
- "Access Manager"

Bruce Hunter

Microsystems, August 1983, p.58

(Retyped by Emmanuel ROCHE.)

Access Manager is a utility for the creation of database-managed systems. A product of Digital Research, this utility is a B+ tree (hierarchical) file system for the building and accessing of data and index files. Access Manager (or AM-80 for short) is designed to interface with Digital Research's languages such as the CBASIC Compiler, Pascal/MT+, and PL/I-80. It is presently furnished in an 8-bit version, and will be furnished in a 16-bit version. It is both CP/M and MP/M compatible, and has both file and record locking.

The mark of any good database system is its storage capability and its accessing speed. In an MP/M environment, AM-80 will allow file sizes of 32 MB each. The maximum number of open data files is 40. Also, the maximum number of index files is 40. 60 simultaneously open buffers are permissible. The result is well over a gigabyte of open storage. The CP/M 2.2 version is limited to 20 open data files of 8 MB each and 10 open index files. Because of the B+ tree organization, a substantial number of indices are stored at each node, and the number of passes through the index, to find a particular record, is unbelievably small. It is possible to search nearly 200,000 records in no more than four passes. Compare this to a binary tree that would require 18 passes for the same accomplishment. Data entered to index files is sorted sequentially. Once entered, data is never moved. The index files do all the housekeeping.

AM-80 is a series of routines (program files) that are linked to the application programs(s) at the final linkage. For a single-user system, a subroutine library and a buffer are all that need to be linked to the program. For multi-user systems, three components are required: the subroutine interface, the shared multiple-user background server, and multiple-user resident system processes. The codes generated by these files are reasonably small, requiring less than 12 KB for all the functions. Multi-user object code will be increased by no more than 20 KB.

Performance

In practice, Access Manager is relatively easy to use. A header file is included into the object code at run-time by way of a preprocessor command. For example, in PL/I the command:

```
%INCLUDE 'AM80EXTR.PLI';
```

will cause the header file to be read into the compiler as part of the object code. It saves about half a page of coding by declaring all the files and

variables external to the program.

There are only four functions necessary to call for system initialization. Data and index files are opened or closed with a single function call. The most intriguing part of Access Manager is the set of index search functions. The functions available are:

- GETKEY -- Searches for an exact match.
- SERKEY -- Searches for a match that will be equal or greater.
- BEFKEY -- Finds the key before the target key.
- AFTKEY -- Finds the key following the target key.
- FRSKEY -- Returns the first key in the index file.
- LASKEY -- Finds the last key in the index file.
- PRVKEY -- Finds the key immediately preceding the last entry found.
- NXTKEY -- Finds the key after the last key found.

(The last 2 functions are not recommended under MP/M.)

The distinction between GETKEY and SERKEY is that GETKEY expects an exact match, as in looking for an ID number, while SERKEY is less exact, as in searching for a name that could be subtly different, like "Alfred Newman" vs. "Alfred E. Newman".

The system not only allows for the deletion of data records, but also allows for recovery of the space taken up by the record by returning that space back to the system for re-use. Statistical functions are available, to show the number of records in use versus the actual number of records. Functions are also available for the setting and unlocking of file locks.

Keys are stored as unsigned two-byte integers for most applications but, for unusually large files, a four-byte integer can be used, with some additional housekeeping, to bring capacity to the gigabyte range. Data files may be of any reasonable length, as long as they are at least four bytes in length. The CP/M restriction of 128-byte logical records has been nicely overcome by the very sophisticated housekeeping performed by AM-80.

In practice, Access Manager is a pleasure to use. Once the preprocessor inclusion command has been put into the code, the variable names that are to be used in the code are declared (this not being necessary in CBASIC, of course). The number of buffers, keys, and index files are defined, and the LOCK request is or is not made. Then, the system is initialized. In all, the last few tasks will account for less than a dozen lines of code. Had this been done in BT-80 (the predecessor of AM-80) instead, a couple of pages would have to be ground out before seeing an executable line of code (something like programming in COBOL). In AM-80, you enter executable code at the bottom of the first page of code. Opening a couple of data files and a couple of index files takes a bit less than half a page more of code.

In my code, index and data records are written simultaneously, at the cost of a half dozen lines per write. Error codes are returned on nearly every AM-80 routine, and error reporting functions can be created with reasonable ease by the programmer, to keep the program from "falling down". The search routines are straightforward, and easily implemented. On the average, they take about four arguments each. For example:

record-number = NXTKEY (act-key, fil-no, lock, idxval);

A mailing list or a chart of accounts can be written with Access Manager in about 10 pages of code, using almost all of AM-80's facilities, including a procedure to rebuild the index files (should they become "corrupted" by a premature closing).

The documentation for Access Manager is a well-written 178-page manual. It tends to get just a bit obscure on the handling of numeric key values, and on the index value of the output argument, IDXVAL, but, other than that, it reads easily. There is a new manual coming out for the 16-bit version, AM-86, that is supposed to be even better, and it will also be the new manual for AM-80. Digital Research is making a distinct effort to have common manuals for both 8-bit and 16-bit versions of just about everything coming out of the Language Division, to demonstrate the near-total compatibility of the versions. Like most of Digital Research's recent documentation, an index is provided. There are a number of programming examples included. The appendix has a program to create a database. It is written in CBASIC, Pascal/MT+, and PL/I.

Unfortunately, the Pascal/MT+ and PL/I versions are direct translations of the CBASIC program. I say "unfortunately", because CBASIC and Pascal/MT+, of necessity, must be programmed bottom-up, as opposed to the top-down programming that PL/I thrives on. The result is that the code is a bit difficult to follow. An entire recreation program is furnished as well, to rebuild an index that has been inadvertently destroyed. It is educational and downright handy.

I have written a number of business programs with AM-80, and to date have found no bugs. Correspondence with a beta test site of AM-80 has revealed no bugs either. It speaks well for Digital Research's beta testing.

There are few utilities comparable to AM-80. I have been informed that there is a system called "Micro B+" written by Bill Fairman of FairCom that is reputed to be a close match. The closest that I have worked with personally is Digital Research's BT-80. I brought up BT-80 about a year earlier, and found that the amount of pointer handling involved, and the enormous amount of initialization required, made the utility difficult to use. In addition, BT-80 is compatible only with PL/I-80 (which was no hardship for me, as PL/I is my language of choice, but it is not everyone's choice). AM-80 is compatible with a minimum of three languages. There are few DBMSs (Data Base Management Systems) available that provide a programming language interface. Among them are MDBMS, TIM, CCA, and IDM-M2. Of these, the closest to AM-80 would be MDBMS (at \$900). The majority of DBMSs, like Ashton Tate's dBase II, are not programmatically accessible. Access Manager is priced at \$300, and is available from both Digital Research directly and through the normal distribution channels. From a consumer's point of view, I should mention that software purchased directly from Digital Research, although not discounted, does have the advantage of being refundable if returned within a reasonable period of time (for a nominal restocking charge).

Programs written using AM-80 to create a database-managed system could and should rival mainframe systems. Retrieval speeds are phenomenal. A common

database built by these programs is accessible from any other program using AM-80, even if it is written in another language. Additional benefits are extremely compact data and index files. AM-80 allows the programmer to assign only as little space as required by the records. In spite of the fact that it still has fixed-length records, the removal of the restriction of multiples of the logical record length allows fantastic economy of storage. Couple this with gigabyte capacities, and you have the capability of writing an application program to run under CP/M that would otherwise need a \$10,000 computer running MP/M. In performance, it may rival programs run on a much larger machines.

Most of the programs I have written and linked with AM-80 routines produced under 30 KB of machine code, including high-speed storage allocation. The need for using large amounts of storage for sorting data is eliminated, because AM-80 has kept the keys in ascending sort order automatically by its built-in housekeeping routines. With these routines, the writing of search and sort routines become academic. For any programmer involved in the writing of database systems in one of Digital Research's 3 programming languages (CBASIC, Pascal/MT+, and PL/I), AM-80 will be a welcome aid.

EOF

AM86PGPC.WS4 (Access Manager-8086 Programmer's Guide for PC DOS)

(Retyped by Emmanuel ROCHE.)

Access Manager 8086
Productivity Tool
Programmer's Guide
for the IBM Personal Computer
Disk Operating System (= PC DOS, not CP/M-86)

COPYRIGHT

Copyright (C) 1983 by Digital Research. All rights reserved. No part of this publication may be reproduced, transmitted, transcribed, stored in a retrieval system, or translated into any language or computer language, in any form or by any means, electronic, mechanical, magnetic, optical, chemical, manual or otherwise, without the prior written permission of Digital Research, Post Office Box 579, Pacific Grove, California, 93950.

Readers are granted permission to include the example programs, either in whole or in part, in their own programs.

DISCLAIMER

Digital Research makes no representations or warranties with respect to the contents hereof and specifically disclaims any implied warranties of merchantability or fitness for any particular purpose. Further, Digital Research reserves the right to revise this publication and to make changes from time to time in the content hereof without obligation of Digital Research to notify any person of such revision or changes.

TRADEMARKS

CBASIC is a registered trademark of Digital Research. Access Manager, ASM-86, CB86, Concurrent CP/M, Digital Research C, MP/M II, MP/M-86, Pascal/MT+86, PL/I-86, RASM-86, and TEX are trademarks of Digital Research. IBM is a registered trademark of International Business Machines. Z-80 is a registered trademark of Zilog, Inc.

The "Access Manager Productivity Tool Programmer's Guide for the IBM Personal Computer Disk Operating System" was prepared using the Digital Research TEX Text Formatter, and printed in the United States of America.

First Edition: July 1983

Foreword

This Programmer's Guide contains information and instructions for implementing your application programs with Access Manager. It is specifically directed at implementations using the IBM Personal Computer Disk Operating System (IBM PC DOS).

The "Access Manager Productivity Tool Programmer's Guide for the IBM Personal Computer Disk Operating System" was designed and written as the companion manual to the "Access Manager Productivity Tool Reference Manual". You need both publications to make full and proper use of Access Manager.

This guide contains information and instructions for implementing your applications programs with Access Manager. Section 1 describes the general guidelines and restrictions you must observe when you use Access Manager with PC DOS. The remaining sections describe how to link Access Manager with one of the Digital Research programming languages supported by Access Manager. Each of these sections includes example programs.

Table of Contents

oooo (to be done...)

Tables

oooo

Listings

oooo

Section 1: Implementation guidelines

1.1 Main Access Manager components

Your Access Manager distribution disk contains the following files:

- Single-user subroutine libraries

AM86CB86.L86 A complete, binary-relocatable, indexed library of index and data file functions for CB86 application programs.

AM86PLI.L86 A complete, binary-relocatable, indexed library of index and

data file functions for PL/I-86 application programs.

AM86PASC.L86 A complete, binary-relocatable, indexed library of index and data file functions for Pascal/MT+86 application programs.

AM86DRCS.L86 A complete, binary-relocatable, indexed library of index and data file functions for Digital Research C application programs using the SMALL memory model.

AM86DRCC.L86 A complete, binary-relocatable, indexed library of index and data file functions for Digital Research C application programs using the COMPACT memory model.

AM86DRCM.L86 A complete, binary-relocatable, indexed library of index and data file functions for Digital Research C application programs using the MEDIUM memory model.

AM86DRCB.L86 A complete, binary-relocatable, indexed library of index and data file functions for Digital Research C application programs using the BIG memory model.

- Single-user Access Manager buffer modules

AM86BUF.OBJ A relocatable object module containing a prespecified buffer area of 3,840 (decimal) bytes for CB86 and PL/I-86 application programs.

AM86BUF.R86 A relocatable object module containing a prespecified buffer area of 3,840 (decimal) bytes for Pascal/MT+86 application programs.

AM86BUF.A86 The assembly language source code for the buffer modules. Contains entry points (AM8FCB and AM8END) that defines the beginning and end of the buffer area. This module is for use with RASM86.EXE.

AM86BUF.I86 The assembly language source code for the buffer modules. Contains entry points (AM8FCB and AM8END) that defines the beginning and end of the buffer area. This module is for use with ASM86.EXE.

SETAMBUF.EXE A program to change the buffer module sizes without reassembling AM86BUF.

- External procedure declarations

AM86EXTR.BAS External procedure declarations for CB86 application programs.

AM86EXTR.PLI External procedure declarations for PL/I-86 application programs.

AM86EXTR.PSC External procedure declarations for Pascal/MT+86 application

programs.

AM86DRC.H Header file containing definitions of special string structures for use with application programs written in Digital Research C.

- Utility programs

RECREATE A general purpose program for rebuilding index files from existing data files. Source versions of the program are provided in each programming language provided by Access Manager. The program is documented in Section 4 of the "Access Manager Productivity Tool Reference Manual".

DATABASE A complete data base example program. Source code is provided in each programming language supported by Access Manager.

1.2 Memory requirements for Access Manager code

Access Manager's modular design ensures that your application program uses only those parts of Access Manager that are actually required. The calls to Access Manager embedded in the application program determine which modules the linking loader brings into the final file.

Table 1-1 shows Access Manager's code requirements.

Table 1-1. Access Manager code requirements

Module	Memory
-----	-----
Kernel: system initialization and index file setup and searching	4.1 KB
Key value insertion	1.5 KB
Key value deletion	1.6 KB
Data file functions	1.7 KB
Buffer Area	0.8+KB

1.3 Access Manager design constraints

When using Access Manager, the following design constraints apply:

- Data records in a data file must all be the same length.
- Data records must be a minimum of four bytes in length.

- Access Manager reserves the first 128 bytes of every data file for recording status information.
- Key values must not exceed a length of 48 bytes.
- Data records numbers (pointers) associated with key values must not exceed a length of four bytes.
- The length of an index file record must be a multiple of 128 bytes; for example, 128, 256, 512, 1024, and so on.
- There must be a minimum of four key values in any given index file record.
- A minimum of three buffers must be allocated for Access Manager.

You must also observe the design constraints shown in the following table.

Table 1-2. Additional design constraints

Design constraints -----	Limit -----
Maximum number of key values per index file record	254
Maximum index file size	*
Maximum number of index files that can be open at a given time	254
Maximum data file size	**
Maximum number of data files that can be open at a given time	254
Maximum number of Access Manager disk buffers	256

Notes: * = The maximum number of nodes in an index file is 65,535. If you are using 512-byte nodes, the maximum file size is 32 megabytes.

** = Access Manager supports a maximum of 16,777,216 logical records.

1.4 Configuring the single-user Buffer Area

The primary parameters affecting the buffer area size are NNSEC% (the number of sectors per index file record) and NBUFS% (the number of index file buffers). Set NNSEC% to a value of four if compatibility of your application program with other software is a factor. However, if response time is a more critical issue, refer to Table 1-3 for suggested NNSEC% values.

Note that NNSEC% determines the length of the records in the index file. For example, if NNSEC% contains a value of four, the resulting index file record length is 512 bytes, regardless of the physical sector size of the disk.

Table 1-3. Suggested index file record lengths

Physical sector size of disk	Suggested values for NNSEC% (*)
128	2
256	2,4
512	2,4
1024	4,8

Note: (*) NNSEC% specifies the number of 128-byte sectors per index file record.

Within the guidelines of Table 1-3, the selection of a value for NNSEC% is usually based on the length of the keys. Long keys lead to higher values of NNSEC%, to reduce the levels of the B-Tree index structure.

For any specified value of NNSEC%, the more buffers (larger NBUFS%) there are, the fewer node accesses you need to retrieve a key value. However, when processing one index file at a time, the payoff from adding buffers diminishes rapidly when five or six buffers are already in use. If more than one index file is in use at a time, increase the number of buffers beyond six. If the memory space is available, allowing three buffers per index file (in active use at one time) is a reasonable guideline. Buffers are not assigned to individual index files. They are shared according to a least recently used priority scheme, which ensures the active index files make full use of the buffers.

Finally, the amount of available memory determines the size of the buffer area. If there is very little memory available for buffers, you can reduce NBUFS% to the minimum level of three.

The required buffer size for any given specification of the maximum number of index files (NKEY%), the node size (NNSEC%), the number of buffers (NBUFS%), and the number of data file (NDATF%), can be computed as follows:

$$(NKEYS\% * 100) + (NDATF\% * 228) + (NBUFS\% * ((NNSEC\% * 128) + 60))$$

For example, if NKEYS%=3, NBUFS%=6, NDATF%=1, and NNSEC%=4, a buffer size of 3,960 (decimal) bytes is required. Once this buffer space is reserved, however, any combination of the four determining parameters (within 3,960 bytes) can be passed to the SETUP function, that sets up how the buffer area is used.

If you have the RASM-86 assembler, which generates relocatable object files, you can change AM86BUF.A86, then reassemble it to create a new AM86BUF.OBJ. Similarly, you can modify AM86BUF.I86 and reassemble it with ASM-86. Or, you can run the program SETAMBUF.EXE, which is on your distribution disk. Make sure AM86BUF.OBJ and AM86BUF.R86 are on the same disk as SETAMBUF.EXE. SETAMBUF modifies the buffer modules according to your specifications.

1.5 8080 compatibility

In all aspects of use, the index and data files produced by Access Manager in an 8080/Z-80 environment are compatible with use in 8086/8088 environments. Access Manager in an 8086/8088 environment has a greater capacity for key values per node than the 8080 version. Therefore, it is necessary to rebuild index files if the combination of parameters listed in Table 1-4 is used in your application program.

Table 1-4. 8080/8086 compatibility factors

NNSEC%	Critical key length
5	1
6	2
7	3
8	4

Key lengths, less than or equal to the critical key lengths shown in Table 1-4, can cause problems for a given index file record length, as specified by NNSEC%. For example, if NNSEC% is eight, key lengths of four or less might cause problems if the index file is created in an 8080 environment, and then transferred to an 8086 environment.

For a given NNSEC% value, the critical key length is the largest integer value strictly less than:

$$((\text{NNSEC\%} * 128) - 10) / 124) - 4$$

1.6 Multiuser operation

Because PC DOS does not support multiuser or multi-tasking operation, Access Manager for PC DOS does not support multiuser applications. However, you can still code your application programs for multiuser environments. So, if you transport a program to or from a Concurrent CP/M, MP/M-86, or MP/M II operating system environment, it still functions properly.

The remaining sections of this guide explain and demonstrate how to link single-user and multiuser applications. However, because PC DOS only supports single-user applications, the sample multiuser link statements are for informative purposes only. The modules referenced in these sample statements are not included on your distribution disk.

