formatted for printing have individual data fields and summary function values arranged in columns. Multiple data items may be "stacked" in a given column to fit more data into the page width.

Values of Adabas buffers specified in EXTRACT requests are presented in character string format below the columnar data for each record.

All data fields and summary functions have default print formats and column headings. These may be overridden in any request. Scaling factors and number of decimal places may be specified for numeric values.

Reports may be directed to system print spools or written to sequential files. Multiple reports may be written to a given output file. When EXTRACT requests or SUMMARIZE requests with intervals share an output file, print lines may be written with headers that enable interleaved reports to be sorted into proper sequence before printing.

Terminal Displays

When EXTRACT or SUMMARIZE requests specify output for display via the Natural user interface - Insight, the data to be displayed is formatted like it is for printed reports.

- SUMMARIZE requests always send the specified number of lines from the top of the report as of the time requested.
- EXTRACT requests use the buffer in a "ring" fashion so that the most recent lines are always displayed in chronological sequence.

Sequential Data File

Machine-readable output files are designed for later processing by Unicenter CA-APAS or other reporting, statistical, billing, or history programs. Data fields and summary functions from EXTRACT and SUMMARIZE requests have default output formats and lengths; these may be overridden in any request.

Output records have standard headers that identify the request, database, CPU, date and time, etc. Output files from SUMMARIZE requests with intervals are "check-pointed" at the end of each interval to ensure that no completed interval of data is lost due to system failure. This is particularly helpful in avoiding loss of charge-back information.

Automated Dual (X/Y) File Handling

Any COPY file, print report or sequential data file may be directed to a pair of physical files known as dual or X/Y files. Dual files offer the following advantages:

- Capture of data from extended Adabas sessions without dedicating tape drives or excessively large disk files
- Availability of data for viewing or processing each time a switch occurs between the files of a dual set

Unicenter CA-APAS writes to one file of a dual set until either that file is full or a Unicenter CA-APAS user executes a FLIPXY command to force a switch to the other file. When the switch from one file to the other occurs, an automatic process copies the previously active file to tape or disk and then resets it to an empty status. The reset file is then ready to receive data when the next switch occurs.

This automated process normally requires no reply, thus minimizing the involvement of operations personnel in dual file processing.

Chapter

2

General Syntax Information

This section describes the rules of statement syntax and use that apply to the various types of Unicenter CA-APAS Insight Monitor for Adabas (Unicenter CA-APAS) requests.

Statement Records

All Unicenter CA-APAS control and request statements for batch input streams are in the form of 80-byte card image records. Columns 1 through 80 of the record are considered meaningful input.

Statements may be stored in Natural source program libraries for dynamic submission to Unicenter CA-APAS. While in the Natural source library, these statements have conventional Natural statement numbers, and their statement lengths are as specified by the Natural "SL" parameter.

The Statement Batch Input Stream

The statement batch input stream for a Unicenter CA-APAS component consists of job control statements and optional comments. Each batch input stream to the Data Collector follows this sequence:

- Any number of comment records placed in any position
- A GLOBALS statement describing the global operating environment
- Optionally, one or more DECLARE statements that define fields to be developed during processing
- One or more EXTRACT, SUMMARIZE, and/or COPY statements defining the desired reports and/or output files

When the Data Collector is executed as part of an MPM session, the above input stream is preceded by control statements for the Unicenter CA-APAS User Exit 4 routine. The User Exit 4 statements are explained in *Unicenter CA-APAS Systems Guide.*

Comments

Comments may be inserted anywhere in the input stream, and may contain arbitrary text. Comments are identified in one of two ways:

An asterisk, *, in the first column of an input record identifies the entire record as a comment.

The characters /* or // identify the rest of any statement record as a comment and cause the statement interpreter to skip to the next record.

Note: The characters /* and // may not be used in columns 1 and 2 to indicate a comment. In columns 1 and 2, these characters are interpreted as JCL and mark the end of the input stream.

The following are examples of comments.

```
* THIS ENTIRE RECORD IS A COMMENT
IF DATN >= '20020101'     /* HAPPY NEW YEAR    these are
AND     <= '20020131';    /* END OF THE MONTH  trailing comments
```

Statements

Statements begin with a label, which is optional for some statement types, followed by a statement identifier keyword. Statements are terminated by any of the following:

- A semi-colon, ;
- The next statement identifier keyword
- The end of the input file

Data Collector statements are further described in the section Statement Descriptions of this document. Unicenter CA-APAS User Exit 4 routine statements are described in *Unicenter CA-APAS Systems Guide*.

End of Input

The end of a set of statement input records is indicated by either of the following:

- An end-of-file on the statement input file.
- A statement input separator record containing the characters ## in the first two columns; this enables input statements for two Unicenter CA-APAS components to be provided in the same input file.

Statement Elements

Statement elements are the components of statements and are separated by spaces. Place commas between statement elements only where specified.

Labels

Labels are strings of one to eight characters without surrounding quotes. They are distinguished from other character strings by an immediately suffixed colon.

```
EXTLONG:  /* is a label
EXTLONG : /* is not a label, and will result in a syntax error
```

Keywords

Keywords, or reserved words, are a predefined set of words, each 32 characters or less in length. Keywords are reserved; that is, they may not be entered as labels, values or field names you supply. Never surround these words with quotation marks or apostrophes. Keywords are listed below.

ALL	AND	ANYREC	ASC
ASCENDING	ASSO	BOTH	BPL
BY	CIPS	CMDLOG	COLSPC
COPY	CPU-COST	CPU-FACTOR	CPU-ID
CPU-INSTRS-SEC	DATA	DAY	DBID
DBNAME	DDCLOG	DE	DECLARE
DEMAND	DESC	DESCENDING	DEVICE
DISPLAY	DISK	DOS-FILES	DUAL
DUAL-COPY-EXIT	DUAL-FULL-EXIT	ELSE	ELSEIF
ENSU	ENSU-BUFFER	EQ	EXTRACT
GE	GLOBALS	GMT-ADJUSTMENT	GRAND-TOTAL-TABLE
GT	GTT	HALF-HOUR	HC
HD	HOUR	HOURS	IF
IFEND	IGNORE-INIT-ERRORS	INIT-JOB	INPUT-EXIT
INSIGHT-LINES	INTERVAL	IO-COST	IOLOG
IOLOG-SWITCH	LE	LEAVE	LINESIZE
LOG	LOGONLY	LOG-DEFAULT	LOG-GLOBALS
LP	LS	LT	LWKP2

LZ	MAXLINES	MAXRECS	MAXTSIZE
MINUTE	MINUTES	ML	MODIFY
MR	MT	NATURAL-SYSTEM-FILES	NE
NL	NO	NONE	NOT
OF	OFF	ON	OR
OREQ	OUTPUT-EXIT	OUTPUT-FILE	OUTPUT-INSIGHT-LINES
OWNER	PAGESIZE	PC	PC-FILE
PF	PRH	PRINT	PRINT-INSIGHT-LINES
PRINT-REC-HEAD	PRODUCT-CODE	PROD-CODE	PS
QTR-HOUR	RABN-RANGES	RECTYPE	REDEFINES
RESULTIS	REWIND	SECOND	SECONDS
SECS	SECURITY	SF	SK
SL	SORT	SPM-GLOBALS	ST
SUBTASK	SUBTOT	SUMMARIZE	TAPE
THEN	TITLE	TO	TP-TRANS-GAP
UEX4	UNLOAD	UNLOG	UQE-CPU-COST
UQE-CPU-FACTOR	UQE-IO-COST	USER-EXIT	VALOF
VSE-FILES	WHERE	WORK	YES

Statement Identifier Keywords

The statement identifier keyword specifies the type of statement being written, for example, GLOBALS, EXTRACT, DECLARE. These keywords must be written exactly as and where shown in the statement syntax descriptions.

Field Names and Summary Functions

Names of Command Log and derived data fields, along with summary functions of applicable fields, are listed in the section "Unicenter CA-APAS Data Fields and Summary Functions." An incorrectly spelled field or function name causes the syntax analyzer to treat it as a character string, leading to corresponding error messages.

Field and Summary Function Lists

A field list is a list of data field names. A summary function list is a list of summary functions. Valid field and summary function names are shown in the "Unicenter CA-APAS Data Fields and Summary Functions" section of this document. The list follows the request statement identifier keyword. A list ends with the first following parameter keyword. An example of a summary function list is:

```
COUNT PERCENT SUM(DATA-IO) MEAN(DATA-IO)
```

An example of a field list is:

```
CMD FILE RESPONSE-CODE NAT-MOD-ID
```

Parameter-Value Lists

Certain parameter keywords take as values a single word, a list of single words, or a group of parameter lists. Each simple entry is equivalent to a character string but does not require surrounding apostrophes. Quoted and hexadecimal character strings are acceptable as list entries.

Parameter list entries are separated by commas, ",". The list ends with the entry that is not followed by a comma or the entry followed by the next keyword. You may surround the parameter list with parentheses ( ). Parentheses improve clarity when writing groups of parameter lists. The following example specifies:

```
EXTRACT ASSO-IO (HD='A-IO','D-IO','W-IO')
```

Parameter value lists may be continued across statement record (card image) boundaries by ending the portion of the list on the first record with a comma and continuing the list in any position on the next record.

```
EXTRACT ASSO-IO (HD= 'A-IO',        /* beginning of list
                 'D-IO','W-IO')      /* continuation of list
```

Constants

Constants may be numeric or character string. Both are defined below.

Numeric Constants

Numeric constants may be fixed integer or decimal. For either type, negative values are indicated by prefixing the numeric digits with a minus sign, -. Commas are not allowed within numeric constants.

Fixed Integer Constants

Fixed integer constants do not have decimal points. The largest acceptable integer is 2,147,483,647; the smallest acceptable integer is -2,147,483,648.

```
1234 -57 50000        /* valid fixed integer constants
1,234 -57. 50,000.00  /* invalid fixed integer constants
```

Decimal Constants

Decimal constants may have up to 15 numeric digits total, with a decimal point at any position.

```
.005 8. 1000.7 12345.1234512345 /* valid decimal constants
1,000.7 12345.12345123456       /* invalid decimal constants
```

Character String Constants

Two types of character string constants may be used: alphanumeric and hexadecimal.

Alphanumeric Character Strings

An alphanumeric character string may contain any character in the collating sequence for which a data entry key is available. Character string constants are usually entered between apostrophes, " ' ". However, single words (no embedded spaces) can be treated as character string constants if they meet the following criteria. They must:

- Begin with a non-numeric character
- Not be keywords
- Not be enclosed between apostrophes

The following character strings are equivalent:

```
LB102TST
'LB102TST'
```

To be treated as character strings, numeric literals must be enclosed in apostrophes.

```
930901 /* is a fixed integer numeric constant
'930901' /* is a character string constant
```

The use or omission of apostrophes with a numeric literal causes significant differences in how the given value is used when specified for an alphanumeric parameter. For instance, the entry CPU-ID=4341 results in the numeric constant 4341 being right-justified within the eight-byte alphanumeric field CPU-ID.

On the other hand, CPU-ID='4341' results in the character string constant 4341 being left-justified within the receiving alphanumeric field.

Use of the apostrophe or quote character as a member character within an alphanumeric character string is subject to the following restrictions:

- If an apostrophe is to be included, then the entire string must be enclosed by quotes.
- If a quote is to be included, then the entire string must be enclosed by apostrophes.
- No individual alphanumeric character string constant may contain both apostrophe and quote characters as member characters. This effect may be achieved, however, either within a single hexadecimal character string constant or by concatenating character string constants as explained later.

```
"ABC'XYZ" 'ABC"XYZ'                /* valid strings
ABC' "ABC '"'ABC ABC'"'XYZ ABC'XYZ /* invalid strings
```

Hexadecimal Character String

A hexadecimal character string is an even number of hexadecimal digits, 0 through F, preceded by the letter H and enclosed in apostrophes, " ' ". Any character of the collating sequence, any binary number or any packed decimal number may be represented within a hexadecimal character string.

```
H'C1C2C3'   H'FFFFFFFF'    H'0000'
```

Continuing Character Strings

Character string constants may be continued by using the concatenation operator dash "-" between them. For example, the following character strings are equivalent:

```
'THE QUICK ' - 'BROWN'
* and
'THE QUICK BROWN'
'ABC' - H'C4' - 'EF'
* and
'ABCDEF'
'THIS IS A CHARACTER STRING' -
'SPANNING INPUT RECORDS'                                  /* (really)
* and
'THIS IS A CHARACTER STRING SPANNING INPUT RECORDS' /* (not really)
```

Date Format

Enter dates between apostrophes in the format *yymmdd* or *yyyymmdd* where *yy* is the last two digits of the year, *yyyy* is the year including the century, *mm* is the month, and *dd* the day.

```
'990322' '20010621'
```

Time Format

Enter time between apostrophes in the format *hhmmss*, where *hh* is the hour of the day (24 hour clock), *mm* is the minutes, and *ss* is the seconds of the time.

```
'123000' '123059' '235959' '000000'
```

Logical Expressions

Logical expressions in WHERE clauses specify command selection criteria. Logical expressions are composed of at least one logical comparison. Several comparisons may be made if connected by either the AND or the OR operator. A logical expression is evaluated to give a result of TRUE or FALSE. The syntax is:

```
WHERE
[NOT] field_name[(length,offset)] [NOT] operator constant
[ {AND|OR}
[NOT] field_name [(length,offset)] [NOT] operator constant
] ...
```

Parentheses may be used as required to assign priority in evaluating expressions.

Operators

Meaning	Keyword	Character
Equal to	EQ	=
Less than	LT	<
Greater than	GT	>
Not equal to	NE	¬=
Less than or equal to	LE	<=
Greater than or equal to	GE	>=
And	AND	&
Or	OR	|
Not	NOT	¬

Note: The NOT operator can precede all other operators, both keyword and character, with the exception of the AND, OR, and NOT operators. It can follow the AND and OR operators.

Logical expressions conform to standard evaluation rules for logical conditions in languages such as COBOL and Natural. Comparisons may be combined with AND and OR operators. Expressions in parentheses are combined first, innermost to outermost, followed by the NOT operator, followed by the AND operator, and lastly the OR operator. For example,

```
WHERE FILE = 3 AND TERM = S183 OR RSP > 3
                AND NOT DATE-TIME < '971212040000'
* is equivalent to
WHERE ((FILE = 3 AND TERM = S183) OR (RSP > 3
                 AND NOT (DATN-TIME < '19971212040000')))
```

Field References

Field references specify fields in the Command Log record, derived fields or summary functions of fields. Field and summary function names are found in the section "Unicenter CA-APAS Data Fields and Summary Functions." The Adabas buffers (record buffer, format buffer, etc.) are not valid fields for reference in logical expressions. They may be accepted by syntax checking, but execution results will not be as desired. A field name may optionally be followed by a length and offset in parentheses. This allows reference to a part of a field. For example,

```
WHERE JOBNAME(4) = 'CICS'
```

selects commands from all jobs with the first four characters = CICS. '4' is the length to compare.

This example:

```
WHERE JOBNAME(3,4) = '106'
```

selects commands from all jobs where the fourth through sixth characters equals 106. '3' is the length to compare and '4' is the offset where the comparison begins.

Numeric comparisons take into account the default scaling factor for a field. The number to which the field is being compared may contain decimal digits.

```
WHERE CPU > 0.010 SECS  /* over 10 milliseconds
```

Date and Time Comparisons

Dates and times are handled as character string variables. To distinguish date and time constants from numbers, the values must be enclosed in apostrophes. Dates must be entered in yymmdd or yyyymmdd format, and times in hhmmss format.

Entering a time range alone, without qualification by date, causes all times between the given values to be selected, regardless of the date. To obtain a true range of dates and times, you must use the 12-character DATE-TIME field (yymmddhhmmss) or the 14-character DATN-TIME field (yyyymmddhhmmss). You may omit trailing digits of the time portion. The following example selects all Command Log records between 0200 AM of one day and 0100 AM of the next day:

```
WHERE DATN-TIME >= '200209010200'
AND <= '200209020100'
```

The following rejects all records:

```
WHERE TIME >= '020000'
AND <= '010000'
```

Duration Comparisons

Duration fields, for example: CPU, DURATION, and ENQ-TIME, must be compared to a duration value, usually in seconds. The following selects commands with estimated CPU time greater than 200 milliseconds:

```
WHERE CPU > 0.200 SECS
```

Valid units of time duration are:

- SECOND
- SECONDS
- SECS
- MINUTE
- MINUTES
- QTR-HOUR
- HALF-HOUR
- HOUR
- HOURS
- DAY

Omitting the Left Side of Comparisons

Unicenter CA-APAS allows you to omit the left side of a comparison, as shown in some of the above examples. If the left field is omitted, the field most recently referred to is assumed. Note that arithmetic operators (the equal sign, "=" in the following example) may not be omitted.

```
WHERE FILE = 31 OR = 33 OR = 36
```

Inadmissible Comparisons

Unicenter CA-APAS does not support comparisons between two fields, or referral to buffer or I/O lists as fields to be compared. Constants may not be expressed as ranges of values. The following example illustrates invalid constants:

```
WHERE RSP = 100-230     /* value range constants are invalid
WHERE FILE = 25 THRU 49 /* THRU is not an APAS operator
```

Command Notation

This guide uses the following conventions and command notation.

`MIXEd CAse`	Identifies command abbreviations. The uppercase letters are the minimum abbreviation; lowercase letters are optional, for example: `Execute the EDit command.`
`Italics`	Identifies values supplied by the user and values supplied programmatically as well as published books or manuals. For example: `Execute the command: Enter` *`filename`*`.`
`|`	Items separated by the vertical bar indicate a choice between several mutually exclusive items, for example: `Yes|No.`
`[ ]`	Items within brackets are optional. Variations can include the following. `[a|b|c]` Indicates that you can choose none or any one value. `[a][b][c]` Indicates that you can choose none, one, several, or all values.
`{ }`	Items within braces are required. Variations can include the following. `{a|b|c}` Indicates that you must choose one and only one value. `{[a][b][c]}` Indicates that you must choose at least one value, but can choose more than one.
`Yes`	Underline indicates the default value, for example: `Yes|No.`
`,...,`	Ellipses indicate that the preceding item or group of items can be repeated more than once, for example: `file1,file2,...,filen.`

Chapter

3

Statement Descriptions

This section is a reference for writing statements to be processed by the Unicenter CA-APAS Insight Monitor for Adabas (Unicenter CA-APAS) Data Collector. These statements may direct the Data Collector to produce any combination of:

- Reports to be printed and/or displayed at terminals
- Output files containing data extracted from individual commands or summarized across multiple commands
- Output files containing Command Log records

Statement Overview

Each type of statement is described by a figure summarizing the syntax of the statement, a general description of the statement and its use, and an explanation of parameters and their uses. The types of statements used are:

1. GLOBALS statement establishes global parameters for the execution. It is best to use only one GLOBALS statement per Data Collector input stream.
2. DECLARE statement defines the format and computation of a user-defined field. If used in a batch input stream, DECLARE statements should precede all other types of statements.
3. EXTRACT statement requests a report and/or output file of information from individual Adabas commands.
4. SUMMARIZE statement requests a report and/or output file of information summarized over multiple Adabas commands
5. COPY statement writes all or selected Adabas Command Log records plus Unicenter CA-APAS derived fields and ACBX data information to an output file. The copied records may be processed in greater detail in a later batch execution of the Data Collector.
6. LOG and UNLOG statements are used to control the logging of Adabas Command Log records by Adabas to the Adabas Command Log file, DDLOG/DDCLOGRn.

All statement types are supplied in card image format according to conventions explained in the chapter "General Syntax Information."

GLOBALS Statement

Each Data Collector statement input stream should include a single GLOBALS statement. It defines the processing environment and values to be used in developing data for reports and output files. The GLOBALS statement does not require a statement label.

Syntax

The GLOBALS statement takes the following syntax.

```
GLOBALS
  [COLSPC = number]
  [CPU-COST = number]
  [CPU-FACTOR = number]
  [CPU-ID = name]
  [CPU-INSTRS-SEC = number]
  [DBID = number]
  [DBNAME = name]
  [DUAL-COPY-EXIT = module_name]
  [DUAL-FULL-EXIT = module_name]
  [ENSU = {ON|OFF}]
  [ENSU-BUFFER = number]
  [GMT-ADJUSTMENT = number]
  [GRAND-TOTAL-TABLE = {YES|NO}]
  [IGNORE-INIT-ERRORS = {YES|NO}]
  [INIT-JOB = jclname]
  [IO-COST = number]
  [IOLOG-SWITCH = {ON|OFF}]
  [LINESIZE  = number]
  [LOG-DEFAULT = {ON|OFF}]
  [MAXRECS = number]
  [MAXTSIZE = {number in k|800}]
  [PAGESIZE  = number]
  [PRINT-REC-HEAD = {YES|NO}]
  [RABN-RANGES = ( [ ASSO=(from_rabn,to_rabn,rabn_name,...)
                   | DATA=(from_rabn,to_rabn,rabn_name,...)
                   | WORK=(from_rabn,to_rabn,rabn_name,...)
                   ] )
  [TP-TRANS-GAP = n {SECONDS|SECS}]
  [UEX4 = module_name]
  [UQE-CPU-COST = number]
  [UQE-CPU-FACTOR = number]
  [UQE-IO-COST = number]
  [USER-EXIT = module_name]
  [WHERE logical-expression]
;
```

The following parameters apply to the VSE operating system only.

```
[DDCLOG = {DISK|DUAL|TAPE}]
[SUB-TASK = {YES|NO}]
[VSE-FILES = (filename [,logical_unit_number]
             [,blocksize][,record_format]
             [,logical_record_length][,tape_label_type]
             [,tape_disposition][,device_class])]
```

Example

The following is an example of the GLOBALS statement.

```
GLOBALS PAGESIZE=23
        LOG-DEFAULT=OFF
        CIPS=5000000      /* 5.0 MIPS
        VSE-FILES=(APASPRI,4,ODISK3,8)
        CPU-ID=SYSA
        DBID=2
        DBNAME=DEVELOP1;
```

GLOBALS Parameters

The following parameters are available in the GLOBALS statement.

COLSPC = *number*

Synonym: SF

The COLSPC parameter establishes a default number of spaces between columns in reports. This value can be overridden by the SF format parameter in individual requests.

```
COLSPC=3
```

CPU-COST = *number*

The COST of a command is calculated by multiplying the computed CPU, including CPU-FACTOR, in seconds by the CPU-COST and adding the TOTAL-IO multiplied by the IO-COST. CPU-COST is entered as, for example, dollars per CPU second. Up to five significant decimal places are supported.

For example, if you specify:

```
GLOBALS CPU-COST = 0.75  IO-COST = 0.00010;
```

then COST = (0.75 * CPU) + (0.0001 * TOTAL-IO).

CPU-FACTOR = *number*

The estimated command EST-CPU-TIME (synonym CPU), calculated with the Software AG formulas makes no provision for operating system or Adabas nucleus overhead. System accounting reports, however, include this overhead.

The computed command CPU time is adjusted using the CPU-FACTOR. This is intended to give a total CPU use closer to the Adabas job-step CPU time reported by system accounting routines. Values greater than 1.0 prorate Adabas nucleus and operating system overhead across commands in proportion to the estimated CPU instruction count given by the Software AG formulas. Up to three (3) significant decimal places may be specified.

```
CPU-FACTOR=1.06
```

CPU-ID = *name*

CPU-ID specifies a one to eight character alphanumeric name that identifies a CPU.

This value is printed out in all reports and is available in the header portion of output records. It is a major key field for the Performance History System. If a numeric literal is used for CPU-ID, it should be enclosed in apostrophes as shown in the example below.

```
CPU-ID=CPUA
CPU-ID='4341' /* apostrophes for left justification
```

CPU-INSTRS-SEC = *number*

Synonym: CIPS

This parameter is a numeric estimate of the processing power of the computer on which Adabas session was run.

This number is used to derive the EST-CPU-TIME (synonym CPU) field for commands. CPU is derived by calculating the estimated number of instructions for a command based on formulas supplied by Software AG, dividing by this CIPS parameter, and multiplying by the CPU-FACTOR parameter. CPU, in turn, is used to calculate the COST field for commands.

Note: The value specified for CIPS must be in whole numbers rather than a factor. Commas and decimal points are not allowed.

```
CIPS=4700000 /* CORRECT FOR 4.7 MIPS
CIPS=4.7     /* WRONG - SPECIFY A VALUE OF CIPS, NOT OF MIPS
```

DBID = *number*

DBID specifies a number from zero to 65535 which identifies an Adabas database. It is printed out in all reports and is available in the header portion of output records. It is a major key field for the Performance History System.

```
DBID=27
```

DBNAME = *name*

DBNAME specifies an eight-character name which identifies an Adabas database. This value is printed out in all reports and is available in the header portion of output records. It is a key field for the Performance History System.

```
DBNAME=TEST2
```

DUAL-COPY-EXIT = *module_name*

This parameter specifies the name of an exit module that submits batch jobs to copy dual (X/Y) files when a switch occurs between files.

A switch occurs either because the active file becomes full or because a FLIPXY command is issued. The submitted job copies the formerly active, and usually now full, file to another sequential file. For information on the use of dual files, see the *Unicenter CA-APAS Systems Guide*.

DUAL-FULL-EXIT = *module_name*

This parameter specifies the name of an exit module that automates the selection of the action to take when both of the dual files are full. For complete information on the use of dual files, see the *Unicenter CA-APAS Systems Guide*.

ENSU = {ON|OFF}

This parameter is used to turn on the Data Collector. The default value is OFF.

```
ENSU=ON
```

ENSU-BUFFER = *number*

This parameter specifies the size of the buffer used to pass records from the APASUEX4 main task to the Unicenter CA-APAS Data Collector sub-task. A value between 16000 and 32000 is normally sufficient.

```
ENSU-BUFFER=32000
```

GMT-ADJUSTMENT = *number*

This numeric parameter is provided for installations whose computer time-of-day clock and adjusted operating system time are set to Greenwich Mean Time rather than to local time. It is used to convert the adjusted operating system time to local time.

The value entered is the correction in hours. Installations east of Greenwich enters a positive value and installations west of Greenwich enters a negative value.

```
GMT-ADJUSTMENT=-5 /* correction for US East Coast
GMT-ADJUSTMENT= 1 /* correction for Italy
```

GRAND-TOTAL-TABLE = {YES|NO}

Synonym: GTT

This parameter affects the building of grand total tables for SUMMARIZE requests that contain an INTERVAL specification. A virtual storage table to hold the totals for the current interval is always maintained while processing such a request. The GTT parameter determines whether an additional table for the total across all intervals is used.

Grand total tables for certain kinds of requests can grow very large and sometimes cause out-of-storage problems. Therefore, the default is to not use grand total tables.

An additional consideration is that no GT (grand total) type records are written to output files of interval summary requests if the grand total table is not used.

To override the default for all interval requests, use the GTT parameter as follows:

```
GTT=YES
```

To override the default for only a single request, use the GTT parameter within that SUMMARIZE statement.

IGNORE-INIT-ERRORS = {YES|NO}

This keyword is used during initialization of the Data Collector. If a value of NO is specified or the default is used, the MPM session abnormally terminates with an error code of User 26 if there is a syntax error in the control cards read at initialization of the Data Collector. If a value of YES is specified, the MPM session does not abnormally terminate.

If a value of YES is specified and a syntax error is detected, the request in error is not added. Therefore, the Data Collector does not accumulate the information that was expected from that request.

INIT-JOB = *jcl_name*

This parameter specifies the name of a JCL job stream in the APASJCL data set that automatically gets submitted during the initialiazation of Unicenter CA-APAS.

IO-COST = *number*

This figure is entered as, for example, dollars per EXCP or SIO. Up to five significant decimal places are supported.

```
IO-COST=.0012 /* $0.0012 PER I/O
```

IO-COST is used in computing the COST of a command, as explained above in the discussion of CPU-COST.

IOLOG-SWITCH = {ON|OFF}

IOLOG-SWITCH is used to control whether or not CA-SpaceMan I/O Log records are created in DBGIOR5. If IOLOG-SWITCH is set to ON, DBGIOR5 passes I/O Log records to the Data Collector. The records are not passed if IOLOG-SWITCH is set to OFF.

IOLOG-SWITCH can also be dynamically controlled using the IOLOGSW Insight command.

When IOLOG-SWITCH is set to OFF, DBGIOR5 performs no work and therefore imposes no overhead on ADAIOR. One can install DBGIOR5, set the IOLOG-SWITCH to a value of OFF in the GLOBALS statement, then use the IOLOGSW Insight command to collect IOLOG records only when desired. This allows users to limit the sampling overhead on ADAIOR to those times when they want I/O Log data.

LINESIZE = *number*

Synonym: LS

The LINESIZE parameter establishes a default line length for reports to be printed or displayed at terminals. This value is the maximum line length for all reports unless overridden by the LS parameter in individual requests. There is no system limit to the logical line length of a report.

```
LS=132
```

LOG-DEFAULT = {ON|OFF}

This parameter is used to control whether Adabas Command Log records are written to the DDLOG/DDCLOGRn.

LOG-DEFAULT=ON enables the writing of Adabas Command Log records.

LOG-DEFAULT=OFF suppresses the writing of Adabas Command Log records. The default value is OFF.

```
LOG-DEFAULT=ON
```

Even if this value is set to OFF, this function can be performed using the Unicenter CA-APAS LOG request statement.

MAXRECS = *number*

Synonym: MR

The MAXRECS parameter limits the number of Command Log records read during a stand-alone batch execution of the Data Collector. When the specified number of records have been read, the utility ends as if end-of-file had been reached on the Command Log or COPY input file.

This parameter allows short trial runs to verify that statements produce desired results without having to process an entire input file. MAXRECS is ignored when the Data Collector is being executed within an MPM session.

```
MAXRECS=1000
```

MAXTSIZE = {*number in k*|*800*}

Synonym: MT

The MAXTSIZE parameter, expressed as a number of kilobytes, limits the size of a SUMMARIZE request's summarization table, the TSIZE of the request. When the TSIZE of a SUMMERIZE request exceeds the value of MAXTSIZE, the Data Collector treats it as an Interval break summarizing the data collected, producing the required output, releasing the summarization table space, and then continues collecting data. This could result in multiple interval reports or records for any given interval period. Use of this parameter is to avoid excessive processing time occationally experienced due to processing huge amounts of collected data.

```
MAXTSIZE=800
```

To override the default for only a single request, use the MT parameter within that SUMMARIZE statement.