1.7 Additional error codes for PC DOS version 2.0

Access Manager can be used with both PC DOS versions 1.1 and 2.0. When used with version 2.0, there are eight user-error codes to consider, in addition to those listed in Section 4 of your "Access Manager Productivity Tool Reference

Manual".

Table 1-5. Error codes

Error code	Error description
EN/94	There are no file handles available while opening an index file. Refer to the section on CONFIG.SYS in your PC DOS 2.0 documentation, to change the maximum number of files that can be used simultaneously.
EO/95	A bad file handle occurred during an index file operation. This indicates the Access Manager buffer area (AM86BUF) has been overwritten. Check to see if you have exceeded the bounds of an array.
FA/97	Access has been denied during an index file operation.
FB/98	A directory patch could not be located during an index file operation.
GL/124	There are no file handles available while opening a data file. (See error code EN/94.)
GM/125	A bad file handle occurred during a data file operation. (See error code EO/95.)
GN/126	Access has been denied during a data file operation.
GO/127	A directory patch could not be located during a data file operation.

1.8 Path names

In the PC DOS version 2.0 environment, Access Manager accepts filenames that include a directory path. The filenames with a directory path have the format:

```
d:\level 1\level 2\...\level n\filename.typ
```

where "d:" is the disk drive indicator. "Level" is the nth level directory name. "Typ" can be up to three characters. For use with Access Manager, the string directory path and filename must not exceed 64 characters.

1.9 Compatibility between PC DOS versions 1.1 and 2.0

Index files created with Access Manager are fully compatible between both versions. Data files created with Access Manager under version 1.1 are always upward compatible with version 2.0. However, data files created under version 2.0 with record lengths that are not a multiple of 128 bytes might report an erroneous end-of-file condition when run under version 1.1.

Section 2: CBASIC Compiler (CB86) applications

This section contains instructions for implementing Access Manager with application programs coded in CB86.

Two examples are provided in this section. The first shows the use of many Access Manager functions described in Section 3 of the "Access Manager Productivity Tool Reference Manual", and how to use CB86 strings for data file buffer areas. The second example illustrates a multiple index and data file application.

2.1 Linking Access Manager to your CB86 program

This section discusses a CB86 application program that you write and compile to produce a binary, relocatable file called "MYPROG".

2.1.1 Linking single-user CB86 applications

You must link your compiled application program to the appropriate Access Manager subroutine library and index file buffer module. The following command line can be used to create an executable version of MYPROG:

```
LINK86 MYPROG,AM86CB86.L86[S],AM86BUF
```

Before linking, ensure that AM86BUF is large enough to contain your buffers (as specified in the SETUP function). You can use SETAMBUF.EXE to create a correctly sized buffer module.

2.1.2 Linking multiuser CB86 applications

Because PC DOS does not support multiuser or multi-tasking operation, Access Manager for PC DOS does not support multiuser applications. However, you can still code your application programs for multiuser environments. So, if you transport a program to or from a Concurrent CP/M, MP/M-86, or MP/M II operating system environment, it still functions properly.

This section explains how to link multiuser applications. However, because PC DOS only supports single-user applications, the sample multiuser link statements are for informative purposes only. The modules referenced in these sample statements are not included on your distribution disks.

If your single-user version of "MYPROG" is coded with appropriate data-locking procedures, you do not have to recompile it to create a multiuser version. You just relink the program.

You must link your compiled application program to the appropriate Access Manager multiuser interface. The interface makes the queue calls to the shared code in the background server. The background server resides in its own memory

segment.

To create an EXE file that calls the Access Manager background server, use the following command line:

```
LINK86 MYPROG,AMQ6CB86.L86
```

No buffer area module (such as AM86BUF) is permitted in the multiuser link statement. Whereas AM86CB86 contains the actual Access Manager code, AMQ6CB86 simply contains the message handler necessary to get the shared Access Manager code to perform the necessary actions.

2.2 External declaration of Access Manager routines

CB86 requires that external routines (those not coded in the program module, but referenced by it) be explicitly declared. The file AM86EXTR.BAS contains the external function declarations for the entire set of Access Manager functions. Use the %INCLUDE feature of CB86 to make these external declarations a part of your application program.

2.3 Coding numeric key values

See the ADDKEY function description in Section 3 of the "Access Manager Productivity Tool Reference Manual" for a discussion of "Coding Numeric Key Values".

2.4 Using the RECREATE.BAS utility program

RECREATE.BAS contains the CB86 source code for the RECREATE utility program. You can change the source code however you want.

To create RECREATE.EXE, compile RECREATE.BAS using CB86, and then link as follows:

```
LINK86 RECREATE,AM86CB86.L86[S],AM86BUF
```

The buffer area for RECREATE is 4,904 bytes, based on the following SETUP parameter values:

- NNSEC% = 4
- NBUFS% = 8
- NDATAF% = 1
- NKEYS% = 1

Note that only one data file and one index file are open at the same time RECREATE is running. Use SETAMBUF to configure AM86BUF.OBJ.

Table 2-1 shows the layout and content of records in a Recreate Parameter

File. This example file can be used to reconstruct DATABASE.BAS (as shown in Listing 2-2).

Table 2-1. Example CB86 Recreate Parameter File

Record type	Contents
Header	1,4
Data file	CUSTOMER.DAT,100,3,0
Index file	NAME.IDX,10,0,1,1,Y
Key part	22,8
Index file	NUMB.IDX,4,0,0,1,N
Key part	2,4
Index file	ZIPC.IDX,11,0,1,1,Y
Key part	84,9

If you want to change the capacity of the RECREATE program (hence, its memory requirements) note the following parameters and associated DIMENSION statements.

- MAX.NO.KEYS% and MAX.NO.KEY.PARTS% specify the maximum number of index files associated with a data file, and the maximum number of fields comprising a key value, respectively. If the value for either of these parameters is increased in RECREATE.BAS, the following dimension statements must be modified to reflect the changes.

```
DIM INDEX.NAMES$ (...  
DIM AUTO.SUFFIX% (...
```

- MAX.SORT% determines the maximum number of key values that are buffered by RECREATE.BAS before being sorted and added to the index file being recreated. If MAX.SORT% is increased, the following dimension statement must be changed.

```
DIM KEYVAL$ (...
```

The routine SORT.SETUP in RECREATE.BAS uses the FRE and MFRE functions of CB86 to determine the amount of available memory for buffered key values. The actual number of key values buffered is stored in NO.SORT%. For long key lengths, the memory space available limits NO.SORT%. To make the value for NO.SORT% more conservative, reduce the values of G.SORT and M.SORT before NO.SORT% is computed.

2.5 CB86 data file example listing

In the following listing, notice how the SADD function determines the value of the buffer pointer (Access Manager parameter BUFFER%) for the READAT and WRDAT functions. The result of SADD is increased by two, because each string variable in CB86 has a two-byte header that contains the length of the string.

When Access Manager fills in the input buffer (INP.BUFFER\$) during the READAT

function, the two-byte length header is not affected. Therefore, you can use one such string input buffer for all the data files, if it is long enough to accommodate the longest record length.

Conversely, the output buffer (OUT.BUFFER\$) is constantly adjusted, because it is reconstructed for each WRDAT function. Therefore, it is not advisable to use the input buffer for output. Reserve the input buffer for input only, and create the output buffer strings as needed.

0000

Listing 2-1. CB86 data file example

2.6 CB86 DATABASE source code

Your Access Manager distribution disks contain sample code for building and maintaining a data base in CB86. The code is designed so you can add or substitute your own key attributes as required. The sample code is on your distribution disk in a file named "DATABASE.BAS".

DATABASE.BAS demonstrates the integration of Access Manager with CB86 applications. It builds a name and address data base, and provides facilities for examining, updating, and/or listing the information contained therein. You might also want to use routines from DATABASE.BAS directly in your application programs.

[SINGLE] To create DATABASE.EXE, compile DATABASE.BAS with CB86.EXE, and link as follows:

```
LINK86 DATABASE,AM86CB86.L86,AM86BUF
```

Prior to the preceding link command, make sure AM86BUF.OBJ can accommodate the SETUP parameters of DATABASE.BAS. You can use SETAMBUF.EXE to modify AM86BUF.OBJ.

[MULTI] In the multiuser environment, enter your link statement as follows:

```
LINK86 DATABASE,AMQ6CB86.L86
```

Note that the listing of DATABASE.BAS, in Listing 2-2, might not include recent changes. Always treat the copy on your distribution disk as the definitive version.

0000

Listing 2-2. DATABASE.BAS source code listing

Section 3: Using Access Manager with PL/I-86 applications

This section contains instructions for implementing Access Manager with

application programs coded in PL/I-86.

Two examples are provided in this section. The first shows the use of many Access Manager functions described in the "Access Manager Productivity Tool Reference Manual", and, specifically, how to use the data file functions in your PL/I-86 applications. The second example provides an extensive illustration of using Access Manager to construct and maintain a data base.

3.1 Linking Access Manager to your application program

This section discusses a PL/I-86 application program that you write and compile to produce a binary, relocatable file called "MYPROG".

3.1.1 Linking single-user PL/I-86 applications

You must link your compiled application program to the appropriate Access Manager subroutine library and index file buffer module. The following command line can be used to create an executable version of MYPROG:

```
LINK86 MYPROG,AM86PLI.L86[S],AM86BUF
```

Before linking, ensure that AM86BUF.OBJ is large enough to contain your buffers (as specified in the SETUP function). You can use SETAMBUF.EXE to create a correctly sized buffer module.

AM86BUF contains the buffer area, beginning with entry point AM8FCB, and ending with AM8END.

3.1.2 Linking multiuser PL/I-86 applications

Because PC DOS does not support multiuser or multi-tasking operation, Access Manager for PC DOS does not support multiuser applications. However, you can still code your application programs for multiuser environments. So, if you transport a program to or from a Concurrent CP/M, MP/M-86, or MP/M II operating system environment, it still functions properly.

This section explains how to link multiuser applications. However, because PC DOS only supports single-user applications, the sample multiuser link statements are for informative purposes only. The modules referenced in these sample statements are not included on your distribution disks.

If your single-user version of "MYPROG" is coded with appropriate data-locking procedures, you do not have to recompile it to create a multiuser version. You just relink the program.

You must link your compiled application program to the appropriate Access Manager multiuser interface. The interface makes the queue calls to the shared code in the background server. The background server resides in its own memory

segment.

To create an EXE file that calls the Access Manager background server, use the following command line:

```
LINK86 MYPROG,AMQ6PLI.L86
```

3.2 External declaration of Access Manager routines

PL/I-86 requires that external routines (those not coded in the program module, but referenced by it) be explicitly declared. The file AM86EXTR.PLI contains external function declarations for the entire set of Access Manager routines. Use the %INCLUDE feature of PL/I-86 to make these external declarations a part of your application program.

Note that AM86EXTR.PLI expects two compile-time constants to be defined with the %REPLACE macro of PL/I-86. NAME_LEN (the maximum length of index and data filenames) and MAX_KEY_LEN (the maximum key value length) must be set to appropriate values before AM86EXTR.PLI is included. For example, you can use the code segment:

```
%REPLACE  
    NAME_LEN by 14,  
    MAX_KEY_LEN by 48;  
  
%INCLUDE 'AM86EXTR.PLI';
```

It is not necessary to set MAX_KEY_LEN at the Access Manager maximum of 48. Any value, less than or equal to 48, that is sufficient for your particular application is valid.

Note: Consider the following points concerning the passing of parameters between Access Manager and PL/I-86:

- Access Manager requires all string-valued parameters (FILENAME, IDXNAME, KEYVAL, and IDXVAL) to be declared as CHARACTER() VARYING. CHARACTER VARYING strings in PL/I-86 reserve the leading byte for a length counter that Access Manager uses to determine the actual length of a string-valued parameter.
- The output string parameter IDXVAL must be declared by reference, as opposed to value. Therefore, the actual parameter passed to Access Manager (for IDXVAL) must be declared with exactly the same attributes as the formal IDXVAL parameter in AM86EXTR.PLI. This implies the actual variable used for IDXVAL, ACTUAL_IDXVAL must be declared as follows:

```
DCL ACTUAL_IDXVAL CHAR (MAX_KEY_LEN) VAR;
```

3.3 Coding numeric key values

For a general discussion of coding numeric key values, refer to the ADDKEY function description in Section 3 of the "Access Manager Productivity Tool Reference Manual".

In the PL/I environment, the easiest approach to represent numeric key values is to use FIXED DECIMAL quantities. Because FIXED DECIMAL quantities are stored in BCD (Binary Coded Decimal) format with the least-significant byte first and the sign bit set in the last byte, KEYTYP% should be one. Access Manager requires a CHARACTER VARYING value for the KEYVAL\$ and IDXVAL\$ parameters. Therefore, FIXED DECIMAL variables should be overlaid or based on KEYVAL\$ and IDXVAL\$ string variables.

The key length is based on the number of bytes required to store the FIXED DECIMAL quantities. A FIXED DECIMAL with "p" digits requires

$$\text{INT} ((p + 2) / 2)$$

bytes, where "INT" returns the integer portion of its argument. For example, a FIXED DECIMAL quantity with eight digits requires five bytes of storage.

The following declarations and assignments permit the use of FIXED DECIMAL quantities as key values in Access Manager.

```
%REPLACE
    MAX_KEY_LEN by 48,
    NAME_LEN by 14,
    P by 8,      /* Example precision */
    Q by 2,      /* Fractional places */
    KEYLEN by 5; /* INT ( (P + 2) / 2 */

%INCLUDE 'AM86EXTR.PLI';

DCL
    (KEYVAL, IDXVAL) CHAR (MAX_KEY_LEN) VAR,
    (BCDINP_PTR, BCDOUT_PTR) POINTER;

DCL
    1 BCDINP BASED (BCDINP_PTR),
    2 LEN FIXED BINARY (7),
    2 VAL FIXED DECIMAL (P, Q);

    1 BCDOUT BASED (BCDOUT_PTR),
    2 LEN FIXED BINARY (7),
    2 VAL FIXED DECIMAL (P, Q);

BCDINP_PTR = ADDR (KEYVAL); /* Overlay BCD on string */
BCDINP.LEN = KEYLEN;      /* Set length byte of string */

BCDOUT_PTR = ADDR (IDXVAL);
BCDOUT.LEN = KEYLEN;
```

When you use a numeric quantity with an Access Manager function, use KEYVAL for input values, and IDXVAL for output values. To manipulate the key as a

numeric quantity, refer to BCDINP.VAL and BCDOUT.VAL for input and output key values, respectively. For example,

```
BCDINP.VAL = 123.45;
DRN = SERKEY (KEY_NO, DFILE, DLOCK, KEYVAL, IDXVAL);
IF DRN ~= 0
| DATVAL() ~= 0
  THEN PUT SKIP LIST (BCDOUT.VAL);
```

prints the numeric value of the first key value in the index greater than, or equal to 123.45, unless no such key exists.

3.4 Using the RECREATE.PLI utility program

RECREATE.PLI contains the PL/I-86 source code for the RECREATE utility program. You can change the source code however you want.

To create RECREATE.EXE, compile RECREATE.BAS using PL/I-86, and then link as follows:

```
LINK86 RECREATE,AM86PLI.L86[S],AM86BUF
```

The buffer area for RECREATE is 4,904 bytes, based on the following SETUP parameter values:

- NNSEC% = 4
- NBUFS% = 8
- NDATA% = 1
- NKEYS% = 1

Note that only one data file and one index file are open at the same time that RECREATE is running. Use SETAMBUF to configure AM86BUF.OBJ.

Table 3-1 shows the layout and content of records in a Recreate Parameter File. This particular example file can be used to reconstruct DATABASE (see Listing 2-2).

Table 3-1. Example PL/I-86 Recreate Parameter File

Record type	Contents
-----	-----
Header	1,4
Data file	CUSTOMER.DAT,100,3,0
Index file	NAME.IDX,10,0,1,1,Y
Key part	22,8
Index file	NUMB.IDX,4,0,0,1,N
Key part	2,4
Index file	ZIPC.IDX,11,0,1,1,Y
Key part	84,9

If you want to change the capacity of the RECREATE program (hence, its memory requirements) note the key constants:

- MAX_NO_KEYS specifies the maximum number of index files associated with a data file, MAX_KEY_PARTS indicates the maximum number of fields comprising a key value.
- MAX_SORT is the maximum number of key values that can be buffered by RECREATE.PLI before being sorted and added to the index file being recreated.
- MAX_SPACE specifies the actual number of bytes available for the buffered key values. Each key value requires one more byte than its key length.

The actual number of buffered key values depends on the key length. For short key lengths, MAX_SORT is the limiting factor. MAX_SPACE is the limiting factor for long key lengths.

Increase the constant MAX_REC_LEN if your applications require data files with record lengths exceeding 1024 bytes.

3.5 PL/I-86 data file example listing

The following listing illustrates the use of the primary Access Manager functions to update records in a data file.

0000

Listing 3-1. Example of PL/I-86 data file

3.6 PL/I-86 DATABASE source code

Your Access Manager distribution disk contains sample code for building and maintaining a data base in PL/I-86. The code is designed so you can add or substitute your own key attributes, as required. The sample code is on your distribution disk, in a file called "DATABASE". DATABASE is comprised of these three separate components:

- DATABAS1.PLI
- DATABAS2.PLI
- DATABASE.DCL

DATABASE demonstrates the integration of Access Manager with PL/I-86 applications. It builds a name and address data base, and provides facilities for examining, updating, and/or listing the information contained therein. You might also want to use routines from DATABASE directly in your application programs.

[SINGLE] To create DATABASE.EXE, compile DATABAS1.PLI and DATABAS2.PLI with PL/I-86, and link as follows:

LINK86 DATABASE=DATABAS1,DATABAS2,AM86PLI.L86[S],AM86BUF

[MULTI] In a multiuser environment, enter the link statement as follows:

LINK86 DATABASE=DATABAS1,DATABAS2,AMQ6PLI.L86

Note that the listing of DATABASE (see Listing 3-2) might not include recent changes. Always treat the copy on your distribution disk as the definitive version.

0000

Listing 3-2. DATABASE source code

Section 4: Using Access Manager with Pascal/MT+86 applications

This section contains instructions for implementing Access Manager with application programs coded in Pascal/MT+86.

Two examples are provided in this section. The first illustrates the use of many Access Manager functions described in the "Access Manager Productivity Tool Reference Manual", and, in particular, how to use the data file functions in your Pascal/MT+86 applications. The second example illustrates the use of Access Manager to create and maintain a data base.

4.1 Linking Access Manager to your Pascal/MT+86 program

This section discusses a Pascal/MT+86 application program that you write and compile to produce a binary, relocatable file called "MYPROG".

4.1.1 Linking single-user Pascal/MT+86 applications

You must link your compiled application program to the appropriate Access Manager subroutine library and index file buffer module. The following command line can be used to create an executable version of MYPROG:

LINKMT MYPROG,AM86PASC/S,AM86BUF,PASLIB/S

Before linking, ensure that AM86BUF.R86 is large enough to contain your buffers (as specified in the SETUP function). You can use SETAMBUF.EXE to create a correctly sized buffer module. For further details, see the example link statements for DATABASE.SRC and RECREATE.SRC in this section.

4.1.2 Linking multiuser Pascal/MT+86 applications

Because PC DOS does not support multiuser or multi-tasking operation, Access

Manager for PC DOS does not support multiuser applications. However, you can still code your application programs for multiuser environments. So, if you transport a program to or from a Concurrent CP/M, MP/M-86, or MP/M II operating system environment, it still functions properly.

This section explains how to link multiuser applications. However, because PC DOS only supports single-user applications, the sample multiuser link statements are for informative purposes only. The modules referenced in these sample statements are not included on your distribution disks.

If your single-user version of "MYPROG" is coded with appropriate data-locking procedures, you do not have to recompile it to create a multiuser version. You just relink the program.

You must link your compiled application program to the appropriate Access Manager multiuser interface. The interface makes the queue calls to the shared code in the background server. The background server resides in its own memory segment.

To create an EXE file that calls the Access Manager background server, use the following command line:

```
LINKMT MYPROG,AMQ6PASC,PASLIB/S
```

4.2 External declaration of Access Manager routines

Pascal/MT+86 requires that external routines (those not coded in the program module, but referenced by it) be explicitly declared. The file AM86EXTR.PSC contains external function declarations for the entire set of Access Manager routines. Use the Include File compiler toggle of Pascal/MT+86 to make these external declarations a part of your application program.

```
[$I AM86EXTR.PSC]
```

includes the external declarations as required.

All Access Manager string-valued parameters (FILENAME, IDXNAME, KEYVAL, and IDXVAL) must be declared as type STRING. Strings, as compared to character arrays, reserve the leading byte for a length counter that Access Manager needs to determine the actual length of a string-valued parameter.

4.3 Coding numeric key values

For a general discussion of coding numeric key values, refer to the ADDKEY function description in Section 3 of the "Access Manager Productivity Tool Reference Manual".

In a Pascal/MT+86 environment, the most straightforward use of numeric keys is with the BCD REAL variables that store numeric quantities with the most significant digits in the first byte position, the least significant digits in

the ninth byte, and the sign indicator in the tenth byte. Because the most significant byte comes first, negative quantities are not properly handled, and you should avoid them. Also, set the KEYTYP parameter to zero. The BCD REALS provide eighteen digits, including four decimal places.

The following declarations and assignments overlay BCD REALS onto the string variables that must be passed to the Access Manager functions.

```
CONST
  KEYLEN by 10;      (* BCD REAL uses ten bytes *)

TYPE
  BCDOVL = RECORD;
    LEN : BYTE;      (* Use compiler B switch *)
    VAL : REAL;

VAR
  KEYVAL, IDXVAL : STRING [KEYLEN];
  BCDINP, BCDOUT : ^BCDOVL;

{$I AM86EXTR.PSC}

BCDINP := ADDR (KEYVAL);      (* Overlay BCD on string *)
BCDINP^.LEN := KEYLEN;      (* Set length byte of string *)

BCDOUT := ADDR (IDXVAL);
BCDOUT^.LEN := KEYLEN;
```

Use KEYVAL and IDXVAL, respectively, to pass key values to and from Access Manager. Use BCDINP^.VAL and BCDOUT^.VAL to manipulate the key values as numeric quantities. For example,

```
BCDINP^.VAL := 123.4567;
DRN := BEFKEY (KEY_NO, DFILE, DLOCK, KEYVAL, IDXVAL);
IF (DRN <> 0)
OR (DATVAL <> 0)
  THEN WRITELN (BCDOUT^.VAL);
```

prints the numeric value of the index entry that immediately precedes 123.4567, unless no such entry exists.

The space savings for this approach with Pascal/MT+86 is only meaningful if numbers with more than ten digits are involved, because BCD REALS are forced to use ten bytes, and the key length must be set to ten bytes. Note that you can accomplish the same kind of overlaying with INTEGER variables. If you overlay integers instead of reals, the key length must be set as necessary (that is, two bytes for regular integers, and four bytes for long integers). Then, set KEYTYP to one, because the least significant byte comes first. The key values are treated as signed integers.

4.4 Using the RECREATE.SRC utility program

RECREATE.SRC contains the Pascal/MT+86 source code for the RECREATE utility program. You can change the source code however you want.

To create RECREATE.EXE, compile RECREATE.SRC using MT+86.EXE, and then link as follows:

```
LINKMT RECREATE,AM86PASC/S,AM86BUF,FPREALS/S,RANDOMIO/S,PASLIB/S
```

The buffer area for RECREATE is 4,904 bytes, based on the following SETUP parameter values:

- NNSEC% = 4
- NBUFS% = 8
- NDATA% = 1
- NKEYS% = 1

Note that only one data file and one index file are open at the same time that RECREATE is running. Use SETAMBUF to configure AM86BUF.R86.

Table 4-1 shows the layout and content of records in a Recreate Parameter File. This particular example file can be used to reconstruct DATABASE.SRC (see Listing 4-2).

Table 4-1. Example Pascal/MT+86 Recreate Parameter File

Record type	Contents
Header	1 4
Data file	CUSTOMER.DAT
Data file	100 3 0
Index file	NAME.IDX
Index file	10 0 1 1
Index file	Y
Key part	22 8
Index file	NUMB.IDX
Index file	4 0 0 1
Index file	N
Key part	2 4
Index file	ZIPC.IDX
Index file	11 0 1 1
Index file	Y
Key part	84 9

If you want to change the capacity of the RECREATE program (hence, its memory requirements) note the following key Pascal constants:

- MAX_NO_KEYS specifies the maximum number of index files associated with a data file, MAX_KEY_PARTS indicates the maximum number of fields comprising a key value.
- MAX_SORT is the maximum number of key values that can be buffered by RECREATE.SRC before being sorted and added to the index file being recreated.

- MAX_SPACE specifies the actual number of bytes available for the buffered key values. Each key value requires one more byte than its key length.

The actual number of buffered key values depends on the key length. For short key lengths, MAX_SORT is the limiting factor. MAX_SPACE is the limiting factor for long key lengths.

4.5 Pascal/MT+86 data file example listing

The following listing illustrates the use of the primary Access Manager functions to update records in a data file.

0000

Listing 4-1. Pascal/MT+86 data file example

4.6 Pascal/MT+86 DATABASE source code

Your Access Manager distribution disk contains sample code for building and maintaining a data base in Pascal/MT+86. The code is designed so you can add or substitute your own key attributes, as required. The sample code is on your distribution disk, in a file called "DATABASE.SRC".

DATABASE.SRC demonstrates the integration of Access Manager with Pascal/MT+86 applications. It builds a name and address data base, and provides facilities for examining, updating, and/or listing the information contained therein. You might also want to use routines from DATABASE.SRC directly in your application programs.

[SINGLE] To create DATABASE.EXE, compile DATABASE.SRC with MT+86.EXE, and link as follows:

```
LINKMT DATABASE,AM86PASC/S,AM86BUF,PASLIB/S
```

[MULTI] In a multiuser environment, enter the link statement as follows:

```
LINKMT DATABASE,AMQ6PASC,PASLIB/S
```

Note that the listing of DATABASE.SRC (see Listing 4-2) might not include recent changes. Always treat the copy on your distribution disk as the definitive version.

0000

Listing 4-2. DATABASE.SRC source code

Section 5: Using Access Manager with Digital Research C applications

This section contains instructions for implementing Access Manager with application programs coded in Digital Research C for the 8086 family of processors.

Two examples are provided in this section. The first shows the use of many Access Manager functions described in the "Access Manager Productivity Tool Reference Manual", and, specifically, how to use the data file functions in your C applications. The second example provides an extensive illustration of using Access Manager to construct and maintain a data base.

5.1 Linking Access Manager to your application program

This section discusses a C application program that you write and compile to produce a binary object file called "MYPROG.OBJ".

5.1.1 Linking single-user C applications

You must link your compiled application program to the appropriate Access Manager subroutine library and index file buffer module. The following command line can be used to create an executable version of MYPROG:

```
LINK86 MYPROG,AM86DRCx.L86[S],AM86BUF
```

where "x" equals S, C, M, or B. The appropriate choice for "x" is based on the memory model in use when MYPROG is compiled, according to the following table:

Table 5-1. Memory models

Memory model	"x"
SMALL	S
COMPACT	C
MEDIUM	M
BIG	B

Note: If you link an application program compiled under one memory model with an Access Manager library designed for another memory model, run-time errors are not reported, but serious and unpredictable run-time errors can occur.

AM86BUF contains the single-user buffer area, beginning with entry point AM8FCB, and ending with AM8END.

Before linking, ensure that AM86BUF.OBJ is large enough to contain your buffers (as specified in the SETUP function). You can use SETAMBUF.EXE to create a correctly sized buffer module.

5.1.2 Linking multiuser C applications

Although PC DOS does not provide multiuser or multi-tasking capabilities, it is available to see how a single-user application is converted to a multiuser application. If your single-user version of MYPROG is coded with appropriate data locking procedures, you do not have to recompile it to create a multiuser version. You only have to relink the program.

You must link your compiled application program to the appropriate multiuser interface. The interface makes the queue calls to the shared code in the background server. The background server resides in its own memory segment.

To create an EXE file that calls the Access Manager background server, use LINK86.EXE as follows:

```
LINK86 MYPROG,AMQ6DRCx.L86
```

where "x" equals S, C, M, or B, based on the memory model used to compile MYPROG, as explained earlier in this section.

5.2 Access Manager parameters

Access Manager uses two types of parameters : two-byte integers and strings. Either int (WORD) or unsigned int (UWORD) variables can be used whenever an integer parameter is required. In C, strings are usually represented as null ('0') terminated character arrays. Because Access Manager expects the first byte of a string to specify the actual length of the string, and does not interpret a null byte as a string delimiter, a more general data structure is required.

The header file AM86DRC.H, supplied on your distribution disks, contains the following structure definition:

```
struct am86strg {
    BYTE len;
    BYTE str [MAX_KEY_LEN+1];
};
```

This structure assumes your application program specifies the compile time constant MAX_KEY_LEN in a #define macro. For example, the following code sequence specifies the maximum key length (which can range from 1 to 48), and sets up the extended string definition required by Access Manager:

```
#define MAX_KEY_LEN 48
#define <am86drc.h>
```

With the previous two statements in your code, you can define extended strings and/or their pointers as follows:

```
struct asm86strg keyval,*ptrkey;
```

You must pass all string parameters (that is, file names and key values) to Access Manager using a pointer to an am86strg structure, or its equivalent.

Note: Make certain the length byte of the am86strg structure is valid before passing a string parameter to Access Manager.

The following two examples show how to set the length byte, assuming the key value ends with a null byte. This example shows how to place a constant value into an extended string structure, and set its length byte:

```
struct am86strg keyval,filename;
```

This example shows how to place a value into the character array, and set the length byte:

```
filename.len = strlen (strcpy (filename.str, "TEST.IDX" ) );
scanf ("%10s", keyval.str);
keyval.len = strlen (keyval.str);
```

Note: You can only use the strlen library function when the string ends with a null; otherwise, you must use some other approach to set the length byte.

With the previous example in mind, the following are more examples that show how to pass extended strings to Access Manager:

```
WORD file_no, drn;
file_no = OPNIDX (-1, &filename, 22, 0, 1);
drn = GETKEY (file_no, 0, 0, &keyval);
```

Notice how the & operator is used to pass a pointer to the extended string structure.

When you pass pointers to extended string structures (am86strg) to your own routines, you can declare the pointers as BYTE-type. Therefore, you can increment or decrement the pointers to move up or down the structure. If you declare them as pointers to am86strg structures, incrementing the pointer causes the size of the am86strg structure to be added to the pointer, instead of simply adding one.

In addition to the definition of am86strg, AM86DRC.H contains the following compile-time constant definitions:

```
#define N_LOCK 0      /* No lock/unlock request      */
#define S_LOCK 1      /* Shared record lock/unlock   */
#define X_LOCK 2      /* Exclusive record lock/unlock */
#define S_FILE 3      /* Shared file lock/unlock     */
#define X_FILE 4      /* Exclusive file lock/unlock   */
#define R_LOCK 5      /* Shared/exclusive record unlock */
#define A_FILE 6      /* Unlock all locks on a file   */
                      /* for calling user.           */
#define A_USER 7      /* Unlock all locks on all file */
                      /* for calling user.           */
```

5.3 Coding numeric key values

The ADDKEY function description in Section 3 of your "Access Manager Productivity Tool Reference Manual" contains a general discussion of coding numeric key values.

In Digital Research C programs, the most direct approach to represent numeric key values is to use long, 4-byte integers. Four-byte integers accommodate key values ranging from -2147483648 to 2147483647. Because long integers are stored with the least significant byte first and the sign bit set in the last byte, set the key type parameter KEYTYP% to one.

Even though the key values are numeric, they must appear as extended strings (see Section 5.2, "Access Manager Parameters", earlier in this section). The length of these keys is four bytes. You can use the following declarations and assignments with long integer key values.

```
#define MAX_KEY_LEN 48
#include <am86drc.h>

struct am86strg keyval,idxval;
long *linp,*lout;

linp = keyval.str; /* Overlay long integer on string */
keyval.len = 4; /* Set key length */

lout = idxval.str;
idxval.len = 4;
```

When using a numeric quantity with an Access Manager function, use &keyval for input values, and &idxval for output values. To manipulate the key as a numeric quantity, use *linp and *lout for input and output values, respectively. For example,

```
*linp = 123456789; /* Length bytes already set above */
drn = SERKEY (key_no, dfile, dlock, &keyval, &idxval);
if (drn || DATVAL() )
    printf ("\nSERKEY output value = %d", *lout);
else
    printf ("\nNo value found.");
```

5.4 Using the RECREATE.C utility program

Section 5 in your "Access Manager Productivity Tool Reference Manual" contains a detailed description of the RECREATE utility program.

RECREAT1.C, RECREAT2.C, and RECREATE.H contain the source code for the RECREATE utility program. You can change the source code however you want.

To create RECREATE.EXE, compile RECREATE.SRC using DRC.EXE, and then link as follows:

```
LINK86 RECREATE=RECREAT1,RECREAT2,AM86DRCS.L86[S],AM86BUF
```

The buffer area for RECREATE is 4,904 bytes, based on the following SETUP parameter values:

- NNSEC% = 4
- NBUFS% = 8
- NDATAF% = 1
- NKEYS% = 1

Note that only one data file and one index file are open at the same time that RECREATE is running. Use SETAMBUF to configure AM86BUF.OBJ.

Table 5-2 shows the layout and content of records in a Recreate Parameter File. This particular example file can be used to reconstruct DATABASE (see Listing 5-2).

Table 5-2. Example C Recreate Parameter File

Record type	Contents
Header	1 4
Data file	CUSTOMER.DAT 100 3 0
Index file	NAME.IDX 10 0 1 1 Y
Key part	22 8
Index file	NUMB.IDX 4 0 0 1 N
Key part	2 4
Index file	ZIPC.IDX 11 0 1 1 Y
Key part	84 9

If you want to change the capacity of the RECREATE program (hence, its memory requirements) note the following key C constants:

- MAX_KEYS specifies the maximum number of index files associated with a data file, MAX_KEY_PARTS indicates the maximum number of fields comprising a key value.
- MAX_SORT is the maximum number of key values that can be buffered by RECREATE before being sorted and added to the index file being recreated.
- MAX_SPACE specifies the actual number of bytes available for the buffered key values. Each key value requires one more byte than its key length.

The actual number of buffered key values depends on the key length. For short key lengths, MAX_SORT is the limiting factor. MAX_SPACE is the limiting factor for long key lengths.

Increase the value of MAX_REC_LEN if the length of the records in the file exceeds 1024 bytes.

5.5 C data file example listing

The following listing illustrates the use of the primary Access Manager functions to update records in a data file.

0000

Listing 5-1. C data file example

5.6 C DATABASE source code

Your Access Manager distribution disk contains a sample program coded in Digital Research C for creating and maintaining a data base. You can make changes to the code. The sample code is on your distribution disk, in these three files:

- DATABAS1.C
- DATABAS2.C
- DATABASE.H

DATABASE demonstrates the integration of Access Manager with Digital Research C applications. You can use this program to build a name and address database, and to examine, update, and list the information contained therein.

[SINGLE] To create DATABASE.EXE, compile DATABAS1.C and DATABAS2.C with DRC.EXE, and link as follows:

```
LINK86 DATABASE=DATABAS1,DATABAS2,AM86DRCS.L86[S],AM86BUF
```

[MULTI] In a multiuser environment, enter the link statement as follows:

```
LINK86 DATABASE=DATABAS1,DATABAS2,AMQ6DRCS.L86
```

Note that the following listing of DATABASE might not include recent changes. Always treat the copy on your distribution disk as the definitive version.

0000

Listing 5-2. C DATABASE source code

Index

0000 (to be done...)

EOF

(Retyped by Emmanuel ROCHE.)

Digital Research
Access Manager 8086
Productivity Tool
Reference Manual

COPYRIGHT

Copyright (C) 1983 by Digital Research. All rights reserved. No part of this publication may be reproduced, transmitted, transcribed, stored in a retrieval system, or translated into any language or computer language, in any form or by any means, electronic, mechanical, magnetic, optical, chemical, manual or otherwise, without the prior written permission of Digital Research, Post Office Box 579, Pacific Grove, California, 93950.

This manual is, however, tutorial in nature. Thus, the reader is granted permission to include the example programs, either in whole or in part, in his own programs.

DISCLAIMER

Digital Research makes no representations or warranties with respect to the contents hereof and specifically disclaims any implied warranties of merchantability or fitness for any particular purpose. Further, Digital Research reserves the right to revise this publication and to make changes from time to time in the content hereof without obligation of Digital Research to notify any person of such revision or changes.