PAGESIZE = *number*

Synonym: PS

The PAGESIZE parameter establishes a default number of lines per page for printed reports. This value applies to all reports except when overridden by the PS parameter in individual requests.

```
PS=60
```

PRINT-REC-HEAD = {YES|NO}

Synonym: PRH

This parameter sets the default for PRINT TO file record headers. If set to YES, all PRINT TO files have record headers unless suppressed in individual request parameter lists. Print Record Headers are useful only when multiple reports are directed to a single PRINT TO file for deferred printing.

The print header can be used to sort records in the file before they are printed. Print Record Headers prevent the data for an EXTRACT or SUMMARIZE statement from being properly displayed at a terminal.

Use Print Record Headers when one or more reports being directed to a single PRINT TO file is an EXTRACT or is a SUMMARIZE that contains an INTERVAL specification.

```
PRH=YES
```

Print Record Headers are unnecessary when a given PRINT TO file receives only multiple summarize reports, none of which contains an interval specification. More information on this parameter is provided in *Unicenter CA-APAS Systems Guide.*

```
RABN-RANGES = ( [ ASSO=(from_rabn,to_rabn,rabn_name,...)
                | DATA=(from_rabn,to_rabn,rabn_name,...)
                | WORK=(from_rabn,to_rabn,rabn_name,...) ] )
```

Note: This parameter should not be specified with Unicenter CA-SpaceMan, which builds its own RABN-RANGE table.

The RABN-RANGES parameter establishes a table of names for Relative Adabas Block Numbers (RABNs). This list is used to process the RABN list generated when LOGIO is specified in the Adabas ADARUN parameters. For further information about using RABN-RANGES, see the RABN-RANGE field description in the chapter "Unicenter CA-APAS Data Fields." Also see the *Adabas DBA Reference Manual.*

Three tables of RABN-RANGES are established, one each for Associator, Data, and Work RABNs. Each range is specified by a from-RABN-number, a to-RABN-number, and the name of the range, up to 10 characters. The maximum number of entries in each table is 100.

```
RABN-RANGES = ( ASSO = (1,10000,AVOL001,
                        10001,20000,AVOL002)
                DATA = (1,20000,DVOL001,
                        20001,50000,DVOL002,
                        60001,60500,DFILE023)
                WORK = (1,1000,WORKPART1,
                        1001,4000,WORKPART2,
                        4001,6000,WORKPART3)
              )    /*close rabn-range paren
```

TP-TRANS-GAP = *n* {SECONDS|SECS}

This parameter is used to derive the TP-TRANS-COUNT field. Any interval between commands larger than the TP-TRANS-GAP duration is assumed to represent user-think time, TP monitor processing time, and line transmission time that occurs during a TP interaction with the user.

Ideally, the value is set shorter than the best response time for TP transactions and longer than the worst Adabas command duration. Some compromise is probably necessary. For example,

```
TP-TRANS-GAP = 4 SECONDS
```

counts any interval longer than four seconds between Adabas commands for a particular user as a TP transaction. For example, when you press the ENTER key, wait for a new screen display, then study the new screen display before initiating the next transaction.

UEX4 = *module_name*

The UEX4 parameter specifies the name of an additional local User Exit 4 to be invoked after Unicenter CA-APAS User Exit 4. It is called before Adabas has written the Command Log record to the Command Log file. Specification and use of the UEX4 module are described in the *Unicenter CA-APAS Systems Guide*.

UQE-CPU-COST = *number*

The UQE-COST of a command is calculated by multiplying UQE-CPU-TIME by the UQE-CPU-FACTOR and the UQE-CPU-COST and adding the TOTAL-IO multiplied by the UQE-IO-COST. UQE-CPU-COST is entered, for example, as dollars per CPU second. Up to five significant decimal places are supported.

For example, if you specify:

```
UQE-CPU-COST = 0.75          UQE-IO-COST = 0.00010
UQE-CPU-FACTOR = 1.06
```

then

```
UQE-COST = (UQE-CPU-TIME * 1.06 * 0.75) + (0.00010 * TOTAL-IO).
```

UQE-CPU-FACTOR = *number*

The CPU time accumulated by Software AG in the UQE for each user omits some operating system and Adabas overhead. Operating system account reports include this overhead.

The UQE-CPU-TIME is adjusted using the UQE-CPU-FACTOR to give a closer measure of actual CPU time used by each command. Up to three significant decimal places may be specified.

```
UQE-CPU-FACTOR = 1.06
```

UQE-IO-COST = *number*

This figure is entered, for example, as dollars per EXCP or SIO. Up to five significant decimal place are supported.

```
UQE-IO-COST = 0.0012          /*$0.0012 per I/O
```

UQE-IO-COST is used in computing UQE-COST of a command, as explained above in the discussion of UQE-CPU-COST.

USER-EXIT = *module_name*

This parameter specifies the name of a user-written module that functions somewhat like an Adabas User Exit 4 routine. The module is dynamically loaded at execution time and called before each Command Log record is processed. Specifications and use of the USER-EXIT module are described in *Unicenter CA-APAS Systems Guide.*

WHERE = *logical_expression*

The WHERE parameter specifies command selection criteria that are to be applied before any command selection by EXTRACT, SUMMARIZE, or COPY statements. Command records excluded by the WHERE criteria of a GLOBALS statement are excluded from processing by all other statements regardless of the selection criteria of those statements.

```
GLOBALS WHERE DATE-TIME >= '930901130330'
                       AND <= '930901130400'
        AND TERMINAL-ID = 'S183'
```

VSE Parameters

DDCLOG = {DISK|DUAL|TAPE}

This parameter is used only for stand-alone batch execution of the Data Collector under VSE. It specifies the input Command Log file device type and format. The default is TAPE. The filename is DDCLOG. Valid specifications are the following:

DISK specifies the input file is a disk file with standard variable blocked format.

DUAL specifies the input file is an Adabas disk dual Command Log file.

TAPE specifies the input file is on tape.

```
DDCLOG=DISK
```

SUBTASK = {YES|NO}

When the SUBTASK parameter is set to YES, the Data Collector works as in MVS. APASUEX4 puts Command Log data into the ENSU-BUFFER, the Data Collector processes this data when Adabas is not busy.

When SUBTASK is set to NO, the Data Collector is called every time a command is processed in MPM, without waiting for an period of low use, thereby potentially slowing Adabas processing. Specifying NO however reduces memory requirements because no ENSU-BUFFER or sub-task stack areas are used.

```
VSE-FILES = (filename [,logical_unit_number] [,blocksize]
             [,record_format] [,logical_record_length]
             [,tape_label_type] [,tape_disposition] [,device_class] )
```

The VSE-FILES parameter enables the VSE user to override default characteristics of Computer Associates standard VSE filenames or to add new filenames and characteristics if no defaults exist. Parameters specified internally by the program cannot be overridden and user-specified values for those parameters are ignored. All parameters up to the last desired overridden parameter must be specified for any given filename. Any of these intervening parameters may be a null parameter except for filename. Parameters are listed as follows.

filename

A filename of one to seven characters must be specified.

logical_unit_number

Specifies the SYSnnn name for the specified filename. This defaults as described in the appendix "Logical Unit Definitions" of the *Unicenter CA-APAS Systems Guide*; if it is not defaulted to one of those, the default is 0. The reassignments must not cause conflicts in logical unit use or assign incorrect physical devices to output units.

blocksize

Specifies maximum blocksize associated with the filename. The blocksize defaults as described in the appendix "Logical Unit Definitions" of the *Unicenter CA-APAS Systems Guide*. Value specified must not be larger than device maximums.

record_format

Specifies the OS record format associated with the file. The default OS format is VB. Valid specifications are the following:

F specifies fixed format unblocked records

FA specifies fixed format unblocked records with ASA control characters

FB specifies fixed format blocked records

V or VB specifies variable format records is always blocked

logical_record_length

Specifies the maximum logical length of each record in the file. Logical record length defaults as described in the appendix "Logical Unit Definitions" of the *Unicenter CA-APAS Systems Guide*. May be reset internally to blocksize -4 bytes for variable format records.

tape_label_type

Specifies type of labels associated with a tape file. The default is SL. Valid specifications are the following:

BLP specifies bypass label processing; treated as non-labeled

NL specifies non-labeled tapes

SL specifies IBM-standard labeled tapes

tape_disposition

Specifies tape handling disposition when closing a tape file. The default is unload. Valid specifications are the following:

LEAVE, which leaves tape positioned as is after close label processing

REWIND, which rewinds tape to load point

UNLOAD, which rewinds and unloads the tape

device_class

Specifies the device class for a filename. One specifies either TAPE or DISK.

device_class is used where logical unit is not assigned until OPEN, and where the default logical unit would assume the wrong device type. If not specified, the device class is determined by the type of device assigned to the specified or defaulted *logical_unit_number*. If no *logical_unit_number* is specified or defaulted to by *filename* as indicated in the appendix "Logical Unit Definitions" of the *Unicenter CA-APAS Systems Guide*, it defaults to unit 9 and unit record device type.

The following example specifies that filename 0DISK3X is SYS035 with a blocksize of 19000 bytes and is a tape device, while the filename MYTAPE is logical unit SYS047 and defaults to tape unit with VB records of 32760 bytes.

```
VSE-FILES ((0DISK3X,35,19000,,,,,TAPE),
(MYTAPE,47))
```

DECLARE Statement

The DECLARE statement defines a new field and the means of computing its value from existing fields. Two forms of definition are supported: static redefinition of an existing field, and assignment of values based on logical criteria.

To redefine or reference another user-defined field, the redefined or referenced field must be declared first. In other words, a field referenced in a DECLARE must itself have already been declared. The DECLARE statement does not require a statement label.

Syntax

The DECLARE statement takes the following syntax.

```
DECLARE
[CMDLOG | IOLOG]
field_name ({A|B|H|N|P}length [.decimal_places [/scale]]
              [HD 'header_string[,…]']
            )
{REDEFINES field_name [(length [,offset])]
|VALOF
    IF logical_expression
    [ELSE {conditional_statement | RESULTIS constant}]
    [ELSEIF {conditional_statement | RESULTIS constant}
      [ELSEIF {conditional_statement | RESULTIS constant}]…]
    [THEN {statement | RESULTIS constant}]
    IFEND
} ;
```

Keep the following considerations in mind when using DECLARE statements:

- Declared fields are global in nature. Once defined in a single DECLARE statement, a field may be referenced in any subsequent statements.
- DECLARE statements may be added but not deleted. Once in effect, the field definition from a DECLARE continues to be available through the rest of the current execution of the Data Collector.
- A given field must not be declared in more than one DECLARE statement in any given execution of the Data Collector.

Example

The following is an example of the DECLARE statement.

```
DECLARE JOBCODE (A4 HD='JOB','CODE')
        REDEFINES JOBNAME(4,3); /* CHARS 3-6
DECLARE DEPT (N3 HD='DEPT','NO') VALOF
        IF JOBCODE <= 'C104' RESULTIS '201'
           ELSEIF JOBCODE > 'C104' AND <= 'H210' RESULTIS '460'
           ELSEIF JOBCODE > 'H210' AND <= 'Q360' RESULTIS '600'
           ELSEIF JOBCODE = 'CICS' THEN
                  IF TERMID(1) = 'K' THEN RESULTIS '801'
                     ELSEIF TERMID(1) = 'L' THEN RESULTIS '802'
                     ELSE RESULTIS '899'
                  IFEND
           ELSE RESULTIS '999'
        IFEND
```

DECLARE Parameters

The parameters of the DECLARE statement are described below.

```
DECLARE
[CMDLOG | IOLOG]
field_name ({A|B|H|N|P}length [.decimal_places [/scale]]
      [HD 'header_string[,...]']

           )
```

Log Record Type

```
CMDLOG | IOLOG
```

The log record type, either CMDLOG or IOLOG, declares which type of input record and request is applicable for the declared field.

If no log record type is specified or if CMDLOG is coded, the field may be used only in EXTRACT and SUMMARY requests processing CMDLOG fields. If IOLOG is specified, the field may be used only in processing Unicenter CA-SpaceMan ADAIOR log records.

Format Parameters

A declared field may be defined as any one of the following formats.

A*n* (Alphanumeric):

The length *n* specifies the number of characters.

B*n* (Binary):

The length *n* should be 2 for half-word or 4 for full-word. The *decimal_places* and *scaling_factor* are used in printing the field. A PF= format is necessary in requests to specify the number of digits to be printed.

H*n* (Hexadecimal):

The field is printed in hex format. The field column width is the number of hexadecimal digits and must be a multiple of 2.

N*n* (Numeric Character):

Character format numbers. A field defined with this format is treated as a character field and must not be used in arithmetic expressions or summary functions.

P*n* (Packed):

The length *n* gives the number of digits. If an even value is specified for length, it is incremented by one. The *scaling_factor* and *decimal_places* are used in printing.

Header Parameter

HD = '*character string*',...

This parameter supplies column headers that override the default headers in a report. This is especially useful when more than one field is written in a column.

REDEFINES/VALOF Expressions

After having specified the DECLARE *field_name* and any format specifications, the user must specify either the REDEFINES or VALOF keywords and their respective parameters.

REDEFINES Declaration

The parameters of the REDEFINES are described below.

```
REDEFINES field_name [(length [,offset])]
```

A declared field may redefine all or a portion of an existing field. This is particularly useful with the USER-AREA of the Adabas Control Block, which is often used to pass a transaction-id or other data for use in analysis or billing.

```
DECLARE CICS-TRAN(A4 HD='CICS','TRAN')
REDEFINES USER-AREA;
```

The format of the field being declared must be compatible with that of the area being redefined.

length

Specifies the length of the portion of that field that is being redefined. This should be the same as the length of the declared field.

offset

Specifies the offset within the field being redefined at which the redefinition starts. Offset is always given in bytes, not digits, from left to right with an offset value of one (1) equating to the first, that is, the leftmost byte. One (1) is the default offset value if no offset is specified.

VALOF Declaration

The parameters of the VALOF are described below.

```
VALOF
    IF logical_expression
    [ELSE {conditional_statement | RESULTIS constant}]
    [ELSEIF {conditional_statement | RESULTIS constant}
      [ELSEIF {conditional_statement | RESULTIS constant}]…]
    [THEN {statement | RESULTIS constant}]
    IFEND
```

When the VALOF specification is used, the statement following the VALOF keyword is executed. This is always a conditional statement. When a RESULTIS keyword is encountered, the constant following the RESULTIS is assigned to the declared field and evaluation of the statement stops.

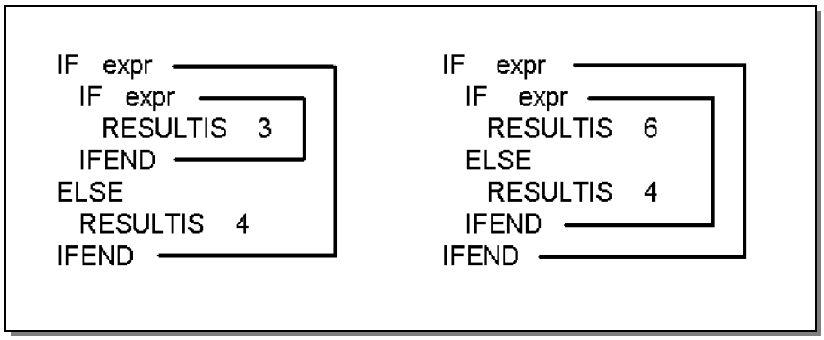

Conditional statements have a standard IF, THEN, ELSE structure. Unicenter CA-APAS syntax for this structure has the following attributes:

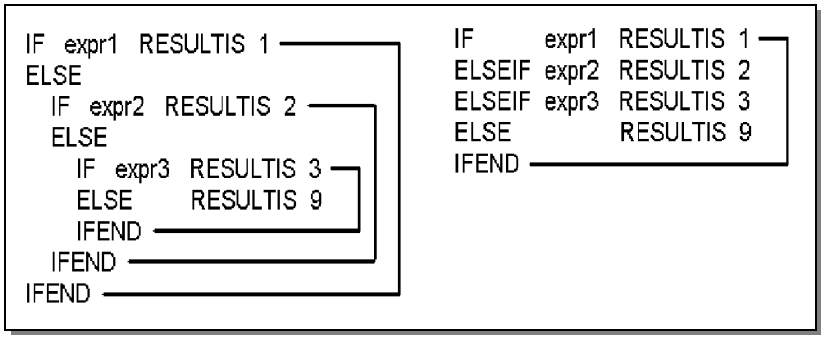

- An IF statement must terminate with an IFEND. This makes the matching of ELSE to IF statements unambiguous. The lines in the diagrams above represent this matching.
- An ELSEIF is provided. This permits easy construction of CASE statements. ELSEIF is almost equivalent to ELSE IF except a matching IFEND is not required for each ELSEIF.

An IF condition without an ELSE statement gives unpredictable results when the IF condition is not satisfied. This can lead to program exceptions for declared binary or packed fields; an ELSE should set the value to zero if no other value is appropriate. In the following example, the value of FILENAME is unpredictable for values of FILE other than 10 or 11.

```
DECLARE FILENAME (A8 HD='FILE NAME') VALOF
        IF FILE = 10        RESULTIS 'SALARY'
          ELSEIF FILE =11   RESULTIS 'TIME CARD'
        IFEND;
```

EXTRACT Statement

The EXTRACT statement requests detailed output from individual commands for exception reporting, debugging, or other purposes. As many EXTRACT statements as desired may be used in a given execution of the Data Collector.

Syntax

The EXTRACT statement takes the following syntax.

```
label: EXTRACT
       [(format_parameter_list)]
       [CMDLOG | IOLOG]
       extract_field_list
       [INSIGHT-LINES = number]
       [OUTPUT-EXIT = module_name]
       [OUTPUT-FILE {ddname|(DUAL,ddname)}]
       [OUTPUT-INSIGHT-LINES = {YES|NO}]
       [OWNER = {userid|'*'}]
       [PRINT-INSIGHT-LINES = {YES|NO}]
       [PRINT TO {ddname|(DUAL,ddname)}]
       [SECURITY = {DISPLAY|MODIFY|NONE}]
       [TITLE = 'character_string']
       [WHERE logical_expression]
;
```

The EXTRACT statement specifies the following:

- Output fields and their report/output format
- Report titles
- Selection criteria for accepting commands to be processed
- Output destinations

Example

The following is an example of the EXTRACT statement.

```
SIGRESPX: EXTRACT(SF=2 SK=1)
RSPCL RSP CMDSEQ DATE / TIME
JOBNAME (HD='JOBNAME','NATID') / Natural-LOGON / NAT-PROGRAM
TERM (HD='TERMINAL') / TERM (PF=H8)
COMMAND-CODE CID (HD='CID') / CID (PF=H8)
COPT1 COPT2 FB-ERR-FIELD
FILE ISN ISNQ ISNLL CPU TOTAL-IO
ASSO-IO (HD='A-IO','D-IO','W-IO') / DATA-IO / WORK-IO
ECBS
WHERE RSPCL GT 2
TITLE = 'EXCEPTION REPORT OF SIGNIFICANT ' -
'RESPONSE CODES'
INSIGHT-LINES=150
PRINT TO SIGRESP
PRINT-INSIGHT-LINES=YES
SECURITY=DISPLAY
OWNER=MSCROOGE;
```

EXTRACT Parameters

The parameters of the EXTRACT statement are described below.

Label

label:

A label is required and must be unique. It is written to the report generated and in the record header on the output file. It may be referenced from Unicenter CA-APAS in ADD, PAUSE, RESUME, and DELETE commands.

Log Record Type

CMDLOG | IOLOG

The log record type, either CMDLOG or IOLOG, specifies the type of log record the request processes. The fields named in the request must all apply to this record type. If no record type is specified, CMDLOG, the default, is applied.

Format Parameter List

The following optional parameters format the request and override the corresponding parameter stated in the GLOBALS statement. If used, they are entered immediately after the EXTRACT keyword and are enclosed in a single pair of parenthesis. The format parameters available are the following.

COLSPC = *number*

Synonym: SF

Specifies the number of spaces placed between each output column. The default is one (1).

LINESIZE = *number*

Synonym: LS

Specifies the line length for this request. This is the maximum line size; actual line size is determined by the request field list. A syntax error is generated if actual line size exceeds the specified maximum line size.

MAXLINES = *number*

Synonym: ML

Specifies the maximum number of commands to be reported on a printed report. This is not actually the number of print lines if multiple lines per command have been specified.

MAXRECS = *number*

Synonym: MR

Specifies the maximum number of records to be written to an OUTPUT-FILE.

PAGESIZE = *number*

Synonym: PS

Specifies number of lines per page for this request.

PRH = {YES|NO}

YES specifies that header data be placed at the beginning of all PRINT TO output records. Print Record Headers are useful only in cases where lines from multiple reports are being interleaved on a single PRINT TO file. Print headers allow the records of the PRINT TO file to be sorted to regroup them into separate and distinct reports prior to printing. They should be used only when an EXTRACT uses the same PRINT TO ddname that is used by another EXTRACT or SUMMARIZE statement.

Warning: Coding the PRH parameter in an EXTRACT statement overrides the GLOBALS statement value of PRH that would otherwise have been in effect for the request-- and for all other requests with print output to the same ddname. That is, if any request specifies PRH=YES, then all requests using that PRINT TO ddname have print headers, regardless of the value of PRH in the GLOBALS statement or in the individual requests.

SK = *number*

Specifies the number of lines to skip following each set of lines printed out for a given command. Multiple line output per command is specified by using slashes, /, between fields.

Extract Field List

The *extract_field_list* specifies the fields to be reported and/or written to output files and optionally, override formats. Fields are written in the order specified in the request.

Syntax

field_specification **[/] [*extract_field_list*]**

where ***field_specification*** is:
field_name **[*field_format*]**

where ***field_format*** is:
([*header_parm*][*print_format*][*output_format*][*justification*])

Field Specification

The *field_specification* is any Command Log or Unicenter CA-APAS data field listed in the chapter "Unicenter CA-APAS Data Fields."

A slash, /, between fields prints them in the same column with the value of the second field "stacked" under the value of the first field. The longest field in the column or the column header, whichever is greater, establishes the column width. The default column heading is that of the first field in the column. When fields are stacked, it is usually desirable to override the heading with the HD *field_format* parameter following the first field in the stacked set.

Field Format

The *field_format* is a parenthesized parameter list that allows you to override default field formats for reporting and output files. Available parameters are:

Header Parm

HD = '*character string*',...

This parameter supplies column headers that override the default headers in a report. This is especially useful when more than one field is written in a column. In this case, the HD parameter must be specified for the first (top) field in the column.

```
EXTRACT ASSO-IO (HD='A-IO','D-IO',
    'W-IO') / DATA-IO / WORK-IO
```

Print Format

PF={A|H|N|O|D}length [.decimal_places]

The print format specifies the report format, including type, length and decimal places for the field. Valid print formats are:

A*n* (Alphanumeric):

The field column width is the specified length (*n*).

```
PF=A10
```

H*n* (Hexadecimal):

The field is printed in hex format. The field column width is the number of hexadecimal digits and must be a multiple of 2. The following prints only the first two bytes of the COMMAND-ID.

```
CID(PF=H4)
```

N*n* (Numeric):

In the following example, the number is printed with five digits preceding the decimal place and three digits following.

```
PF=N5.3
```

A specification of NC prints the number without commas. The default numeric format includes commas after every third digit.

The following formats are used only in printing buffers. These "out of line" formats must be the last fields in the *extract_field_ list.*

O*n* (Out-of-Line):

The buffer is printed in character format following the main report line. This is the default format for the FB and SB. If a length of 0 is specified, the entire buffer is printed in a variable length format. If a non-zero length is specified, it will be the maximum number of characters printed.

D*n* (Dump Format):

This is an out-of-line format with hexadecimal display of the buffers. This is the default format for the RB, VB, IB, UEXB, and IO-LIST. The length of 0 causes printing of the entire buffer, no matter what length, except for UEXB, which will print a maximum of 200 bytes.

Output Format

```
OF = {A|B|N|P}length [.decimal_places]
```

The output parameter, OF, overrides the default output file format for fields.

```
EXTRACT  FILE (OF=P5)
         ASSO-IO (OF=P5)
         WORK-IO (OF=P5)
         DATA-IO (OF=P5) ;
```

Valid output formats are as follows:

A*n* (Alphanumeric):

Fields are copied from input to output up to the length *n* specified; length indicates the length of the output field. This is the only valid output format for buffers.

B*n* (Binary):

Numbers are converted to 2 or 4 byte binary fields. Overflow is indicated by a field that is the maximum negative number (e.g. X'8000'). The length *n* indicates the number of bytes in the output field.

N*n* (Numeric Character):

Numbers are converted to zoned decimal. Decimal places are implicit. The length *n* is the number of digits.

P*n* (Packed):

Numbers are converted to packed format. The length *n* is the number of digits, rather than bytes. An even value for length is incremented by one.

The length value, *n*, must be supplied for all of the above formats. The length is in bytes for all formats except packed, in which case it specifies number of digits. The length for numbers is the number of digits to the left of the decimal point. *decimal_places* gives the number of digits to the right of the decimal point. All decimal points are implicit; no character is ever written to indicate the position of the decimal point. The length may be positive for buffers, in which case a fixed format field is output. A buffer output with a fixed length always includes the original buffer length at the beginning of the field.

A zero may be specified for the output length of a buffer, in which case it is written with a variable length. This is the default. Variable length buffers must be written at the end of the output record.

Justification

HC = {L|C|R}

This parameter specifies header justification: Left, Center, or Right. The default is Center for alpha fields and Right for numeric fields. Enter the letter L for left justification.

Insight Lines

INSIGHT-LINES = *number*

The INSIGHT-LINES parameter is applicable only when the Data Collector is being executed within an Adabas region. It specifies the number of output lines to be kept in a buffer for display at Unicenter CA-APAS Insight terminals. Lines stored in this buffer are transferred to Unicenter CA-APAS Insight terminals when requested.

The maximum amount of data that can be transferred to Unicenter CA-APAS Insight users is limited by the maximum record buffer length allowed by Natural. Therefore, the useful maximum for INSIGHT-LINES is approximately (32750-150)/request-line-length. The minimum request line length is 79. The buffer for an EXTRACT request always contains the most recent lines of the report.

Output-Exit

```
OUTPUT-EXIT = module_name
```

An OUTPUT-EXIT may be specified in addition to or instead of an output file. All output records generated by the request is passed to the specified module. Standard linkage conventions are used. On entry to the module, Register 15 contains the entry point of the module, Register 14 contains the return address, Register 13 points to a standard 18-word OS save area which can be used by the module. Register 1 points to a two-word parameter list:

- Offset 0: Address of the output record. 0 on the final call.
- Offset 4: Address of a word that may be used by the exit module. The contents of this word are preserved between calls to the exit module, so it can be used to store the address of a work area for the exit module. The value of the word is 0 on the first call to the exit module. The exit is called once for each output record and a final time when Unicenter CA-APAS ends or the MPM shuts down.

A return code may be inserted into Register 15 before returning from the exit module. A return code of 16 (decimal) bypasses writing of the output record to the output file. Any other value results in the output record being written if OUTPUT-FILE parameter was specified.

When the Data Collector is run under the Adabas User Exit 4, the exit must be written to accept control from the Data Collector in the addressing mode being used by Adabas. It must also return control to the Data Collector in the same addressing mode in which it received control. If the exit must switch to a different addressing mode; for example, into 24 bit mode to do non-VSAM I/O, the exit must preserve the addressing mode that it had on entry and restore that addressing mode before returning control to the Data Collector. Failure to adhere to these addressing requirements results in addressing errors that causes the Data Collector to terminate abnormally.

Output-File

```
OUTPUT-FILE = {ddname|(DUAL,ddname)}
```

The optional OUTPUT-FILE parameter specifies a file to which output records are written. Multiple output extracts may be written to the same file. Output records from different extracts are intermixed, but can be separated by a post-processor program on the basis of the extract names (label). For more details of the formats of headers and data areas for output files, see the appendix "OUTPUT-FILE Formats."

The keyword DUAL may be used as follows:

```
OUTPUT-FILE=(DUAL,OUTFILA)
```

DUAL directs output alternately to two files, an "X" and a "Y" file. The actual ddnames appearing in the JCL for the two files are the ddname coded suffixed by the letter "X" for one and the letter "Y" for the other. To allow suffixes, the ddname coded in the request statement cannot exceed seven characters (six characters for VSE). For more information about dual files, see the *Unicenter CA-APAS Systems Guide* .

Output-Insight-Lines

```
OUTPUT-INSIGHT-LINES = {YES|NO}
```

When you specify both the INSIGHT-LINES parameter and either the OUTPUT-FILE or OUTPUT-EXIT parameter, or both, the OUTPUT-INSIGHT-LINES parameter specifies how the Data Collector processes a DISPLAY command from Unicenter CA-APAS for this request.

The default value, NO, specifies that the record generated for the DISPLAY command are not written to the OUTPUT-FILE data set or passed to the OUTPUT-EXIT. A value of YES specifies that Unicenter CA-APAS writes a record to the OUTPUT-FILE data set and/or passed to the OUTPUT-EXIT whenever an Unicenter CA-APAS Insight DISPLAY command is processed for this request.

Note: This parameter has no effect on calls to OUTPUT-EXIT or records written to the OUTPUT-FILE data set at the end of an interval or at the end of an MPM session. In addition, setting the value of this parameter to YES when the request does not contain an INSIGHT-LINES parameter and either the OUTPUT-FILE or OUTPUT-EXIT parameter, or both, has no meaning and is treated as a syntax error.

Owner

```
OWNER = {userid | '*'}
```

The OWNER parameter specifies the eight-character identifier of the "owner" of the request.

It is used to control who is authorized to use Unicenter CA-APAS Insight commands to modify the request, that is, to use the DELETE, PAUSE, RESUME, or RESET commands, or to issue a DISPLAY command to display its INSIGHT-LINES.

Typically, the OWNER value is the userid of the person who creates the request. The default value of an asterisk, `'*'`, indicates a "public" request; a request that can be accessed by all Unicenter CA-APAS users.

Note: The single quotation marks shown around the asterisk are required because an asterisk is not an alphanumeric character. The quotation marks do not become part of the value. For a more complete explanation of Unicenter CA-APAS security, see the *Unicenter CA-APAS User Guide.*

Print-Insight-Lines

```
PRINT-INSIGHT-LINES = {YES|NO}
```

When you specify both the PRINT TO and INSIGHT-LINES parameters, the PRINT-INSIGHT-LINES parameter specifies how the Data Collector processes a DISPLAY command from Unicenter CA-APAS for this request.

The default value, NO, specifies that the lines generated for the DISPLAY command are not written to the PRINT TO data set also. A value of YES specifies that Unicenter CA-APAS writes all request data lines to the PRINT TO data set whenever an Unicenter CA-APAS Insight DISPLAY command is processed for this request.