TRADEMARKS

CBASIC and CP/M are registered trademarks of Digital Research. Access Manager, CB80, CB86, MP/M, Pascal/MT+, Pascal/MT+86, PL/I-80, and PL/I-86 are trademarks of Digital Research. Z-80 is a registered trademark of Zilog, Inc.

The "Access Manager Reference Manual" was prepared using the Digital Research TEX Text Formatter, and printed in the United States of America.

First Edition: March 1983

Foreword

Access Manager is a general purpose, keyed, file accessing package for 8080 or 8086 compatible microprocessors which run under CP/M or MP/M operating systems. It is designed for programmers creating application packages that must access data records on the basis of identifying key values, such as a name or a number.

The Access Manager hardware environment must include an 8080, 8085, 8088, 8086, or Z-80 microprocessor, at least 48K bytes of Random Access Memory (RAM), at least one random access secondary storage device such as a floppy or hard disk drive, and a console.

Access Manager is a set of routines called from programs written in CB80, CB86, PL/I-80, PL/I-86, Pascal/MT+, and Pascal/MT+86; Access Manager is not a stand-alone data base management system. The routines require approximately 5 KB to 11 KB, excluding buffer areas. In a multiple-user environment, Access Manager requires a separate memory segment of at least 20K bytes, a minimum of three pages (768 bytes) of common memory for special Access Manager queues, and 2K bytes of memory in each user's memory segment. Up to eight users can share Access Manager under MP/M II, if there is enough queue space.

Access Manager documentation set

This manual describes Access Manager, a programming productivity tool from Digital Research. It is important that you understand how this manual is organized and what it contains, for it to serve you well.

There are three manuals in this documentation set. The first is the "Access Manager Reference Manual". Contained here are descriptions of the product, instructions for using it, and numerous examples of how it can be applied in your application programs.

The second manual is the Access Manager Programmer's Guide for either the 8080 or 8086 implementations. Your set contains one of the following:

- "Access Manager Programmer's Guide for the CP/M Family of Operating Systems". (8080 implementation)
- "Access Manager Programmer's Guide for the CP/M-86 Family of Operating Systems". (8086 implementation)

Here, you will find information for using Access Manager with a specific operating system and programming language. Your Programmer's Guide also contains facts concerning basic requirements and design constraints for Access Manager. Finally, there are several examples showing how to organize your application program for the various languages that are supported.

The third manual in your documentation set is the "Access Manager Function Summary". The summary is a compact, abbreviated version of information you need when using Access Manager. As you become familiar with the Access Manager functions, you will rely less on your Reference Manual and Programmer's Guide, and more on the Function Summary.

Prerequisite

These manuals assume you are knowledgeable about keyed file accessing methods. If you are not familiar with these methods, the following reading is recommended:

- Knuth, D.E., "The Art of Computer Programming, Vol.3, Sorting and Searching", Reading, Massachusetts: Addison-Wesley Publishing Company, 1973.
- Wiederhold, G., "Database Design", New York, New York: McGraw-Hill Book Company, 1977.

How to use this documentation

Your Access Manager documentation contains several forms of indexing, to help you find information quickly and efficiently.

- Your Reference Manual and Programmer's Guide each contain a Table of Contents and subject index.
- Appendix A of your Reference Manual contains a function index. All the functions you can perform with Access Manager are listed alphabetically, and cross-referenced to the mnemonic function name.
- The function descriptions in Section 3 are arranged alphabetically by their mnemonic function names.

Take a few moments to understand how the indexing facilities are constructed in this manual, and you will be able to find the information you need when you need it.

Conventions used in this documentation

Access Manager is designed for use in two distinct environments: single-user and multiple-user. A single-user environment consists of a single microcomputer, a single terminal, and one person to operate them. A multiple-user environment (which we hereafter shorten to multiuser), consists of a single microcomputer, two or more terminals, and two or more people simultaneously running programs on the computer.

With a few exceptions, the rules for using Access Manager are the same in single-user and multiuser environments. When a distinction is necessary in the procedure, the text is preceded by [SINGLE] or [MULTI], whichever is applicable.

Hexadecimal values are noted by the letter h appended to the number. For example, 20h represents a hexadecimal value of 20; 0FFh represents hexadecimal

FF.

All examples shown in your "Access Manager Reference Manual" are coded in CBASIC Compiler language. Similarities to PL/I and Pascal/MT+ should be apparent as the logic is the same in all cases.

Table of Contents

0000

Tables

0000

Figures

0000

Listings

0000

Section 1: Product Description

1.1 What is Access Manager?

Think about your public library for a moment. There is a staggering volume of information located there. Yet, the card indexes make it possible to find specific information quickly and easily. Access Manager is something like the card index files in the public library. But, Access Manager provides a way to find specific information on computer disks, rather than between the covers of a book.

The librarian prepares the card index in the public library based on specific subject matter, with Access Manager you decide how to index the information in your computer files.

Access Manager is a general purpose, keyed, file accessing method for microcomputers. It is designed for use with application programs that need to access data records based on identifying key values, such as a name or a number.

1.2 Access Manager architecture

Access Manager is comprised of several different functions that you can call from your application program. These functions are categorized according to basic purpose. This categorization has no effect on how you can use Access Manager, but does help explain its logical structure.

Access Manager functions are categorized into the following groups:

- system initialization and maintenance
- index file setup and maintenance
- index file updates
- index file search
- data file setup and maintenance
- data file updates
- data file and record locking

1.2.1 System initialization and maintenance

These are commonly known as housekeeping functions. They prepare Access Manager to work with your application program by indicating how you want to handle errors, how many data files will be used by the program, and more. Normally, your program will only use these functions at the beginning.

1.2.2 Index file setup and maintenance

These functions ready your index files for use by Access Manager and your application program. They are used to open a file, close it, or save updates made to it. There are also functions for determining the size of the file, and erasing it when desired.

1.2.3 Index file updates

Update functions are used to add information to, or remove it from, an index file. You can also change the information the index uses to point to its associated data records. These pointers are called data record numbers.

1.2.4 Index file search

To locate a data record in a data file, you must first find its corresponding entry in the matching index file. Index File Search functions help you find the index file entry. They can be used to find specific entries in the index or entries relative to some other value. For example, with these functions, you can find the first entry in an index, the last entry, a specific entry, an entry that is greater than or less than one you specify, and more.

1.2.5 Data file setup and maintenance

These functions ready a data file for use by Access Manager and your application program. For example, they can be used to open a data file, close it, or save updates made to it at critical points in your program. There are also functions to aid in counting the entries (or records) in a data file, or to erase it when necessary.

1.2.6 Data file updates

You can use these functions to make actual changes in a data file. For example, there is a function to read a record from the file, and another to write a record into it.

1.2.7 Data file and record locking

In multiuser environments, it is often necessary to protect information in a data file to prevent its use at critical time; such as while a customer's balance is being updated. There are Access Manager functions to prevent anyone from using a data file or data record at these critical points.

The wide variety of Access Manager functions provides you with simple methods to accomplish the following:

- create and/or maintain data files
- create and/or maintain indexes for your data files
- retrieve information from your data files
- protect the integrity of your data files when they are being used simultaneously by two or more people

1.3 Access Manager benefits

There are numerous benefits to using Access Manager. Here are just a few:

- simplified programming
- fast data retrieval
- language portability

1.3.1 Simplified programming

Standardizing access methods for keyed data files across single-user and multiuser environments can be a complex programming task. This is further complicated by the variety of programming languages in use. Access Manager overcomes these complications by being readily compatible with various operating system environments and most programming languages.

Furthermore, you will find application programs can be developed faster, and are much less prone to error, when Access Manager is used.

1.3.2 Fast data retrieval

Most keyed access methods for retrieving information from a data file are slow, compared to the computer's processing speed. Consequently, many otherwise efficient application programs slow to a snail's pace when retrieving information under these conditions. Access Manager overcomes this problem with its unique indexing structure. The structure (known as a B-Tree, and described later in this section) minimizes the number of times an index file has to be read to locate a data record. Under typical circumstances, Access Manager can locate a specific record in a file of one-half million records with a maximum of four disk accesses. This is not only fast, it is efficient use of your hardware.

With Access Manager, you can reduce disk accesses by properly allocating disk buffers. Each buffer holds an index file record, and is shared by all the index files. A least-recently-used priority scheme manages the assignment of index records to buffers. Of course, Access Manager checks to see if the required index record is in a buffer before needlessly accessing the disk.

1.3.3 Language portability

Access Manager can be used with a variety of programming languages. And because the functions are standardized, application programs written in different languages are able to access a common data base. For example, in a multiuser environment, one user can be accessing the data base via an application program written in PL/I-80, another user accesses the same data base through a Pascal/MT+ program, and yet another via a CBASIC Compiler program.

1.4 Concepts and facilities

Access Manager uses the following concepts and facilities to efficiently create, index, share, and recreate your index and data files:

- file structures
- B-Tree structure
- data locking
- file recreation

1.4.1 Files structures

Access Manager provides the necessary functions to create and manage index files and data files on secondary storage devices, such as floppy and hard

disks.

Data files

Every records in a data file used with Access Manager must be the same length. That is, every record must contain the same number of bytes, and the record length must be at least four bytes.

Data records

Unlike data base systems, Access Manager treats a data record as a single entity, not a collection of fields. This means your application program must parse each data record into the required fields when necessary. One or more fields in a data record can be used as the key to that particular record. Normally, the key is a name, number, or other value uniquely identifying the contents of the data record.

Index files

Index files contain the key value for each record, and its assigned data record number. With just a key value, the index files make it possible for an application program to locate the associated data record (even allowing for duplicate keys) without a lengthy search of the data file. Access Manager creates index files using a height-balanced, multiway tree structure known as a B-Tree. This index structure guarantees the least number of disk accesses to search an index file. Besides being speedy, B-Trees eliminate the need to reorganize the index files.

Access Manager assumes nothing about the relationships between index and data files. Your application program must interpret the contents of the index files and how, if at all, they relate to the data files. Because Access Manager does not explicitly link index and data files, you can create key values in any way that suits your application.

Access Manager's separate index and data files permit flexibility in your program design. For example, several index files can reference one data file, or an index file can be referenced without a companion data file. For exceptionally large data files, one index file can reference many volumes. However, this flexibility leaves you with the responsibility for maintaining consistency between your index and data files.

Index and data file concepts used by Access Manager are discussed in detail under the ADDKEY function description in Section 3.

1.4.2 B-Tree structure

A B-Tree is a height-balanced, multiway tree structure. Under this structure, the tree is inverted, with the root at the top and the nodes at the bottom. Access Manager uses a variant of the B-Tree structure known as a B+Tree, where all key values are stored at the bottommost level of the tree.




Figure 1-1. B+Tree structural diagram

Figure 1-1 shows part of a simple B+Tree structure. To locate the data record with a key value of 180, the following steps are taken:

- 1) Access the root node.
- 2) Because 180 lies between 150 and 270, the middle branch is taken to access node 12.
- 3) 180 is greater than 178, so the right branch is taken to access node 4.
- 4) Because a match is found in node 4, the data record corresponding to key value 180 is record number 9 in the data file.

Consider these key points regarding Access Manager's index structures:

- In a multiway tree structure, many branches can emanate from each node of the tree. By comparison, in a binary tree there are at most two branches from each node. The advantage of additional branches is the height of the tree decreases as the number of branches per node increases; and the height of the tree determines the maximum number of

disk accesses to search the tree.

For example, if Access Manager is set up with 512-byte nodes, and the length of the key values is ten bytes, up to 34 key values can be stored in each node. This ensures that a tree structure with no more than four levels can store 192,780 key values under worst case conditions. By contrast, the same number of key values in a binary tree would require 18 levels.

- A height-balanced tree structure ensures that all nodes at the bottom of the tree are equidistant from the root node. Therefore, no key value ever has a long access path. A height-balanced tree eliminates the problems of overflow areas with unpredictable access times.
- Placing all key values in the leaf nodes speeds and simplifies sequential key values accesses, because Access Manager links the nodes in both ascending and descending key value order.

Any time you add a key to, or delete on from, an index file using Access Manager, the required changes in the B+Tree structure are made automatically. This guarantees that key searches always remain efficient. Reorganization of the index structure is not necessary, even after thousands of updates to the index file.

If you want further information on B-Tree structures, the following reading material is recommended:

- Comer, D. "The Ubiquitous B-Tree", "ACM Computing Surveys", Vol.11, No.2, June 1979, pp.121-137.
- Knuth, D.E., "The Art of Computer Programming, Vol.3, Sorting and Searching", pp.451-479, Reading, Massachusetts: Addison-Wesley Publishing Company, 1973.

1.4.3 Data locking

To ensure that data files maintain their integrity in multiuser environments, Access Manager provides the ability to lock data records and/or data files. The locking facilities make it possible for the same or different application programs running concurrently to share and/or update index and data files at the same time.

The locking facilities take two forms: shared locks and exclusive locks. Shared locks permit separate users simultaneous access to the same data file or data record. An exclusive lock can only be held by one user at a time, and bars all other users from having access to the locked data file or data record.

Data record locks

An application program can request a shared or exclusive record lock on individual data records. Any number of users can hold shared locks on the same data record at the same time. However, only one user can hold an exclusive record lock on a particular data record. Used correctly, the lock allow several users to access a data record, but only one to update it. Access Manager record locks are set at the logical record level, not the physical or logical sectors of the operating system or storage medium.

Data file locks

Besides locking data records, Access Manager can set locks on individual data files. Once a user holds an exclusive data file lock, all other requests for file or record locks on that file are refused. A shared data file lock signals your intent to use and possibly update a data file without the need to block other users from it.

Note: The locking facilities pertain only to data files and data records. The only way to exclude other users from accessing an index file is by using passwords.

1.4.4 File recreation

Access Manager contains a utility program for quickly recreating index and data files when file integrity has been lost. You can set up a special file containing parameters which tell the utility program precisely how to reconstruct the files. Thus, file reconstruction becomes a simple task when necessary. Section 5 contains a complete description of the RECREATE utility program.

1.5 Application program structure

This subsection discusses common structures for application programs that use Access Manager. However, you should not conclude from this discussion that other program structures or uses of Access Manager are inappropriate.

You will find the suggested program structures in Figure 1-2 and 1-3 more meaningful after you have studied the individual function descriptions in Section 3.

```
+-----+
| Initialize Access Manager |
| (INTUSR, SETUP)         |
+-----+
||
```