Note: This parameter has no effect on the records written to the PRINT TO data set at the end of an interval or at the end of an MPM session. In addition, setting the value of this parameter to YES when the request does not contain both the INSIGHT-LINES and PRINT TO parameters has no meaning and is treated as a syntax error.

Print To

```
PRINT TO = {ddname|(DUAL,ddname)}
```

This parameter specifies the file to which a report should be written. Each report is ordinarily written to a different file to avoid interspersing lines from different reports.

However, it is possible to write any number of reports to the same PRINT TO file if the records are sorted prior to printing. Refer to the previous discussion of the PRH parameter. The keyword DUAL may be used as described earlier for OUTPUT-FILE.

```
PRINT TO = (DUAL,OUTFILA)
```

When DUAL is specified with PRINT TO, a Print Record Header is automatically written as part of each output record. For more information about Dual files, see the *Unicenter CA-APAS Systems Guide*.

Security

```
SECURITY = {DISPLAY|MODIFY|NONE}
```

The SECURITY parameter defines the level of Unicenter CA-APAS request security for the request. See the *Unicenter CA-APAS User Guide* for a complete description of Unicenter CA-APAS request security.

The default value, NONE, allows all users complete access to the request. A value of MODIFY allows all users to display data collected by the request but allows only authorized user to modify it, that is, to use Unicenter CA-APAS Insight DELETE, PAUSE, RESUME, or RESET commands against this request. A value of DISPLAY allows only authorized users to display or modify the request.

Title

TITLE = '*character_string*'

TITLE is a list of character strings that are used as title lines for a report. They are centered at the top of every page immediately following the standard report heading lines.

Where

WHERE = *logical_expression*

The logical expression is used to select commands to be included in this request. It is ordinarily used to:

- Create exception reports of commands requiring investigation
- Select commands of certain applications for debugging or audit-trail purposes

You may select commands on the basis of general logical expressions involving any actual or derived Command Log fields except the buffers. Any commands that do not satisfy the WHERE criteria of a GLOBALS statement are excluded from processing by all requests, regardless of the WHERE criteria in each request.

SUMMARIZE Statement

The SUMMARIZE statement summarizes information over multiple commands. Summarization is based on control break fields. If an INTERVAL parameter is supplied, summary tables are output at breaks of the specified interval. If no INTERVAL parameter is coded, a summary table is output at the end of processing.

Syntax

The SUMMARIZE statement takes the following syntax.

```
label: SUMMARIZE
       [(format_parameter_list)]
       [CMDLOG | IOLOG]
       summary_function_list
       [BY control_break_field [(BY number)]
          [BY control_break_field [(BY number)]]...]
       [INSIGHT-LINES = number]
       [INTERVAL = duration_value]
       [OUTPUT-EXIT = module_name]
       [OUTPUT-FILE {ddname|(DUAL,ddname)}]
       [OUTPUT-INSIGHT-LINES = {YES|NO}]
       [OWNER = {userid|'*'}]
       [PRINT-INSIGHT-LINES = {YES|NO}]
       [PRINT TO {ddname|(DUAL,ddname)}]
       [SECURITY = {DISPLAY|MODIFY|NONE}]
       [TITLE = 'character_string']
       [WHERE logical_expression]
```

Any number of SUMMARIZE requests may be included in an execution of the Data Collector. The SUMMARIZE statement specifies the following:

- The fields to be summarized and their report/output formats
- The control break fields for summarization
- Any special sort order for the report
- Optionally, the periodic time interval at which summary reports are to be created
- Titles for the report
- Selection criteria for including command records in the summarization
- Output destinations

Example

The following is an example of the SUMMARIZE statement.

```
QTRHRSUM: SUMMARIZE(SF=3) COUNT PERCENT CMD-RATE
          SUM(TOTAL-IO) PCT(TOTAL-IO)
          MEAN(TOTAL-IO) MAX(TOTAL-IO)
          SUM(CPU) PCT(CPU) MEAN(CPU) MAX(CPU)
          PCT(DURATION) MEAN(DURATION) MAX(DURATION)
          BY FILE BY COMMAND-CODE
          INTERVAL 15 MINUTES
          TITLE = ('QUARTER-HOUR SUMMARY'
          'BY FILE / COMMAND CODE')
          OUTPUT-FILE QTRHRFIL
          PRINT TO QTRHRSUM;
```

SUMMARIZE Parameters

The parameters of the SUMMARIZE request are described below.

Label

label:

The label specified for the SUMMARIZE request must be unique. It is written both to the report generated from the request and in the record header on the output file. It may be referenced from Unicenter CA-APAS to ADD, PAUSE, RESUME, and DELETE commands.

Log Record Type

CMDLOG | IOLOG

The log record type, either CMDLOG or IOLOG, specifies the type of log record the request processes. The fields named in the request must all apply to this record type. If no record type is specified, CMDLOG, the default is applied.

Format Parameter List

The following optional parameters format the report and override the corresponding parameter stated in the GLOBALS statement. If used, they are entered within a single pair of parentheses immediately after the keyword SUMMARIZE. The format parameters available in the SUMMARIZE statement include:

COLSPC = *number*

Synonym: SF

This number of spaces is placed between each output column on the report. The default is one (1).

GRAND-TOTAL-TABLE ={YES|NO}

Synonym: GTT

Specifies whether a grand total is developed if the request includes an INTERVAL parameter. A grand total table can require additional virtual storage. A GRAND-TOTAL-TABLE value coded here overrides, for this individual SUMMARIZE statement, a GRAND-TOTAL-TABLE on the GLOBALS statement.

LINESIZE= *number*

Synonym: LS

Specifies line length for this request. This is the maximum linesize; actual linesize is determined by the request field list. A syntax error is generated if the actual linesize exceeds the specified linesize.

MAXLINES = *number*

Synonym: ML

Maximum number of lowest break level items to be printed on a report. This is not actually the number of print lines if multiple lines per BY entity have been specified.

MAXRECS = *number*

Synonym: MR

Specifies the maximum number of records to be written to an OUTPUT-FILE.

MAXTSIZE = *number in k*

Synonym: MT

Limits the size of the summarization table or TSIZE of the request. The number is expressed in kilobytes.

PAGESIZE = *number*

Synonym: PS

Specifies the number of lines per page for this request.

PRH = {YES|NO}

YES specifies that header data be placed at the beginning of all PRINT TO output records. Print Record Headers are useful only in cases where lines from multiple reports are being interleaved on a single data set. Print headers allow the lines to be sorted to regroup them into separate and distinct reports prior to printing.

Print headers are unnecessary if a PRINT TO file is shared only by multiple SUMMARIZE requests, none of which contain an INTERVAL specification. Print headers should be used in any other cases of multiple reports sharing the same PRINT TO file.

Coding this parameter in a SUMMARIZE statement overrides the GLOBALS statement value of PRH. However, if any request specifies PRH=YES, then all requests using the same PRINT TO ddname have print headers, regardless of the value of PRH in the GLOBALS statement or in the individual requests. Unicenter CA-APAS users adding a request with PRH=YES must ensure that the PRINT TO ddname in the request has not been used by any other request or that the ddname is already receiving Print Record Headers.

SK = number

Specifies the number of lines skipped following each set of lines printed out for any given combination of BY values (lowest break level). Multiple-line output for each BY entity is specified for a report by using slashes, /, between fields. To make such reports easier to read, skip one or more lines.

SUBTOT = OFF

Suppresses the generation of all levels of subtotal and total lines and output records for a SUMMARIZE request. The Data Collector ordinarily generates a subtotal record for each level of break of the control break fields. SUBTOT=OFF generates report lines and output records only for the lowest level of break. This may be useful if the output is being passed to another report writer, since the higher level control records would serve no purpose.

Note: SORT always suppresses subtotals.

Summary Function List

The *summary_function_list* specifies the report/output data and its format. Functions are output in the order specified in the request.

Syntax

```
summary_specification
    [SORT [ASC|ASCENDING|DESC|DESCENDING] ]
    [/] [summary_function_list]

    where summary_specification is:
        summary_field [field_format] |
        summary_function(field) [field_format]

        where summary_field is:
            COUNT|PERCENT|CMD-RATE|JOB-START|JOB-END|
            NAT-LOAD-COUNT|TP-TRANS-COUNT

        where summary_function is:
            SUM|PCT|MAX|MIN|MEAN|RATE

        where field_format is:
            ([header_parm][print_format][output_format]
            [justification])
```

Summary Specification

The *summary_fields* and *summary_functions* are listed in the chapter "Unicenter CA-APAS Data Fields."

A slash, /, between any two specifications prints them in the same column with the value of the second being "stacked" under the value of the first. Column width is the width of the widest specifications in the column or the width of the column heading whichever is greater. The default column heading is that of the first specification in the column. To override the heading, specify a heading with the HD *field_format* parameter following the first specification.

Note: Instantaneous *summary_fields* and *summary_functions* cannot be the first field requested in a Unicenter CA-APAS SUMMARY request. They must be preceded by at least one session or interval-related *summary_field* or *function*.

The SORT specification immediately follows the *summary_specification* to which it applies. If no sort specification is given, the output table will automatically be in control-break field sequence with default sub-totals. Only one sort specification may be given in a SUMMARIZE statement. Most commonly, it is used to sort tables in descending order of magnitude or resource usage.

```
SUMMARIZE SUM (TOTAL-IO) SORT DESCENDING
          MEAN(TOTAL-IO)
          MAX (TOTAL-IO)
          BY   CMD
```

Note: Sorting on the MEAN function is not allowed.

Field Format

The *field_format* is a parenthesized parameter list which allows you to override the default field or summary function formats for both reports and output files. Available parameters are:

Header Parm

HD = '*character string*',...

This parameter supplies column headers that override the default headers in a report. This is especially useful when more than one field is written in a column. In this case, the HD parameter must be specified for the first (top) field in the column.

```
SUMMARIZE SUM(ASSO-IO)
  (HD='TOT A-IO','TOT D-IO','TOT W-IO')
  / SUM(DATA-IO) / SUM(WORK-IO)
```

Print Format

PF = {{A|G|H|N}length [.decimal_places]}

The print format specifies the report format, including type, length, and decimal places for the field. Valid print formats are:

A*n* (Alphanumeric):

The field column width is the specified length (*n*).

```
PF=A10
```

G*n* (Graph):

Prints histogram horizontally in the column width specified. One column is used for the base of the histogram and the rest are used for building bars. By default, the length of bars is normalized to the maximum value for the summary function in the table. This is overridden by specifying a normalizing value such as:

```
SUM(TOTAL-IO) (PF=G11/100)
```

which prints a histogram with each character representing 10 I/O's That is, G11 results in one column being used as the base of the histogram and the remaining ten columns being used for bars. Since the number of bars is based on the normalizing value of 100, each bar represents 10 I/O's.

Note: This format cannot be used for the MEAN function.

H*n* (Hexadecimal):

The field is printed in hex format. The field column width is the number of hexadecimal digits and must be a multiple of 2. Hex format is normally not used for summary values; it use in a SUMMARIZE request should be limited to fields that are specified by the BY *control-break-field* parameters. In the example shown below, the hex format displays terminal-id values. This avoids the problem of values in default character format being non-displayable.

```
SUMMARIZE COUNT SUM(TOTAL-IO)
 BY TERMINAL-ID (PF=H8)
 WHERE USER-TYPE = 'TSO'
```

N*n* (Numeric):

In the following example, the number is printed with five digits preceding the decimal place and three digits following.

```
PF=N5.3
```

A specification of NC prints the number without commas. The default numeric format includes commas after every third digit.

The value NL specifies the default input/output length of a numeric field. If NL is less than the field length, the value is truncated. If it is greater, values are expanded with blanks.

Output Format

```
OF = {A|B|N|P}length [.decimal_places]
```

The default output format for fields is overridden using the OF parameter. This parameter is coded with other field-format parameters, PF and HD, immediately following the field or summary function specification.

```
SUMMARIZE SUM(TOTAL-IO) (OF=B4)
          RATE(TOTAL-IO) (OF=N9)
          BY FILE (OF=P5)
```

Valid output formats are:

A*n* (Alphanumeric):

Fields are copied from input to output up to the length specified. The length *n* indicates the length of the output field.

B*n* (Binary):

Numbers are converted to 2 or 4 byte binary fields. Overflow is indicated by a field that is the maximum negative number, for example: X'8000'. The length *n* indicates the number of bytes in the output field.

N*n* (Numeric):

Numbers are converted to zoned decimal. Decimal places are implicit. The length *n* is the number of both digits and bytes.

P*n* (Packed):

Numbers are converted to packed format. The length *n* is the number of digits, rather than bytes. If an even value is specified for length, it is incremented by one.

The length value must be supplied with all of the above formats. The length is in bytes for all formats except packed; it specifies number of digits. The length for numbers is the number of digits to the left of the decimal point. *Decimal_places* gives the number of digits to the right of the decimal point. All decimal points are implicit; no character is ever written to indicate the position of the decimal point.

Justification

HC = {L|C|R}

This parameter specifies header placement: Left, Center, or Right. Default placement is Center for alpha fields and Right for numeric fields. The letter L specifies left justification.

Control Break Field

BY *field_name* [BY *number*] ...

The control break fields are listed after the BY keyword. They are listed from major to minor break level. These fields are always printed in front of the list of summary functions. *Field_formats* may be given for the fields, but slashes are not allowed.

Report subtotals and subtotal records for each specified break level are automatically generated, unless suppressed by SUBTOT=OFF or a SORT specification. SUBTOT=OFF in the format parameter list immediately following the SUMMARIZE keyword suppresses all levels of subtotals. SUBTOT=OFF may be used following one or more of the control break fields to selectively suppress individual levels of subtotaling.

The BY parameter can specify only fields with a numeric format. It allows grouping of numeric values by rounding them down to the next lower multiple of the BY value.

```
INSUM: SUMMARIZE COUNT BY FILE
       BY ISN (BY 500)
       WHERE CMD = 'L1'
```

This report accesses by ISN ranges 0-499, 500-599, etc.

Insight-Lines

```
INSIGHT-LINES = number
```

The INSIGHT-LINES parameter is applicable only when the Data Collector is being executed within an Adabas region. It specifies the number of output lines kept in a buffer for display at Unicenter CA-APAS Insight terminals. Print format lines are stored in this buffer while data are being collected and then transferred to Unicenter CA-APAS Insight terminals when requested by users.

The maximum amount of data that can be transferred to Unicenter CA-APAS Insight users is limited by the maximum record buffer length allowed by Natural. Therefore, the useful maximum for INSIGHT-LINES is approximately (32750-150)/request-line-length. The minimum request line length is 79. The buffer for a SUMMARIZE request always contains the first lines of the generated summary table.

Interval

```
INTERVAL = n {HALF-HOUR|HOUR|HOURS|MINUTE|MINUTES|QTR-HOUR|
              SECONDS|SECS}
```

The INTERVAL parameter generates interval reports and gives the length of the interval.

Durations are calculated beginning at regular time points. This allows the data computed during these time intervals to be coordinated with data collected by various system monitoring functions, usually on a ten or fifteen minute basis. A complete summary table is calculated at the end of every interval. A grand total table is printed at the end of the entire run if GTT=YES was included in this SUMMARIZE statement or in the GLOBALS statement.

The Data Collector prints the beginning and ending time of each interval, as well as the time of the first and last command within the interval. If no commands were logged within the interval, subtotal records with data of zero are written.

A value of 0 SECONDS specifies a “demand interval”. These interval breaks occur only when data is output by the request as the result of a Unicenter CA-APAS Insight DISPLAY command. Because the minimum time in the Command Log record is 1.048576 seconds, a non-zero interval of less than that is not meaningful and produces unpredictable results.

Output-Exit

OUTPUT-EXIT = *module_name*

An OUTPUT-EXIT may be specified in addition to or instead of an output file. All output records generated by the request is passed to the specified module. Standard linkage conventions are used. On entry to the module, Register 15 contains the entry point of the module, Register 14 contains the return address, Register 13 points to a standard 18-word OS save area which can be used by the module. Register 1 points to a two-word parameter list:

- Offset 0: Address of the output record. 0 on the final call.
- Offset 4: Address of a word that may be used by the exit module. The contents of this word are preserved between calls to the exit module, so it can be used to store the address of a work area for the exit module. The value of the word is 0 on the first call to the exit module. The exit is called once for each output record and a final time when Unicenter CA-APAS ends or the MPM shuts down.

A return code may be inserted into Register 15 before returning from the exit module. A return code of 16 (decimal) bypasses writing of the output record to the output file. Any other value results in the output record being written if the OUTPUT-FILE parameter was specified.

When the Data Collector is run under the Adabas User Exit 4, the exit must be written to accept control from the Data Collector in the addressing mode being used by Adabas. It must also return control to the Data Collector in the same addressing mode in which it received control. If the exit must switch to a different addressing mode; for example, into 24-bit mode to do non-VSAM I/O, the exit must preserve the addressing mode that it had on entry and restore that addressing mode before returning control to the Data Collector. Failure to adhere to these addressing requirements results in addressing errors that causes the Data Collector to terminate abnormally.

Output-File

OUTPUT-FILE = {*ddname*|(DUAL,*ddname*)}

The optional OUTPUT-FILE parameter specifies a file to which output records are written. Multiple output summaries may be written to the same file. Records from different requests may be intermixed; they can be separated by a post-processor program on the basis of the request names (labels). For details of the formats of headers and data access for output files, see the appendix “OUTPUT-FILE Formats.”

The keyword DUAL may be used as follows:

```
OUTPUT-FILE (DUAL,OUTFILA)
```

DUAL directs output alternately to two files, an "X" and a "Y" file. The actual ddnames appearing in the JCL for the two files are the ddname coded suffixed by the letter "X" for one and the letter "Y" for the other. To allow suffixes, the ddname coded in the request statement cannot exceed seven characters (six characters for VSE). For a discussion of dual files, see the *Unicenter CA-APAS Systems Guide*.

Output-Insight-Lines

```
OUTPUT-INSIGHT-LINES = {YES|NO}
```

When you specify the INSIGHT-LINES parameter and either the OUTPUT-FILE or OUTPUT-EXIT parameter, or both, the OUTPUT-INSIGHT-LINES parameter specifies how the Data Collector processes a DISPLAY command from Unicenter CA-APAS for this request.

The default value, NO, specifies that the record generated for the DISPLAY command are not written to the OUTPUT-FILE data set or passed to the OUTPUT-EXIT. A value of YES specifies that Unicenter CA-APAS writes a record to the OUTPUT-FILE data set and/or passed to the OUTPUT-EXIT whenever an Unicenter CA-APAS DISPLAY command is processed for this request.

Note: This parameter has no effect on calls to the OUTPUT-EXIT or records written to the OUTPUT-FILE data set at the end of an interval or at the end of an MPM session. In addition, setting the value of this parameter to YES when the request does not contain the INSIGHT-LINES parameter and either the OUTPUT-FILE or OUTPUT-EXIT parameter, or both, has no meaning and is treated as a syntax error.

Owner

```
OWNER = {userid|'*'}
```

The OWNER parameter specifies the eight-character identifier of the "owner" of the request.

It is used to control who is authorized to use Unicenter CA-APAS Insight commands to modify the request, that is, to use the DELETE, PAUSE, RESUME, or RESET commands or to issue a DISPLAY command to display its INSIGHT-LINES.

Typically, the OWNER value is the userid of the person who creates the request. The default value of an asterisk, `'*'` indicates a "public" request; a request that can be accessed by all Unicenter CA-APAS users.

Note: The single quotation marks shown around the asterisk are required because an asterisk is not an alphanumeric character. The quotation marks do not become part of the value. For a more complete explanation of Unicenter CA-APAS security, see the *Unicenter CA-APAS User Guide.*

Print-Insight-Lines

```
PRINT-INSIGHT-LINES = {YES|NO}
```

When you specify both the PRINT TO and INSIGHT-LINES parameters, the PRINT-INSIGHT-LINES parameter specifies how the Data Collector processes a DISPLAY command from Unicenter CA-APAS for this request.

The default value, NO, specifies that the lines generated for the DISPLAY command are not written to the PRINT TO data set also. A value of YES specifies that Unicenter CA-APAS write all request data lines to the PRINT TO data set whenever an Unicenter CA-APAS DISPLAY command is processed for this request.

Note: This parameter has no effect on the records written to the PRINT TO data set at the end of an interval or at the end of an MPM session. In addition, setting the value of this parameter to YES when the request does not contain both the INSIGHT-LINES and PRINT TO parameter has no meaning and is treated as a syntax error.

Print To

```
PRINT TO = {ddname|(DUAL,ddname)}
```

This parameter specifies the ddname where a report is written. Summaries without interval specifications may share a single ddname. Each summary report begins on a new page.

Each interval report is ordinarily written to a different output ddname. However, it is possible to print any number of reports to the same PRINT TO ddname if Print Record Headers are specified. The keyword DUAL may be used as described above for OUTPUT-FILE.

```
PRINT TO (DUAL,PRTFILA)
```

When DUAL is specified with PRINT TO, a Print Record Header is automatically written as part of each output record. For more information about Dual files, see the *Unicenter CA-APAS Systems Guide*.

Security

```
SECURITY = {DISPLAY|MODIFY|NONE}
```

The SECURITY parameter defines the level of Unicenter CA-APAS request security for the request. For a complete description of Unicenter CA-APAS request security, see the *Unicenter CA-APAS User Guide*.

The default value, NONE, allows all users complete access to the request. A value of MODIFY allows all users to display data collected by the request but allows only authorized user to modify it, that is, to use Unicenter CA-APAS DELETE, PAUSE, RESUME, or RESET commands against this request. A value of DISPLAY allows only authorized users to display or modify the request.

Title

`TITLE = 'character_string'`

TITLE is a list of character strings that are used as title lines for a report. They are centered at the top of every page immediately following the standard report heading lines.

Where

`WHERE = logical_expression`

The logical expression selects command records to be summarized. You may select commands on the basis of general logical expressions involving any fields except the buffers.

COPY Statement

The COPY statement writes Command Log records to a file for later batch processing by the Data Collector. Multiple COPY statements may be used in a given execution of the Data Collector. COPY statements may be ADDED, DELETED, PAUSED, and RESUMED from Unicenter CA-APAS Insight terminals. However, appropriate JCL for the ddnames must have been included when starting the Adabas nucleus session.

Unicenter CA-APAS derived fields, which are listed below, are not calculated and inserted into the COPY record unless those fields are referenced by one or more EXTRACT or SUMMARIZE requests. One may include such a request in either the Unicenter CA-APAS input stream or add it later to force the fields onto the COPY records. Once started, the request maybe deleted and the field continues to be included in the COPY records.

COMMAND-COST is included in the COPY file only if it is referenced by a request. EST-CPU-TIME and EST-INSTRUCTIONS are included if either field is referenced. Before using the COPY statement, review the discussion of COPY files in *Unicenter CA-APAS Systems Guide*.

Syntax

The COPY statement takes the following syntax.

```
label: COPY
   [CMDLOG | IOLOG]
   [buffer_list]
   OUTPUT-FILE {ddname|(DUAL,ddname)}
   [OUTPUT-EXIT = module_name]
   [OWNER = {userid| '*'}]
   [SECURITY = {DISPLAY|MODIFY|NONE}]
   [WHERE logical_expression]
   ;
```

The COPY statement specifies the following:

- Buffers which are output in addition to the standard Command Log fields and fields derived by Unicenter CA-APAS.
- Selection criteria for command records to be output.
- The ddname for sequential file output.

Example

The following is an example of the COPY statement.

```
COPYTEST: COPY SB VB
          WHERE NAT-LOGON = 'SYSTEST'
          OUTPUT-FILE = (DUAL,DDCOPY1) ; /* X/Y file
```

COPY Parameters

The parameters of the COPY statement are described below.

Label

label:

The label specified for the COPY request must be unique. It may be referenced from Unicenter CA-APAS to ADD, PAUSE, RESUME and DELETE the request.

Log Record Type

CMDLOG | IOLOG

The log record type, CMDLOG or IOLOG, specifies the type of log record the request processes. The fields named in the request must all apply to this record type. If no record type is specified, the default, CMDLOG, is applied.

Buffer List

[FB [,IB] [,IOL] [,RB] [,SB] [,UEXB] [,VB]]

Any buffers specified are logged for selected commands.

The buffers specified are copied in addition to the standard Adabas Command Log fields and the data derived by Unicenter CA-APAS. However, only those buffers that are logged by Adabas can be copied, that is, ADARUN LOGSB=YES must be specified before the search buffer is available in the Adabas Command Log record to be copied to an Unicenter CA-APAS COPY file.

The *buffer_list* may be omitted, in this case, only the fixed length portion of the command record and the Unicenter CA-APAS derived fields are written to the COPY file.

The buffer UEXB is available only when the value of LNUINFO is greater than zero (0) in the Adabas link routine. UEXB data is included with the other Adabas buffers if the ADARUN LOGGING parameter is set to YES and the LOGUX parameter is ON.

Output-File

```
OUTPUT-FILE = {ddname|(DUAL,ddname)}
```

This OUTPUT-FILE parameter specifies the ddname to which records are written.

The keyword DUAL may be used as follows:

```
OUTPUT-FILE (DUAL,CPYFIL)
```

DUAL directs output alternately to two files, an "X" and a "Y" file. The actual ddnames appearing in the JCL for the two files are the ddname coded suffixed by the letter "X" for one and the letter "Y" for the other. To allow suffixes, the ddname coded in the request statement cannot exceed seven characters or six characters for VSE. For a discussion of dual files, see the *Unicenter CA-APAS Systems Guide*.

Output-Exit

```
OUTPUT-EXIT = module_name
```

An OUTPUT-EXIT may be specified in addition to or instead of an output file. All output records generated by the request is passed to the specified module. Standard linkage conventions are used. On entry to the module, Register 15 contains the entry point of the module, Register 14 contains the return address, Register 13 points to a standard 18-word OS save area which can be used by the module. Register 1 points to a two-word parameter list:

- Offset 0: Address of the output record. 0 on the final call.
- Offset 4: Address of a word that may be used by the exit module. The contents of this word is preserved between calls to the exit module, so it can be used to store the address of a work area for the exit module. The value of the word is 0 on the first call to the exit module. The exit is called once for each output record and a final time when Unicenter CA-APAS ends or the MPM shuts down.

A return code may be inserted into Register 15 before returning from the exit module. A return code of 16 (decimal) bypasses writing of the output record to the output file. Any other value results in the output record being written if an OUTPUT-FILE parameter was specified.

When the Data Collector is run under the Adabas User Exit 4, the exit must be written to accept control from the Data Collector in the addressing mode being used by Adabas. It must also return control to the Data Collector in the same addressing mode in which it received control. If the exit must switch to a different addressing mode; for example, into 24 bit mode to do non-VSAM I/O, the exit must preserve the addressing mode that it had on entry and restore that addressing mode before returning control to the Data Collector. Failure to adhere to these addressing requirements results in addressing errors that causes the Data Collector to terminate abnormally.

Owner

OWNER = {*userid*|'*'}

The OWNER parameter specifies the eight-character identifier of the "owner" of the request.

It is used to control who is authorized to use Unicenter CA-APAS commands to modify the request, that is, to use the DELETE, PAUSE, RESUME, or RESET commands against the request.

Typically, the value specified here is the userid of the person who creates the request. The default value of an asterisk, '*' indicates a "public" request; a request that can be accessed by all Unicenter CA-APAS users.

Note: The single quotation marks shown around the asterisk are required because an asterisk is not an alphanumeric character. The quotation marks do not become part of the value. For a more complete explanation of Unicenter CA-APAS security, see the *Unicenter CA-APAS User Guide.*

Security

SECURITY = {DISPLAY|MODIFY|NONE}

The SECURITY parameter defines the level of Unicenter CA-APAS request security for the request. For a complete description of Unicenter CA-APAS request security, see the *Unicenter CA-APAS User Guide.*

The default value, NONE, allows all users complete access to the request. A value of MODIFY allows all users to display data collected by the request but allows only authorized user to modify it, that is, to use the Unicenter CA-APAS DELETE, PAUSE, RESUME, or RESET commands against this request. A value of DISPLAY allows only authorized users to display or modify the request.

Where

WHERE = *logical_expression*

This logical expression selects commands to be copied. Some typical uses include:

- Selecting exceptional commands needing investigation
- Selecting commands from an application for debugging

You may select commands on the basis of general logical expressions involving fields that are either derived or from the Command Log (except the buffers).

Note: If no WHERE clause is present, a COPY file record is written for every command.

LOG Statement

The LOG statement specifies selection criteria for record logging by Adabas and specifies which buffers are to be logged. LOG statements are optional; any number of them may be included in the input stream and/or submitted through Unicenter CA-APAS Insight by means of the ADD command. The format of this statement is shown below.

Syntax

The format of the LOG statement is as follows:

```
label: LOG
  [ FB   [(length)] ]
  [ IB   [(length)] ]
  [ IOL  [(length)] ]
  [ RB   [(length)] ]
  [ SB   [(length)] ]
  [ UEXB [(length)] ]
  [ VB   [(length)] ]
  [ WHERE logical_expression ] ;
```

The LOG statement specifies the following:

- Buffers that are output in addition to the standard Command Log record.
- Selection criteria for command records to be output.

Example

The following is an example of the LOG statement.

```
LOGRSPG2: LOG SB VB(32)
          WHERE RSPCL GT 2;
```

The LOG statement requests logging of specific records and, optionally, buffers based on logical selection criteria. As many LOG statements as desired may be included in the input command stream.

LOG Parameters

The elements of the LOG statement are described below.

Label

```
label: LOG
```

The *label* specified for the LOG request must be unique. It is used when deleting the LOG request with an UNLOG request through Unicenter CA-APAS.

Buffer List

```
FB    [(length)]
IB    [(length)]
IOL   [(length)]
RB    [(length)]
SB    [(length)]
UEXB  [(length)]
VB    [(length)]
```

The *buffer_list* specifies the buffers that are to be logged if the WHERE selection criteria is satisfied. Each buffer optionally is followed by a length representing the maximum length of data for that buffer that is logged. The buffer is truncated to the specified inclusive length.

If no buffers are specified, the fixed length portion of the log record is logged if the selection criterion is true. This is useful only if LOG-DEFAULT=OFF.

Note: When using CLOG layout 4, the buffer list is ignored and all buffers turned on in the ADARUN parameters are included when the Command Log record is written to the Adabas Command Log file.

Where

```
WHERE = logical_expression
```

This logical expression is used to select records to be logged. Some typical uses include:

- Selecting exceptional commands needing investigation
- Selecting commands from an application for debugging

You may select commands on the basis of general logical expressions involving fields that are either derived or from the Command Log (except the buffers).