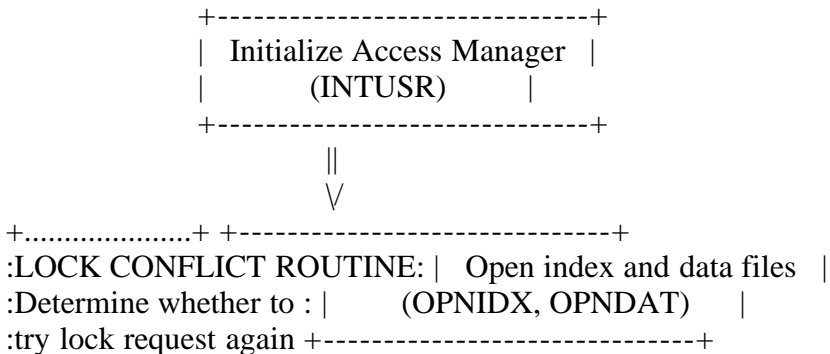


Figure 1-2. Single-user program flow



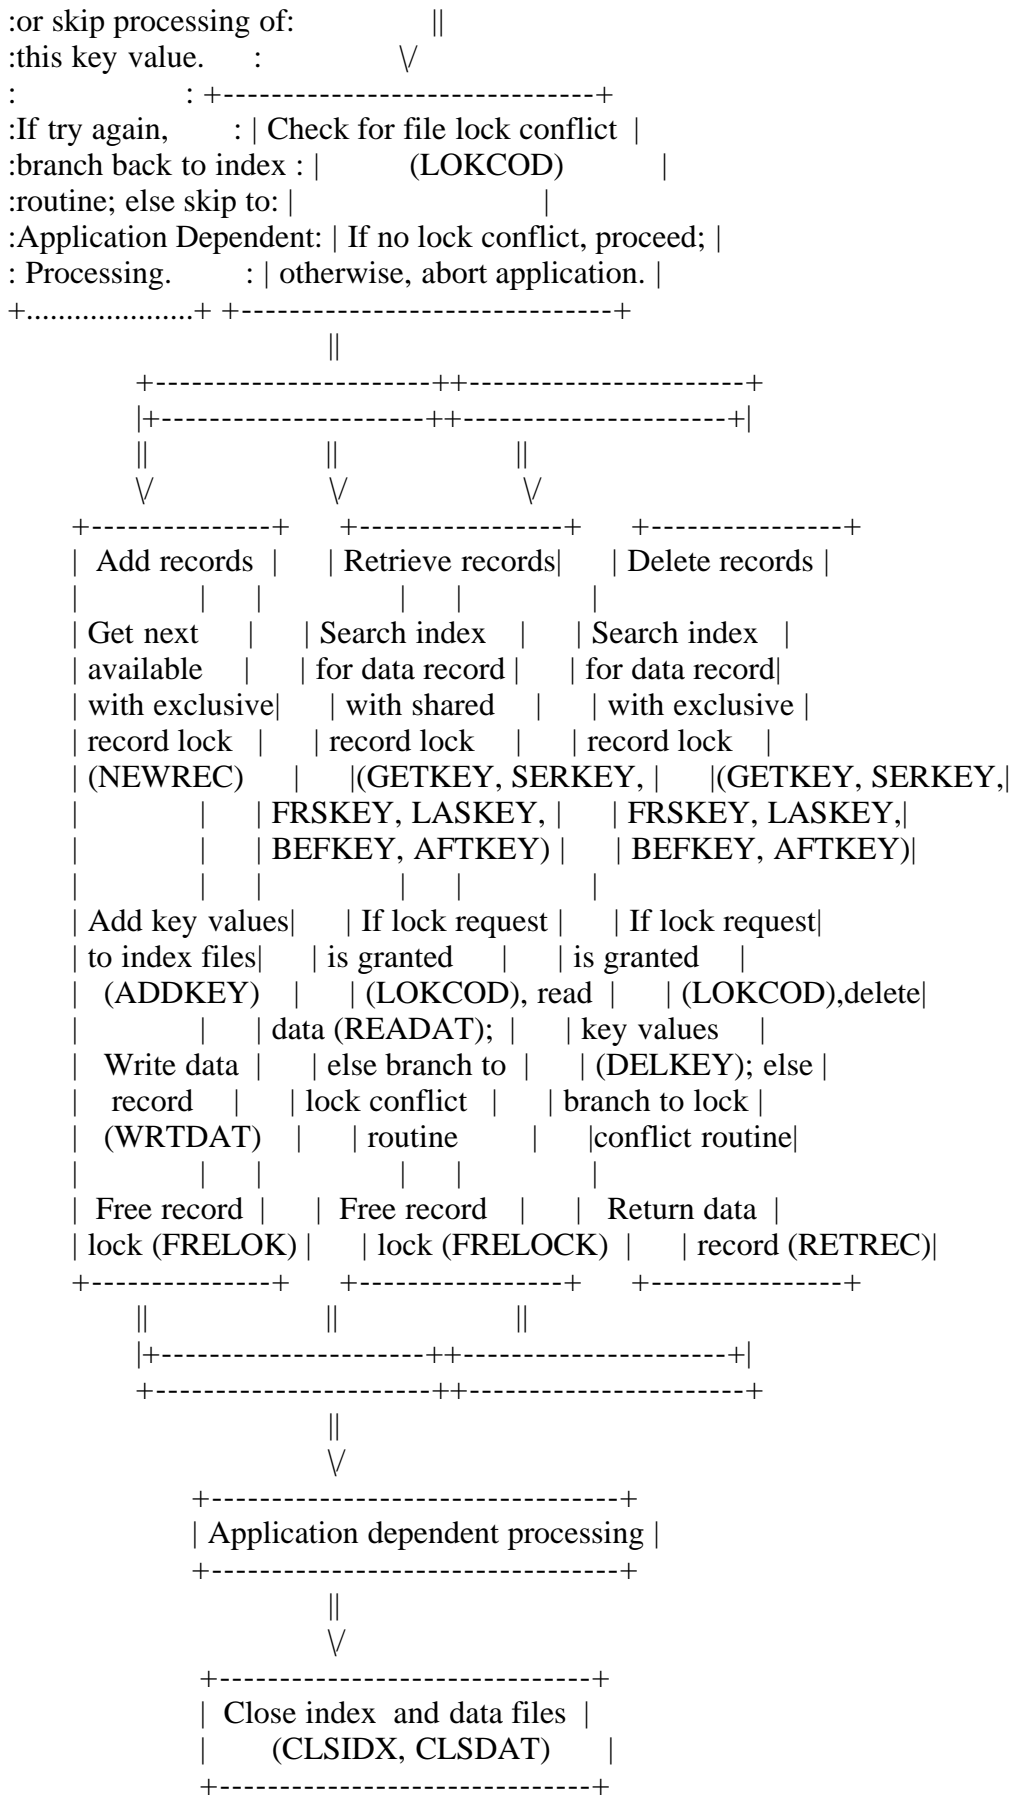


Figure 1-3. Multiuser program flow

Figure 1-2 outlines the program flow in a single-user environment. If your application is intended for a multiuser environment, the program flow in Figure 1-3 is more appropriate. Remember, an application program that runs in a multiuser environment also runs, without modifications, in a single-user environment. For this reason, you might want to consider designing all of your application programs for a multiuser environment.

If an application program requires more than 65,535 entries in an index file or records in a data file, you will want to modify the program structures outlined in Figures 1-2 and 1-3 to call the DATVAL function after any other Access Manager function that returns a data record number. Call the SETDAT function before any other Access Manager function that accepts a data record number as an input parameter.

To simplify the program structures, error handling routines have been omitted from Figures 1-2 and 1-3. If you do not want error trapping, no changes are required to the program structures shown. This structure causes user error messages to display or print on the console, and then return control to your operating system. If you want to control program actions when user errors occur, call the INTUSR function with a non-zero ERROPT% parameter. In this case, you should call the ERRCOD function after each call to the other Access Manager functions. If ERRCOD returns with a non-zero value, you can either branch to your error handling routines to recover, or perform an orderly shutdown of the application.

Note the following points about the multiuser program flow shown in Figure 1-3.

- Your application program can request data record locks as part of the index search or get-next-available-data-record function. This eliminates the need to get a data record number and then request a lock as a separate action.
- If your application program opens a data file with a shared or exclusive data file lock, a call to NEWREC with an exclusive data record lock is always accepted.
- If a lock request is denied as part of an index search, the search function still returns the associated data record number and the key value found in the index (via the IDXVAL\$ parameter). Therefore, your application program can use the before and after key functions (BEFKEY and AFTKEY) to skip the locked item.
- The RETREC function, which deletes data records, automatically releases the data record lock if the user holds one. Therefore, do not call the FRELOK function after RETREC.

The importance of data record locking in multiuser environments cannot be overemphasized. If you set shared data record locks when you scan or review a data record and set exclusive record locks when updating, your application program will operate smoothly in multiuser environments.

Remember, Access Manager makes it possible to override the lock settings. Unless your experience and requirements demand you make such overrides, you

should strictly adhere to the lock settings. In particular, if a LOKCOD value indicates you did not get the requested lock, be sure your application program can follow an appropriate course of action. Do not ignore the LOKCOD results!

Your Programmer's Guide contains examples illustrating the use of Access Manager to create and maintain a simple data base. Your distribution disk contains the complete source code for these examples.

Section 2: Function parameters

Most Access Manager functions are called with parameters that restrict or determine what the function will do. Before reading the function descriptions in Section 3, you should have a basic understanding of these parameters.

2.1 Parameter types

There are three types of parameters used with Access Manager functions: two-byte integers, characters strings, and pointers.

2.1.1 Two-byte integers

Access Manager treats two-byte integers as unsigned quantities. Most application languages treat them as signed quantities ranging from -32,768 to +32,767. Therefore, if you print a data record number returned from an Access Manager function, you might get a negative result. For example, the largest, unsigned, two-byte quantity is 65,535, which corresponds to -1 as a signed quantity. If you want to display the actual, unsigned value of a two-byte integer, add 65,536 to the quantities in the negative range.

Data record numbers and attributes measuring the number of data records and index file entries are stored in the files as four-byte integer quantities. Two-byte integers span the range from 0 to 65,535. Because most application programs find this range of data record numbers sufficient, Access Manager is designed to use two-byte parameters. Whenever an application program requires the four-byte capacity of Access Manager, the SETDAT and DATVAL functions allow an additional two bytes to be passed to or returned from Access Manager. For example, if your program references multiple data files by a single index file, the two-byte integer returned by an index search function might represent the data record number, while the additional two bytes represent the data file number. SETDAT is called before, and DATVAL after, other Access Manager functions.

2.1.2 Character strings

Key values and filenames are always represented as character strings.

2.1.3 Pointers

The READAT and WRDAT functions use pointers as one of their parameters. Different operating systems sometimes require different pointer lengths. Fortunately, your application language automatically sizes pointer parameters correctly.

2.2 Parameter descriptions

Descriptions of the parameters used with Access Manager functions are summarized in Table 2-1. Specific information on the use and content of each parameter is provided with the individual function descriptions in Section 3. Note that parameter names are suffixed with a dollar ("\$\$") or percent ("%") sign to indicate the type of value it must contain:

\$ = character string value

% = two-byte integer

Table 2-1. Access Manager parameters

Format: Parameter
Description

BUFFER%

This parameter is used only in conjunction with the READAT and WRDAT functions. It contains a pointer to the input/output buffer area in your application program. When your program reads a data record, Access Manager places it in this buffer area; when it writes a data record, it is written from the buffer.

DLOCK%

[MULTI] A code is passed in this parameter to request a lock on, or release a lock from, a data file or data record. See the SETLOK function for lock request codes; the FRELOK function for lock release codes.

DRN%

All records in a data file are assigned a unique data record number. Entries in an index file contain the key value of the corresponding record in the data file, and its assigned data record number. Thus, the data record number points to the record in the data file with an equal key values. Note that a data record number of zero is considered an error by Access Manager.

DUPKEY%

The value passed in this parameter tells Access Manager how duplicate key values in an index file are to be handled. DUPKEY% is used in conjunction with the KEYTYP% parameter. If DUPKEY% is one and KEYTYP% is zero, Access Manager assigns unique sequence numbers to each key value. If KEYTYP% is other than zero, DUPKEY% has no effect.

ERROPT%

The value passed in this parameter tells Access Manager whether or not to trap user errors for handling by your application program. A zero value disables the error-trapping facility; a non-zero value causes Access Manager to trap user errors, and make them available to your application program. See the INTUSR function in Section 3 and in Section 4, "Error Codes", for further information.

FILE%

With Access Manager, your application program can process several data files simultaneously. Hence, every data file is assigned a unique file number. The value passed in this parameter tells Access Manager which data file you want to access. Data file numbers used by Access Manager are separate from those used with your application language. Assigning a particular number to an Access Manager data file does not preclude your using the same file number with an application language file.

FILNAMES\$

Before opening a data file, Access Manager looks in the directory of the disk for the name of the file. The value passed in this parameter must match exactly the name of the data file as it is recorded in the directory.

IDXNAMES\$

Before opening a data file, Access Manager looks in the directory of the disk for the name of the file. The value passed in this parameter must match exactly the name of the data file as it is recorded in the directory.

IDXVAL\$

This is the only function parameter that passes a value to your application program. It is used in conjunction with those functions that locate a key value in an index file. At the conclusion of one of these functions, Access Manager places the value of the key it locates in this parameter. Your application program can then use subsequent functions to process that key as required.

KEY%

With Access Manager, your application program can process several index files simultaneously. Hence, every index file is assigned a unique file number. The value passed in this parameter tells Access Manager which index file you want to access. Index file numbers are independent of data file numbers assigned by Access Manager or the application language.

KEYLEN%

To open an index file, Access Manager needs to know the length of the key values in the file. This parameter is used to indicate the number of bytes contained in each key value.

KEYTYP%

To open an index file, Access Manager must know whether the key values are stored in alphanumeric or numeric order. Pass a zero in this parameter to indicate alphanumeric order; a one indicates that key values are stored in numeric order.

KEYVAL\$

Entries in an index file contain a key value, and associated data record

number. To add, delete, or retrieve an entry from an index file, your program must specify the exact value of the key for that entry. This parameter is used to pass the key value to Access Manager.

NBUFS%

[SINGLE] This parameter specifies the number of input/output buffers allocated to your application program. Note that there must be at least three such buffer areas.

NDATF%

[SINGLE] This parameter specifies the maximum number of data files that will ever be open by your program at any given time.

NKEYS%

[SINGLE] This parameter specifies the maximum number of index files that will ever be open by your program at any given time.

NNSEC%

[SINGLE] This parameter determines or specifies the length of the records in an index file. Specifically, it refers to the number of 128-byte disk sectors in each index file record.

PROGID%

[MULTI] This parameter contains the unique identifier assigned to each program or task.

RECLLEN%

The value of this parameter tells Access Manager the length (that is, number of bytes) of each record in a data file.

TIMOUT%

[MULTI] This parameter specifies the number of seconds within which the background server must respond to the INTUSR function call.

Table 2-2. Parameter use table

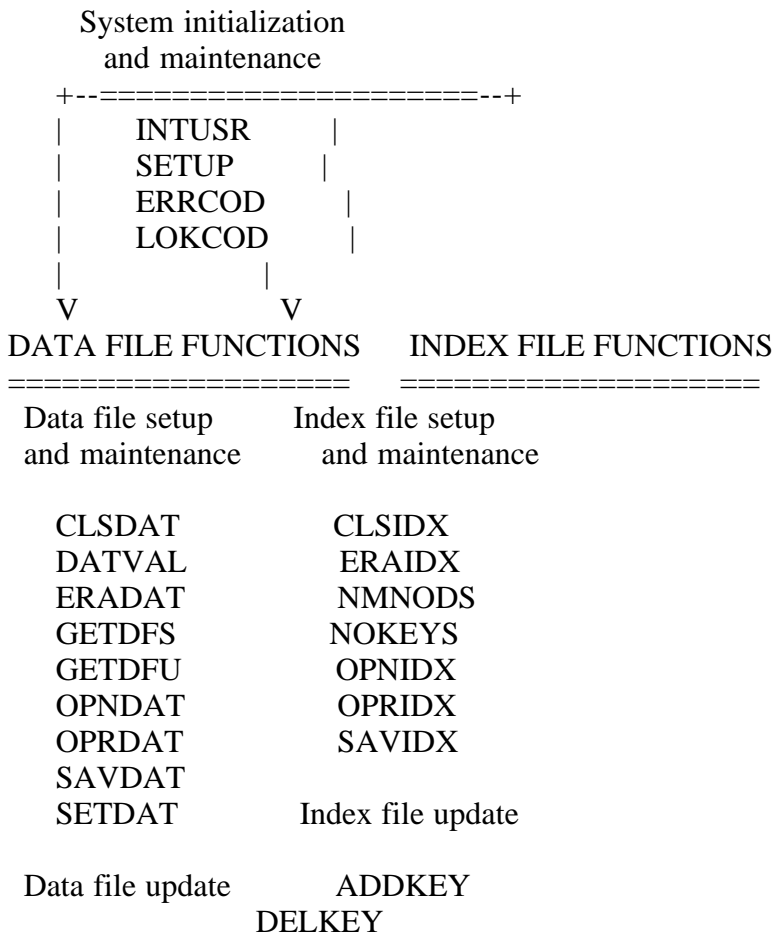
Parameter	B	D	D	D	E	F	F	I	I	K	K	K	N	N	N	N	P	R	T		
---->	U	L	R	U	R	I	I	D	D	E	E	E	E	E	B	D	K	N	R	E	I
	F	O	N	P	R	L	L	X	X	Y	Y	Y	Y	U	A	E	S	O	C	M	
Function	F	C	.	K	O	E	N	N	V	.	L	T	V	F	T	Y	E	G	L	O	
	E	K	.	E	P	.	A	A	A	.	E	Y	A	S	F	S	C	I	E	U	
	R	.	Y	T	.	M	M	L	.	N	P	L	.	.	.	D	N	T			
V	E	E	

- A list of the parameters required in the command syntax.
- When appropriate, a section of additional comments regarding uses and restrictions placed on the function.
- A descriptive list of the error codes that can be returned by the function. Note that only a brief explanation of the problem is provided. A complete description of each can be found in Section 4, "Error Codes".
- When appropriate and useful, a segment of program code is included as an example of how to use the function in a program.

Access Manager functions return values to your program. The returned value can be an error code, a code indicating the results of the function, or a data record number. Your program must, in all cases, examine the returned value to determine exactly what took place.

Some functions, for example, ADDKEY, DELKEY, and UPDPTR, do not return error codes. The ERRCOD function should be called following any of these functions, to test for errors (unless error trapping is not enabled by the ERROPT% parameter). ERRCOD returns a non-zero value if a user error occurs.

Figure 3-1 shows the way Access Manager functions are categorized. The categories are explained in Section 1.



NEWREC	UPDPTR
READAT	
RETREC	Index file search
WRDAT	
	AFTKEY
Data locking	BEFKEY
	FRSKEY
FRELOCK	GETKEY
SETLOCK	LASKEY
	NXTKEY
	PRVKEY
	SERKEY

Figure 3-1. Access Manager functions by category

ADDKEY function Index file update function

Syntax: ADDKEY (key%, file%, dlock%, keyval\$, DRN%)

Explanation:

ADDKEY adds a new key and associated data record number to an index file and, optionally, assigns sequence numbers to duplicate key values. At the conclusion of the function, a value is returned indicating its success or failure (see "Additional Comments" below).

Parameters:

KEY%

The number of the index file where the key value is added.

FILE%

[MULTI] This parameter is ignored in a single-user environment. If you want to ignore it in a multiuser environment, set the DLOCK% parameter to zero. It is the number of data file associated with the index file specified by the KEY% parameter.

DLOCK%

[MULTI] This parameter is ignored in a single-user environment. If you want to ignore it in a multiuser environment, assign a value of zero. It is the type of data record lock requested for the data file specified by the FILE% parameter. See the SETLOK function for a list of acceptable lock codes.

You might encounter situations in a multiuser environment where the associated data record, DRN% parameter, should already have an exclusive record lock before calling ADDKEY (for example, by a previous call to NEWREC). In this case, ADDKEY can be called with a zero DLOCK%, or DLOCK% can be set to two (the code for an exclusive record lock). The lock request should always be

granted because, presumably, the same program already holds an exclusive lock. The advantage of the latter approach is that, if the requested exclusive record lock is not granted, you might have detected a flaw in your program logic.

KEYVAL\$

The key value to be added to the index file. Your application program must ensure that the data record number, DRN% parameter, corresponds to KEYVAL\$. If the associated data record number requires more than two bytes, call the SETDAT function immediately before ADDKEY with the higher-order two-byte value as its argument.

If you pass a null string in KEYVAL\$, the ADDKEY function takes no action but return a one, unless the associated lock request is denied. A string is null if it has a zero length. Access Manager handles null key values in this manner, to simplify coding for multiple key applications. For example, if a data file is indexed by name, number, and zip code, a loop that adds the key values to their respective index files can ignore whether or not a key value is actually present for an index, by representing missing values as null strings.

DRN%

The data record number associated with KEYVAL\$.

Additional comments:

The ADDKEY function returns the following values to indicate its success or failure:

Table 3-1. ADDKEY function values

Value Meaning

- 0 KEYVAL\$ has been added to the index but the automatic, duplicate-key sequence numbers are exhausted.
- 1 KEYVAL\$ has been successfully added to the index.
- 2 KEYVAL\$ already exists in the index. The index file update has not taken place.
- 4 The DLOCK% request for the data record (DRN%) in the data file (FILE%) cannot be granted. KEYVAL\$ was not added to the index file.

The only assumption Access Manager makes regarding the content of an index file entry is that zero is never used for the data record number. If this happens, user error 36/BD occurs.

In most applications, you should include the key values in both the data file and index file records. This redundancy provides the best protection when or if you have to reconstruct your index files. You should only omit key values

from data files if you have an extreme secondary storage space constraint and there are other ways to determine the key values associated with each data record.

Key value padding

ADDKEY pads KEYVAL\$'s that are less than KEYLEN% bytes with blank spaces (20H) on the right, and truncates KEYVAL\$'s that are too long. However, to ensure proper handling of numeric keys, all numeric KEYVAL\$'s must be passed to the Access Manager functions with the exact KEYLEN%.