Note: If no WHERE clause is present, the request defaults to true and the specified buffers are logged by the Adabas Command Log file for every command.

UNLOG Statement

The UNLOG statement specifies a LOG statement is to be deleted from Data Collector processing. UNLOG statements are not to be included in the Data Collector input stream; they are only to be submitted through Unicenter CA-APAS Insight.

Syntax

The format of the UNLOG statement is as follows:

```
label: UNLOG;
```

The *label* specified for the UNLOG statement must match the label of the LOG statement that is to be deleted.

Example

The following is an example of the UNLOG statement.

```
LOGRSPG2: UNLOG;
```

This UNLOG statement instructs the Data Collector to delete the LOG statement named LOGRSPG2.

UNLOG Parameters

The elements of the UNLOG statement are described below.

Label

```
label: UNLOG
```

The *label* specified for the LOG request to be deleted from Data Collector processing.

Chapter

4

Unicenter CA-APAS Data Fields

This section describes the data fields and summary functions that may be referenced in Unicenter CA-APAS Insight Monitor for Adabas (Unicenter CA-APAS) requests.

Key to Data Field and Summary Function Descriptions

Descriptions of fields and summary functions follow a standard format. The standard headings of the descriptive items are explained here.

Heading	Description
Adabas Version	Availability of some fields and functions differs according to which version of Adabas Unicenter CA-APAS is reporting. Several conditions about availability are stated separately for different versions of Adabas.
Gvn	Indicates whether the field or function is given for the indicated version of Adabas. "No" indicates that meaningful values are never provided for the Adabas version, even though the field or function name may pass syntax checking when it appears in requests that are being used with that Adabas version. "Yes" indicates that meaningful values are provided for the Adabas version, subject however to any limitations that may be stated in the description.

Heading	Description
ACBX	Indicates whether the optional "Adabas Control Block Extension" (ACBX) facility of Unicenter CA-APAS is required in order to report meaningful values of the field or function for the indicated version of Adabas. For some fields, the ACBX facility is a necessary but not sufficient condition; the Unicenter CA-APAS User Exit 4 routine may also be required.
UX4	Indicates whether the Unicenter CA-APAS User Exit Four routine must be executed with the indicated version of Adabas if meaningful values for the field or function are to be reported. For some fields, this is a necessary but not sufficient condition; the ACBX feature may also be required. Note that execution of the Unicenter CA-APAS User Exit Four routine is dependent on certain Adabas command logging options being ON.
Log	Indicates whether the field or function is available from Adabas Command Logs being processed in stand-alone batch executions of Unicenter CA-APAS. In some cases, availability from Command Logs may be dependent on the Unicenter CA-APAS ACBX and/or UEX4 having been in use with the Adabas session that produced the Command Log.
COPY	Indicates whether the field or function is available from Unicenter CA-APAS COPY files being processed in stand-alone batch executions of Unicenter CA-APAS. In some cases, availability from COPY files may be dependent on the Unicenter CA-APAS ACBX and/or UEX4 having been in use with the Adabas session with which the COPY file was produced.
File Fmt	This is the default output format for the field or function in the OUTPUT-FILE. This is also the field format as defined in the input Command Log record or the derived field area. The output format may be overridden using the OF parameter in the field specification.
Rept Fmt	This is the default print format for the field in print or terminal reports. This print format may be overridden using the PF parameter in the field specification.
Sum Fld	Indicates whether a summary function such as MIN, MAX, SUM or MEAN, etc. is allowed on this field. In general, summary functions may be applied only to numeric fields.

Heading	Description
Acc Size	This is applicable only to fields that may be summarized. It indicates the type of the accumulator for the field. This is where the sum of the field is accumulated for later calculation of SUM, MEAN, PCT, and RATE functions. It is either a full-word (F) or a double-word (D). Fields that tend to have large values use double-word accumulators to avoid overflow as the sums increase.
Column Heading	The default column heading for the field in reports. A heading element following a "/" character is stacked below the element preceding the "/" character, on the next line of the report or display. This column heading may be overridden using the HD parameter in the field specification.
Values	The possible values for a field and their meanings, also the units in which values are reported.
Derivation	For all fields computed or modified by Unicenter CA-APAS, the method used to derive the field. Any potential problem areas in deriving the field are noted.
Use	Indicates what uses that the field has in interpreting Adabas command processing. This explanation is included to give you a general indication of how the field is used.

Note: Whenever the character "*" appear in an actual description rather than "Yes" or "No", conditional factors affecting the item are explained in the text that follows for the field or summary function.

Summary Fields and Functions: Session or Interval Related

The values assigned to these functions reflect activity over an entire Adabas session when the functions are referenced in requests that do not contain an INTERVAL parameter. When requests do contain an INTERVAL parameter, the values reflect activity which occurred during the specified intervals.

These summary fields and functions may be referenced ONLY within the Summary Function list of SUMMARIZE requests.

Field	Description
ASSO-READS-INTERVAL	Indicates the number of Associator reads during the SUMMARIZE interval, that is, the number of physical I/O operations performed by Adabas during the interval to read blocks from the Associator. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No Yes No No B4 N7 No F **Derivation** This value is computed by subtracting the value of ASSO-READS-SESSION at the end of the interval from the value at the start of the interval. A request added after the start of an MPM session reports all Associator reads occurring before the end of the first interval as occurring in that interval. **Column Heading** ASSO/READS/INTERVAL
ASSO-READS-SESSION	Indicates the number of Associator reads since the start of the current MPM session, that is, the number of physical I/O operations performed by Adabas to read blocks from the Associator. The amount is for the entire session. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No Yes No No B4 N7 No F **Column Heading** ASSO/READS/SESSION

Field	Description
ASSO-WRITES-INTERVAL	Indicates the number of Associator writes during the SUMMARIZE interval, that is, the number of physical I/O operations performed by Adabas during the interval to write blocks to the Associator. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No Yes No No B4 N7 No F **Derivation** This value is computed by subtracting the value of ASSO-WRITES-SESSION at the end of the interval from the value at the start of the interval. A request added after the start of an MPM session reports all Associator writes occurring before the end of the first interval as occurring in that interval. **Column Heading** ASSO/WRITES/INTERVAL
ASSO-WRITES-SESSION	Indicates the number of Associator writes since the start of the MPM session, that is, the number of physical I/O operations performed by Adabas to write blocks to the Associator. The amount is for the entire session. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No Yes No No B4 N7 No F **Column Heading** ASSO/WRITES/SESSION
BFLI	See BUFFER-FLUSHES-INTERVAL.
BFLS	See BUFFER-FLUSHES-SESSION.

Field	Description
BUFFER-FLUSHES-INTERVAL	Indicates the number of buffer flushes during the SUMMARIZE interval, that is, the Adabas count of buffer flushes during the current request interval. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No Yes No No B4 N5 No **Derivation** Adabas increases its buffer flush count before completing the buffer flush. This field does not tell you which Adabas command caused a buffer flush to occur. **Use** The values of BFLI and BFLS can be used to determine times when buffer flush activity is high. A detailed examination of the processing going on at those times may allow you to determine what is actually causing the buffer flush activity. Since this requires "after the fact" analysis of Adabas command processing, we recommend that you use COPY files for this detailed analysis. **Alias** BFLI **Column Heading** BUFF/FLUSH/INTER

Field	Description
BUFFER-FLUSHES-SESSION	Indicates the number of buffer flushes during the current MPM session, that is, the Adabas count of buffer flushes during the current MPM session. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No Yes No No B4 N7 No **Derivation** Adabas increases its buffer flush count before completing the buffer flush. This field does not tell you which Adabas command caused a buffer flush to occur. **Use** The values of the BFLI and BFLS fields can be used to determine times when buffer flush activity is high. A detailed examination of the processing going on at those times may allow you to determine what is actually causing the buffer flush activity. Since this requires analysis of Adabas command processing "after the fact", you should use COPY files for this detailed analysis. **Alias** BFLS **Column Heading** BUFF/FLUSH/SESS

Field	Description
CLOG-READS-INTERVAL	Indicates the number of Command Log reads during the SUMMARIZE interval, that is, the number of physical I/O operations performed by Adabas during the interval to read blocks from the Command Log. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No Yes No No B4 N7 No F **Derivation** This value is computed by subtracting the value of CLOG-READS-SESSION at the end of the interval from the value at the start of the interval. A request added after the start of an MPM session reports all Command Log reads occurring before the end of the first interval as occurring in that interval. **Column Heading** CLOG/READS/INTERVAL
CLOG-READS-SESSION	Indicates number of Command Log reads since the start of the current MPM session, that is, the number of Adabas physical I/O operations to read blocks from the Command Log. The amount is for the entire MPM session. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No Yes No No B4 N7 No F **Column Heading** CLOG/READS/SESSION

Field	Description
CLOG-WRITES-INTERVAL	Indicates number of Command Log writes during the SUMMARIZE interval, that is, the number of physical I/O operations performed by Adabas during the interval to write blocks to the Command Log. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No Yes No No B4 N7 No F **Derivation** This value is computed by subtracting the value of CLOG-WRITES-SESSION at the end of the interval from the value at the start of the interval. A request added after the start of an MPM session reports all Command Log writes occurring before the end of the first interval as occurring in that interval. **Column Heading** CLOG/WRITES/INTERVAL
CLOG-WRITES-SESSION	Indicates number of Command Log writes since the start of the current MPM session, that is, the number of physical I/O operations performed by Adabas to write blocks to the Command Log. The amount is for the entire MPM session. **----- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No Yes No No B4 N7 No F **Column Heading** CLOG/WRITES/SESSION

Field	Description
CMD-RATE	Indicates commands per second, that is, the average number of commands per second issued over the interval of the report. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No Yes Yes P8.3 N5.3 Yes F **Derivation** Computed by taking the number of commands in a given break (COUNT field) and dividing by the length of the interval in seconds. The interval can be the specified interval or the time from the first to the last command if no interval was specified. **Use** Indicates relative saturation of the nucleus during the interval. **Values** In commands per second. **Column Heading** CMDS/PER/SEC

Field	Description
COUNT	Indicates the number of commands in a summary group. This is usually the number of commands at a given break level as categorized by a SUMMARIZE request. In certain special cases, it is the number of times that a field name was found in format or search buffers and may be larger than the number of commands involved. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No Yes Yes P11 N8 Yes F **Derivation** Where this function is a count of commands, it is incremented by one each time a command passed to the Unicenter CA-APAS User Exit 4 routine meets the WHERE selection criteria of the request and falls into the indicated summary group. Where a value of this function represents a total command count for an entire Adabas session, the value reported by Unicenter CA-APAS may be less than the session total command count reported by Adabas. Differences between the value Unicenter CA-APAS reports and the Adabas command count may be due to one of the following factors: ■ Unicenter CA-APAS does not report the "dummy" S1 commands which Unicenter CA-APAS issues solely for communication with the Unicenter CA-APAS Data Collector ■ Adabas does not pass any of its "internal" commands to User Exit 4 ■ Some versions of Adabas do not pass commands from Adabas utilities to User Exit 4 **Use** Basic information for SUMMARIZE requests. **Column Heading** CMD/COUNT

Field	Description
DATA-READS-INTERVAL	Indicates the number of Data Storage reads during the SUMMARIZE interval, that is, the number of physical I/O operations performed by Adabas during the interval to read blocks from Data Storage. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No Yes No No B4 N7 No F **Derivation** This value is computed by subtracting the value of DATA-READS-SESSION at the end of the interval from the value at the start of the interval. A request added after the start of an MPM session reports all Data Storage reads occurring before the end of the first interval as occurring in that interval. **Column Heading** DATA/READS/INTERVAL
DATA-READS-SESSION	Indicates the number of Data Storage reads since the start of the current MPM session, that is, the number of physical I/O operations performed by Adabas to read blocks from Data Storage. The amount is for the entire MPM session. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No Yes No No B4 N7 No F **Column Heading** DATA/READS/SESSION

Field	Description
DATA-WRITES-INTERVAL	Indicates the number of Data Storage writes during the SUMMARIZE interval, that is, the number of physical I/O operations performed by Adabas during the interval to write blocks to Data Storage. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No Yes No No B4 N7 No F **Derivation** This value is computed by subtracting the value of DATA-WRITES-SESSION at the end of the interval from the value at the start of the interval. A request added after the start of an MPM session reports all Data Storage writes occurring before the end of the first interval as occurring in that interval. **Column Heading** DATA/WRITES/INTERVAL
DATA-WRITES-SESSION	Indicates the number of Data Storage writes since the start of the current MPM session, that is, the number of physical I/O operations performed by Adabas to write blocks to Data Storage. The amount is for the entire MPM session. **----- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No Yes No No B4 N7 No F **Column Heading** DATA/WRITES/SESSION

Field	Description
FB-OVERWR	Indicates the number of format overwrites, that is, the number of times Adabas had to overwrite an internal format during the process of translating another external format buffer to an internal format. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No Yes No No B4 N9 No **Use** Helpful in analyzing the adequacy of the size of the internal format pool (ADARUN LFP parameter). Note: An INTERVAL is required in the SUMMERIZE request before the data field is computed. **Column Heading** FB/OVERWRITES
FB-TRANS	Indicates the number of format translations, that is, the number of times Adabas has translated an external format buffer into an internal format within the internal format pool. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No Yes No Yes B4 N9 No **Use** Helpful in analyzing the adequacy of the size of the internal format pool (ADARUN LFP parameter), detecting possibility of external format buffers which translate into excessively long internal formats and the potential or actual payback of using Global Formats. Note: An INTERVAL is required in the SUMMERIZE request before the data field is computed. **Column Heading** FORMAT/BUFF/TRANS

Field	Description
FB-TRAN-%OVRWRT	Indicates relative format overwrite activity, that is, format overwrites are presented as a percentage of format translations. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No Yes No No B2 N3.2 No **Derivation** FB-TRAN -%OVRWRT = (FB-OVERWR/FB-TRANS) X 100. Note that multiple overwrites can result from a single format translation. **Use** Helpful in analyzing efficiency of the format translation process. The smaller the value of this item, the better. Improvement may result from: ■ Increasing the value of the ADARUN parameter, LFP ■ Reducing the lengths of internal format buffers Note: An INTERVAL is required in the SUMMERIZE request before the data field is computed. **Column Heading** FB-TRAN/%OVRWRT

Field	Description
FB-%TRAN	Indicates relative format translation activity, that is, format translations are presented as a percentage of commands that use format buffers. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No Yes No No B2 N3.2 No **Derivation** FB-%TRAN = (FB-TRANS/SUM(FB-CMD)) X 100 **Use** Helpful in analyzing efficiency of the format translation process. The smaller the value of this item, the better. Improvement may result from: ■ Increasing the value of the ADARUN LFP parameter. ■ Using Adabas Global Formats. ■ Reducing the lengths of internal format buffers. Note: An INTERVAL is required in the SUMMERIZE request before the data field is computed. **Column Heading** FB/%TRAN
IO-BUFF-EFF	Indicates buffering efficiency for the SUMMARIZE interval as the ratio of the number of read accesses to the Adabas I/O buffer pool to the number of physical reads to Associator and Data Storage blocks. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No Yes No No B4 N6.1 No F **Derivation** IO-BUFF-EFF= 10 * LOGICAL-IO/(ASSO-READS-INTERVAL + DATA-READS-INTERVAL) **Use** Helpful in optimizing the size of the Adabas I/O buffer pool (ADARUN LBP parameter). **Column Heading** IO/BUFF/EFF

Field	Description
IO-BUFF-EFF-SESSION	Indicates I/O buffering efficiency for the MPM session as the ratio of the number of read accesses to the Adabas I/O buffer pool to the number of physical reads to Associator and Data Storage blocks. All counts are since the start of the current MPM session. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No Yes No No B4 N6.1 No F **Derivation** IO-BUFF-EFF-SESSION = 10 * LOGICAL-IO (for the session) /(ASSO-READS-SESSION+DATA-READS-SESSION) **Use** Helpful in optimizing the size of the Adabas I/O buffer pool (ADARUN LBP parameter). **Column Heading** IO BUFF/EFF/SESSION
JOB-END	Indicates the time of the user's last Adabas command. The field is identical to specifying MAX(TIME), except for the label. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No Yes Yes N6 A6 No **Derivation** MAX(TIME) **Values** hhmmss (hour - minute - second) **Column Heading** JOB/END/TIME

Field	Description
JOB-START	Indicates the time of the user's first Adabas command. This field is identical to specifying MIN(TIME), except for the label. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No Yes Yes N6 A6 No **DERIVATION** MIN(TIME) **Values** hhmmss (hour - minute - second) **Column Heading** JOB/START/TIME
LOGICAL-IO	Indicates the number of Adabas logical I/Os, that is, the number of times Adabas found a needed Associator or Data Storage block already in the I/O buffer pool and thereby avoided a physical I/O during the interval. Adabas logical I/Os and physical I/Os are mutually exclusive. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No Yes No No B4 N9 No **Use** Helpful in analyzing optimum size of the ADARUN LBP parameter; can also be helpful in analyzing optimum mixes of application jobs to minimize physical I/O and elapsed processing time. **Column Heading** BUFFER/ACCESSES

Field	Description
LOGICAL-IO-SESSION	Indicates the number of Adabas logical I/Os, that is, the number of times Adabas found a needed Associator or Data Storage block already in the I/O buffer pool and thereby avoided a physical I/O for the entire MPM session. Adabas logical I/Os and physical I/Os are mutually exclusive. All counts are since the start of the current MPM session. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No Yes No No B4 N9 No **Use** Helpful in analyzing optimum size of the ADARUN LBP parameter; can also be helpful in analyzing optimum mixes of application jobs to minimize physical I/O and elapsed processing time. **Column Heading** BUFFER/ACCESSES/SESSION
MAX(field)	The MAXIMUM value summary function, which is, the maximum value of the summarized field for the current break (BY) level. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No Yes Yes No **Derivation** The largest value in this break group of records. **Use** Use to search for extreme values of a performance variable. **Values** The file and print formats are identical to the summarized field. **Column Heading** MAX/'field heading'

Field	Description
MEAN(field)	The MEAN value summary function, which is, the mean (average) value of the summarized field for the current break (BY) level. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No Yes Yes Pn+.2 Nn+.2 No **Derivation** The sum of the field in this break group of records divided by the number of records (COUNT). **Use** Use to look for an indication of typical performance or load. Both the SORT specification and the graph print format (PF=Gnn) give invalid results for this function. **Values** The default print format is equal to the digits of the summarized field plus two decimal places (column width is three larger). The output field format is packed, with two digits of decimal precision. **Column Heading** MEAN/'field heading'

Field	Description
MIN(field)	The MINIMUM value summary function, which is, the minimum value of the summarized field for the current break (BY) level. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No Yes Yes No **Derivation** The smallest value in this break group of records. **Use** Use to search for extreme values of a performance variable. **Values** The file and print formats are identical to the summarized field. **Column Heading** MIN/'field heading'

Field	Description
NAT-LOAD-COUNT	Indicates the number of times a Natural module is invoked, that is, the count of the number of times a Natural module has been invoked by an EXECUTE or FETCH. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes Yes Yes No Yes B4 N4 Yes F **Derivation** This count is incremented every time the current module name changes for a given user. The count is not incremented if a module fetches itself. If the ACBX feature is in use, then changes in module name are detected based on names that are set by Natural. This includes invocations from the Natural buffer pool. Note, however, that experience seems to suggest that the names set by Natural do not always accurately reflect actual operations; in these cases, values for this field usually are useful, though not perfect. **Use** Used only in SUMMARIZE statements by NAT-MODULE to show how often the module was executed. This can be used to estimate commands or I/O per invocation of the module. **Column Heading** NAT/LOAD/CNT

Field	Description
PCT(field)	The PERCENT summary function, which is, the percent of the summarized field for this break group by the total for the entire log or interval. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No Yes Yes P4.1 N3.1 No **Derivation** The percent of the summarized field is computed by taking the sum of the field value for the current break level and dividing it by the sum of the field for the entire interval (or log, if no interval is specified) and multiplying by 100. **Use** As an indication of the importance of a given break level. **Column Heading** PCT/'field heading'
PERCENT	Indicates percent of total commands in summary group, that is, the percent of all commands in the current SUMMARIZE table that fall in the current break level. This is equivalent to PCT(COUNT). **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No Yes Yes P4.1 N3.1 Yes F **Use** Basic SUMMARIZE information. **Column Heading** PCT/TOTAL/CMDS

Field	Description
RATE(field)	The RATE summary function, which is, the average rate per second of the summarized field during the reported time period. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No Yes Yes Pn+.2 Nn+.2 No **Derivation** The sum of the field in this break group of records divided by the number of seconds in the period being reported. This gives a per-second value for the field. For requests that contain an INTERVAL parameter, the period being reported is the length of the specified interval. For requests with no INTERVAL parameter, the period reported is the time between the first and latest Command Log records selected for the request. **Use** Indicates the intensity of use of resources. **Values** The default print format is equal to three more digits than the summarized field plus two decimal places (column width is six larger). The output file format is packed, with the precision of the summarized field plus three, and two digits of decimal fraction. **Column Heading** RATE/'field heading'

Field	Description
SIBA-READS-INTERVAL	Indicates number of Protection Log reads during the SUMMARIZE interval, that is, the number of physical I/O operations performed by Adabas during the interval to read blocks from the Protection Log. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No Yes No No B4 N7 No F **Derivation** This value is computed by subtracting the value of SIBA-READS-SESSION at the end of the interval from the value at the start of the interval. A request added after the start of an MPM session reports all Protection Log reads occurring before the end of the first interval as occurring in that interval. **Column Heading** SIBA/READS/INTERVAL
SIBA-READS-SESSION	Indicates number of Protection Log reads since the start of the current MPM session, that is, the number physical I/O operations performed by Adabas to read blocks from the Protection Log. The amount is for the entire MPM session. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No Yes No No B4 N7 No F **Column Heading** SIBA/READS/SESSION

Field	Description
SIBA-WRITES-INTERVAL	Indicates the number of Protection Log writes during the SUMMARIZE interval, that is, the number of physical I/O operations performed by Adabas during the interval to write blocks to the Protection Log. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No Yes No No B4 N7 No F **Derivation** This value is computed by subtracting the value of SIBA-WRITES-SESSION at the end of the interval from the value at the start of the interval. A request added after the start of an MPM session reports all Protection Log writes occurring before the end of the first interval as occurring in that interval. **Column Heading** SIBA/WRITES/INTERVAL
SIBA-WRITES-SESSION	Indicates the number of Protection Log writes, that is, the number of physical I/O operations performed by Adabas to write blocks to the Protection Log. The amount is for the entire MPM session. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No Yes No No B4 N7 No F **Column Heading** SIBA/WRITES/SESSION
SUM(field)	The SUM summary function, which is, the sum of the summarized field for the current break level. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No Yes Yes No **Values** The default print and output file formats vary depending on the format of the field being summarized. **Column Heading** SUM/'field heading'

Field	Description
THREAD-CHANGES	The number of Adabas switches between threads, that is, the number of times Adabas has changed from processing a command in one of its threads to processing a different command in a different thread. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No Yes No No B4 N9 No **Column Heading** THREAD/CHANGES
THREADnn-CMDS-INTERVAL	Indicates the number of commands processed in thread nn during the SUMMARIZE interval. There are multiple summary function names of this form, where nn ranges from 01 to 56 as well as A1, A2, A3, and A4. Threads A1, A2, A3, and A4 are threads used by internal Adabas commands. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No Yes No No B4 N9 No **Column Heading** THREADnn/CMDS/INTERVAL
THREADnn-CMDS-SESSION	Indicates the number of commands processed in thread nn during the current MPM session. There are multiple summary function names of this form, where nn ranges from 01 to 56 as well as A1, A2, A3, and A4. Threads A1, A2, A3, and A4 are used by internal Adabas commands. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No Yes No No B4 N9 No **Column Heading** THREADnn/CMDS/SESSION

Field	Description
THROW-BACKS	The number of times Adabas has accomplished some amount of processing of a command, found that some resource required for continued processing is unavailable, and has discontinued processing of the command, in effect, "throwing the command back into the Command Queue" to later re-process from the beginning. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No Yes No No B4 N7 No **Derivation** In Adabas V7, this value is the sum of nucleus counters for ISN and space related throwbacks at the end of the interval from the sum at the start of the interval. A request added after the start of an MPM session reports all throwbacks occurring before the end of the first interval as occurring in that interval. **Column Heading** THROW/BACKS
THROW-BACKS-ISN	The number of times Adabas has accomplished some amount of processing of a command and had to throwback the command because of ISN contention. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** * No Yes No No B4 N7 No **Derivation** This value is computed by subtracting the nucleus counter of ISN related throwbacks at the end of the interval from the value at the start of the interval. A request added after the start of an MPM session reports all ISN related throwbacks occurring before the end of the first interval as occurring in that interval. Note: Only available with Adabas V7. **Column Heading** THROW/BACKS/ISN

Field	Description
THROW-BACKS-SPACE	The number of times Adabas has accomplished some amount of processing of a command and had to throwback the command because of work pool space (LWP) not available. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** * No Yes No No B4 N7 No **Derivation** This value is computed by subtracting the nucleus counter of space related throwbacks at the end of the interval from the value at the start of the interval. A request added after the start of an MPM session reports all space related throwbacks occurring before the end of the first interval as occurring in that interval. Note: Only available with Adabas V7. **Column Heading** THROW/BACKS/SPACE

Field	Description
TP-TRANS-COUNT	Indicates the number of TP transactions, that is, the approximate number of TP transactions that issued Adabas commands. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes * No * Yes B4 N5 Yes F **Derivation** With the ACBX facility installed, this field is incremented by 1 each time the TP-TRANS-ID changes for a given TP user. In CICS it is not valid if the CICS application is running in conversational mode; otherwise, values appear to be satisfactory. Experience has shown that this approach does not yield accurate values for versions of COM-PLETE that do not generate suitable values of TP-TRANS-ID. Without the ACBX facility, this field is incremented by 1 each time the length of time specified for TP-TRANS-GAP is exceeded between two consecutive Adabas commands from a given user. This approach is imprecise, at best; you may or may not find that it yields useful values in your environment. **Use** Use this field to estimate the number of commands that are issued by a TP transaction program between interactions with the terminal/user. If this number is large, it may indicate applications that are poorly designed for interactive systems. **Values** Values reported for this field are usually less than the actual numbers of times users pressed ENTER or RETURN keys. This is primarily because transactions that do not involve calls to Adabas are not subject to monitoring by Unicenter CA-APAS. **Column Heading** TP/TRANS/CNT

Field	Description
WORK-READS-INTERVAL	Indicates the number of Work data set reads during the SUMMARIZE interval, that is, the number of physical I/O operations performed by Adabas during the interval to read blocks from the Work data set. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No Yes No No B4 N7 No F **Derivation** This value is computed by subtracting the value of WORK-READS-SESSION at the end of the interval from the value at the start of the interval. A request added after the start of an MPM session reports all Work data set reads occurring before the end of the first interval as occurring in that interval. **Column Heading** WORK/READS/INTERVAL
WORK-READS-SESSION	Indicates the number of Work data set reads during the current MPM session, that is, the number of physical I/O operations performed by Adabas to read blocks from its Work data set. The amount is for the entire MPM session. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No Yes No No B4 N7 No F **Column Heading** WORK/READS/SESSION

Field	Description
WORK-WRITES-INTERVAL	Indicates the number of Work data set writes during the SUMMARIZE interval, that is, the number of physical I/O operations performed by Adabas during the interval to write blocks to the Work data set. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No Yes No No B4 N7 No F **Derivation** This value is computed by subtracting the value of WORK-WRITES-SESSION at the end of the interval from the value at the start of the interval. a request added after the start of an MPM session reports all Work data set writes occurring before the end of the first interval as occurring in that interval. **Column Heading** WORK/ WRITES/INTERVAL
WORK-WRITES-SESSION	Indicates the number of Work data set writes during the current MPM Session, that is, the number of physical I/O operations performed by Adabas to write blocks to its Work data set. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No Yes No No B4 N7 No F **Column Heading** WORK/WRITES/SESSION

Summary Fields and Functions: Instantaneous

The values assigned to these functions reflect conditions at a specific instant of time, and not over an interval or session. In fact, the values of some of these functions never change during an Adabas session. Values are computed at the ends of intervals for the requests in which these functions are referenced.

None of these instantaneous fields and functions have logical significance for any BY fields for which they may be reported since they are not maintained in summary table entries. This should particularly be remembered when reporting BY time periods. At each output point in time, a new current value is developed for each instantaneous field or function; the new value is then reported for all BY entries, even if they represent past time periods for which different values of the instantaneous field or function have been reported earlier.

The availability of these summary functions may vary from release to release of Adabas.

Instantaneous summary fields and functions cannot be the first field requested in an Unicenter CA-APAS summary request! At least one session or interval-related summary field or function must precede them.

Note: These functions may be specified ONLY within the Summary Function lists of Unicenter CA-APAS SUMMARIZE requests.