If you want to store alphanumeric KEYVAL\$'s in right-justified form, you must ensure that the KEYVAL\$'s are properly justified in view of the truncation of oversized keys.

Duplicate key values

There are many situations, such as building an index based on last names, where the key values are not unique. When this occurs, you must append a unique identifier to the key, possibly after truncating the original key value to a prespecified length. This way, the KEYVAL\$'s are still stored in the expected order, but there is no conflict between like-valued entries.

If DUPKEY% is one when the index file is opened, Access Manager automatically places a unique, binary sequence number in the last two bytes of the key value. For example, if KEYLEN% is ten and KEYVAL\$ equals the string 12345, Access Manager adds three spaces of padding and a two byte sequence number to the end of KEYVAL\$ before adding it to the index file. Each set of duplicate key values maintains a separate set of sequence numbers. Thus, there can be up to 65,535 entries in each set of duplicates. Each time another duplicate is added to an existing set of duplicates, the new entry gets the next higher sequence number. Therefore, Access Manager stores duplicate key values in entry order.

You can exhaust the sequence numbers if more than 65,535 identical entries are made. However, if the last entry in a set of duplicates is deleted, this sequence number is reused before the sequence numbers are incremented. When the last available sequence is used for a set of duplicates, ADDKEY returns a value of zero. At this time, the index file must be rebuilt or it will, either start to reject duplicates with a return code of two, or enter the duplicates out of sequence. Which of these possibilities takes place depends on the pattern of sequence numbers still in the set of duplicates.

Note: Read the DELKEY function description for important information concerning the adverse effect of automatic duplicate keys on the time required to delete a duplicate key value.

Coding numeric key values

If your application requires integer key values, Access Manager provides three alternatives for representing them. In all cases, however, the KEYVAL\$ parameter must be a string valued variable. Only the last alternative uses signed, integer key values; the first two alternatives require alphanumeric key values.

The simplest approach to represent an integer key value is to create a string variable which equals the ASCII representation of the integer. With the CBASIC Compiler, the STR\$ function automatically performs this conversion. In PL/I, a simple assignment statement of a numeric variable to a CHAR VAR string accomplishes the conversion. In Pascal/MT+, redirected I/O can be used to convert a numeric quantity to a string.

If you use this approach, the resulting strings must be carefully right-justified and padded with blank spaces or zeros on the left, and KEYTYP% should be zero (for an alphanumeric key type). The main disadvantage to this approach is that the key length must be set to the maximum number of digits in the number, as opposed to the number of bytes required to store the number in internal binary format.

The two remaining approaches for representing integer key values do take advantage of the compact representation of integers in binary format. One approach treats the key values as signed quantities, while the other applies to unsigned integers. In both cases, it is necessary to create a string with characters actually equal to the bytes that comprise the integer quantity. The resulting string is probably an unprintable image, because any bit pattern from 00h to 0FFh might result in each element of the string.

For unsigned integer key values, the CBASIC Compiler function UNSIGNED.INT.KEY\$ (see example below) converts an integer quantity into a string equivalent. The function assumes NUMBER is a real, positive quantity; NUMBER can be represented in KEY.LEN% bytes. CHR\$ is capable of converting an arbitrary byte value to a string, whether or not the byte corresponds to a valid ASCII character.

```
DEF Unsigned.Int.Key$ (number, Key.Len%)
  STRING temp$
  INTEGER i%, byte%
  REAL factor

  temp$ = ""
  FOR i% = 1 TO Key.Len%
    factor = INT (number / 256.)
    byte% = number - 256. * factor
    temp$ = CHR$ (byte%) + temp$
    number = factor
  NEXT i%
  Unsigned.Int.Key$ = temp$
  RETURN
FEND
```

The preceding function example creates a string whose individual bytes correspond to the bytes necessary to represent NUMBER as an integer, with the

most significant byte first and the least significant last. As with the first approach, ensure KEYTYP% is zero (for alphanumeric keys).

If you want to treat the integers as signed quantities with negative values preceding positive ones, set KEYTYP% to one, and revise the UNSIGNED.INT.KEY\$ function so the least significant byte is first and the most significant byte (which includes the sign bit) is last. You must also convert NUMBER to a positive quantity before conversion. With these changes, the conversion function looks like this:

```
DEF Signed.Int.Key$ (number, Key.Len%)
  STRING temp$
  INTEGER i%, byte%
  REAL factor

  IF number < 0 THEN number = number + constant
  temp$ = ""
  FOR i% = 1 TO Key.Len%
    factor = INT (number / 256.)
    byte% = number - 256. * factor
    temp$ = CHR$ (byte%) + temp$
    number = factor
  NEXT i%
  Signed.Int.Key$ = temp$
  RETURN
FEND
```

In the preceding example, "constant" should be replaced by 256 raised to the KEY.LEN% power. Some example values for constant follow:

Table 3-2. Constant values

KEY.LEN%	Constant
2	65536.0
3	16777216.0
4	4294975296.0

The preceding values for "constant" also indicate the maximum value plus one for unsigned integer values. The signed integer key values range from $(-\text{constant}/2)$ to $(\text{constant}/2)-1$.

For example, let us compare the first and second approaches for integer keys as applied to social security numbers. If SOC.SEC.NO is a real quantity, then an example of using the first approach would be:

```
KEYVAL$ = RIGHT$ (" " + STR$ (SOC.SEC.NO), 9)
```

and using the second approach:

```
KEYVAL$ = UNSIGNED.INT.KEY$ (SOC.SEC.NO, 4)
```

The first approach is faster, but the index file is about 62.5% larger, because each entry requires thirteen bytes; nine for the social security

number, and four for the data record number. The second approach requires only eight bytes.

Note: If signed integer quantities are implemented by setting KEYTYP% to one in OPNIDX, the key value string passed to the index functions must have the least significant byte as the first byte of the string, followed by increasingly significant bytes, until the most significant byte is in the last position. Further, the most significant bit of the last byte is treated as the sign; zero if positive, one if negative.

If either the second or third approach is used to compact integer key values, it might be necessary to unpack these strings when they are returned to your program in the output parameter IDXVAL\$. The following function unpacks key values prepared according to the second approach:

```
DEF Unpack.Unsigned.Int (idxval$, Key.Len%)
  REAL temp, power
  INTEGER i%

  power = 1
  temp = 0
  FOR i% = Key.Len% TO 1 STEP -1
    temp = temp + power * ASC (MID$ (idxval$, i%, 1) )
    power = power * 256.
  NEXT i%
  Unpack.Unsigned.Int = temp
  RETURN
FEND
```

The function MID\$ returns the string, beginning at the position specified by the second parameter, with a length given by the third parameter. The function ASC returns the byte value of the first character of its string parameter.

See your Programmer's Guide for information concerning the use of numeric key values with PL/I or Pascal/MT+.

Large data files

An Access Manager index file must be contained in one physical disk file. However, this is not true for the data file, because the two files are separate entities. Hence, the data file can be spread over more than one disk file. This segmentation of the data file is represented in the index file by the data record numbers associated with the key values.

To use the associated data record numbers to keep track of segmented data files, set the two high-order bytes of the data record number to the segment number, and the two low-order bytes to the relative data record number in the segment. For example, the following code inserts the key value corresponding to the ten-thousandth data record of the fifth data file segment:

```
Data.Segment% = 5
DRN% = 10000
```

```
CALL setdat (Data.Segment%)
Ret.Code% = addkey% (key%, Data.Segment%, X.Lock%, keyvalue$, DRN%)
```

An important point concerning segmented data files is setting record locks during index file searches. Because the target segment (that is, data file) is unknown until after the index search, it seems unfeasible to set a record lock at the time the index is searched. However, you can avoid this by always setting the lock on the first segment, regardless of the actual segment where the associated data record is found. This approach is effective, because the lock function uses all four bytes (two for the actual segment, and two for the record number) to identify the locked record.

The following examples show how to manage a very large data file and its associated index file. To simplify the discussion, we assume only positive, two-byte logical record numbers (0 through 32,767) are required within each segment.

First, the physical segments comprising the large logical file are opened with shared file locks. For example,

```
FOR Seg.No% = 1 TO Total.Seg%
  File.No% (Seg.No%) = opndat (-1, 3, filename$ (Seg.No%), rlen%)
  IF errcod <> 0 THEN CALL Error.Handler (1)
  IF lokcod <> 0 THEN CALL Lock.Conflict (1)
NEXT Seg.No%
```

To add a new entry, you must first determine the correct physical segment for its placement. The simplest approach is to assume a fixed, upper limit on the active number of records in each segment. The following function returns the segment number to use for the new data file entry:

```
DEF Get.Segment.No% (Active.Rec.Limit%, No.Of.Segement%)
  INTEGER Seg.No%

  FOR Seg.No% = 1 TO No.Of.Segments%
    IF getdfu (File.No% (Seg.No%) ) < Active.Rec.Limit% \
      THEN Get.Segment.No% = Seg.No% : \
        RETURN
  NEXT Seg.No%
  CALL Error.Handler (2)  REM No more space
FEND
```

The next code example shows how to set locks and add the new entry to the data and index files. Assume up to 10,000 records can be active in each physical segment:

```
Null.Lock% = 0
Share.Rec.Lock% = 1
Xclsv.Rec.Lock% = 2
```

REM Get next available data record number.

```
Seg.No% = Get.Segment.No% (10000, Total.Seg%)
DRN% = newrec (File.No% (Seg.No%), Null.Lock%)
```

REM Always set lock on the first segment, with a 4-byte DRN
REM that includes the actual segment number.

```
CALL setdat (Seg.No%)  
IF setlok% (File.No% (1), Xclsv.Rec.Lock%, DRN%) <> 0 \  
  THEN CALL Lock.Conflict (3)
```

REM Add key value with upper portion of DRN = SEG.NO

```
CALL setdat (Seg.No%)  
IF addkey (key%, File.No% (1), Xclsv.Rec.Lock%, keyval$, DRN%) <> 1 \  
  THEN CALL Error.Handler (4)
```

REM Write data record to actual physical segment.
REM Note that SETDAT is not used, since DRN% is
REM the actual record number.

```
IF wrtdat (File.No% (Seg.No%), DRN%, Buffer.Ptr%) <> 0 \  
  THEN CALL Error.Handler (5)
```

REM Free exclusive record lock.

```
CALL setdat (Seg.No%)  
IF frelok (File.No% (1), Xclsv.Rec.Lock%, DRN%) <> 0 \  
  THEN CALL Lock.Conflict (6)
```

The next code example searches for a data file entry based on the specified
KEYVAL\$, and applies a shared record lock if it is found.

REM Find, lock, and read data record referenced by KEYVAL%.

Wait.Loop:

```
DRN% = getkey (key%, File.No% (1), Share.Rec.Lock%, keyval$)  
Seg.No% = datval  
IF lokcode <> 0 \  
  THEN PRINT "Lock request denied." :\  
    INPUT "Enter 'W' to wait, or press RETURN to skip."; LINE wait$ :\  
    IF wait$ = "" \  
      THEN GOTO Skip.Data \  
      ELSE GOTO Wait.Loop  
IF readat (File.No% (Seg.No%), DRN%, Buffer.Ptr%) <> 0 \  
  THEN CALL Error.Handler (7)
```

REM Process data (no updates, since shared lock).

```
CALL Process.Data
```

REM Free shared lock.

```
CALL setdat (Seg.No%)  
IF frelok (File.No% (1), Share.Rec.Lock%, DRN%) <> 0 \  
  THEN CALL Lock.Conflict (8)
```

Skip.Data:

The final code example demonstrates a deletion from the segmented data file.

REM Find the data file entry by key value.

```
DRN% = getkey (key%, File.No% (1), Xclsv.Rec.Lock%, keyval$)
Set.No% = datval
```

REM Test if entry is found.

```
IF DRN% = 0 \
  THEN GOTO Skip.Delete
```

REM Test if exclusive lock obtained.

```
IF lokcode <> 0 \
  THEN PRINT "Cannot delete "; keyval$; ". Currently in use!" :\
  GOTO Skip.Delete
```

REM Delete key value.

```
CALL setdat (Seg.No%)
IF delkey (key%, File.No% (1), Xclsv.Rec.Lock%, keyval$, DRN%) <> 1 \
  THEN CALL Error.Handler (9)
```

REM Free exclusive lock.

```
CALL setdat (Seg.No%)
IF frelok (File.No% (1), Xclsv.Rec.Lock%, DRN%) <> 0 \
  THEN CALL Lock.Conflict (10)
```

REM Return record to available pool.

```
IF retrec (File.No% (Seg.No%), Null.Lock%, DRN%) <> 0 \
  THEN CALL Error.Handler (11)
```

We have assumed that all access to the data file is through the index file. If you must scan the data file without reference to the index file, it is best to perform such as scan with an exclusive file lock, which ensures no unexpected updates to the file during your processing.

Error Codes:

The ADDKEY function can cause these error codes:

Table 3-3. ADDKEY error codes

Value	Code	Explanation
21	AE	No directory or disk space available.
30	AN	Index file number (KEY%) is out of range.

36 BD Key value has a zero data record number.
46 BN Index file is not open.
147 IC Lock code (DLOCK%) is out of range.
153 II Bad parameter value.

Example:

```
DEF exception (locale, Return.Code)
  INTEGER locale, Return.Code

  PRINT "Exception at location "; locale
  PRINT "  Error Code: "; errcod; \
    "  Lock Code: "; lokcod
  PRINT "  Return Code: "; Return.Code
  STOP
FEND
```

```
INPUT "Part Name: "; Key.Value$
DRN% = newrec (File.No%, X.Lock%)
IF errcod <> 0 \
OR lokcod <> 0 \
  THEN CALL exception (1, 1)
Ret.Code% = addkey (key%, Inv.File%, X.Lock%, Key.Value$, DRN%)
IF errcod <> 0 \
OR Ret.Code% <> 1 \
OR lokcod <> 0 \
  THEN CALL exception (2, Ret.Code%)
```

EXCEPTION provides a concise method to communicate unexpected results during program development. EXCEPTION is not designed to handle on-line resolution of lock conflicts or error conditions.

AFTKEY function Index file search function

Syntax: AFTKEY (key%, file%, dlock%, keyval\$, idxval\$)

Explanation:

AFTKEY returns the data record number associated with the first entry, in key-sequential order, immediately following (that is, strictly greater than) a specific key value. AFTKEY also places the key value it finds in the IDXVAL\$ parameter.

Use AFTKEY to move forward through an index file sequentially in key-ascending order if one of the following situations apply:

- your application operates in a multiuser environment

- you want to update a data file while moving through its index file

If either of the preceding situations apply, you will find it efficient to use the NXTKEY function. However, using AFTKEY ensures compatibility between single-user and multiuser environments.

Parameters:

KEY%

The number of the index file where the key value is added.

FILE%

[MULTI] This parameter is ignored in a single-user environment. It is the number of the data file referenced by the KEY% parameter.

DLOCK%

[MULTI] This parameter is ignored in single-user environments. It specifies the type of data record lock requested for the data file referenced by the FILE% parameter. If the requested lock for the data record that AFTKEY found cannot be granted, LOKCOD returns a non-zero value. See the SETLOK function for a list of acceptable lock codes.

KEYVAL\$

The key value that precedes the one being sought. Note that KEYVAL\$ does not have to exist in the index for AFTKEY to function. KEYVAL\$ serves only as a reference value for seeking the actual index entry.

IDXVAL\$

This is an OUTPUT parameter. Access Manager places the key value found in the index file in IDXVAL\$ at the conclusion of the function. If no index entry is found with a key value greater than KEYVAL\$, AFTKEY returns zero, and IDXVAL\$ is set to all blank spaces (which is not the same as a null string!...).

Before calling the AFTKEY function, IDXVAL\$ must be at least as long as the key length specified for the KEY% file. Normally, it is only necessary to initialize IDXVAL\$ once at the beginning of a program. Notice that KEYVAL\$ and IDXVAL\$ must not be the same variable.

Additional Comments:

To determine the two high-order bytes of the associated data record number, call the DATVAL function immediately following AFTKEY.

Error Codes:

The AFTKEY function can cause the error codes listed in Table 3-4.

Table 3-4. AFTKEY error codes

Value	Code	Explanation
30	AN	Index file number (KEY%) is out of range.
35	BC	IDXVAL\$ too short to contain key value.
46	BN	Index file is not open.
147	IC	Lock code (DLOCK%) is out of range.
153	II	Bad parameter value.

Example:

The following example listing shows many of the index and data file functions. It creates a physically ordered data file from an existing indexed data file. Also, the index file is updated in place to reflect the new data record positions.

REM External AM80 declarations.

%INCLUDE am80extr.bas

REM Error handling routine.

DEF Error.Handler (locale)

INTEGER locale

PRINT "Fatal ERROR at Locale "; locale

PRINT " ERROR Code "; errcod

PRINT " Return Code "; Ret.Code%

STOP

FEND

DEF Temp.Lock

PRINT "Temporary Sort File "; Fil.Name\$:\

" cannot be locked!" :\

DUMMY% = clsdat (Old.File%) :\

STOP

FEND

REM System parameters.

yes% = -1

no% = 0

nbuf% = 6

nkeys% = 1

nnsec% = 4

ndatf% = 2

X.File% = 4

N.Lock% = 0

progid% = -1


```
Error.Trap% = yes%
Time.Out.Test% = 3
```

```
REM System initialization.
```

```
progid% = intusr (progid%, Error.Trap%, Time.Out.Test%)
```

```
IF errcod <> 0 \
```

```
    THEN CALL Error.Handler (1)
```

```
IF setup (nbuf%, nkeys%, nnsec%, ndatf%) <> 0 \
```

```
    THEN CALL Error.Handler (2)
```

```
REM Data file set up.
```

```
INPUT "Enter data filename & record length:"; Fil.Name$, Rec.Len%
```

```
Old.File% = -1
```

```
New.File% = -1
```

```
Old.File% = opndat (Old.File%, X.File%, File.Name$, Rec.Len%)
```

```
IF errcod <> 0 \
```

```
    THEN CALL Error.Handler (3)
```

```
IF lokcod <> 0 \
```

```
    THEN PRINT "Data File "; Fil.Name$; " cannot be locked at"; \
```

```
        " this time. Try later!" : \
```

```
    STOP
```

```
Fil.Name$ = "SORT.$$" + RIGHT$ (STR$ (progid%), 1)
```

```
New.File% = opndat (New.File%, X.File%, Fil.Name%, Rec.Len%)
```

```
REM Check for left over SORT.$$. If lock granted,
```

```
REM erase then reopen file, to ensure fresh data file.
```

```
REM See Section 4 for description of User Error Codes.
```

```
IF errcod <> 0 \
```

```
AND errcod <> 53 \
```

```
    THEN CALL Error.Handler (4)
```

```
IF lokcod <> 0 \
```

```
    THEN CALL Temp.Lock
```

```
IF eradat (New.File%, N.Lock%) <> 0 \
```

```
    THEN CALL Error.Handler (14)
```

```
New.File% = -1
```

```
New.File% = opndat (New.File%, X.File%, X.File%, Fil.Name$, Rec.Len%)
```

```
IF errcod <> 0 \
```

```
    THEN CALL Error.Handler (15)
```

```
IF lokcod <> 0 \
```

```
    THEN CALL Temp.Lock
```

```
REM I/O buffer set up.
```

```
filler$ = "....+....1....+....2....+....3....+....4....+...." : \ 48 long
```

```
IO.Buffer$ = ""
```

```
No.Fillers% = INT% ( (Rec.Len% + 47) / 48)
```

```
FOR i% = 1 TO No.Fillers%
```

```

IO.Buffer$ = IO.Buffer$ + filler$
NEXT i%
Buffer.Ptr% = SADD (IO.Buffer$) + 2
REM "+2" because each string has a 2-byte length header.

REM Index file set up.

INPUT "Enter index name, key length, & type:"; Idx.Name$, \
    Key.Len%, Key.Type%
key% = -1
key% = opnidx (key%, Idx.Name$, Key.Len%, Key.Type%, 0)
IF errcod <> 0 \
    THEN CALL Error.Handler (4)
Idx.Entry$ = LEFT$ (filler$, Key.Len%)

REM Get the first key value.

drn% = frskey (key%, Old.File%, N.Lock%, Idx.Entry$)
drn2% = datval
IF errcod <> 0 \
    THEN CALL Error.Handler (5)

REM Loop over key values until index is exhausted,
REM writing data to new file in key value order.

WHILE drn% <> 0
    OR drn2% <> 0

    REM Read old data file.

    CALL setdat (drn2%)
    IF readat (Old.File%, drn%, Buffer.Ptr%) <> 0 \
        THEN CALL Error.Handler (6)

    REM Get next record in new file.

    Sort.Drn% = newrec (New.File%, N.Lock%)
    Sort.Drn2% = datval
    IF errcod <> 0 \
        THEN CALL Error.Handler (7)

    REM Write out new record (in key value order).

    CALL setdat (Sort.Drn2%)
    IF wrtdat (New.File%, Sort.Drn%, Buffer.Ptr%) <> 0 \
        THEN CALL Error.Handler (8)

    REM Update pointer in index file to reflect sorted order.

    CALL setdat (Sort.Drn2%)
    Ret.Code% = updptr (key%, New.File%, N.Lock%, Idx.Entry$, Sort.Drn%)
    IF errcod <> 0 \
        THEN CALL Error.Handler (9)

```

REM Get next key value.

```
target$ = LEFT$ (Idx.Entry$, Key.Len%)
drn% = aftkey (key%, Old.File%, N.Lock%, target$, Idx.Entry$)
drn2% = datval
IF errcod <> 0 \
    THEN CALL Error.Handler (10)
```

WEND

REM Close files.

```
IF clsdat (New.File%) <> 0 \
    THEN CALL Error.Handler (11)
IF clsdat (Old.File%) <> 0 \
    THEN CALL Error.Handler (12)
IF clsidx (key%) <> 0 \
    THEN CALL Error.Handler (13)
```

REM Sign-off message.

```
PRINT "Last record written: "; Sort.Drn2%; Sort.Drn%
STOP
```

Listing 3-1. AFTKEY function program code

BEFKEY function Index file search function

Syntax: BEFKEY (key%, file%, dlock%, keyval\$, idxval\$)

Explanation:

BEFKEY returns the data record number assigned to the index entry immediately preceding (that is, strictly less than) a key value you provide. BEFKEY also places the value of the key it finds in the IDXVAL\$ parameter.

Use BEFKEY to move backward through an index file sequentially in key-ascending order if either of the following situations apply:

- 1) your application operates in a multiuser environment
- 2) you want to update the index file while moving through it

If neither of the preceding situations apply, you will find it more efficient to use the PRVKEY function. Use BEFKEY to maintain compatibility between single-user and multiuser environments.

Parameters:

KEY%

The number of the index file where the search takes place.

FILE%

[MULTI] This parameter is ignored in a single-user environment. It is the number of the data file referenced by the KEY% parameter.

DLOCK%

[MULTI] This parameter is ignored in single-user environments. It specifies the type of data record lock requested for the data file referenced by the FILE% parameter. If the requested lock for the data record that BEFKEY found cannot be granted, LOKCOD returns a non-zero value. See the SETLOK function for a list of acceptable lock codes.

KEYVAL\$

The key value immediately following the one being sought. Note that KEYVAL\$ does not have to exist in the index for BEFKEY to function. KEYVAL\$ serves only as a reference value for seeking the actual index entry.

IDXVAL\$

This is an OUTPUT parameter. Access Manager places the key value found in the index file in IDXVAL\$ at the conclusion of the function. If no index entry is found with a key value less than KEYVAL\$, BEFKEY returns zero, and IDXVAL\$ is set to all blank spaces (which is not the same as a null string!...).

Before calling the BEFKEY function, IDXVAL\$ must be at least as long as the key length specified for the KEY% file. Normally, it is only necessary to initialize IDXVAL\$ once at the beginning of a program. Notice that KEYVAL\$ and IDXVAL\$ must not be the same variable.

Additional Comments:

To determine the two high-order bytes of the associated data record number, call the DATVAL function immediately after BEFKEY.

Error Codes:

Table 3-5 lists the BEFKEY function error codes.

Table 3-5. BEFKEY error codes

Value	Code	Explanation
30	AN	Index file number (KEY%) is out of range.
35	BC	IDXVAL\$ too short to contain key value.
46	BN	Index file is not open.
147	IC	Lock code (DLOCK%) is out of range.
153	II	Bad parameter value.

Example:

In this example, BEFKEY is used to determine the sequence number automatically appended to the key value by the ADDKEY function. This assumes the DUPKEY% parameter is set to one in the call to OPNIDX. Further assumptions are that this is a single-user environment, and the function SPACE\$ has already been defined and returns a string of blank spaces.

```
INPUT "Enter key value: "; keyval$
IF addkey (key%, 0, 0, keyval$, drn%) <> 1 \
  THEN CALL Error.Handler (1)

IF lokcod <> 0 \
  THEN CALL Error.Handler (2)

target$ = LEFT$ (keyval$ + space$ (keylen%), keylen%-2) + \
  CHR$ (0FFH) + CHR$ (0FFH)

idxval$ = space$ (keylen%)

IF befkey (key%, 0, 0, target$, idxval$) <> drn% \
  THEN CALL Error.Handler (3)

IF errcod <> 0 \
  THEN CALL Error.Handler (4)

Sequence.No$ = RIGHT$ (idxval$, 2)
```

The approach used in this example works, because the last key added to the index has the largest sequence number less than 0FFFFH.

CLSDAT function Data file setup and maintenance function

Syntax: CLSDAT (file%)

Explanation:

CLSDAT closes a data file. If the data file is closed successfully, CLSDAT returns a zero at the conclusion of the function; otherwise, a non-zero user code is returned.

Parameters:

FILE%

The number of the closed data file.

Additional Comments:

If you make any changes to a data file, your application program must call the CLSDAT or SAVDAT function to force the updates to the disk file. If you do not do this, the integrity of the data file can be disrupted, because the Access Manager header record could be incorrect. Further, the operating system might be holding data in internal buffers. CLSDAT and SAVDAT force the updated header record to the disk, and flush the operating system buffers.

[MULTI] If other programs are using the data file, CLSDAT actually performs a save operation. The data file is still available to the other users. Because Access Manager opens files in a locked mode, a file is not accessible to users outside of Access Manager, unless all programs that opened the file subsequently closed it.

If CLSDAT is called prior to program termination, you should also call the FRELOK function to release whatever type of data file lock you requested at the time the file was opened.

Error Codes:

If a data file is not properly closed or saved after it is updated, Access Manager issues user error 70/DF the next time that data file is opened. In such cases, the data file must be reconstructed.

Table 3-6 lists the CLSDAT function error codes.

Table 3-6. CLSDAT error codes

Value	Code	Explanation
55	CG	Data filename not in directory.
60	CL	Data file number (FILE%) is out of range.
74	DJ	Data file (FILE%) is inactive.
177	KA	Lock code (DLOCK%) is out of range.
183	KG	Bad parameter value.

Example:

```
IF clsdat (File.No%) <> 0 \
  THEN PRINT "Error while closing data file."
```

CLSIDX function Index file setup and maintenance function

Syntax: CLSIDX (key%)

Explanation:

CLSIDX closes an index file. If the index file is closed successfully, CLSIDX returns a zero at the conclusion of the function; otherwise, a non-zero user error code is returned.

Parameters:

KEY%

The number of the closed index file.

Additional Comments:

If you make any changes to an index file, your application program must call the CLSIDX or SAVIDX function to force the updates to the disk file. If you do not do this, the integrity of the index file can be disrupted, because some updated nodes might still be in an I/O buffer and/or the header record could be incorrect.

[MULTI] If other programs are still using the index file, CLSIDX actually performs a save operation. The index file is still available to the other users. Because Access Manager opens files in a locked mode, a file is not accessible to users outside of Access Manager, unless all programs that opened the file subsequently closed it.

Error Codes:

If an index file is not properly closed or saved after it is updated, Access Manager issues error 40/BH the next time that index file is opened. In such cases, the index file must be reconstructed.

Table 3-7 lists the CLSIDX function error codes.

Table 3-7. CLSIDX error codes

Value	Code	Explanation
----	----	-----
21	AE	Disk or directory full.
25	AI	Index filename not found in directory.
30	AN	Index file number (KEY%) is out of range.
44	BL	Index file (KEY%) is inactive.
46	BN	Index file (KEY%) is not open.
153	II	Bad parameter value.

Example:

IF clsidx (Key.No%) <> 0 \
THEN CALL Error.Handler

DATVAL function Data file setup and maintenance function

Syntax: DATVAL

Explanation:

DATVAL returns the high-order two bytes of the data record number. DATVAL can be called following any Access Manager function that returns a data record number as a function value (such as NEWREC) or a size attribute of the data file (GETDFS, GETDFU, or NOKEYS).

Parameters:

The DATVAL function is called without parameters.

Additional Comments:

If your application program does not require more than 65,535 data records or index file entries, there is no need to call DATVAL. However, there is no harm in doing so under these circumstances, because a value of zero is returned.

Error Codes:

The DATVAL function does not cause user errors.

Example:

Note that DATVAL is called immediately after the NEWREC function in the following example. DRN2% contains the most significant two bytes of the data record number, and DRN% contains the least significant two bytes.

```
dlock% = 2    REM Exclusive record lock  
drn% = newrec (File.No%, dlock%)  
drn2% = datval  
IF errcod <> 0 \  
THEN CALL Error.Handler (3)  
IF lokcod <> 0 \  

```


THEN CALL Lock.Conflict (3)

DELKEY function Index file update function

Syntax: DELKEY (key%, file%, dlock%, keyval\$, drn%)

Explanation:

DELKEY removes a key and its data record number from a specified index file. At the conclusion of the function, DELKEY returns a value indicating its success or failure (see "Additional Comments" below).

Parameters:

KEY%

The number of the index file where the key value and data record number are removed.

FILE%

[MULTI] This parameter is ignored in a single-user environment. It is the number of the data file referenced by the KEY% parameter.

DLOCK%

[MULTI] This parameter is ignored in single-user environments. It contains the code to specify the type of data record lock requested for the data file referenced by the FILE% parameter. See the SETLOK function for a list of acceptable lock codes.

To ignore locking protocols in a multiuser environment, assign DLOCK% a value of zero. This procedure is not normally recommended.

KEYVAL\$

The value of the key record to be deleted from the index file.

If you pass a null string in the KEYVAL\$ parameter, the DELKEY function takes no action but returns a value of one, unless the requested DLOCK% is denied.

This simplifies coding in multiple key applications. A loop designed to delete all key values for a given data record from their respective index files does not have to test for nonexistent keys, as long as they appear as null strings.

DRN%

The data record number associated with the key value contained in KEYVAL\$. Specifically, this is the number of the record to be deleted from the index file. Access Manager checks the data record number associated with KEYVAL\$ before making the deletion.

Additional Comments:

DELKEY returns one of the following values at its conclusion, to indicate its success or failure:

Table 3-8. DELKEY function values

Value	Meaning
-------	---------

- | ----- | ----- |
|-------|---|
| 0 | KEYVAL\$ was not found in the index. |
| 1 | KEYVAL\$ was successfully deleted from the index file. |
| 2 | The data record number in the index file containing KEYVAL\$ does not agree with the DRN% parameter. The index entry was not deleted. |
| 4 | The DLOCK% request on the data record given by DRN% for the data file specified by FILE% could not be granted. The index entry was not deleted. |

If an index file is opened with a DUPKEY% parameter of one (implying automatic suffixing of key values), DELKEY ignores the last two bytes of KEYVAL\$, and searches for an entry matching the first KEYLEN%-2 bytes, and for which the data record number equals DRN%. If a match is found, the entry is deleted from the index file. If the key value in KEYVAL\$ is less than KEYLEN% bytes in length, it is padded with blank spaces before the last two bytes are truncated.

Because Access Manager must search a set of duplicates for the one with a matching data record number, the number of disk accesses required to delete a duplicate key value can be excessive for a very large set of identical keys.

If you foresee large sets of duplicate key values (several hundred or thousands of identical values) and a regular need to delete members of these sets, you can avoid longer delete times by not using the Access Manager automatic duplicate feature. Instead, you should append your own unique suffix to these keys. For example, you might append a data record number to the end of the key value, or an associated unique key value. If you append your own suffix, set the DUPKEY% parameter in the OPNIDX function to zero. When you delete one of these key values, you can construct the exact composite key (identical portion plus unique suffix) to use in a regular deletion operation. No special Access Manager search of the duplicate set is required.

Error Codes:

The DELKEY function error codes are listed in Table 3-9.

Table 3-9. DELKEY error codes

Value	Code	Explanation
-------	------	-------------

-----	-----	-----
-------	-------	-------

30 AN Index file number (KEY%) is out of range.
46 BN Index file is not open.
147 IC Lock code (DLOCK%) is out of range.
153 II Bad parameter value.

Example:

In this example, SETDAT passes the two high-order bytes of the associated data record number.

```
CALL setdat (drn2%)  
IF delkey (key%, File.No%, lock%, Key.Value$, drn%) <> 1 \  
  THEN PRINT "Unsuccessful Deletion"
```

ERADAT function Data file setup and maintenance function

Syntax: ERADAT (file%, dlock%)

Explanation:

ERADAT erases a data file. Note that the data file must have been previously opened using either the OPNDAT or OPRDAT function. At the conclusion of the ERADAT function, a zero value is returned if the erasure was successful; otherwise, a non-zero user error code results.

Parameters:

FILE%

The number of the erased data file. If the data file is not open when ERADAT is called, an error results.

DLOCK%

[MULTI] Specifies the type of data file lock operation requested for the data file to be erased. If the value is set to zero when ERADAT is called, all data file lock operations are ignored. Use extreme care when calling ERADAT without an exclusive data file lock. See the SETLOK function for a list of acceptable lock codes.

If a file lock is requested and granted, the data file is erased. If the file lock request cannot be granted, the data file is not erased. Your application program must examine the contents of LOKCOD immediately after calling ERADAT, to determine if the lock request was granted.

Additional Comments:

Use the ERADAT function with care, to prevent accidental loss of a data file.

If your application program creates temporary data files, ERADAT can be used to delete them when necessary.

If your application program requires that a data file be new, successive calls to OPRDAT, ERADAT, and then OPNDAT ensure this.

Error Codes:

Access Manager returns a value in ERADAT to indicate the results of the operation. If no user errors occur, a zero is returned. The ERADAT function can cause any of the following user errors:

Table 3-10. ERADAT error codes

Value	Code	Explanation
60	CL	Data file number (FILE%) is out of range.
175	JO	The data file (FILE%) is inactive.
176	K@	File name not found in directory.
177	KA	Lock code (DLOCK%) is out of range.
183	KG	Bad parameter value.

Example:

In the following example, the data file is assumed to be open. The arguments of the ERROR.HANDLER and LOCK.CONFLICT functions indicate from where the functions were called.

```
dlock% = 4    REM Exclusive data file lock
IF eradat (File.No%, dlock%) <> 0 \
  THEN CALL Error.Handler (2)
IF lokcod <> 0 \
  THEN CALL Lock.Conflict (2)
```

ERAIDX function Index file setup and maintenance function

Syntax: ERAIDX (key%)

Explanation:

ERAIDX erases an index file. The index file must have been opened previously using the OPNIDX or OPRIDX function. ERAIDX returns a zero at the conclusion of the function if the erasure was successful; otherwise, a non-zero user

error code results.

Parameters:

KEY%

The number of the erased index file. If the index file is not open when ERAIDX is called, an error results.

Additional Comments:

Use the ERAIDX function with care to prevent accidental loss of an index file.

If your application program creates temporary index files, ERAIDX can be used to delete them when necessary.

If your application program requires that an index file be new, successive calls to OPRIDX, ERAIDX, and OPNIDX will initialize the file correctly.

Error Codes:

ERAIDX returns the error code as its function value. If ERAIDX returns zero, erasure of the index file is successful; otherwise, a user error has occurred. The ERAIDX function can cause the user errors listed in Table 3-11.

Table 3-11. ERAIDX error codes

Value	Code	Explanation
30	AN	Index file number (KEY%) is out of range.
46	BN	Index file is not open.
145	IB	The index file (KEY%) is inactive.
146	IC	Index file name not found in directory.
153	II	Bad parameter value.

ERRCOD function System initialization and maintenance function

Syntax: ERRCOD

Explanation:

ERRCOD returns the current value of the user error code.

ERRCOD returns a zero if the function completes execution without a user error occurring. On the other hand, if a user error DOES occur and the ERROPT%

parameter (see INTUSR function) was passed with a non-zero value, ERRCOD returns the appropriate error code value.

Parameters:

The ERRCOD function is called without parameters.