Field	Description
AB-ALLOC	Indicates the size of the Adabas Attached Buffer area as bytes of storage allocated. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No Yes No No B4 N7 No **Column Heading** AB/SPACE/ALLOC
AB-USED	Indicates the high-water mark for Attached Buffer usage, that is, the maximum number of bytes of Adabas Attached Buffers used at any point so far during the Adabas session. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No Yes No No B4 N7 No **Column Heading** AB/SPACE/USED

Field	Description
BUFF-POOL-SIZE	Indicates the size of the Adabas I/O Buffer Pool as bytes of storage allocated. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No Yes No No B4 N9 No **Column Heading** BUFF/POOL/SIZE
BUFF-POOL-%ACT	The percentage of Adabas I/O Buffer Pool containing active blocks. Adabas marks blocks as active while they are involved in certain stages of updating activity. Access by other users to active blocks is restricted. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No Yes No No B4* N3.1 No * The file output is multiplied by 10 to give tenths of a percent. For example, x'00000121' equals 28.9. **Column Heading** BUFF/POOL/%ACT
BUFF-POOL-%ASSO	The percentage of Adabas I/O Buffer Pool containing Associator blocks. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No Yes No No B4* N3.1 No * The file output is multiplied by 10 to give tenths of a percent. For example, x'00000121' equals 28.9. **Column Heading** BUFF/POOL/%ASSO

Field	Description
BUFF-POOL-%AVAIL	The percentage of Adabas I/O Buffer Pool available. The percentage of the Adabas I/O Buffer Pool that consists of elements available to have new Associator or Data blocks read into them. Elements containing blocks that have been updated but not yet written back to disk cannot be overlaid by new blocks, so any such elements are not yet available. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No Yes No No B4* N3.1 No * The file output is multiplied by 10 to give tenths of a percent. For example, x'00000121' equals 28.9. **Column Heading** BUFF/POOL/%AVAIL
BUFF-POOL-%DATA	The percentage of Adabas I/O Buffer Pool containing Data blocks. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No Yes No No B4* N3.1 No * The file output is multiplied by 10 to give tenths of a percent. For example, x'00000121' equals 28.9. **Column Heading** BUFF/POOL/%DATA
BUFF-POOL-%USED	The percentage of Adabas I/O Buffer Pool in use. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No Yes No No B4* N3.1 No * The file output is multiplied by 10 to give tenths of a percent. For example, x'00000121' equals 28.9. **Derivation** The number of bytes of ASSO or DATA blocks/BUFF-POOL-SIZE. Always less than 100 percent because of unused bytes due to fit of blocks into elements and because of bytes used for control information. **Column Heading** BUFF/POOL/%USED

Field	Description
BUFF-POOL-%WRIT	The percentage of the Adabas I/O Buffer Pool which contains ASSO or DATA blocks which are marked for writing to disk before they can be overlaid by new blocks that need to be read in to the Buffer Pool. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No Yes No No B4* N3.1 No * The file output is multiplied by 10 to give tenths of a percent. For example, x'00000121' equals 28.9. **Column Heading** BUFF/POOL/%WRIT
CQE-ALLOC	Indicates the size of the Adabas Command Queue as bytes of storage allocated. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No Yes No No B4 N7 No **Column Heading** CQE/SPACE/ALLOC
CQE-USED	Indicates high-water mark for Command Queue usage, that is, the maximum number of bytes of the Adabas Command Queue used at any point so far during the Adabas session. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No Yes No No B4 N7 No **Column Heading** CQE/SPACE/USED
CURR-ECBS	Indicates the number of ECBS currently posted. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No Yes No No B4 N3 No **Column Heading** POSTED/ECBS

Field	Description
HELD-RECORDS	Indicates the number of records held for update. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No Yes No No B4 N5 No **Derivation** The number of active entries in the Adabas Hold Queue. **Use** Helpful in evaluating the adequacy of the size of the Hold Queue. Possibly helpful in detecting excessive holding of records by application programs. **Column Heading** TOTAL/RECORDS/HELD
HQ-ALLOC	Indicates the size of the Adabas Hold Queue as bytes of storage allocated. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No Yes No No B4 N7 No **Column Heading** HQ/SPACE/ALLOC
HQ-USED	Indicates high-water mark for Hold Queue usage as the maximum number of bytes of the Adabas Hold Queue used at any point so far during the Adabas session. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No Yes No No B4 N7 No **Column Heading** HQ/SPACE/USED

Field	Description
IF-POOL-SIZE	Indicates the size of the Adabas Internal Format Pool. This is the equivalent to the LFP summary function described below. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No Yes No No B4 N9 No **Column Heading** IFB/SIZE
IF-POOL-%USED	Indicates the percentage of the Adabas Internal Format Pool that contains active internal formats. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No Yes No No B4* N3.1 No * The file output is multiplied by 10 to give tenths of a percent. For example, x'00000121' equals 28.9. **Column Heading** IFB/%USED
IF-USED	Indicates the high-water mark for Internal Format Pool usage as the maximum number of bytes of the Adabas Internal Format Pool used at any point so far during the Adabas session. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No Yes No No B4 N7 No **Column Heading** IF/USED
LFP	Indicates the size of the Adabas Internal Format Pool as bytes of storage allocated. This is the equivalent to the IF-POOL-SIZE summary function described above. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No Yes No No B4 N7 No **Column Heading** LFP

Field	Description
MAX-HELD-RECORDS	Indicates the maximum number of records held by any user. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No Yes No No B4 N5 No **Column Heading** MAX/RECORDS/HELD
MAX-HOLDER-COMMID	Indicates the 28-byte Communication ID field of the user who has more records held for update than any other user. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No Yes No No A28 A28 No **Column Heading** COMMID/HOLDER/MAX
MAX-HOLDER-CPUID	Indicates the CPU ID of the user who has more records held for update than any other user. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No Yes No No A8 A8 No **Column Heading** CPU-ID/HOLDER/MAX
MAX-HOLDER-ID	Indicates the Adabas ID of the user who has more records held for update than any other user. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No Yes No No A8 A8 No **Column Heading** USERID/HOLDER/MAX
MAX-HOLDER-ID-V5	Obsolete. See MAX-HOLDER-COMMID.

Field	Description
MAX-HOLDER-IDPT1	Indicates the leftmost four (4) bytes of MAX-HOLDER-ID. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No Yes No No A4 A4 No **Column Heading** USERID/HOLDER/MAX PT1
MAX-HOLDER-IDPT2	Indicates the rightmost four (4) bytes of MAX-HOLDER-ID. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No Yes No No A4 A4 No **Column Heading** USERID/HOLDER/MAX PT2
MAX-HOLDER-JOB	Indicates the jobname of the user who has more records held for update than any other user. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No Yes No No A8 A8 No **Column Heading** JOBNAME/HOLDER/MAX
MAX-HOLDER-OSID	Indicates the OS ID of the user who has more records held for update than any other user. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No Yes No No A4 A4 No **Column Heading** OS-ID/HOLDER/MAX
MAX-HOLDER-TID	Indicates the TID of the user who has more records held for update than any other user. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No Yes No No A8 A8 No **Column Heading** TID/HOLDER/MAX

Field	Description
MAX-HOLDER-VMID	Indicates the VM ID of the user who has more records held for update than any other user. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No Yes No No A8 A8 No **Column Heading** VM-ID/HOLDER/MAX
NO-USERS	Indicates the current number of Adabas users, using the number of entries in the Adabas User Queue at the end of the interval. Includes entries for active users and any entries being kept only for possibly sending response code 9 to inactive users who have exceeded time limits. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No Yes No No B4 N4 No **Column Heading** NO/USERS
SC-ALLOC	Indicates the size of the Adabas Security Pool as the bytes of storage allocated. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No Yes No No B4 N7 No **Column Heading** SC/SPACE/ALLOC
SC-USED	Indicates the high-water mark for Security Pool usage as the maximum number of bytes of the Adabas Security Pool used at any point so far during the Adabas session. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No Yes No No B4 N7 No **Column Heading** SC/SPACE/USED

Field	Description
SESS-MAX-CMD-COUNT	Indicates the maximum command count for any user at the present time during the Adabas session. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No Yes No No B4 N9 No **Column Heading** SESSION/MAX CMD/COUNT
SESS-MAX-CMD-USERID	Indicates the ID of user with max command count, that is, the ID of the user with the maximum number of commands so far during the Adabas session. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No Yes No No A8 A8 No **Column Heading** SESSION/MAX CMD/USERID
SESS-MAX-CPU-MINUTES	Indicates the CPU minutes for the user with maximum CPU time, that is, the number of minutes of CPU time used by the user who has used the maximum amount of CPU time within Adabas so far during the Adabas session. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No Yes No No B4 N8 No **Column Heading** SESSION/MAX CPU/MINUTES
SESS-MAX-CPU-SECONDS	Indicates CPU seconds for the user with maximum CPU time, that is, the number of seconds, in addition to minutes, of CPU time used by the user who has used the maximum amount of CPU time within Adabas so far during the Adabas session. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No Yes No No B4 N3 No **Column Heading** SESSION/MAX CPU/SECONDS

Field	Description
SESS-MAX-CPU-USERID	Indicates ID of user with maximum CPU time, that is, the ID of the user with the maximum amount of CPU usage within Adabas so far during the Adabas session. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No Yes No No A8 A8 No **Column Heading** SESSION/MAX CPU/USERID
SESS-MAX-IO-COUNT	Indicates maximum I/O count for any user, that is, the maximum number of I/Os for any user so far during the Adabas session. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No Yes No No B4 N9 No **Column Heading** SESSION/MAX IO/COUNT
SESS-MAX-IO-USERID	Indicates ID of user with maximum I/O count, that is, the ID of the user with the maximum number of I/Os so far during the Adabas session. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No Yes No No A8 A8 No **Column Heading** SESSION/MAX IO/USERID
SESS-USER-COUNT	Indicates total number of users, that is, the total number of users so far during the Adabas session. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No Yes No No B4 N8 No **Column Heading** SESSION/USER/COUNT

Field	Description
TBI-ALLOC	Indicates the size of the Adabas Table of ISN Lists as bytes of storage allocated. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No Yes No No B4 N7 No **Column Heading** TBI/SPACE/ALLOC
TBI-USED	Indicates the high-water mark for Table of ISN Lists usage as the maximum number of bytes of the Adabas Table of ISN Lists used at any point so far during the Adabas session. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No Yes No No B4 N7 No **Column Heading** TBI/SPACE/USED
TBS-ALLOC	Indicates the size of the Adabas Table of Sequential Commands as bytes of storage allocated. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No Yes No No B4 N7 No **Column Heading** TBS/SPACE/ALLOC
TBS-USED	Indicates the high water mark for the Table of Sequential Commands as the maximum number of bytes of the Adabas Table of Sequential Commands used at any point so far during the Adabas session. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No Yes No No B4 N7 No **Column Heading** TSB/SPACE/USED

Field	Description
UQ-ALLOC	Indicates the size of the Adabas User Queue as bytes of storage allocated. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No Yes No No B4 N7 No **Column Heading** UQ/SPACE/ALLOC
UQ-USED	Indicates high-water mark for User Queue usage as the maximum number of bytes of the Adabas User Queue used at any point so far during the Adabas session. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No Yes No No B4 N7 No **Use** SUM (UQE-COST) may be used to report charges for Adabas usage. **Column Heading** UQ/SPACE/USED
WORK-ALLOC	Indicates the size of the Adabas Work Pool as bytes of storage allocated. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No Yes No No B4 N7 No **Column Heading** WORK/SPACE/ALLOC
WORK-POOL-SIZE	Indicates the size of the Adabas Work Pool (ADARUN LWP parameter) as bytes of storage allocated. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No Yes No No B4 N9 No **Column Heading** WORK/POOL/SIZE

Field	Description
WORK-POOL-%USED	Indicates percentage of the Adabas Work Pool that is currently involved in processing of active commands. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No Yes No No B4* N3.1 No * The file output is multiplied by 10 to give tenths of a percent. For example, x'00000121' equals 28.9. **Column Heading** WORK/POOL/%USED
WORK-USED	Indicates high-water mark for Work Pool usage as the maximum number of bytes of the Adabas Work Pool used at any point so far during the Adabas session. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No Yes No No B4 N7 No **Column Heading** WORK/SPACE/USED

Selected ADARUN Parameters

Unicenter CA-APAS collects and reports some of the Adabas ADARUN initialization parameters. In this section we list the parameters we currently report. See the appropriate Adabas operations manual for detailed explanations of the ADARUN parameters listed below. The Unicenter CA-APAS names for all of these parameters begin with the prefix AP-. Some have aliases.

Like the instantaneous summary functions described immediately previous, these values are assigned to the various parameters at the ends of intervals specified in requests in which the parameters appear.

Parameter	-------- Adabas V7 ------					File	Rept	Sum	Acc
	Gvn	ACBX	UX4	Log	COPY	Fmt	Fmt	Fld	Size
AP-ASTAR	Yes	No	Yes	No	No	A1	A1	No	
AP-CLOGV	Yes	No	Yes	No	No	A1	A1	No	
AP-CT	Yes	No	Yes	No	No	B4	N8	No	
AP-DEVICE	Yes	No	Yes	No	No	A4	A4	No	
AP-IGNDIB	Yes	No	Yes	No	No	A1	A1	No	
AP-LBP	Yes	No	Yes	No	No	B4	N9	No	
AP-LCP	Yes	No	Yes	No	No	B4	N8	No	
AP-LFIOP	Yes	No	Yes	No	No	B4	N8	No	
AP-LFP	Yes	No	Yes	No	No	B4	N8	No	
AP-LGBP	Yes	No	Yes	No	No	A8	A8	No	
AP-LI	Yes	No	Yes	No	No	B4	N8	No	
AP-LOGCB	Yes	No	Yes	No	No	A1	A1	No	
AP-LOGFB	Yes	No	Yes	No	No	A1	A1	No	
AP-LOGGING	Yes	No	Yes	No	No	A1	A1	No	
AP-LOGIB	Yes	No	Yes	No	No	A1	A1	No	
AP-LOGIO	Yes	No	Yes	No	No	A1	A1	No	
AP-LOGRB	Yes	No	Yes	No	No	A1	A1	No	

Parameter	Adabas V7 Gvn	ACBX	UX4	Log	COPY	File Fmt	Rept Fmt	Sum Fld	Acc Size
AP-LOGSB	Yes	No	Yes	No	No	A1	A1	No	
AP-LOGUX	Yes	No	Yes	No	No	A1	A1	No	
AP-LOGVB	Yes	No	Yes	No	No	A1	A1	No	
AP-LP	Yes	No	Yes	No	No	B2	N5	No	
AP-LQ	Yes	No	Yes	No	No	B4	N8	No	
AP-LS	Yes	No	Yes	No	No	B4	N8	No	
AP-LU	Yes	No	Yes	No	No	B4	N8	No	
AP-LWKP2	Yes	No	Yes	No	No	B2	N5	No	
AP-LWP	Yes	No	Yes	No	No	B4	N8	No	
AP-MODE	Yes	No	Yes	No	No	A1	A1	No	
AP-NAB	Yes	No	Yes	No	No	B4	N5	No	
AP-NC	Yes	No	Yes	No	No	B2	N5	No	
AP-NH	Yes	No	Yes	No	No	B4	N6	No	
AP-NISNHQ	Yes	No	Yes	No	No	B4	N5	No	
AP-NONDES	Yes	No	Yes	No	No	A1	A1	No	
AP-NQCID	Yes	No	Yes	No	No	B4	N8	No	
AP-NSISN	Yes	No	Yes	No	No	B2	N6	No	
AP-NT	Yes	No	Yes	No	No	B2	N3	No	
AP-NU	Yes	No	Yes	No	No	B4	N8	No	
AP-OPENRQ	Yes	No	Yes	No	No	A1	A1	No	
AP-PLOGNO **Alias:** PLOG, PLOG-NO, PLOG-NUMBER	Yes	No	Yes	No	No	B2	N5	No	
AP-PLOGR	Yes	No	Yes	No	No	A1	A1	No	

Parameter	Adabas V7 Gvn	ACBX	UX4	Log	COPY	File Fmt	Rept Fmt	Sum Fld	Acc Size
AP-PREFETCH	Yes	No	Yes	No	No	A1	A1	No	
AP-PREFSBL	Yes	No	Yes	No	No	B2	N5	No	
AP-PREFTBL	Yes	No	Yes	No	No	B4	N8	No	
AP-READONLY	Yes	No	Yes	No	No	A1	A1	No	
AP-SMPID **Alias:** SMPID	Yes	No	Yes	No	No	B2	N5	No	
AP-SVC	Yes	No	Yes	No	No	B1	N3	No	
AP-TLSCMD	Yes	No	Yes	No	No	B4	N8	No	
AP-TNAA	Yes	No	Yes	No	No	B4	N8	No	
AP-TNAE	Yes	No	Yes	No	No	B4	N8	No	
AP-TNAX	Yes	No	Yes	No	No	B4	N8	No	
AP-TT	Yes	No	Yes	No	No	B4	N8	No	
AP-UEX1	Yes	No	Yes	No	No	A8	A8	No	
AP-UEX2	Yes	No	Yes	No	No	A8	A8	No	
AP-UEX3	Yes	No	Yes	No	No	A8	A8	No	
AP-UEX4	Yes	No	Yes	No	No	A8	A8	No	
AP-UEX5	Yes	No	Yes	No	No	A8	A8	No	
AP-UEX6	Yes	No	Yes	No	No	A8	A8	No	
AP-UEX7	Yes	No	Yes	No	No	A8	A8	No	
AP-UEX8	Yes	No	Yes	No	No	A8	A8	No	
AP-UEX9	Yes	No	Yes	No	No	A8	A8	No	

Unicenter CA-APAS Data Fields

Values for these fields pertain to the processing of individual commands, as opposed to the summary functions described earlier that usually reflect activity across the processing of multiple commands. Many of these field values are taken directly from the Command Log record image that Adabas creates for each command when the appropriate Adabas logging parameters are ON. Some fields are derived from other sources.

Field	Description
ACB12	Indicates the first two bytes of the Adabas Control Block (ACB). **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No Yes No Yes A2 A2 No **Derivation** Directly from the ACB. **Use** X'00' in byte 1 indicates a logical database call. Certain values in byte 2 indicate that the call came from Natural. **Column Heading** ACB/BYTES/1+2
ADABAS-DATE	Indicates the assembly date of the Adabas nucleus, which can be used to differentiate between releases of Adabas. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No Yes No No A6 A6 No **Derivation** The assembly date is extracted from the Adabas nucleus, ADANC0, and put into the character string in the form YYMMDD. This could be useful during migration of applications between release levels of Adabas. **Column Heading** ADABAS/DATE

Field	Description
ADABAS-VERSION	The Adabas version number. This field identifies the Adabas version of the MPM. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No Yes No No A3 A3 No **Derivation** Determined at Unicenter CA-APAS initialization by seeing which Adabas nucleus is running. As the Data Collector is being initialized under the Adabas User Exit 4, it sets this field to either a C'71x', C'72x', or C'74x'. **Use** This field can be used to identify the version of MPM being used. This could be useful during migration of applications between versions of Adabas. **Column Heading** ADABAS/VERSION

Field	Description

ADDITIONS-1	The ADDITIONS-1 field from the Adabas Control Block contains different kinds of information for different command codes; see the *Adabas Command Reference* manual.
	
	------- Adabas V7 ----- File Rept Sum Acc
	Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size
	Yes No No Yes Yes A8 A8 No
	
	Use
	
	Important in detailed analysis and exception reporting of S2 and S9 commands (see below).
	
	Any analysis of descriptor use must take these values into account in addition to analyzing search buffers (see DESCRIPTOR) and logical sequential reading (see L3-SEQUENCE). Use of descriptors for sorting may be reported by a request of the following form:
	`SUMMARIZE COUNT`
	`BY FILE BY ADDITIONS-1`
	`WHERE (CMD = 'S2' OR = 'S9') . . . ;`
	
	The application external user-id must be determined when tracing ET-data use. These IDs may be listed from the Adabas checkpoint file or may be reported by an Unicenter CA-APAS request of the following form:
	`SUMMARIZE COUNT`
	`BY ADDITIONS-1 . . .`
	`WHERE (CMD = 'OP' AND ADDITIONS-1 NE ' ')`
	`. . . ;`
	
	Since the usage of this field differs according to the command code, requests should include WHERE conditions that exclude command codes not desired for a particular interpretation of the field.
	
	Alias
	ADDIT1, APPL-USER-ID
	
	Column Heading
	ADDIT 1

Field	Description
ADDITIONS-2	The ADDITIONS-2 field from the Adabas Control Block contains different kinds of information for different command codes. For more information, see the *Adabas Command Reference* manual.

------- Adabas V7 -----					File	Rept	Sum	Acc
Gvn	ACBX	UX4	Log	COPY	Fmt	Fmt	Fld	Size
Yes	No	No	Yes	Yes	B4	A4	No	

Use

See FB-ERR-FIELD, REC-LENGTH-DEC, and REC-LENGTH-COMP for descriptions of other fields that are derived from ADDITIONS-2.

Since the usage of this field differs according to the command code, requests should include WHERE conditions that exclude command codes not desired for a particular interpretation of the field.

Values

To avoid reporting meaningless values, which are sometimes present in this field in the Adabas Control Block, Unicenter CA-APAS sometimes reports binary zeroes for all or a portion of this field. If the Adabas response code is equal to 9, this field is not reset to zero. Otherwise, the field is set to binary zeroes as explained in the table below. Possible outcomes are posted as:

- A = the entire field is set to zero
- B = the second half of the field is set to zero.

CMD Code	Response Code = 0	Response Code = 1 - 8, 10 - 255
A1/A4	B	A
BT	See Note 1	See Note 1
CL	A	A
C1	A	A
C2	A	A
C3	A	A
C5	A	A

Field	Description		
	ET	A	A
	E1/E4		A
	HI		A
	LF	A	A
	L1/L4		A
	L2/L5		A
	L3/L6		A
	L9	A	A
	N1/N2	B	A
	OP	A	A
	RC	A	A
	RE	A	A
	RI	A	A
	S1/S2/S3	See Note 2	A
	S5	S	A
	S8		A
	S9		A

Note 1: Only set to zero when Command Option 2 is not equal to 'F'.

Note 2: Set to zero only if there is a null format buffer (one that begins with a period, .) in the Command Log record.

Note 3: The reader is reminded that if the response code is nine (9), the field is not reset to zero (0).

Alias
ADDIT2

Column Heading
ADD 2

Field	Description
ADDITIONS-2-REPORTED	The ADDITIONS-2 field from the Adabas Control Block contains different kinds of information for different command codes. For more information, see the *Adabas Command Reference Manual*. The difference between this field and the ADDITIONS-2 field is that this field is not altered as described in the ADDITIONS-2 field description. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No Yes Yes B4 A4 No **Use** See the ADDITIONS-2 field description. **Alias** ADDIT2-REPORTED **Column Heading** ADD 2/REPTD

Field	Description
ADDITIONS-3	The ADDITIONS-3 field from the Adabas Control Block contains the Adabas password for the command. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No Yes Yes A8 A8 No **Use** The security administrator may want to know who is using which passwords. This may help assess the impact of password changes. Be aware of security implications with the presence of passwords in the Command Log file, COPY files, reports and data files. In some versions of Adabas, the ADDITIONS-3 field that is written to the Command Log contains only blanks to prevent disclosure of the password. **Values** The Adabas password for the command. **Alias** ADDIT3 **Column Heading** ADDIT 3

Field	Description
ADDITIONS-4	The ADDITIONS-4 field in the Adabas Control Block contains the Adabas cipher code. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No Yes Yes N8 N8 No **Use** Be aware of security implications with cipher codes in Command Logs, COPY files, reports and data files. **Value** The eight-digit cipher code. **Alias** ADD4, ADDIT4 **Column Heading** ADDIT 4
ADDITIONS-5	The ADDITIONS-5 field of the Adabas Control Block (ACB). **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No Yes Yes A8 A8 No **Use** May contain a Format Buffer ID different from the CID, or a Global Format ID, for some command codes. See the *Adabas Command Reference Manual*. **Alias** ADDIT5 **Column Heading** ADDIT 5
ADDIT1	See ADDITIONS-1
ADDIT2	See ADDITIONS-2

Field	Description
ADDIT2-REPORTED	See ADDITIONS-2-REPORTED
ADDIT3	See ADDITIONS-3
ADDIT4	See ADDITIONS-4
ADDIT5	See ADDITIONS-5
ADD4	See ADDITIONS-4
APPL-USER-ID	See ADDITIONS-1
ASSO-IO	Indicates the Associator I/O count, that is, the number of physical I/O operations Adabas reported it performed to read and/or write Associator blocks for the processing of the command. The field may include a buffer flush if ADARUN LFIOP=0. See the BUFFER-FLUSHES field. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No Yes Yes B2 N4 Yes F **Use** Critical in analyzing performance, particularly search and sort (Sx) commands. Large amounts of I/O to the Associator for single Sx commands may indicate poor use of descriptors. **Column Heading** ASSO/IO
BFL	See BUFFER-FLUSHES.

Field	Description
BUFFER-FLUSHES	Represents the buffer flush indicator, which estimates whether an Adabas I/O buffer flush occurred as part of the processing of the command. The field is set to 1 if a buffer flush has occurred, to zero if none has. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No Yes Yes B2 N1 Yes F **Derivation** Unicenter CA-APAS marks a command as having caused a buffer flush if the command caused more than one I/O operation to Data Storage and at least one I/O to Work. **Use** Useful to analyze the buffer flush control parameter, ADARUN TFLUSH and the buffer pool size. This field may not provide any meaningful information under Adabas Version 5 or 6. This is due to changes in the Adabas buffer flush logic. See summary fields BUFFER-FLUSHES-INTERVAL and BUFFER-FLUSHES-SESSION for an alternative source of buffer flush activity reporting. **Values** 0 = no buffer flush 1 = buffer flush occurred **Alias** BFL **Column Heading** B/F/L
CC	See CENTURY.

Field	Description
CENTURY	Contains the first 2 digits of the year, the century, from the command time stamp. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No Yes Yes A2 A2 No **Alias** CC **Column Heading** CC
CICS-TERMID	See CICS-TERMINAL-ID.
CICS-TERMINAL-ID	This field indicates the last four bytes of the GLOBAL-USER field. For CICS users, this is the four-character terminal-id. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No Yes No Yes A8 A8 No **Alias** CICS-TERMID, TERM-ID2 **Column Heading** CICS/TERM/ID
CID	See COMMAND-ID.

Field	Description
CL-TOT-IO	Adabas stores the count of the number of I/O operations done by a user in the Internal Sequence Number (ISN) field for a CL command. **Note:** Adabas keeps a separate total in each UQE. Reporting GLOBAL-USERID or one of its sub-fields is necessary to identify which user's I/O count is being reported. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No Yes Yes B4 N8 Yes F **Column Heading** TOTAL/IO
CL-TOT-CMDS	Adabas stores the count of the number of Adabas commands issued by a user in the ISN-Lower Limit field for a CL command. **Note:** Adabas keeps a separate total in each UQE. Reporting GLOBAL-USERID or one of its sub-fields is necessary to identify which user's I/O count is being reported. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No Yes Yes B4 N8 Yes F **Column Heading** TOTAL/COMMANDS
CL-TOT-CPU	Adabas stores an estimate of CPU time used by this user in the ISN Quantity field for a CL command. Experience has shown that this field measures elapsed time of the commands and not CPU time and is a very poor estimate of CPU time. **Note:** Adabas keeps a separate total in each UQE. Reporting GLOBAL-USERID or one of its sub-fields is necessary to identify which user's I/O count is being reported. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No Yes Yes B4 N8 Yes F **Column Heading** TOTAL/CPU

Field	Description
CMD	See COMMAND-CODE.
CMD-TIME	See COMMAND-TIME.
CMD-TYPE	See COMMAND-TYPE.
CMDSEQ	See COMMAND-SEQ-NUMBER.
COMM-CPU	See GLOBAL-CPU-ID.
COMM-ID	See GLOBAL-ID.
COMM-OS	See GLOBAL-OS-ID.
COMM-STCK	See GLOBAL-ID-STCK.
COMM-TOKEN	See GLOBAL-PROCESS-TOKEN.
COMM-USER	See GLOBAL-USER-ID.
COMM-USERID	Indicates the unique internal user-id. The values assigned to this field by Adabas have varied according to SM and optional zaps. This value is the UQETID field and is extracted from the UQE by the Data Collector when running under the Adabas User Exit 4. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No Yes Yes Yes B4 H8 No **Use** This internal user-id is used to trace the activity of a particular user. It may be used to differentiate between batch job steps within a given job name (SUMMARIZE ... BY JOBNAME BY CUID). **Alias** CUID, UQETID **Column Heading** USER ID
COMM-VM	See GLOBAL-VM-ID.

Field	Description
COMMAND-CODE	Indicates the Adabas Command Code. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No Yes Yes A2 A2 No **Use** Critical to performance analysis and exception reporting. **Alias** CMD **Column Heading** C/M/D

Field	Description
COMMAND-COST	This field contains a cost calculated for the individual command. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No Yes Yes B4 N2.5 Yes D **Derivation** Calculated by multiplying the estimated CPU, which includes the CPU-FACTOR, in seconds by the CPU-COST and adding the TOTAL-IO multiplied by the IO-COST. CPU-FACTOR, CPU-COST and IO-COST are parameters on the GLOBALS statement. Computed only if referenced in an EXTRACT or SUMMARIZE request; absent from COPY file records written before computation has begun. Computation continues even after request has been deleted. **Use** The SUM(COST) field may be used to report charges for Adabas usage. **Alias** COST **Column Heading** CMD/COST

Field	Description
COMMAND-ID	Indicates the command-id from the Adabas Control Block. This code identifies a sequence of related commands, held ISN list, internal format, etc. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No Yes Yes A4 A4 No **Use** In detailed analysis of command sequences. Verifies that a command-id is being used in a sequence of commands with the same format buffer. **Values** Helpful to display this field in both alpha and hexadecimal formats. **Alias** CID **Column Heading** CID

Field	Description
COMMAND-OPT1	Indicates Command Option One field from the Adabas Control Block. The Command Option 1 field has a number of meanings to Adabas depending on the command code; see the *Adabas Command Reference Manual* for complete information. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No Yes Yes A1 A1 No **Use** In detailed analysis of commands. Since the use of this field differs according to the command code, requests should include WHERE conditions that exclude command codes not desired for a particular interpretation of the field. **Alias** COPT1 **Column Heading** C/O/1
COMMAND-OPT2	Indicates Command Option Two field from the Adabas Control Block. The Command Option 2 field has a number of meanings to Adabas depending on the command code; see the Adabas *Command Reference Manual* for complete information. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No Yes Yes A1 A1 No **Use** In detailed analysis of commands. Since the usage of this field differs according to the command code, requests should include WHERE conditions that exclude command codes not desired for a particular interpretation of the field. **Alias** COPT2 **Column Heading** C/O/2

Field	Description
COMMAND-SEQ-NUMBER	Indicates the Command Sequence Number, that is, the number assigned by Adabas when the command is placed in the Command Queue, not the sequential number assigned as records written to the Command Log file. Therefore, there may be gaps in the sequence number, in an ALLTRACE request, for example. The missing numbers assigned to internal or utility commands that are never written to the Command Log file. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No Yes Yes B4 N10 No **Use** May be used in batch runs against a Command Log to skip to records of interest identified from earlier processing. **Alias** CMDSEQ **Column Heading** CMD/SEQ/NO

Field	Description
COMMAND-TIME	Indicates the command duration or estimated CPU time stored in the Adabas Command Log record. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No Yes Yes B4 N4.3 Yes D **Derivation** Taken from bytes 73-76 of the Adabas Control Block. **Values** Various duration and estimated CPU time, depending on release and SM of Adabas. **Alias** CMD-TIME **Column Heading** CMD/TIME

Field	Description
COMMAND-TYPE	Indicates if the command is simple, complex or update. The type of command as determined by Adabas. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No Yes Yes A1 A1 No **Use** As a first cut in analyzing the overall balance of commands. Inefficient complex search commands are a frequent cause of performance problems for inexperienced users. Also helps in analysis of ENQ-TIME to determine whether complex searches or update commands are being delayed in the Command Queue. **Values** The values are S for simple, C for complex, and U for update. **Alias** CMD-TYPE **Column Heading** T/Y/P
COPT1	See COMMAND-OPT1.
COPT2	See COMMAND-OPT2.
COST	See COMMAND-COST.
CPU	See EST-CPU-TIME.
CUID	See COMM-USERID.
DA	See DAY.

Field	Description
DATA-IO	Indicates Data Storage I/O count, that is, the number of physical I/O operations Adabas reported it performed to read and/or write Data Storage blocks for the processing of the command. The value may include a buffer flush, see BUFFER-FLUSHES field. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No Yes Yes B2 N4 Yes F **Use** Important in analyzing buffer pool use. Read command (Lx) sequences should have few reads to Data Storage if read in an efficient sequence. **Column Heading** DATA/IO
DATABASE-ID	Indicates the database-id (DBID) from the Command Log record or the MPM. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No Yes No Yes B4 N5 No **Derivation** When running under the Adabas User Exit 4, this field is taken from the Adabas nucleus. When processing COPY records, this field is taken from the derived fields area. When processing Command Log records in batch, this field is taken from the GLOBALS DBID keyword. **Use** If Command Logs from several databases are processed together, this field may be used to distinguish databases. **Alias** DBNO **Column Heading** DBID

Field	Description
DATE	Indicates the date on which Adabas thread processing of a command was completed. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No Yes Yes A6 A6 No **Derivation** Computed from the truncated Store-Clock end-of-thread-processing time stamp for the command. **Values** The values are year excluding century, month, and day in the format: yymmdd. **Column Heading** DATE
DATN	Indicates the date on which Adabas thread processing of a command was completed. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No Yes Yes A8 A8 No **Derivation** Computed from the truncated Store-Clock end-of-thread-processing time stamp for the command. **Values** The values are year including century, month, and day in the format: yyyymmdd. **Column Heading** DATE

Field	Description
DATE-TIME	Indicates the date and time at which Adabas thread processing of a command was completed. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No Yes Yes A12 A12 No **Derivation** Computed from the Store-Clock end-of-thread-processing time stamp for the command. Installations where the system clock is set to Greenwich Mean Time may adjust the date and time to local values using the GMT-ADJUSTMENT parameter of the GLOBALS statement. **Use** This field is used for selecting commands from a given time interval. **Values** The values are year excluding century, month, day, hour, minute, and second in the format: yymmddhhmmss. **Column Heading** DATE/TIME

Field	Description
DATN-TIME	Indicates the date and time at which Adabas thread processing of a command was completed. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No Yes Yes A14 A14 No **Derivation** Computed from the Store-Clock end-of-thread-processing time stamp for the command. Installations where the system clock is set to Greenwich Mean Time may adjust the date and time to local values using the GMT-ADJUSTMENT parameter of the GLOBALS statement. **Use** This field is used for selecting commands from a given time interval. **Values** The values are year including century, month, day, hour, minute, and second in the format: yyyymmddhhmmss. **Column Heading** DATE/TIME

Field	Description
DATEJ	Indicates the Julian date from the command time stamp. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No Yes Yes N5 A5 No **Derivation** Computed from the command Store-Clock time stamp. **Use** To correlate other performance data that uses this type of date. **Values** The value is specified in the format yyddd where yy is the year excluding century and ddd is the day of the year. **Column Heading** JULIAN/DATE
DAT4J	Indicates the Julian date from the command time stamp. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No Yes Yes N7 A7 No **Derivation** Computed from the command Store-Clock time stamp. **Use** To correlate other performance data that uses this type of date. **Values** The value is specified in the format yyyyddd where yyyy is the year including century and ddd is the day of the year. **Column Heading** JULIAN/DATE

Field	Description
DAY	Indicates the day of the month from the command time stamp. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No Yes Yes A2 A2 No **Values** The values can be 1 through 31. **Alias** DA **Column Heading** DY
DAY-OF-WEEK	See WEEKDAY.
DBNO	See DATABASE-ID.
DE-UPD	See DESCR-UPDATED.
DE-UPD-REPORTED	See DESCR-UPDATED-REPORTED.

Field	Description
DESCR-UPDATED	Indicates the number of descriptors that were changed by an update, add or delete command. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No Yes Yes B2 N3 Yes F **Use** Indicates the performance of update commands. Updates that do not update descriptors are less expensive. Updates, adds and deletes which update many descriptors can be very inefficient. Watch for updates which always cause a buffer flush because all descriptor blocks can't be updated in the buffer pool at once. **Values** To avoid reporting meaningless values, Unicenter CA-APAS sets this field to zero for all commands that do not update descriptors. **Alias** DE-UPD **Column Heading** DES/UPD

Field	Description
DESCR-UPDATED-REPORTED	Indicates number of descriptors that were changed by an update, add or delete command. This field differs from the DESCR-UPDATED field in that it is not altered when the command code and response code indicates that no descriptors were updated. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** **Yes** No No Yes Yes B2 N3 Yes F **Use** See DESCR-UPDATED. **Alias** DE-UPD-REPORTED **Column Heading** DES/UPD/RPT

Field	Description

DESCRIPTOR	The name of search buffer fields, that is, the Adabas field names appearing in search buffers.
	
	------- Adabas V7 ----- File Rept Sum Acc
	Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size
	Yes No No Yes Yes A2 A2 No
	
	Derivation
	
	Derived by analyzing search buffers of selected commands and incrementing the count each time a given field name is found.
	
	Search buffers are available for this analysis only if the Adabas ADARUN parameter LOGSB is ON.
	
	Use
	
	May be used to help determine which descriptors in a file are actually used; also whether non-descriptors are being used. The only useful type of request using this field is:
	`SUMMARIZE COUNT BY FILE BY DESCRIPTOR . . .`
	`WHERE (CMD = 'S1' OR = 'S2' OR = 'S4' OR = 'L9') . . . ;`
	
	Requests specifying this field should include only the command codes shown in the preceding example.
	
	Because there are often multiple field names in a search buffer (or the same name several times), the count may total more than the actual number of search commands. Used as indicated in the above example, this field does not reflect descriptor usage for logical sequential reading or sorting. See L3-SEQUENCE and ADDITIONS-1 for these cases.
	
	Alias
	
	SB-FLD
	
	Column Heading
	
	DESC/FLD

DUR	See DURATION.

Field	Description
DURATION	**------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No Yes Yes B4 N4.3 Yes D CLOG Layout 4: Indicates thread processing time, that is, the "wall-clock" time for the execution of a command as measured by Adabas. DURATION runs from the most recent time the command was selected from the Command Queue for execution in a thread until it completed processing. In the case of commands that have been "thrown back" at some point in earlier attempts to process them, time spent in Adabas threads prior to the final, successful initiation is not included in DURATION; instead, it is included in ENQ-TIME. CLOG Layout 5: Indicates the elapsed time for the command in the Command Queue as measured by Adabas. In the case of commands that have been "thrown back" at some point in earlier attempts to process them, time spent in Adabas threads prior to the final, successful initiation is included in DURATION. 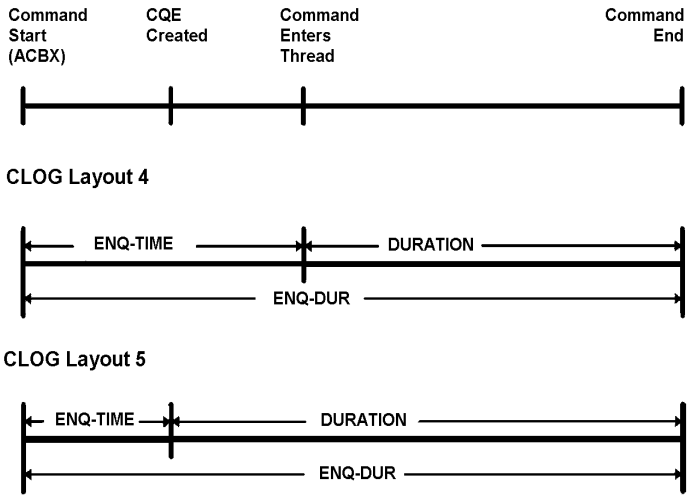

Field	Description
DURATION (cont.)	**Derivation** Computed from the 16-microsecond units given by Adabas. **Use** Duration is usually proportional to the amount of I/O. Watch for unusually long commands. These may indicate that the Adabas task has been interrupted, which should not happen often. This field is a good indication of overall tuning of the system. If duration increases with increasing thread number, CPU saturation is indicated since Adabas processed threads in thread number sequence. **Values** The value is shown in seconds. **Alias** DUR **Column Heading** DURAT/SECS
ECBS	See POSTED-ECBS.
ENQ	See ENQ-TIME.

Field	Description
ENQ-DUR	Indicates sum of ENQ-TIME and DURATION, which is the elapsed time from when the Adabas command was issued by the calling program, until Adabas has finished processing the command in one of its threads. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes Yes Yes No Yes B4 N4.3 Yes D **Derivation** Computed by subtracting a STCK value taken in the ACBX routine from a STCK value taken in the Unicenter CA-APAS User Exit Four routine. The value of this field may be questionable for commands that are issued on one CPU but are processed by Adabas on a different CPU. The Store-Clock values on the two different CPUs may not be synchronized closely enough. **Use** Can be helpful in determining appropriate values for COM-PLETE's ADAROLL parameter. Can also be helpful in analyzing online response times. See the ENQDUR request in the Unicenter CA-APAS Natural library. Also see the explanations of ENQ-TIME and DURATION in this document. **Values** The value is shown in seconds. **Column Heading** ENQ-DUR/SECS

Field	Description
ENQ-TIME	**------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes Yes Yes No Yes B4 N3.4 Yes D CLOG Layout 4: Indicates wait time before thread processing, which is the "wall-clock" time between the time the application issued the command and Adabas started processing the command in one of its threads. Generally, it reflects the amount of time the command was in the Adabas Command Queue before it was selected for thread processing. It also reflects inter-task communication time (SVC) and Adabas throwback(s) of a command, that is, thread processing that was discontinued and eventually duplicated. CLOG Layout 5: Indicates wait time before thread processing, which is the "wall-clock" time between the time the application issued the command and the command was put in the Adabas Command Queue. It reflects inter-task communication time (SVC). 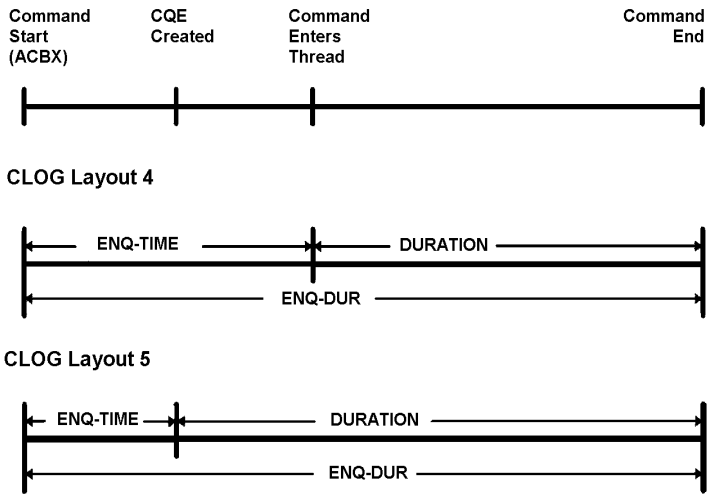

Field	Description
ENQ-TIME (cont.)	**Derivation** The field is computed by determining the total time from when the application called Adabas until the end of processing in the thread, and then subtracting the DURATION field computed by Adabas, to eliminate time spent processing in the thread. The value of this field may be questionable for commands that are issued on one CPU but are processed by Adabas on a different CPU. The Store-Clock values on the two different CPUs may not be synchronized closely enough. **Use** This field is the best indication of saturation of the nucleus. In a well-balanced system, the enqueued time should be much less than the duration, for most commands. Very high values of this field may indicate MVS swapping problems. Other reasons for high values are shortage of memory for starting additional searches marked by Adabas as complex, (LWP space), or Work data set part 2 space (LWKP2 area); Adabas I/O buffer flushes; Adabas being delayed by external events or conditions in the operating environment. **Values** The value is expressed in seconds. **Alias** ENQ **Column Heading** ENQ TIME/SECS

Field	Description
EST-CPU-TIME	Indicates CPU time to process command, that is, the estimated amount of CPU time that the Adabas nucleus spent processing a command. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No Yes Yes B4 N5.3 Yes D **Derivation** This value is derived directly from the EST-INSTRUCTIONS value, as per Software AG formulae, divided by the CPU-INSTRS-SEC parameter, then multiplied by the CPU-FACTOR parameter. The estimated CPU time for a command may sometimes be greater than the duration reported for the command. In such cases, the estimated CPU time is clearly greater than the actual CPU time. These and other instances of overestimation are due to any combination of the following reasons: ■ the CPU-FACTOR value is greater than 1.0 ■ the CPU-INSTRS-SEC parameter is overstated ■ Deficiencies in the Software AG formulae This value is computed only if referenced in an EXTRACT or SUMMARIZE request; absent from COPY file records written before computation has begun. Computation continues even after the request that initiated computation has been deleted. As of the publication date of this document, the prospects of being able to use an improved method of estimating this field for Adabas 7 was uncertain. Until better approaches for Adabas 7 are developed, the method used for Adabas 4 is used for Adabas 7. **Use** This value is a very rough approximation. The CPU-FACTOR parameter may be adjusted in an attempt to get the total of all commands' estimated CPU time to correspond to the CPU time reported by system accounting for the Adabas job step. The formulae for calculating the number of instructions are dominated by the physical I/O, therefore TOTAL-IO itself may be as useful for estimating relative resource usage.

Field	Description
EST-CPU-TIME (cont.)	**Values** The value is expressed in 10's of microseconds. **Alias** CPU **Column Heading** CPU/SECS
EST-INSTRUCTIONS	Indicates the number of instructions to process command, that is, the number of instructions in hundreds to execute a command, estimated using the formulae supplied by Software AG. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No Yes * B4 N5 Yes D **Derivation** Computed using Software AG formulae. Computed only if referenced in an EXTRACT or SUMMARIZE request; absent from COPY file records written before computation has begun. Computation continues even after the request that initiated computation has been deleted. **Use** See the discussion of EST-CPU-TIME above. A relative indicator of resource use. **Alias** INSTRS **Column Heading** PROC/INSTR
FB	See FORMAT-BUFFER.

Field	Description
FB-ERR-FIELD	Indicates format or search buffer field in error. This field is set to the first two bytes of the ADDITIONS-2 field when Adabas detects an error in the format buffer or search buffer. These two bytes are the two-character name of the field in error. When no error occurs, this field is set to blanks. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No Yes Yes A2 A2 No **Use** May be displayed for detailed error analysis and debugging of specific commands selected on the basis of error response codes. This field should be used strictly on an exception reporting basis, not for all commands. **Alias** SB-ERR-FIELD **Column Heading** FB/ERR

Field	Description
FB-FIELD	Indicates name of field in format buffer, the field is used in connection with counts of occurrences of field names in format buffers. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No Yes Yes A2 A2 No **Derivation** Derived by analyzing each format buffer and adding an entry to the summary table for each field name found. Adabas format buffer logging must be active. In the case of Adabas Command Logs written under control of Unicenter CA-APAS UEX4, the command must have satisfied the selection criteria of a LOG statement that specified format buffer logging. **Use** Use this field to determine the relative use of fields in a file. The only useful type of request using this field is: `FLDUSE:SUMMARIZE COUNT BY FILE BY FB-FIELD . . . ;` Because a given field name may appear multiple times in a format buffer, the count may total more than the actual number of commands using format buffers. **Column Heading** FB/FLD
FBL	See FBLENGTH.

Field	Description
FBLENGTH	Indicates format buffer length, that is, the format buffer length specified in the Adabas Control Block. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No Yes Yes B2 N5 Yes F **Use** May be analyzed for excessive length. **Alias** FBL **Column Heading** FBL
FILE	See FILE-NO.

Field	Description
FILE-NO	Indicates Adabas file number from the Adabas Control Block. For commands where the file number has no meaning, this field is set to zero. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No Yes Yes B2 N5 No **Derivation** Some Adabas command codes do not require a file number; in these cases Adabas ignores the contents of the file number field of the ACB. Some applications specify actual file number values with these command codes; others do not. Further, some applications are inconsistent in this respect. To avoid attributing such commands to spurious file number values, Unicenter CA-APAS ignores the actual contents of the file number field in the ACB (as does Adabas) and reports a file number value of zero for the following command codes: BT (unless Command Option 2 = 'F') CL C1 C2 C3 C5 ET OP RC RE For all other command codes, Unicenter CA-APAS reports the value it finds in the ACB. **Use** Basic to performance analysis and capacity planning. **Values** 0 to 32767. **Alias** FILE **Column Heading** FIL/NUM

Field	Description
FILE-NO-REPORTED	Indicates the Adabas file number from the Adabas Control Block without correction for commands where the file number is meaningless. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No Yes Yes B2 N5 No **Use** Use field name FILE unless your desire is the unaltered value reported by Adabas. **Value** The values may be 0 through 32767. **Alias** FILE-REPORTED **Column Heading** FIL/NUM/RPT
FILE-REPORTED	See FILE-NO-REPORTED.

Field	Description
FORMAT-BUFFER	Indicates the contents of Adabas format buffer for the command. Adabas format buffer logging must be active. In the case of Adabas Command Logs written under control of Unicenter CA-APAS UEX4, the command must have satisfied the selection criteria of a LOG statement that specified format buffer logging. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No Yes Yes A Dump No **Use** Detailed command analysis. Analyze for field use. **Alias** FB **Column Heading** FORMAT BUFFER
GLOBAL-CPU	See GLOBAL-CPU-ID.
GLOBAL-CPU-ID	This is a sub-field of the GLOBAL-ID field. The derivation and usage of the field are described in the description of GLOBAL-ID. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No * * Yes A8 A8 No * See the explanation of GLOBAL-ID. **Alias** COMM-CPU, GLOBAL-CPU **Column Heading** COMM/CPU-ID

Field	Description
GLOBAL-ID	This field contains the 28-byte Global Unique Identifier that is created by Adabas for each user. This field can be used to uniquely identify each Adabas user. It is taken from the Communication-ID field of the Adabas Command Log layout 5 record, or from the UQE Global Unique ID field with CLOG layout 4. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No * * Yes A28 A28 No * When using CLOG layout 5, this field is taken from the Command Log record and does not require the Adabas User Exit 4. **Derivation** This is a Unicenter CA-APAS derived field; the field name is DERGLBID, which is extracted from the Adabas UQE when the Data Collector is running under the Adabas User Exit 4. It is not available when processing Adabas Command Logs in batch, but it is part of the Unicenter CA-APAS derived fields area in Unicenter CA-APAS COPY files. **Use** This field can be used to identify uniquely each Adabas user based on the information supplied by the Adabas link modules and the MPM. This might be useful when reporting Adabas usage by user. The field consists of four sub-fields, as follows: Bytes Field Description 1-8 CPU ID An eight-byte CPU identifier 1-8 STCK An eight-byte STCK field for X'48' calls 9-16 VM ID An eight-byte virtual machine identifier 9-12 TOKEN A four-byte process token for X'48' calls 17-20 OS ID A four-byte operating system identifier 21-28 USER ID An eight-byte user identifier Several of these sub-fields, particularly CPU ID, STCK, VM ID, TOKEN, and OS ID, are not necessarily printable characters. Specifying a print format of H for one or more of these sub-fields might be helpful. A field name and one or more aliases have been defined for each of these sub-fields.

Field	Description
GLOBAL-ID (cont.)	**Alias** COMM-ID, GLOBAL-UNIQUE-ID **Column Heading** COMM/ID
GLOBAL-ID-STCK	This is a sub-field of the GLOBAL-ID field in hex'48' calls. The derivation and usage of the field are described in the description of GLOBAL-ID. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No * * Yes A8 A8 No * See the explanation of GLOBAL-ID. **Alias** COMM-STCK **Column Heading** COMM/STCK
GLOBAL-OS	See GLOBAL-OS-ID.
GLOBAL-OS-ID	This is a sub-field of the GLOBAL-ID field. The derivation and usage of the field are described in the description of GLOBAL-ID. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No * * Yes A4 A4 No * See the explanation of GLOBAL-ID. **Alias** COMM-OS, GLOBAL-OS **Column Heading** COMM/OS-ID

Field	Description
GLOBAL-PROCESS-TOKEN	This is a sub-field of the GLOBAL-ID field in hex'48' calls. The derivation and usage of the field are described in the description of GLOBAL-ID. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No * * Yes A4 A4 No * See the explanation of GLOBAL-ID. **Alias** COMM-TOKEN **Column Heading** COMM/TOKEN
GLOBAL-UNIQUE-ID	See GLOBAL-ID.
GLOBAL-USER	See GLOBAL-USER-ID.
GLOBAL-USER-ID	This is a sub-field of the GLOBAL-ID field. The derivation and usage of the field are described in the description of GLOBAL-ID. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No * * Yes A8 A8 No * See the explanation of GLOBAL-ID. **Alias** COMM-USER, GLOBAL-USER, GLOBAL-USERID **Column Heading** COMM/USER
GLOBAL-USERID	See GLOBAL-USER-ID.
GLOBAL-VM	See GLOBAL-VM-ID.

Field	Description
GLOBAL-VM-ID	This is a sub-field of the GLOBAL-ID. The derivation and usage of the field are described in the description of GLOBAL-ID. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No * * Yes A8 A8 No * See the explanation of GLOBAL-ID. **Alias** COMM-VM, GLOBAL-VM **Column Heading** COMM/VM-ID

Field	Description
HOLD-COUNT	Provides an estimate of the number of records held by the current user (TP user or batch job). **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No No Yes B4 N4 Yes F **Derivation** This field is incremented by one every time an L4, L5, L6, S4, N1, N2 or HI command occurs. It is assumed that the record was previously read for Ax and Ex commands. It is decremented by one for RI command with ISN and set to zero for other RI, ET, BT, OP, and CL commands. The count is not accurate for multiple read-with-holds of the same record, which can occur by design or in an unintended program loop, nor for updates and deletes where the record was not previously read with hold. **Use** The MAX of this field helps to identify potential Hold Queue overflow situations or excessive holds that could cause transaction time-outs. This field is an attribute of the user that issued the command, not an attribute of the command itself. This distinction should be borne in mind when writing requests and interpreting outputs from requests. **Alias** REC-HOLD-CNT **Column Heading** REC/HOLD/CNT

Field	Description
HOUR	Indicates the hour of the day from the command time stamp. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No Yes Yes A2 A2 No **Use** May be used in calculating use profiles by hour of day. **Values** The values can be 00 to 23. **Alias** HR **Column Heading** HR
HR	See HOUR.
IB	See ISN-BUFFER.
IBL	See IBLENGTH.
IBLENGTH	Indicates the length of the ISN buffer from the Adabas Control Block. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No Yes Yes B2 N5 Yes F **Use** Compare to ISN-QUANTITY for determining whether size is efficient. **Alias** IBL **Column Heading** IBL
IFBL	See IFB-LENGTH.

Field	Description
IFBL-ALLOC	Indicates the allocated length of the Internal Format Buffer, that is, the length of the area in the format buffer pool that has been allocated to receive the interpreted user format buffer. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No Yes No Yes B2 N5 Yes F **Use** Compare to the actual IFBL. If IFBL-ALLOC is significantly greater than IFBL, you are allocating too much space from the Work Pool for each command. Decrease the initial zap amount. **Column Heading** IFBL/ALLOC
IFB-LENGTH	Indicates the Internal Format Buffer length, this is, the actual length of the fully interpreted format buffer for the command. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No Yes No Yes B2 N5 Yes F **Use** MEAN (IFBL) is useful when applying the Software AG zap to set the initial size of an area of LWP into which the user format buffer is interpreted. Also identify commands with unusually large Internal Format Buffers caused by reading too many fields or too many occurrences of a periodic group or multi-value field. **Alias** IFBL **Column Heading** IFBL
IMS-CDATE	See IMS-CURR-DATE.
IMS-CTIME	See IMS-CURR-TIME.

Field	Description
IMS-CURR-DATE	Indicates the current date from the IMS IOPCB control block. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes Yes Yes No Yes P7 N2.3 No **Derivation** From the IMS IOPCB control block. **Alias** IMS-CDATE, IOPCB-CDATE **Column Heading** IMS/DATE
IMS-CURR-TIME	Indicates the current time from the IMS IOPCB control block. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes Yes Yes No Yes P7 N5 No **Derivation** From the IMS IOPCB control block. **Alias** IMS-CTIME, IOPCB-CTIME **Column Heading** IMS/TIME

Field	Description
IMS-LTERM-ID	Indicates the logical terminal ID from the IMS IOPCB control block. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes Yes Yes No Yes A8 A8 No **Derivation** From the IMS IOPCB control block. **Alias** LTERM, IOPCB-LTERM **Column Heading** LTERM-ID
IMS-STATUS	See IMS-STATUS-CD.
IMS-STATUS-CD	Indicates the IMS status code from the IMS IOPCB control block. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes Yes Yes No Yes A2 A2 No **Derivation** From the IMS IOPCB control block. **Alias** IMS-STATUS, STATUS-CODE, IOPCB-STATUS **Column Heading** IMS/STATUS

Field	Description
IMSN	Indicates the IMS message number from the IMS IOPCB control block. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes Yes Yes No Yes B4 H8 No **Derivation** From the IMS IOPCB control block. **Alias** IOPCB-IMSN **Column Heading** IMSN
INSTRS	See EST-INSTRUCTIONS.
IO-LIST	Indicates the RABNs (Relative Adabas Block Numbers) of physical I/O operations, that is, a list of up to 40 RABNs of blocks that were physically read or written by ADAIOR for the command. Adabas I/O logging must be active. In the case of Adabas Command Logs written under control of Unicenter CA-APAS UEX4, the command must have satisfied the selection criteria of a LOG statement that specified I/O list logging. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No Yes Yes A Dump No **Use** In detailed analysis of patterns of I/O activity; in ADAM files, for example. **Alias** IOL **Column Heading** IO RABN LIST
IOL	See IO-LIST.

Field	Description
IOPCB-CDATE	See IMS-CURR-DATE.
IOPCB-CTIME	See IMS-CURR-TIME.
IOPCB-IMSN	See IMSN.
IOPCB-LTERM	See IMS-LTERM-ID.
IOPCB-STATUS	See IMS-STATUS-CD.
ISN	Indicates the Internal Sequence Number from the Adabas Control Block, altered for the Adabas CL command, as described below. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No Yes Yes B4 N13 No **Use** In detailed analysis of commands. Identifies records accessed very frequently, thereby causing hold conflicts, and redundant accesses to the same record in a sequence of commands. For a CL command, Adabas stores a count of total I/O operations for this user. The Data Collector resets this field to zero for a CL command. Users who want the I/O operation count should use field CL-TOT-IO. Users who want the unaltered ISN field should use field name ISN-REPORTED. **Column Heading** ISN

Field	Description
ISN-BUFFER	Indicates the contents of ISN buffer following execution of command. Adabas ISN buffer logging must be on. In the case of Adabas Command Logs written under control of Unicenter CA-APAS UEX4, the command must have satisfied the selection criteria of a LOG statement that specified ISN buffer logging. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No Yes Yes A Dump No **Use** Not usually of interest. **Alias** IB **Column Heading** ISN-BUFFER
ISN-LOWER-LIMIT	Indicates the ISN Lower Limit field from the Adabas Control Block, altered as described below. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No Yes Yes B4 N13 No **Use** Useful in debugging search command problems. **Values** Non-zero values are reported only for commands where the ISN Lower Limit is meaningful; zero is reported for all commands except for all Sx commands regardless of the Adabas response code. **Alias** ISNLL **Column Heading** ISN/LOWER/LIMIT

Field	Description
ISN-LOWER-LIMIT-REPORTED	This is the ISN Lower Limit field from the Adabas Control Block. It differs from the field ISN-LOWER-LIMIT in that it is not set to zero for commands where a non-zero value is meaningless. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No Yes Yes B4 N13 No **Use** See ISN-LOWER-LIMIT. **Values** Note that this field is used for other purposes such as CL or OP. See the *Adabas Command Reference Manual* before using the values from this field. **Alias** ISNLL-REPORTED **Column Heading** ISN/LOWLIM/RPTD

Field	Description
ISN-QUANTITY	Indicates the number of ISNs found by a search, corresponding to the ISN Quantity field from the Adabas Control Block, altered as described below. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No Yes Yes B4 N13 Yes D **Use** Large values may indicate inefficient searches. Use this to find an application's most efficient ISN buffer size. **Values** Non-zero values are reported only for commands where the ISN Quantity is meaningful; zero is reported for all commands except Sx or L9 commands with a zero response code. **Alias** ISNQ **Column Heading** ISN/QUANT

Field	Description
ISN-QUANTITY-REPORTED	Indicates the number of ISNs found by a search, this is the ISN Quantity field from the Adabas Control Block. It differs from the field ISN-QUANTITY in that it is not set to zero for commands where a non-zero value is meaningless. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No Yes Yes B4 N13 Yes D **Use** See ISN-QUANTITY. **Values** Note that this field is used for other purposes by commands such as CL, OP, or ET. See the *Adabas Command Reference Manual* before using the values from this field. **Alias** ISNQ-REPORTED **Column Heading** ISN/QUANT/RPTD
ISNLL	See ISN-LOWER-LIMIT.
ISNLL-REPORTED	See ISN-LOWER-LIMIT-REPORTED.
ISNQ	See ISN-QUANTITY.
ISNQ-REPORTED	See ISN-QUANTITY-REPORTED.

Field	Description
ISN-REPORTED	Indicates the Internal Sequence Number from the Adabas Control Block. This field differs from ISN in that it is not set to zero for CL commands. For CL commands, this field does not contain an Internal Sequence Number. Instead, it contains a count of I/O operations for this user. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No Yes Yes B4 N13 No **Use** See ISN. **Column Heading** ISN/RPTD
IUB-LENGTH	Indicates the length of Adabas Intermediate User Buffer, that is, the actual IUB length for the command. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No Yes No Yes B4 N6 Yes D **Use** MAX(IUBL) may guide setting the value of the ADARUN parameter LU. MEAN(IUBL) may be useful in setting the value of the ADARUN NABS parameter. **Alias** IUBL **Column Heading** IUB/LENGTH
IUBL	See IUB-LENGTH.
JEDR	See JOB-ENTRY-DATE-REPORTED.
JMR-USER-ID	See USER-ID.
JOB	See JOBNAME.

Field	Description

JOB-ELAPSED-TIME	Indicates the time since user's first Adabas command, that is, the time in seconds elapsed since the first Adabas command issued by this user during the Adabas session.
	
	------- Adabas V7 ----- File Rept Sum Acc
	Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size
	Yes No No No Yes B4 N6 Yes F
	
	Derivation
	
	Current time minus time of first command of the Adabas session for the particular value of CUID. This field is never reset for a given user during the Adabas session.
	
	Use
	
	This field is meaningful only for batch job steps. It is not valid for measuring transaction times of TP users. All requests which reference this field should include:
	
	`". . .WHERE USER-TYPE = 'BT' . . ."`
	
	SUMMARIZE requests should reference MAX(JOB-ELAPSED-TIME) and should include:
	
	`". . .BY JOBNAME BY CUID. . . "`
	
	With the above qualifications, MAX(JOB-ELAPSED-TIME) gives the duration of that portion of a batch job step between the step's first and last Adabas commands.
	
	Values
	
	The values are expressed in seconds.
	
	Column Heading
	
	JOB/ELAPSED/TIME

Field	Description
JOB-ENTRY-DATE	Indicates system job entry date for batch jobs, that is, the Julian date with a 2 digit year when the batch job was submitted. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes Yes Yes No Yes P5 N2.3 No **Derivation** Obtained from the JMREDATE field. **Use** All requests which reference this field should include: `". . .WHERE USER-TYPE = 'BT' . . ."` **Values** The values are expressed as a 5-digit Julian date: yy.ddd. **Column Heading** JOB/ENTRY/DATE

Field	Description
JOB-ENTRY-DATE-REPORTED	Indicates system job entry date for batch jobs, that is, the Julian date when the batch job was submitted. The field corresponds to the equivalent field in SMF accounting records. The difference between this field and the JOB-ENTRY-DATE and JOB-ENTRY-DAT4J fields is that this field is not altered to drop the century flag as with JOB-ENTRY-DATE or to convert the century flag to the proper century as with the JOB-ENTRY-DAT4J field. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes Yes Yes No Yes * H8 No * The field is defined as B4 but the contents is formatted by JES as a packed P7 field. **Derivation** Obtained from the JMREDATE field. **Use** May be used to match job accounting data from Unicenter CA-APAS with system SMF data. All requests which reference this field should include: `". . .WHERE USER-TYPE = 'BT' . . ."` **Values** The values are expressed in the IBM format: 0cyydddf. **Alias** JEDR **Column Heading** JOB/ENTRY/DATE/REPORTED

Field	Description
JOB-ENTRY-DAT4J	Indicates system job entry date for batch jobs, that is, the Julian date including century when the batch job was submitted. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes Yes Yes No Yes P7 N4.3 No **Derivation** Obtained from the JMREDATE field. **Use** All requests which reference this field should include: `". . .WHERE USER-TYPE = 'BT' . . ."` **Values** The values are expressed as a 7-digit Julian date including century: c,cyy.ddd. To eliminate the comma in the century portion, specify in the request an override Print Format of (PF=NC7) for the field. **Column Heading** JOB/ENTRY/DAT4J

Field	Description
JOB-ENTRY-TIME	Indicates system job entry time for batch jobs, that is, the time when the batch job was submitted. The field corresponds to the equivalent field in SMF accounting records. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes Yes Yes No Yes B4 N7 No **Derivation** Obtained from the JMRENTRY field. **Use** May be used to match job accounting data from Unicenter CA-APAS with system SMF data. All requests which reference this field should include: `". . .WHERE USER-TYPE = 'BT' . . ."` **Values** The values indicate 1/100 seconds since midnight. **Column Heading** JOB/ENTRY/TIME
JOB-NO	See JOB-NUMBER.

Field	Description
JOB-NUMBER	Indicates system job number for batch jobs. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes Yes Yes No Yes P7 N7 No **Derivation** Obtained from the SSIBJBID field of the Subsystem Identification Block. **Use** May be used to match job accounting data from Unicenter CA-APAS with system SMF data. All requests which reference this field should include: `". . .WHERE USER-TYPE = 'BT' . . ."` **Alias** JOB-NO **Column Heading** JOB/NUMBER
JOB-STEP	Indicates the system job step number, that is, the number of the step in a batch job issuing the command. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes Yes Yes No Yes B1 N3 No **Derivation** Obtained from the JMRSTEP field. All requests which reference this field should include: `". . .WHERE USER-TYPE = 'BT' . . ."` **Column Heading** JOB/STEP/NO

Field	Description
JOBNAME	Indicates the name of job issuing command, that is, the jobname, as determined by Adabas. It is usually identical to the name from the JOB statement in the JCL for the program issuing the command. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No Yes Yes A8 A8 No **Derivation** Taken directly from the Adabas Command Log record, as provided by Adabas. **Use** Identifies batch jobs and TSO users of the database. **Alias** JOB **Column Heading** JOB NAME
JOBNAME2	Indicates alternative job name field. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes Yes Yes No Yes A8 A8 No **Derivation** From the Adabas Command Queue Element (CQE). **Use** May be used in lieu of the field, JOBNAME, which is taken directly from the Adabas Command Log record, in cases where JOBNAME has incorrect values. **Column Heading** JOBNAME2

Field	Description
LOG-REC-TYPE	Indicates the Command Log record type from the Adabas Control Block. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No * Yes B2 H4 Yes F * Only valid when using Adabas Command Log layout 5. **Derivation** Taken directly from an Adabas CLOG layout 5 Command Log record. **Use** This field can be used to differentiate commands generated by Adabas triggers and stored procedures. The values include: X'0001' for Adabas Basic Log Record X'0005' for Pre-command Trigger Record X'0006' for Post-command Trigger Record **Alias** LREC-TYPE **Column Heading** LOG/REC/TYPE
LNUINFO	Indicates the length of the User Information Area, that is, the UEXB buffer length from the Adabas Control Block. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No Yes Yes B2 N5 Yes F **Use** May be analyzed for excessive LNUIFNO length. **Alias** UEXITBL, UEXBL **Column Heading** LNUINFO

Field	Description
LREC-TYPE	Indicates the Command Log record type from the Adabas Control Block. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No Yes Yes B2 H4 Yes F **Derivation** Taken directly from the Adabas Command Log record. **Use** This field can be used to differentiate commands generated by Adabas triggers and stored procedures. The values include: X'0001' for Adabas Basic Log Record X'0005' for Pre-command Triggers Record X'0006' for Post-command Trigger Record **Alias** LOG-REC-TYPE **Column Heading** LREC/TYPE
LTERM	See IMS-LTERM-ID.

Field	Description

L3-SEQUENCE	Indicates the name of sequencing descriptor, that is, the descriptor used in a logical-sequential read of a file (L3 or L6 command).
	
	------- Adabas V7 ----- File Rept Sum Acc
	Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size
	Yes No No Yes Yes A2 A2 No
	
	Derivation
	
	The first two characters of the ADDITIONS-1 field.
	
	Use
	
	Identifies potential benefits of reordering the records of a file within Data Storage. If the most heavily-used descriptor has a high MEAN(DATA-IO) figure, the file should probably be reordered. Useful only in requests of the following form:
	`SUMMARIZE COUNT MEAN(DATA-IO)`
	`BY FILE BY L3-SEQUENCE`
	`WHERE CMD = 'L3' OR = 'L6' . . .;`
	
	Column Heading
	
	L3/SEQ

MON	See MONTH.
MONTH	Indicates the month (01-12) from the time stamp for the command. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No Yes Yes A2 A2 No **Values** The values may be 01 to 12. **Alias** MON **Column Heading** MN

Field	Description
NAT-CMD	Indicates the second byte of the Adabas Control Block (ACB). **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No Yes No Yes A1 A1 No **Derivation** Directly from the Adabas Control Block. **Use** Certain values in this field indicate that the call came from Natural. **Column Heading** NAT/CMD
NAT-LOAD-DATE	Indicates the date a module was loaded, that is, the date of the FETCH/READ of a Natural module/program. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No Yes Yes A6 A6 No **Derivation** See the NAT-LOAD-DTTM field. **Values** The values are year, month, and day in the format: yymmdd. **Column Heading** NATURAL/LOAD/DATE

Field	Description
NAT-LOAD-DTTM	Indicates the date/time a module was loaded, that is, the date and time of the FETCH/READ of a Natural module/program. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No Yes Yes A12 A12 No **Derivation** Computed from the Store-Clock time stamp of the L3 commands for the Natural FETCH/READ function. See the NAT-MOD-ID field description. **Use** Use this field to differentiate between executions of Natural programs that have been modified during an Adabas session. **Values** The values express the year, month, day, hour, minute, and second in the format: yymmddhhmmss. **Column Heading** NATURAL/LOAD/DATE-TIME
NAT-LOAD-TIME	Indicates the time a module was loaded, that is, the time of the FETCH/READ of a Natural module/program. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No Yes Yes A6 A6 No **Derivation** See the NAT-LOAD-DTTM field. **Values** The values express hour, minute, and second in the format: hhmmss. **Column Heading** NATURAL/LOAD/TIME

Field	Description
NAT-LOGON	See NATRUAL-LOGON.
NAT-MOD-ID	Indicates the ID of source or object module, thereby identifying the Natural object module or source program that issued a command. NAT-MOD-ID contains the logon id, the program name, the mode indicator (the letter P for program, the letter M for module, or the letter A for ADHOC), and system file number. The ACBX logic derives the item type (P, M, or A) based on indicators maintained by Natural. Object modules may be identified erroneously as programs in some cases in which RUN has been executed by the user or a program earlier in the Natural session. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes Yes Yes No Yes A22 A22 No **Derivation** The Natural module-id is derived from fields maintained by Natural. Experience indicates that the values maintained by Natural do not always reflect the program that actually issued the command. In some cases, the value reflects the name of a Natural object that was being read from the Natural system file; in these cases the command was issued by Natural's own routines, though in behalf of the named object. **Use** Basic for analysis of systems written in Natural. **Values** The values include: NAT-LOGON(A8) NAT-PROGRAM(A8) NAT-M-OR-P(A1) NAT-SYSTEM-FILE(N5) **Alias** NAT-MODULE **Column Heading** NATURAL/MODULE/ID

Field	Description
NAT-MODULE	See NAT-MOD-ID.
NAT-PROG	See NAT-PROGRAM.
NAT-PROGRAM	Indicates the name of the source or object module, that is, the name of the module or program that issued the command. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes Yes Yes No Yes A8 A8 No **Derivation** See NAT-MOD-ID. This field is simply a redefinition of bytes 9-16 of NAT-MOD-ID. **Alias** NAT-PROG **Column Heading** NATURAL/PROGRAM

Field	Description
NAT-STMT-NO	Indicates the source line number of the Natural statement that caused the command to be issued, usually. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No Yes Yes B2 H4 No **Derivation** Internal redefinition of the left two bytes of the CID. Note that Natural is not consistent in its use of this field, as mentioned above. **Use** Helpful in finding the Natural source code causing expensive or unnecessary commands. **Values** The values range from 1-9999 for actual statement numbers. Other values do not represent statement numbers. Including the following WHERE clause may be used in some requests. `WHERE NAT-STMT-NO LE H'9999'` **Column Heading** NAT/STMT/NO
NAT-SYSFILE	See NAT-SYSTEM-FILE.
NAT-SYSFILE-DBID	Indicates the database-id of the system file from which the current Natural program was run, read, fetched, or edited. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes Yes Yes No Yes B2 N5 No **Derivation** The DBID for the user system file specified in NATPARMS. This may be incorrect for SYSLIB programs. **Column Heading** NAT/SYSF/DBID

Field	Description
NAT-SYSTEM-FILE	Indicates the number of the system file from which the current Natural program was run, read, fetched, or edited. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** **Y**es Yes Yes No Yes A5 A5 No **Derivation** If the Unicenter CA-APAS ACBX facility is installed, this field is the user system file specified in NATPARMS. This may be incorrect for system (SYSLIB) programs in the system file. **Alias** NAT-SYSFILE **Column Heading** NAT/SYS/FILE
NATURAL-LOGON	Indicates the library or application name, that is, the logon id of the current Natural program. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes Yes Yes No Yes A8 A8 No **Derivation** See the NAT-MOD-ID discussion. This field is a redefinition of bytes 1-8 of NAT-MOD-ID. **Use** Used to analyze performance and load by system. **Alias** NAT-LOGON **Column Heading** NATURAL/LOGON
OPERATOR-ID	See USER-ID.

Field	Description
POSTED-ECBS	Indicates the number of posted Command Queue elements, that is, the total number of commands ready for execution and/or currently executing when Adabas selected the command for execution in a thread. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No Yes Yes B1 N3 Yes F **Use** Used to estimate relative saturation of Adabas nucleus command processing capacity. ENQ-TIME is a better indicator of waits in the Command Queue, but POSTED-ECBS may be used if it is not possible to install the Unicenter CA-APAS ACBX facility and User Exit 4. MAX(ECBS) may be used to estimate the number of CQEs needed (ADARUN NC parameter). **Alias** ECBS **Column Heading** PST/ECB
PROGRAM	Indicates the name of the program issuing the command, that is, the name of the module issuing commands in either the batch or TP environment. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes Yes Yes No Yes A8 A8 No **Derivation** In batch, the name of the program module first loaded for the task or sub-task. For CICS, the name of the module which issued the LINK to the Adabas module. **Column Heading** PROGRAM/NAME
PRTY	See USER-PRIORITY.
QTRHR	See QUARTER-HOUR.

Field	Description
QUARTER-HOUR	Indicates the fifteen-minute interval of the day, that is, the fifteen-minute period within the twenty-four hour day from the command time stamp. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No Yes Yes A2 A2 No **Use** For merging Adabas data with performance data from other systems. **Values** The values range from 00 to 95. **Alias** QTRHR **Column Heading** QTR/HR

Field	Description
RABN-RANGE	Indicates the user-defined RABN interval, that is, the user-defined name from the GLOBALS RABN-RANGE parameter for each physical I/O to Associator, Data Storage, or Work. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No Yes Yes A12 A12 No **Derivation** Each RABN in the IOLIST is compared to the RABN-RANGES defined in the GLOBALS statement. A summary by RABN-RANGE has data accumulated by physical I/O rather than command, so that the total count may be greater or less than the number of commands. The Adabas limit of forty RABNs logged per command may result in fewer I/Os being reported than actually occurred. Adabas I/O logging must be active. In the case of Command Logs written under control of Unicenter CA-APAS UEX4, the command must satisfy the selection criteria of a LOG statement that specified I/O List logging. **Use** For detailed analysis of disk usage by pack or pack area. This field should only be used in a request such as: `SUMMARIZE COUNT BY RABN-RANGE . . .;` **Values** The number of physical reads and writes from the IO-LIST, for named RABN ranges defined by the user in the GLOBALS statement. The field is built from the ten-character name, a space, and the letter W for write or the letter R for read. RABNs that do not fall within a specified interval are given range names of ASSO OTHER, DATA OTHER, or WORK OTHER. **Alias** RABNR **Column Heading** RABN/RANGE

Field	Description
RABNR	See RABN-RANGE.
RB	See RECORD-BUFFER.
RBL	See RBLENGTH.
RBLENGTH	Indicates the record buffer length from the Adabas Control Block. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No Yes Yes B2 N5 Yes D **Use** May be compared to the REC-LENGTH-DEC field to identify excessively long buffers. **Alias** RBL **Column Heading** RBL
REC-HOLD-CNT	See HOLD-COUNT.
REC-LENGTH-COMP	See RECORD-LENGTH.
REC-LENGTH-COMP-REPORTED	See RECORD-LENGTH-REPORTED.

Field	Description
REC-LENGTH-DEC	Indicates the length of the decompressed data. This field is the second two bytes of the ADDITIONS-2 field of the Adabas Control Block. It contains the decompressed data length in the record buffer for read commands, corrected as described in the description of the ADDITIONS-2 field. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No Yes Yes B2 N5 Yes F **Use** Compare to the RBLENGTH to detect excessive record buffer sizes. Requests referencing this field should contain WHERE command selection criteria excluding any response code greater than zero and any command code that does not involve processing of a compressed record. **Values** See the discussion of the ADDITIONS-2 field. **Column Heading** REC/LEN/DECMP
REC-LENGTH-DEC-REPORTED	Indicates the length of the decompressed data. This field is the second two bytes of the ADDITIONS-2 field of the Adabas Control Block. It contains the decompressed data length in the record buffer for read commands. The difference between this field and REC-LENGTH-DEC is that this field is not altered to give a zero value when a decompressed data length is not valid. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No Yes Yes B2 N5 Yes F **Use** See REC-LENGTH-DEC **Column Heading** RECL/DECMP/RPTD

Field	Description
RECORD-BUFFER	Indicates the contents of the record buffer for the command. Adabas record buffer logging must be active. In the case of Adabas Command Logs written under control of Unicenter CA-APAS UEX4, the command must have satisfied the selection criteria of a LOG statement that specified record buffer logging. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No Yes Yes A Dump No **Use** For debugging purposes. **Alias** RB **Column Heading** RECORD BUFFER

Field	Description
RECORD-LENGTH	Indicates the length of the compressed record. This field is the first two bytes of the ADDITIONS-2 field of the Adabas Control Block. It contains the compressed record length for read or update commands, corrected as described in the ADDITIONS-2 field. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No Yes Yes B2 N5 Yes F **Use** Detects records that are close to the maximum size (Data Storage block size). This value could be analyzed to determine how effective blocking is, and whether record types should be broken into separate physical files. MIN, MAX, and MEAN functions of the compressed record length may be of interest. Requests referencing this field should contain WHERE command selection criteria excluding any response code greater than zero and any command code that does not involve processing of a compressed record. **Values** See the discussion of the ADDITIONS-2 field. **Alias** REC-LENGTH-COMP **Column Heading** REC/LEN/COMP

Field	Description
RECORD-LENGTH-REPORTED	Indicates the length of the compressed record. This field is the first two bytes of the ADDITIONS-2 field of the Adabas Control Block. It contains the compressed record length for read or update commands. This field differs from REC-LENGTH-COMP in that it is not set to zero for commands and response codes for which a compressed record length is not valid. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No Yes Yes B2 N5 Yes F **Use** See RECORD-LENGTH. **Alias** REC-LENGTH-COMP-REPORTED **Column Heading** RECL/COMP/RPTD

Field	Description
RESPONSE-CLASS	Indicates the error category. This is a classification of the Adabas response code into categories of errors. See Appendix B for a full discussion. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No Yes Yes B1 N2 No **Derivation** Looked up in the RSPCL table. **Use** Used as a first-level breakdown of errors for exception reporting. Different fields in separate requests can be displayed for different classes of errors. **Values** The values range from 0 to 10. **Alias** RSPCL **Column Heading** R/C/L

Field	Description
RESPONSE-CODE	Indicates the response code returned by Adabas in the Adabas Control Block. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No Yes Yes B2 N3 No **Values** The values range from zero to 255. Unicenter CA-APAS can never report response codes returned by the Adabas SVC or link routines because the call never reaches the Adabas nucleus and therefore is never included in Adabas Command Logging. Response codes that may go unreported for this reason include: 22 146 148 151 152 250-255 **Alias** RSP **Column Heading** RSP
RSP	See RESPONSE-CODE.
RSPCL	See RESPONSE-CLASS.
SB	See SEARCH-BUFFER.
SB-ERR-FIELD	See FB-ERR-FIELD.
SB-FLD	See DESCRIPTOR.
SBL	See SBLENGTH.

Field	Description
SBLENGTH	Indicates the search buffer length from the Adabas Control Block. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No Yes Yes B2 N5 Yes F **Use** May be analyzed for excessive buffer length. **Alias** SBL **Column Heading** SBL
SEARCH-BUFFER	Indicates the contents of the search buffer for the command. Adabas search buffer logging must be active. In the case of Adabas Command Logs written under control of Unicenter CA-APAS UEX4, the command must have satisfied the selection criteria of a LOG statement that specified search buffer logging. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No Yes Yes A Dump No **Use** Important in analyzing badly performing search commands. Look for descriptors that do not have a good distribution of values. Identify superdescriptor candidates. May be used to analyze and identify descriptor use. **Alias** SB **Column Heading** SEARCH BUFFER

Field	Description
SM-LEVEL	This field contains the SM level of the Unicenter CA-APAS Data Collector that is executing under the Adabas User Exit 4. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No Yes No No A4 A4 No **Column Heading** APAS/SM
STATUS-CODE	See IMS-STATUS-CD.
STCK	See STCK-TIME.
STCK-TIME	Indicates the command time stamp placed on the Command Log record by the Adabas nucleus. This represents the time that the command completed execution in an Adabas thread. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No Yes Yes B4 H8 No **Use** Not ordinarily of interest. **Values** See the IBM *Principles of Operation* manual. Note that the value provided by Adabas is only the high-order portion of the STCK-TIME value, not a full STCK-TIME value. **Alias** STCK **Column Heading** STCK/TIME

Field	Description
SUB-CODE	Indicates the sub-code of the response code that is returned by Adabas in the ADDITIONS-2 field for some response codes. This field is zero unless Adabas returns a sub-code for the response code. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No Yes Yes B2 N3 No **Use** Exception reporting and debugging. **Values** See the *Adabas Messages and Codes Manual* from Software AG for a description of the sub-code values. **Alias** SUBCODE **Column Heading** SUBCODE
SUBCODE	See SUB-CODE.

Field	Description
SYSTEM-OVERHEAD	Indicates the amount of time, that is, the "wall-clock" time associated with a command that is not actual execution time in a thread. This would include inter-task communication time (SVC), time spent in the Command Queue prior to selection for execution in a thread, Adabas throwback(s), interrupt and I/O time. Applicable to CLOG layout 5 only. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes Yes Yes No Yes B4 N2.5 Yes D **Derivation** Computed by subtracting the THREAD-TIME from the TOT-CMD-TIME. **Use** Determine what portion of the total command time was spent in overhead. **Values** The value is shown in seconds. **Column Heading** SYSTEM/OVHEAD
TERM	See TERMINAL-ID.
TERM-ID2	See CICS-TERMINAL-ID.

Field	Description
TERMINAL-ID	Indicates the user-id for TP users. This field is the unique identifier given to the user by the link module for the TP monitor. It usually corresponds to the terminal identifier for the system. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No Yes Yes B4 A4 No **Use** To identify the terminal which issued online commands. Usually not suitable for COM-PLETE terminals; use TID instead. **Values** See the terminal identifier scheme for your TP monitor. **Alias** TERM **Column Heading** TERM/ID
THD	See THREAD

Field	Description
THREAD	Indicates the Adabas thread number of command, that is, the number of the nucleus thread in which the command was executed. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No Yes Yes B1 N3 No **Use** Performance data by thread is occasionally of interest. Analyzing mean command durations for similar commands in different threads may reflect CPU-bound versus I/O-bound conditions for Adabas during peak loads. Analyzing relative mix of batch versus TP users by individual thread may reflect degradation of online response by concurrent batch processing. **Values** 1 to Number of threads (ADARUN NT parameter). **Alias** THD **Column Heading** T/H/D

Field	Description
THREAD-TIME	Indicates command time in thread excluding interrupt and I/O time. That is, the "wall-clock" time for the execution of a command as measured by Adabas. Applicable to CLOG layout 5 only. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No Yes Yes B4 N2.5 Yes D **Derivation** Computed from the 16-microsecond units given by Adabas. **Use** Determine what portion of the total command time was spent in actual command execution. **Values** The value is shown in seconds. **Column Heading** THREAD/TIME

Field	Description
TID	Indicates the COM-PLETE terminal id. This field is the right half-word of the COMM-USERID. The COM-PLETE link module sets these two bytes to the COM-PLETE terminal number. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No Yes Yes B2 N4 No **Use** To identify the terminal in a COM-PLETE environment. Note that with dynamic terminal allocation, the TID value varies for a given USER-ID or LU-NAME. **Values** 1 to Maximum TID **Column Heading** TID
TIME	Indicates the time the command was completed from the command time stamp. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No Yes Yes A6 A6 No **Derivation** See DATE-TIME **Use** Shows when thread processing of a command ended. **Values** The values are the hour, minute, and second in the format: hhmmss. **Column Heading** TIME

Field	Description
TOT-CMD-TIME	Indicates the total amount of time of a command from the time the application called Adabas until the end of processing in the thread. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes Yes Yes No Yes B4 N2.5 Yes D **Derivation** Computed from the time stamp generated by the ACBX facility and the time stamp given by Adabas as when the request was logged. **Use** Watch for unusually long commands. **Values** The value is shown in seconds. **Column Heading** TOTAL/CMD
TOT-IO	See TOTAL-IO.

Field	Description
TOTAL-IO	Indicates the total physical I/O count. This is the sum of the DATA-IO, ASSO-IO and WORK-IO fields. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No Yes Yes B2 N4 Yes F **Derivation** Sum of other I/O fields. **Use** The best indication of resource use by a particular command. A major determinant of Adabas performance. **Alias** TOT-IO **Column Heading** TOT/IO
TP-TRANS-ID	Indicates CICS transaction identification number, that is, the CICS task number assigned to this transaction. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes Yes Yes No Yes P7 N5 No **Use** Can be used to identify a single interaction with the user when pseudo-conversational programs such as Natural are used. **Column Heading** TP/TRANS/ID

Field	Description
TP-TRANS-NAME	Indicates the CICS transaction name. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes Yes Yes No Yes A4 A4 No **Use** Identifies the current transaction as invoked by the user. **Column Heading** TP/TRANS/NAME
TRANSACTION-DURATION	This field contains the transaction duration in seconds. This field is non-zero only for ET commands. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No Yes Yes B4 N5.2 Yes D **Derivation** Adabas sets the ISN Quantity Field to transaction duration in 16 microsecond units for ET commands. This field is converted to seconds. **Alias** TRANS-DUR **Column Heading** TRANS/DURAT/SECS
TRANS-DUR	See TRANSACTION-DURATION.
TRIG-ADD2-FLD1	See TRIGGER-ADD2-FLD1.
TRIG-ADD2-FLD2	See TRIGGER-ADD2-FLD2.
TRIG-CMD	See TRIGGER-COMMAND-TYPE.
TRIG-FLD	See TRIGGER-FIELD.
TRIG-FLAG	See TRIGGER-FLAG.

Field	Description
TRIG-PART	See TRIGGER-PARTICIPATION.
TRIG-POINT	See TRIGGER-POINT.
TRIG-PROC-NAME	See TRIGGER-PROC-NAME.
TRIG-RSP-CODE	See TRIGGER-RESPONSE-CODE.
TRIG-SUBCODE	See TRIGGER-SUBCODE.
TRIGGER-ADD2-FLD1	This field contains the first half of the ADDITIONS-2 field. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No * Yes B2 H4 No * Only valid when using Adabas Command Log layout 5. **Derivation** From the first two bytes of the Adabas Control Block ADDITIONS-2 field. **Use** All requests which reference this field should include: `". . .WHERE LOG-REC-TYPE = X'0005' OR EQ X'0006' . . ."` **Alias** TRIG-ADD2-FLD1 **Column Heading** TRIG/ADD2/FLD1

Field	Description
TRIGGER-ADD2-FLD2	This field contains the last two bytes of the ADDITIONS-2 field. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No * Yes B2 H4 No * Only valid when using Adabas Command Log layout 5. **Derivation** From the last two bytes of the Adabas Control Block ADDITIONS-2 field. **Use** All requests which reference this field should include: `". . .WHERE LOG-REC-TYPE = X'0005' OR EQ X'0006' . . ."` **Alias** TRIG-ADD2-FLD2 **Column Heading** TRIG/ADD2/FLD2

Field	Description
TRIGGER-COMMAND-TYPE	This field contains the type of command issued by the trigger command or procedure call. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No * Yes A1 A1 No * Only valid when using Adabas Command Log layout 5. **Derivation** From the fifth byte of the Adabas Control Block ADDITIONS-4 field. **Use** All requests which reference this field should include: `". . .WHERE LOG-REC-TYPE = X'0005' OR EQ X'0006' . . ."` **Values** Values include: "R" for read "U" for update "S" for store/add "D" for delete "P" for procedure call **Alias** TRIG-CMD **Column Heading** TRIG/CMD

Field	Description
TRIGGER-FIELD	This field contains the name of the field associated with the trigger. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No * Yes A2 A2 No * Only valid when using Adabas Command Log layout 5. **Derivation** From the seventh and eighth bytes of the Adabas Control Block ADDITIONS-4 field. **Use** All requests which reference this field should include: `". . .WHERE LOG-REC-TYPE = X'0005' OR EQ X'0006' . . ."` **Alias** TRIG-FLD **Column Heading** TRIG/FLD

Field	Description
TRIGGER-FLAG	This field contains information about the CALLNAT parameters category and the record buffer access status. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No * Yes B1 H2 No * Only valid when using Adabas Command Log layout 5. **Derivation** From the sixth byte of the Adabas Control Block ADDITIONS-4 field. **Use** All requests which reference this field should include: `". . .WHERE LOG-REC-TYPE = X'0005' OR EQ X'0006' . . ."` **Values** Values include: "1" for no parameters "2" for response code only "4" for control information "8" for special stored procedure parameters "10" for record buffer access "80" for record buffer update **Alias** TRIG-FLAG **Column Heading** TRIG/FLG

Field	Description
TRIGGER-PARTICIPATION	This field indicates of the trigger is asynchronous or synchronous in relation to the initiating Adabas command and if the trigger is synchronous, whether it is participating or non-participating in the logic of the initiating Adabas command. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No * Yes A1 A1 No * Only valid when using Adabas Command Log layout 5. **Derivation** From the forth byte of the Adabas Control Block ADDITIONS-4 field. **Use** All requests which reference this field should include: `". . .WHERE LOG-REC-TYPE = X'0005' OR EQ X'0006' . . ."` **Values** Values include: "A" for asynchronous "N" for non-participating "P" for participating **Alias** TRIG-PART **Column Heading** TRIG/PART

Field	Description
TRIGGER-POINT	This field indicates if the command was created by a pre-trigger or post-trigger, or from a stored procedure. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No * Yes A1 A1 No * Only valid when using Adabas Command Log layout 5. **Derivation** From the third byte of the Adabas Control Block ADDITIONS-4 field. **Use** All requests which reference this field should include: `". . .WHERE LOG-REC-TYPE = X'0005' OR EQ X'0006' . . ."` **Values** Values include: "P" for pre-trigger "R" for procedure call "S" for post-trigger **Alias** TRIG-POINT **Column Heading** TRIG/POINT

Field	Description
TRIGGER-PROC-NAME	This field contains the name of the procedure invoked. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No * Yes A8 A8 No * Only valid when using Adabas Command Log layout 5. **Derivation** From the ADDITIONS-3 field of the Adabas Control Block. **Use** All requests which reference this field should include: `". . .WHERE LOG-REC-TYPE = X'0005' OR EQ X'0006' . . ."` **Alias** TRIG-PROC-NAME **Column Heading** TRIG/PROC/NAME

Field	Description
TRIGGER-RESPONSE-CODE	This field contains the response code returned from the Natural trigger driver. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No * Yes A2 A2 No * Only valid when using Adabas Command Log layout 5. **Derivation** From the first and second bytes of the Adabas Control Block ADDITIONS-4 field. **Use** All requests which reference this field should include: `". . .WHERE LOG-REC-TYPE = X'0005' OR EQ X'0006' . . ."` **Alias** TRIG-RSP-CODE **Column Heading** TRIG/RSP/CODE

Field	Description
TRIGGER-SUBCODE	This field contains a subcode that indicates the type of trigger that was fired if the trigger procedure return code is non-zero. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No * Yes B2 N3 No * Only valid when using Adabas Command Log layout 5. **Derivation** From the third and fourth bytes of the Adabas Control Block ADDITIONS-4 field. **Use** All requests which reference this field should include: `". . .WHERE LOG-REC-TYPE = X'0005' OR EQ X'0006' . . ."` **Values** Values include: "15" indicates a pre-command trigger "16" indicates a post-command trigger **Alias** TRIG-SUBCODE **Column Heading** TRIG/SUB/CODE
TSO-JOB-ID	See TSO-JOBID.

Field	Description
TSO-JOBID	This field is the last 5 bytes of the GLOBAL-USER field. For TSO users, this is the JES job number that follows the three characters "TSO", which is inserted by Adabas. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No * * Yes A5 A5 No * See the explanation of GLOBAL-ID. **Alias** TSO-JOB-ID **Column Heading** TSO/JOBID
UA	See USER-AREA.
UEXB	See USER-INFO-AREA.
UEXBL	See LNUINFO.
UEXITB	See USER-INFO-AREA.

Field	Description
UEXITB-USERA	Indicates the User Exit B User Field A. An alphanumeric field that is available for users to use in a User Exit B link routine. The data is passed to Adabas with the Unicenter CA-APAS User Information Data Area, the ACBX data. The field can be used like the predefined fields. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes Yes Yes No Yes A10 A10 No **Derivation** By a User Exit B routine. **Use** Possible uses are for computing charges, classifying applications, or computer special resource use function. **Values** User derived. **Column Heading** UEXITB/USERA

Field	Description
UEXITB-USERB	Indicates the User Exit B User Field B. An alphanumeric field that is available for users to use in a User Exit B link routine. The data is passed to Adabas with the Unicenter CA-APAS User Information Data Area, the ACBX data. The field can be used like the predefined fields. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes Yes Yes No Yes A10 A10 No **Derivation** By a User Exit B routine. **Use** Possible uses are for computing charges, classifying applications, or computer special resource use function. **Values** User derived. **Column Heading** UEXITB/USERB

Field	Description
UEXITB-USER1	Indicates the User Exit B User Field 1. This is a numeric field that is available for users to use in a User Exit B link routine. The data is passed to Adabas with the Unicenter CA-APAS User Information Data Area, the ACBX data. The field can be used like the predefined fields. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes Yes Yes No Yes B4 N7 Yes F **Derivation** By a User Exit B routine. **Use** Possible uses are for computing charges, classifying applications, or computer special resource use function. **Values** User derived. **Column Heading** UEXITB/USER1

Field	Description
UEXITB-USER2	Indicates the User Exit B User Field 2. This is a numeric field that is available for users to use in a User Exit B link routine. The data is passed to Adabas with the Unicenter CA-APAS User Information Data Area, the ACBX data. The field can be used like the predefined fields. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes Yes Yes No Yes B4 N7 Yes F **Derivation** By a User Exit B routine. **Use** Possible uses are for computing charges, classifying applications, or computer special resource use function. **Values** User derived. **Column Heading** UEXITB/USER2
UEXITBL	See LNUINFO.
UINFO	See USER-INFO-AREA.

Field	Description
UQE-COMMAND-COST	This field contains a cost calculated for the individual command. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No Yes No Yes B4 N2.5 Yes D **Derivation** Calculated by multiplying the UQE-CPU-TIME, which includes the UQE-CPU-FACTOR, in seconds by the UQE-CPU-COST and adding the TOTAL-IO multiplied by the UQE-IO-COST. UQE-CPU-FACTOR, UQE-CPU-COST and UQE-IO-COST are parameters in the GLOBALS statement. This field in computed only if referenced in an EXTRACT or SUMMARIZE request. It does not appear in COPY file records written before computation has begun. Computation continues after the request has been deleted. **Use** SUM (UQE-COST) may be used to report changes for Adabas usage. **Alias** UQE-COST **Column Heading** UQE-CMD/COST
UQE-COST	See UQE-COMMAND-COST.
UQE-CPU	See UQE-CPU-TIME.

Field	Description
UQE-CPU-TIME	Contains the estimated amount of CPU time that the Adabas nucleus spent processing a command. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No Yes No * B4 N2.3 Yes D **Derivation** Adabas keeps the total CPU time for each user in that user's UQE. For each command this field is the increase in the cumulative CPU time kept in the UQE. This field in computed only if it is referenced in an EXTRACT or SUMMARY request. It is absent from COPY file records before computation of the field has begun. Computation of the field continues even after all requests referencing it have been deleted. **Use** This value is only as accurate as the calculations Adabas performs, which might omit a portion of the total CPU time spent by the operating system processing the command. The UQE-CPU-FACTOR parameter may be adjusted in an attempt to get the total of all commands' estimated CPU time to correspond to the CPU time reported by the operating system accounting for the Adabas session. Before using this value, users are advised to contact their Software AG representative to determine how accurately it reflects actual CPU time. In initial release of Adabas 5, this value was actually elapsed time, which is a highly inaccurate estimate of CPU time. **Values** The values are expressed in seconds with three (3) decimal places; that is, the B4 value is in milliseconds. **Alias** UQE-CPU **Column Heading** UQE-CPU/SECS
UQETID	See COMM-USERID.

Field	Description
USER-AREA	Represents the Adabas Control Block field for users, that is, the four-byte area in the Adabas Control Block available to the installation for its own use. Usually filled in by local modifications to the link modules. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No Yes Yes A4 A4 No **Use** Depends on installation's use of the field. **Alias** UA **Column Heading** USER/AREA
USER-COMMAND-COST	This field may be computed in the Derived Fields User Exit. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** * No No Yes Yes B4 N3.3 Yes D * This field is present only if computed in a user-supplied Derived Fields User Exit. **Derivation** A user written Derived Fields User Exit is used to compute this field. The Derived Fields User Exit should insert an estimate of the cost of the command into the derived fields area DERUSRCS. **Values** The field has units of 1/1000 of a cent. **Alias** USER-COST **Column Heading** USER-CMD/COST

Field	Description
USER-COST	See USER-COMMAND-COST.
USER-CPU	See USER-CPU-TIME.
USER-CPU-TIME	This field may be computed in the Derived Fields User Exit. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** * No No Yes Yes B4 N2.3 Yes D * This field is present only if computed in a user-supplied Derived Fields User Exit. **Derivation** A user written Derived Fields User Exit is used to compute this field. The Derived Fields User Exit should insert an estimate of the command's CPU time into VALUES: or the derived fields area DERUSRCP. **Values** The CPU time is measured in units of tens of microseconds, that is, each unit will equal $1x10^{-5}$ seconds. **Alias** USER-CPU **Column Heading** USER-CPU/SECS

Field	Description
USER-ID	Contains the user-id from the system accounting fields for batch programs or operator-id for TP systems. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes Yes Yes No Yes A8 A8 No **Derivation** The &ACCOUNT option in the DBGLNxxx link routine controls the contents of the field. Options include RACF, JMR, JOBC, and NATLOGON for batch jobs; RACF, TCA, TCT, NATLOGON, and INITID for CICS; and IMSUSER and NATLOGON for IMS. For COM-PLETE, the &USERIDL option is used to control the length of the COM-PLETE user id to be used. **Use** Identifies the terminal user; may contain installation account codes for batch jobs. **Alias** OPERATOR-ID, JMR-USER-ID **Column Heading** USER/ID

Field	Description
USER-INFO-AREA	Contains the contents of the User Information Area, the UEXB buffer. The User Information Area is created by the Adabas link module. It contains data from the UEXITB routine provided it has been installed in the link module. It can also contain data from any other UEXITB routine and from any Adabas user exits called prior to the User Exit 4. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No Yes Yes A Dump No **Use** Detailed command analysis. Extract data stored by UEXITB and Adabas user exits. **Alias** UINFO, UEXITB, UEXB **Column Heading** USER INFO AREA
USER-PRIORITY	Contains the selection priority of the command for execution by the nucleus. This is usually set by the SVC to the system priority value of the task that issued the command. Special priorities may be assigned to particular user-ids by the DBA. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No Yes Yes B1 N3 No **Use** ENQ-TIME versus priority may indicate whether adjusting priorities would be helpful. **Alias** PRTY **Column Heading** USER/PRTY

Field	Description
USER-TYPE	This field is set by the Unicenter CA-APAS User Exit B to indicate the type of user that issued the command. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes Yes Yes No Yes A3 A3 No **Derivation** This field is set by the Unicenter CA-APAS ACBX link routine. **Use** System use may be broken down by USER-TYPE. Classification and analysis of TSO use helps to identify additional interactive use. **Values** The values can be "BT" for Batch, "TP" for TP user (CICS or COMPLETE), or "TSO" for TSO user. **Column Heading** USER/TYPE
USERA	User Field A is an alphanumeric field that may be computed in a user exit routine and then used like the predefined fields. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No No Yes A8 A8 No **Derivation** By a user-exit routine. **Use** Possible uses are for computing charges, classifying applications, or computing special resource-use functions. **Values** Values are user-defined. **Column Heading** USERA

Field	Description
USERB	User Field B is an alphanumeric field that may be computed in a user exit routine and then used like the predefined fields. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No No Yes A8 A8 No **Derivation** By a user-exit routine. **Use** Possible uses are for computing charges, classifying applications, or computing special resource-use functions. **Values** The values are user-defined. **Column Heading** USERB
USERC	User Field C is an alphanumeric field that may be computed in a user exit routine and then used like the predefined fields. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No No Yes A20 A20 No **Derivation** By a user-exit routine. **Use** Possible uses are for computing charges, classifying applications, or computing special resource-use functions. **Values** The values are user-defined. **Column Heading** USERC

Field	Description
USER1	User Field 1 is a numeric field that may be computed in a user exit routine and then used like the predefined fields. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No No Yes B4 N7 Yes F **Derivation** By a user-exit routine. **Use** Possible uses are for computing charges, classifying applications, or computing special resource-use functions. **Values** The values are user-defined. **Column Heading** USER1
USER2	User Field 2 is a numeric field that may be computed in a user exit routine and then used like the predefined fields. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No No Yes B4 N7 Yes F **Derivation** By a user-exit routine. **Use** Possible uses are for computing charges, classifying applications, or computing special resource-use functions. **Values** The values are user-defined. **Column Heading** USER2

Field	Description
USER3	User Field 3 is a numeric field that may be computed in a user exit routine and then used like the predefined fields. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No No Yes B4 N7 Yes F **Derivation** By a user-exit routine. **Use** Possible uses are for computing charges, classifying applications, or computing special resource-use functions. **Values** The values are user-defined. **Column Heading** USER3
VALUE-BUFFER	Contains the contents of value buffer for the command. Adabas value buffer logging must be active. In the case of Adabas Command Logs written under control of Unicenter CA-APAS UEX4, the command must have satisfied the selection criteria of a LOG statement that specified value buffer logging. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No Yes Yes A Dump No **Use** Identifies particular values in poorly performing search commands. Use for debugging. **Alias** VB **Column Heading** VALUE BUFFER

Field	Description
VB	See VALUE-BUFFER.
VBL	See VBLENGTH.
VBLENGTH	Contains the value buffer length from the Adabas Control Block. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No Yes Yes B2 N5 Yes F **Use** Used to look for excessive buffer length. **Alias** VBL **Column Heading** VBL
VERSION	This field is the version number of the Unicenter CA-APAS Data Collector that is executing under the Adabas User Exit 4. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No Yes No No A4 A4 No **Column Heading** APAS/VERSION

Field	Description
WEEK	Contains the week of the calendar year from the command time stamp. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No Yes Yes N2 A2 No **Use** Used in Performance History to build up a time series. **Values** The values range from 01 to 53. **Alias** WK **Column Heading** WK

Field	Description
WEEKDAY	Contains the day of the week from the command time stamp. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No Yes Yes N1 A1 No **Use** Used to build a use profile by day of week. **Values** The values can be one of the following: "1" for Sunday "2" for Monday "3" for Tuesday "4" for Wednesday "5" for Thursday "6" for Friday "7" for Saturday **Alias** DAY-OF-WEEK **Column Heading** WK/DY
WK	See WEEK.

Field	Description
WORK-IO	Contains the Work I/O count, that is, the number of physical I/O operations Adabas reported it performed to read and/or write Work blocks for the processing of the command. The value may include a buffer flush. See the BUFFER-FLUSHES field. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No Yes Yes B2 N4 Yes F **Use** I/O to the Work data set is to the protection area (update commands) and to the work area (saved ISN Lists). Excess I/O to Work can be affected by the ADARUN LWKP2 and NSISN parameters and the Work block size. **Column Heading** WORK/IO
YEAR	Contains the year from the command time stamp. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No No Yes Yes A2 A2 No **Alias** YR **Column Heading** YR
YR	See YEAR.

Field	Description
ZAP-MAP	This field contains the update level for updates 00 through 351 for this version and SM of the Unicenter CA-APAS Data Collector that is being executed under the Adabas User Exit 4. **------- Adabas V7 ----- File Rept Sum Acc** **Gvn ACBX UX4 Log COPY Fmt Fmt Fld Size** Yes No Yes No No A352 A352 No **Derivation** As each update is applied to the Data Collector, a field is updated which shows the level of each update. There is space for 352 updates, numbered 00 through 351. Each update is assigned a one-byte binary field whose displacement corresponds to the update number. A value of zero indicates that an update for the number has not been applied. A non-zero value indicates the level of the corresponding update. **Use** Intended only for use in special requests. The one-byte binary values are automatically replaced by the Data Collector with the period, ., (reflecting non-printable characters) in all PRINT TO or INSIGHT -LINES output. To see the zap levels at an Unicenter CA-APAS terminal, do the following: ■ START ZAPDECL ■ START VERZAPS ■ Execute the Unicenter CA-APAS command APASZAPS. **Column Heading** APAS ZAP MAP

Appendix

A

Output Record Header Formats

This appendix documents the format of records output as a result of specifying the OUTPUT-FILE parameter of EXTRACT and SUMMARIZE requests. All records are output in variable-blocked format. The blocksize depends on the device-type of the output device.

OUTPUT-FILE Record Header

Every OUTPUT-FILE record written by the Data Collector contains an OUTPUT-FILE header. The format is shown below. You should examine member OUTHDRA to obtain the latest version of this DSECT.

```
DSECT for OUTPUT-FILE Record Header

OUTRECRD DSECT
* THIS DSECT MAPS THE START OF AN OUTPUT FILE RECORD.
OUTRECDW DS       XL4          RECORD DESCRIPTOR WORD.
OUTEXNAM DS       CL8          EXTRACT NAME.
OUTDBNAM DS       CL8          DATABASE NAME.
OUTDBID  DS       XL2          DATABASE ID.
OUTCPUID DS       CL8          CPU ID.
OUTRDATE DS       CL8          REPORT DATE.
OUTRTIME DS       CL6          REPORT TIME.
* INTERVAL SUMMARY DESCRIPTORS
OUTINSTD DS       CL8          INTERVAL START DATE.
OUTINSTT DS       CL6          INTERVAL START TIME.
OUTINDAY DS       ZL1          DAY OF WEEK(1-7).
OUTINWKY DS       ZL2          WEEK OF YEAR(0-53).
OUTINEND DS       CL8          INTERVAL END DATE.
OUTINENT DS       CL6          INTERVAL END TIME.
OUTFRSDT DS       CL8          1ST RECORD IN INTERVAL; DATE.
OUTFRSTI DS       CL6          1ST RECORD IN INTERVAL; TIME.
OUTLSTDT DS       CL8          LAST RECORD IN INTERVAL; DATE.
OUTLSTTI DS       CL6          LAST RECORD IN INTERVAL; TIME.
         DS       X            RESERVED.
OUTBRKLV DS       XL2          BREAK LEVEL.
* DETAIL RECORD DESCRIPTORS
         ORG OUTINSTD
OUTRECDT DS       CL8          RECORD DATE FOR DETAIL RECORD.
OUTRECTI DS       CL6          RECORD TIME FOR DETAIL RECORD.
OUTRECDY DS       ZL1          DAY OF WEEK(1-7).
OUTRECWK DS       ZL2          WEEK OF YEAR(0-53).
```

```
* SUMMARY DESCRIPTORS           FOR SUMMARIES WITHOUT INTERVALS.
         ORG      OUTINSTD
OUTSMSTD DS       CL8           SUMMARY START DATE.
OUTSMSTT DS       CL6           SUMMARY START TIME.
OUTSMDAY DS       ZL1           DAY OF WEEK(1-7).
OUTSMWKY DS       ZL2           WEEK OF YEAR(0-53).
OUTSMEND DS       CL8           SUMMARY END DATE.
OUTSMENT DS       CL6           SUMMARY END TIME.
         ORG
OUTRECTY DS       CL2           RECORD TYPE.
* IS = INTERVAL SUMMARY
* GT = GRAND TOTAL OF INTERVAL SUMMARY
* SU = SUMMARY WITHOUT INTERVALS
* DE = DETAIL RECORD
OUTHDRLN EQU *-OUTRECRD HEADER LENGTH.
OUTDATA  DS       0X            START OF DATA.
```

The following is a brief explanation of the fields in this header:

Field	Description
OUTRECDW	Standard variable length record header. COBOL and other high-level language programs do not see this word.
OUTEXNAM	Name coded as the label for the request.
OUTDBNAM	Eight-character DBNAME coded as a GLOBALS parameter.
OUTDBID	Two-byte binary database-id taken from the Adabas nucleus.
OUTCPUID	Eight-character CPU-ID coded as a GLOBALS parameter.
OUTRDATE **OUTRTIME**	Date and time of the Data Collector run which created this file.

The following fields apply to interval summaries:

Field	Description
OUTINSTD OUTINSTT	Date and time the interval started.
OUTINDAY	Day of the week the interval started, with Sunday being day 1.
OUTINWKY	Week of the year the interval started, with the first week being week 0.
OUTINEND **OUTINENT**	Date and time the interval ended.
OUTFRSDT **OUTFRSTI**	Date and time of the first record in the interval (zero if no records in interval).

Field	Description
OUTLSTDT **OUTLSTTI**	Date and time of the last record in the interval (zero if no records in interval).
OUTBRKLV	Break level for the record being written. This ranges from 1 for the highest level subtotal record, to n, where n is the number of control-break fields and the most detailed level of break. The grand total record for the interval has a break level of 0. If SUBTOT=OFF is specified, all have a break level of n. Fields 1-4 are zero if no Command Log records are processed for the request. An example of how this could occur is an MPM session during which no Adabas commands are processed.

Note: Fields 1-4 are zero if no Command Log records are processed for the request. An example of how this could occur is an MPM session during which no Adabas commands are processed.

The following fields redefine the header at OUTINSTD for detail records:

Field	Description
OUTRECDT **OUTRECTI**	Detail record date and time.
OUTRECDY	Day of the week, with Sunday being day 1.
OUTRECWK	Week of the year, with the first week being week 0.

The following fields redefine the header at OUTINSTD for summaries with no interval:

Field	Description
OUTSMSTD **OUTSMSTT**	Date and time of first record read for the summary.
OUTSMDAY	Day of the week, with Sunday being day 1.
OUTSMWKY	Week of the year, with the first week of the year being week 0.
OUTSMEND **OUTSMENT**	Date and time of the last record read for the summary.

Field	Description
OUTRECTY	Type of the record. This is a two-byte alpha field with the following values: ■ IS—interval summary record ■ GT—the grand total summary record written at end of run for INTERVAL request ■ SU—table record for SUMMARIZE request with no INTERVAL ■ DE—detail record (EXTRACT request)

Output Fields—Variable Portion of Record

Fields are put in output records in the order that they are specified in the output request, except that control-break (BY...) fields are always placed at the beginning of the variable portion of the record.

Fields are output with a standard format (specified below) unless overridden by the output-format OF parameter of the field format specification.

Default Output Format—Standard Fields

In the absence of other criteria, fields are output in the same format as the input format. This applies to both fields in EXTRACT field lists and control break fields in SUMMARIZE requests. For information about input formats, see the chapter "Unicenter CA-APAS Data Fields".

Default Output Format—Control-Break Fields

Control break fields are output in the same format as standard extract fields except when subtotal records are being written. Subtotal records have the break fields which are totalled set to high-values (hexadecimal FFs). The grand total record for the table has all break fields set to high values.

Default Output Format—Buffers

Buffers (RB, FB, SB, VB, IB, IOL, and UEXB) are output to the record as variable length fields, preceded by a halfword inclusive length. Variable length buffers must be specified at the end of the output record.

Default Output Format—Summary Functions

Summary functions are always output to the file in a default packed format. This includes MAX and MIN functions on numeric fields.

The number of digits output is determined by taking the precision calculated for the output field as shown in "Unicenter CA-APAS Data Fields and Summary Functions" and rounding up to an odd number of digits, including decimal digits.

One is added if any decimal places are needed. For example, the output format for RATE(DURATION) is 7 (digits of DURATION) + 3 + 2 (digits and decimal places of RATE function) + 1 (for the decimals) = 13. No rounding up necessary, so the output format is P11.2 (13 total).

Appendix

B

RSPCLASS Definitions

This appendix explains which Adabas response codes are included in each of the numeric values that are developed by Unicenter CA-APAS Insight Monitor for Adabas (Unicenter CA-APAS) for the field, RSPCL. For more information on use of RSPCL in Unicenter CA-APAS requests, see its description in the chapter "Unicenter CA-APAS Data Fields," the Unicenter CA-APAS Data Fields section.

Physical and Logical Structure of RSPCLASS

The RSPCLASS csect within the APASINTM load module in the Unicenter CA-APAS load library consists of a 256-byte translate table. Each byte of the table, based on its relative position, represents one of the Adabas response code values 0-255. The first byte represents response code "0", the next byte response code "1", and so on, with the last byte representing response code "255".

Each byte in the table contains the value of one of the response code classes defined in the table shown below thus establishing the class to which its associated response code belongs.

The purpose of these class definitions is to allow convenient specification of exception reporting by response class instead of by individual response codes.

The class numbers and meanings reflect severity of actual or potential problems. The higher the response class number, the greater the severity. A few response codes have multiple possible meanings and could logically fall into multiple classes; each of these response codes has been assigned to the highest severity class into which it might fall.

Class	Meaning
0	All normal, non-error responses
1	User logic, data, or specification error
2	User error in format, record, search, or value buffer
3	Contention for record hold

Class	Meaning
4	Transaction or cluster time-out or CHKPT request denied
5	Conflict between sessions or file availability
6	Security violation
7	Inadequate Adabas resources
8	Database or file space problem
9	Possible database integrity problem
10	Undocumented response code

Default Class Assignments

The default Unicenter CA-APAS classification of response codes is described below. These assignments of response codes to classes may have been changed at your site; any such changes are reflected in the current contents of the RSPCLASS csect in the APASINTM load module.

Class 0 - Normal, Non-error Response Codes

This class includes those few response codes that reflect only conditions that involve no errors or significant circumstances worth reporting.

Example of response codes included are:

- 0 - successfully executed call
- 3 - logical end of file
- 148 - Adabas not active for specified database-id
- 207 - SAF Security completed phase 1 of logon

Class 1 - User Logic, Data or Specification Error

This class includes response codes reflecting a wide variety of routine errors stemming from flaws in application program design or coding. Normally these are of no interest for exception reporting as they would have been detected and resolved by the responsible programmer or users.

Class 2 - User Error in Format, Record, Search or Value Buffer

The response codes in this class, like those in Class 1, should typically be of no interest for exception reporting since they should always be resolved by the responsible programmers or users. These codes were isolated from those in Class 1 simply because these constituted a sizable group related specifically to the buffers mentioned.

Class 3 - Contention for Record Holds

Response codes 89, 144, and 145 are included in this category. These response codes may or may not be of interest to the DBA and/or user personnel. Possible benefits of reporting them might include:

- Identification of programs that request response 145, in particular, Natural programs running with WH=OFF
- Awareness of instances where specific ISNs fail to become available
- Awareness of situations where the Hold Queue is full
- Detection of throwbacks related to record holding situations
- Possible program logic errors

Except for applications where multiple concurrent users compete regularly for use of a single control record that must be read and incremented very frequently, the occurrence of these response codes is usually far less frequent than might be expected; so reporting them should normally not produce excessively long listings.

However, when there are user or vendor program bugs in this area, it isn't unusual for program looping to cause considerably inflated reports to be produced. For this reason, it is advisable to dedicate one particular exception report to this class and no others; this prevents other response codes from getting "lost" within a few inches of printout containing mostly 144 and 145 responses.

Class 4 - Transaction or Checkpoint Problems

Response code 9 indicates transaction back out. Depending on how common and acceptable these events are in your environment, you may or may not wish to have these response codes reported.

Even if the DBA does not care to see this type of report, user personnel may be interested in it in order to stay abreast of the frequency of time-outs and to have the opportunity to detect any situations that should be corrected.

Class 5 - Adabas Session or File Use Conflicts

Examples of response codes in this group are:

- 48 - several reasons for rejecting an OP command
- 64 - utility functions cannot be done on an Adabas system file
- 126 - a problem with multi-target communications
- 129 - attempt to perform unsupported cluster functions
- 130 - error in Sysplex cluster environment
- 228 - ASCII formatted user request issued against an EBCDIC formatted database

Class 6 - Security Violations

This group of response codes indicates attempts to access and/or update unauthorized files or fields. If you use password security you may find it useful to refine this type of reporting so as to have separate reports for individual files or applications according to differing data sensitivities or report distribution requirements. If multiple reports are desired, the most efficient approach would be to use a single output file and subsequently split that file into multiple printed reports.

Examples of response codes included in this category are:

- 200 - security violation or invalid cipher code
- 201 - specified password not defined
- 202 - attempt to use unauthorized file
- 203 - attempt to delete record in unauthorized file
- 208-209 - security user exit

Response code 17 is not included because it commonly does not represent a security violation. You may wish to move it to this category if you regard it as a significant potential indicator of a security violation.

Class 7 - Inadequate Adabas Resources

Certain Adabas response codes indicate that commands could not be processed due to unavailability of space in an Adabas work area. These instances indicate that either the DBA should increase the area size or something should be done with the application to reduce the rate of such commands or even eliminate them entirely.

Examples of response codes included in this category are:

- 1 - LS sort area exceeded or security by value
- 2 - NH Hold Queue full
- 6 - LREPL area is full
- 8 - WORK overflow
- 16 - UQE overflow
- 27 - LWP space unavailable or invalid SB, VB lengths
- 42 - Internal Format buffer too small
- 46 - NQCID exceeded
- 47 - NISNHQ exceeded
- 70 - LQ sequential process table full
- 71 - LI list table full
- 72 - NU User Queue full
- 73 - WORK area for resulting ISN lists full
- 74 - WORK area for intermediate ISN lists too small
- 83 - Hypertable overflow
- 87 - buffer pool locked
- 88 - LWP space not available
- 106 - prefetch buffer too small
- 107 - prefetch space error
- 110 - CID pool I is full
- 123 - an error reported by the Adabas cluster messaging service
- 151 - NC Command Queue full
- 152 - LU intermediate user buffer too small
- 154 - a trigger fired but the queue is full
- 160 - more than 20 blocks active in buffer pool for a command
- 162 - shortage of space for buffer pool header blocks

- 164 - more than 30 work areas are allocated
- 197 - unique descriptor pool too small
- 204 - password-pool overflow
- 218 - not enough space for UB
- 219 - no room in Save Area
- 255 - NAB attached buffer space full

Class 8 - Database or File Space Problems

Response codes in this category indicate some sort of space management problem for the database as a whole or for an individual file.

Examples of response codes included in this category are:

- 65 - space calculation error
- 66 - LU too large
- 75 - file has filled five extents of some type
- 76 - index structure overflow (six levels filled)
- 77 - no space to allocate a required extent
- 78 - AC file cannot be increased

Class 9 - Database Integrity Problems

Certain response codes indicate possible inconsistencies in the control information within the database.

Examples of response codes included in this category are:

- 67 - internal error executing Sx command
- 95 - I/O error to WORK LP area
- 96 - error during repair execution
- 97 - I/O error during buffer flush
- 99 - I/O error
- 161 - invalid RABN chain in buffer pool header list
- 163 - attempted duplicate RABN in RABN header chain
- 166 - error in inverted list
- 167 - possible coupling list out of sync
- 168 - internal CID for coupling processing missing

- 170 - a required RABN could not be found
- 171 - error in GCB (missing constant set)
- 172 - ISN was zero or too high for file
- 173 - invalid Data RABN
- 174 - L2/L5 starting Data RABN not found in A/C
- 176 - inverted list error or invalid subroutine call
- 177 - Address Converter and Data inconsistency
- 178 - FDT and an internal format buffer disagree
- 179 - Internal error in WORK Part 3
- 181 - unsuccessful transaction back out
- 182 - needed ET data not found in Work
- 183 - invalid internal ID of an I/O operation
- 184 - a phonetic name could not be found
- 185 - ADAM field could not be found
- 199 - inconsistency in an inverted list
- 210 - logical ID greater than 255
- 211 - invalid ID table index in UB
- 212 - invalid I/O buffer for internal command
- 251 - Adabas cluster processing error
- 252 - SVC post error or illogical ADAMINT call
- 253 - invalid buffer length or illogical ADAMINT call
- 254 - attached buffer overflow or illogical ADAMINT call

Response code 113 could possibly indicate a database integrity problem, but since it usually doesn't it has not been included in this class. (It is in Class 1.)

Class 10 - Undocumented Response Codes

This class includes all response codes that are not documented in the *Adabas Messages and Codes Manual*. If any of these codes ever occur, the DBA should report them to the Adabas vendor to inquire about their meanings.

Index

A

B

C

D

E

F

G

H

I

J

L

M

N

O

P

Q

R

S

T

U

V

W

Y

Z

Computer Associates™