Additional Comments:

The value of ERRCOD does not change until a subsequent index or data file function is called. Consequently, if a function causes a user error, later calls to ERRCOD without intervening calls to other functions return the same error code value. ERRCOD clears to zero only when an Access Manager function executes without a user error.

Access Manager checks for two types of errors: user errors and failures of internal consistency. (Internal consistency errors are explained in Section 4.) When internal consistency errors occur, Access Manager cannot trap them, so it routes an error message to the console, and returns control to the operating system.

Error Codes:

Calls to ERRCOD do not cause user errors.

Example:

```
Ret.Code% = addkey (Key.No%, File.No%, Lock.Req%, Key.Value*, drn%)
IF errcod <> 0 \
  THEN CALL Error.Handler
IF Ret.Code% <> 1 \
  THEN CALL Add.Key.Problem
```

FRELOK function Data locking function

Syntax: FRELOK (file%, dlock%, drn%)

Explanation:

FRELOK releases (or, frees) an existing lock on a data file or data record. If the lock is released successfully, the FRELOK function returns a zero at its conclusion; otherwise, a one is returned. This function has no effect if

called in a single-user environment.

Parameters:

FILE%

The number of the data file from which you want to release the lock.

DLOCK%

Table 3-12 lists the codes you can use to request release of a data file lock. The appropriate code is passed via this parameter.

DRN%

The number of the data record number from which the lock is to be released.

Additional Comments:

FRELOK attempts to remove a specified type of lock held by the calling program (see PROGID% parameter under the INTUSR function). To release a lock, set the DLOCK% parameter to the appropriate value shown in Table 3-12. Note that, if the specified file or record does not have a lock set when FRELOK is called, no action takes place and LOKCOD is set to one.

Table 3-12. Data file/record lock release requests

DLOCK%

value Unlock request

- 0 No action, but LOKCOD is set to zero.
- 1 Release shared lock on DRN%.
- 2 Release exclusive lock on DRN%.
- 3 Release shared file lock on FILE%.
- 4 Release exclusive file lock on FILE%.
- 5 Release either shared or exclusive record lock on DRN%.
- 6 Release all locks held by the calling program (PROGID%) on the specified FILE%.
- 7 Release all locks held by the calling program (PROGID%) on all files with numbers greater than or equal to FILE%.

It is far more efficient to release each record lock after the record is processed than to release all record locks at the same time. While Access Manager has no practical limit on the number of record locks held simultaneously, a large number of locks forces some of the lock information out to disk. This results in slower processing.

Error Codes:

Table 3-13 lists the error codes for FRELOK function.

Table 3-13. FRELOK error codes

Value	Code	Explanation
52	CD	Data record number (DRN%) is zero or beyond end of file.
60	CL	Data file number (FILE%) is out of range.
177	KA	Lock code (DLOCK%) is out of range.
183	KG	Bad parameter value.

Example:

```

IF frelok (File.No%, 5, drn%) <> 0 \
  THEN PRINT "No record lock was held on "; drn% :\
  PRINT "by this user."

```

FRSKEY function Index file search function

Syntax: FRSKEY (key%, file%, dlock%, idxval\$)

Explanation:

The FRSKEY function locates the first entry in an index file. This function returns the data record number assigned to the first entry, and also returns its key value in the IDXVAL\$ parameter.

Parameters:

KEY%

The number of the index file where you want to search for the first index entry.

FILE%

[MULTI] This parameter is ignored in a single-user environment. It is the number of the data file referenced by the index file specified by the KEY% parameter.

DLOCK%

[MULTI] This parameter is ignored in single-user environments. It specifies the type of data record lock requested for the data file referenced by the FILE% parameter. If the requested lock for the data record that FRSKEY found cannot be granted, LOKCOD returns a non-zero value. See the SRTLOK function for a description of lock codes.

IDXVAL\$

This is an OUTPUT parameter. Access Manager places the key value found in the index file in IDXVAL\$. If the index is empty, FRSKEY returns a zero, and

IDXVAL\$ is set to all blanks (which is not the same as a null string).

Before calling the FRSKEY function, IDXVAL\$ must be at least as long as the key length specified for the KEY% file, or user error 35/BC results.

Additional Comments:

To determine the two high-order bytes of the associated data record number, call the DATVAL function immediately after FRSKEY.

Error Codes:

The FRSKEY function can cause the error codes listed in Table 3-14.

Table 3-14. FRSKEY error codes

Value	Code	Explanation
30	AN	Index file number (KEY%) is out of range.
35	BC	IDXVAL\$ too short to contain key value.
46	BN	Index file is not open.
147	IC	Lock code (DLOCK%) is out of range.
153	II	Bad parameter value.

Example:

```
Rec.Lock% = 2      REM Exclusive record lock request
drn% = frskey (Key.No%, file%, Rec.Lok%, idxval$)
Seg.No% = datval
IF errcod <> 0 \
  THEN CALL Error.Handler
IF lokcod <> 0 \
  THEN CALL Lock.Conflict
IF drn% = 0 \
AND Seg.No% = 0 \
  THEN PRINT "Index file is empty."
```

Error Codes:

The FRSKEY function can cause the error codes listed in Table 3-14.

Table 3-14. FRSKEY error codes

Value	Code	Explanation
30	AN	Index file number (KEY%) is out of range.

35 BC IDXVAL\$ too short to contain key value.
46 BN Index file is not open.
147 IC Lock code (DLOCK%) is out of range.
153 II Bad parameter value.

Example:

```
Rec.Lok% = 2      REM Exclusive record lock request
drn% = frskey (Key.No%, file%, Rec.Lok%, idxval$)
Seg.No% = datval
IF errcod <> 0 \
  THEN CALL Error.Handler
IF lokcod <> 0 \
  THEN CALL Lock.Conflict
IF drn% = 0 \
AND Seg.No% = 0 \
  THEN PRINT "Index file is empty."
```

GETDFS function Datafile setup and maintenance function

Syntax: GETDFS (file%)

Explanation:

GETDFS returns a count of the records in a data file. The count includes the number of records in current use, and those available for use.

Parameters:

FILE%

The number of the data file where the record count takes place.

Additional Comments:

The record count returned by the GETDFS function includes the header record and any other records automatically reserved by the NEWREC function. This ensures that the first available data record begins on or after the 129th byte of the data file, as required by Access Manager.

To determine the two high-order bytes of the logical data file size, call the DATVAL function immediately after GETDFS. Also, see the GETDFU function.

Error Codes:

Table 3-15 lists the GETDFS error codes.

Table 3-15. GETDFS error codes

Value	Code	Explanation
60	CL	Data file number (FILE%) is out of range.
183	KG	Bad parameter value.

Example:

If the file is certain to contain less than 65,535 records, the actual number of kilobytes used by the data file can be computed as follows:

```
IF getdfs (File.No%) < 0 \
  THEN totrec = getdfs (File.No%) + 65536. \
  ELSE totrec = getdfs (File.No%)
kilobyte% = INT% ( (totrec * Rec.Len% + 1023.) / 1024.)
```

where REC.LEN% is the record length of the file specified by FILE.NO%. Note how the signed integer quantity is converted to a positive real quantity.

If the file might contain more than 65,535 records, the number of kilobytes is given by:

```
IF getdfs (File.No%) < 0 \
  THEN totrec = getdfs (File.No%) + 65536. \
  ELSE totrec = getdfs (File.No%)
totrec = totrec + (65536. * datval)
kilobyte% = INT% ( (totrec * Rec.Len% + 1023.) / 1024.)
```

In this case, DATVAL returns the two high-order bytes of the data file size. If the size is less than 65,536, DATVAL returns zero.

GETDFU function Data file setup and maintenance function

Syntax: GETDFU (file%)

Explanation:

GETDFU returns a count of records in use in a data file.

The difference between the counts returned by GETDFS and GETDFU represents the data records returned for reuse, plus the header record, plus (for record lengths less than 128 bytes) records that share the first 128 bytes with the header record.

Parameters:

FILE%

The number of the data file where the record count takes place.

Additional Comments:

To determine the two high-order bytes of the in-use record count, call the DATVAL function immediately after GETDFU.

Error Codes:

Table 3-16 lists the GETDFU error codes.

Table 3-16. GETDFU error codes

Value	Code	Explanation
60	CL	Data file number (FILE%) is out of range.
183	KG	Bad parameter value.

Example: GETDFU and GETDFS can be used together to compute the fraction of data records in active use, as shown in this example:

```
IF getdfu (File.No%) < 0 \
  THEN Active.Records = getdfu (File.No%) + 65536. \
  ELSE Active.Records = getdfu (File.No%)
Active.Records = Active.Records + (65536. * datval)
```

```
IF getdfs (File.No%) < 0 \
  THEN Total.Records = getdfs (File.No%) + 65536. \
  ELSE Total.Records = getdfs (File.No%)
Total.Records = Total.Records + (65536. * datval)
```

utilization = Active.Records / Total.Records

GETKEY function Index file search function

Syntax: GETKEY (key%, file%, dlock%, keyval\$)

Explanation:

GETKEY returns the data record number associated with a specific key value in an index file. GETKEY locates the entry in the index file that exactly matches a key value you provide in the KEYVAL\$ parameter.

Parameters:

KEY%

The number of the index file where the key value/data record number can be located.

FILE%

[MULTI] This parameter is ignored in a single-user environment. It is the number of the data file referenced by the KEY% parameter.

DLOCK%

[MULTI] This parameter is ignored in single-user environments. It specifies the type of data record lock requested for the data file referenced by the FILE% parameter. If the requested lock for the data record GETKEY found cannot be granted, LOKCOD returns a non-zero value. See the SETLOK function for a description of lock codes.

KEYVAL\$

The actual key value of the index file record being sought.

Additional Comments:

If your application program takes advantage of the four-byte capacity of the associated data record numbers, your program must call the DATVAL function immediately after GETKEY, to return the value of the two high-order bytes. GETKEY returns the value of the two low-order bytes.

Error Codes:

Table 3-17 lists the error codes for the GETKEY function.

Table 3-17. GETKEY error codes

Value	Code	Explanation
30	AN	Index file number (KEY%) is out of range.
46	BN	Index file is not open.
147	IC	Lock code (DLOCK%) is out of range.
153	II	Bad parameter value.

Example:

In this example, notice GETKEY only finds exact matches. Also, CONV.INT converts the part numbers to numeric information.

```
DEF Conv.Int (intval%)
  STRING Conv.Int
  INTEGER msb, lsb

  msb = INT% (intval% / 256)
  lsb = intval% - msb * 256
  IF lsb < 0 THEN msb = msb - 1
  Conv.Int = CHR$ (lsb) + CHR$ (msb)
FEND

INPUT "Enter Desired Part #:"; Part.No%
Part.No$ = Conv.Int (Part.No%)
lock% = 1    REM Shared data record lock
pointer% = getkey (Part.Key%, Inv.Fil%, lock%, Part.No$)
segment% = datval
IF errcod <> 0 \
  THEN CALL exception (2, 1)
IF pointer% = 0 \
AND segment% = 0 \
  THEN PRINT "Requested Part Number not in data base."
IF Lock.Code% <> 0 \
  THEN PRINT "Shared Record Lock not granted."
```

INTUSR function System initialization and maintenance function

Syntax: INTUSR (progid%, erropt%, timeout%)

Explanation:

INTUSR initializes Access Manager, and must be called by your application program prior to calling any other functions.

Parameters:

PROGID%
[MULTI] This parameter is ignored in a single-user environment. It should contain the identification number assigned to your application program. Each application program calling the background server must have a unique PROGID% number. PROGID% numbers begin at zero. For example, in a three-user environment, the acceptable values for PROGID% are 0, 1, and 2. (See your Programmer's Guide for additional information on system configuration.)

To simplify program development, if PROGID% is set to -1, INTUSR returns the number of the console requesting the application program. This makes it possible to run an application program from different consoles without any

changes to set up a unique program ID, or having to predetermine the PROGID% values.

ERROPT%

A value to indicate how you want Access Manager to handle user errors when they occur.

Set ERROPT% to zero if you want Access Manager to display a two-character error message on the console, and then return control to the operating system.

Set ERROPT% to a non-zero value if you want Access Manager to trap user errors, and return them to your program for processing. In this case, the ERRCOD function detects user errors, making it possible for your application program to determine the reason for the error, and take appropriate action. The value of ERRCOD should be examined following each call to an Access Manager function. Section 4 contains a list of possible ERRCOD values.

TIMOUT%

[MULTI] This parameter is ignored in a single-user environment. Your program can use it to specify the number of seconds within which the background server must respond to the INTUSR call. If the background server fails to respond within this timeframe, the following message appears on the console:

Be sure that ACCESS MANAGER is operational. Run MPMSTAT to see.

Control then returns to your operating system.

If TIMOUT% is zero, no response time check is made. Therefore, if Access Manager is not running, the application program hangs indefinitely at the INTUSR call. A reasonable value for TIMOUT% is three seconds. The time check starts only after the call to INTUSR is sent to the background server. If the call must wait in a queue, the time check is not affected.

Additional Comments:

[SINGLE] Must be called at least once, but can be called as often as necessary to change the error handling technique.

[MULTI] Subsequent calls of the INTUSR function clear the message queue, and remove all pending locks associated with PROGID%. This can produce unpredictable results, and should be avoided.

Error Codes:

[SINGLE] Calls to the INTUSR function should never cause user errors.

[MULTI] User errors might result if PROGID% is outside the allowed range of the background server, or the server is not active when INTUSR is called. The possible errors are listed in Table 3-18.

Table 3-18. INTUSR error codes

Value	Code	Explanation
161	JA	Program number (PROGID%) is out of range.
162	JB	Cannot open Access Manager input queue.
163	JC	Cannot open output queue.

Example:

```
-----
progid% = intusr (-1, 1, 3)
IF errcod <> 0 \
  THEN PRINT "Cannot Initialize User...Error: "; errcod
IF setup (7, 3, 4, 1) <> 0 \
  THEN PRINT "Illegal SETUP parameters."
```

[SINGLE] The above statements imply there are seven buffers, up to three index files can be open at one time, the index file record length is 512 bytes (four 128-byte sectors), and only one data file is opened at a time. Further, it is implied that user errors are trapped by calling the ERRCOD function following all calls to Access Manager functions, although calls to INTUSR should never result in user errors.

[MULTI] Note the call to SETUP is ignored. An application program cannot affect the SETUP parameters. These are established at the time the background server is configured.

Error trapping is enabled by the non-zero ERROPT% parameter. Because the PROGID% parameter is passed as -1, INTUSR returns the value corresponding to the console number from which the application program is run. Therefore, the program uses PROGID% to identify the console. More importantly, the same application program can be executed from different consoles without predetermined PROGID% values.

If the assigned PROGID% is beyond the permissible range, INTUSR returns a user error code. But, if error trapping is not enabled, the following message is displayed:

Is PROG ID out of range?

If the background server has not been initialized, the TIMEOUT% value of three causes the "...Run MPMSTAT..." message to display (see TIMEOUT% parameter description).

If an Access Manager Resident System Process is required but not included at GENSY, INTUSR returns a user error code. If error trapping is not enabled, the following message is displayed:

Has an Access Manager RSP been included at GENSY?

LASKEY function Index file search function

Syntax: LASKEY (key%, file%, dlock%, idxval\$)

Explanation:

LASKEY returns the data record number assigned to the last entry in an index file. LASKEY also places the value of the last key it locates in the IDXVAL\$ parameter.

Parameters:

KEY%

The number of the index file in which you want to find the last key.

FILE%

[MULTI] This parameter is ignored in a single-user environment. It is the number of the data file referenced by the KEY% parameter.

DLOCK%

[MULTI] This parameter is ignored in single-user environments. It specifies the type of data record lock requested for the data file referenced by the FILE% parameter. If the requested lock for the data record LASKEY found cannot be granted, LOKCOD returns a non-zero value. See the SETLOK function for a description of lock codes.

IDXVAL\$

This is an OUTPUT parameter. Access Manager places the last key value found in the index file in IDXVAL\$ at the conclusion of the function. If the index is empty, LASKEY returns a zero, and IDXVAL\$ is set to all blank spaces (which is not the same as a null string!...).

Before calling the LASKEY function, IDXVAL\$ must be at least as long as the key length specified for the KEY% file, or user error 35/BC results. (See KEYLEN% parameter under OPNIDX function.)

Additional Comments:

To determine the two high-order bytes of the associated data record number, call the DATVAL function immediately after LASKEY.

Error Codes:

Table 3-19 lists the error codes for the LASKEY function.

Table 3-19. LASKEY error codes

Value	Code	Explanation
30	AN	Index file number (KEY%) is out of range.
35	BC	IDXVAL\$ too short to contain key value.
46	BN	Index file is not open.
147	IC	Lock code (DLOCK%) is out of range.
153	II	Bad parameter value.

Example:

The following example prints data record numbers in reverse key order.

```

Rec.Lock% = 1      REM Shared record lock

drn% = laskey (Key.No%, File.No%, Rec.Lok%, Last.Key$)
WHILE drn% <> 0 \
  AND errcod <> 0 \
  AND lokcod <> 0
  PRINT Last.Key$, drn%
  target$ = LEFT$ (Last.Key$, Key.Len%)
  drn% = befkey (Key.No%, File.No%, Rec.Lok%, target$, Last.Key$)
WEND

IF errcod <> 0 \
  THEN CALL Error.Handler
IF lokcod <> 0 \
  THEN CALL Lock.Conflict
PRINT "The End"

```

LOKCOD function System initialization and maintenance function

Syntax: LOKCOD

Explanation:

LOKCOD returns a value indicating whether or not a request to lock a data file or data record was granted.

[SINGLE] Because data file and record locks are unnecessary in a single-user environment, all calls to LOKCOD return zero, indicating a successful lock operation.

When a lock/unlock request is successful, LOKCOD returns a value of zero; otherwise, it returns a non-zero value.

Parameters:

The LOKCOD function is called without parameters.

Additional Comments:

When a data record lock is not granted, LOKCOD returns a one to signify a conflict with another record lock, or a two to signify the data file (not just an individual record) is locked exclusively.

When a request for an exclusive data file lock is not granted, LOKCOD returns a two if another user already holds the exclusive file lock, or a one if another user holds a data record lock in the specified data file.

A request to unlock a data file or individual data record fails only if the expected lock is not found. In this case, LOKCOD returns a one.

Error Codes:

Calls to LOKCOD do not cause user errors.

NEWREC function Data file update function

Syntax: NEWREC (file%, dlock%)

Explanation:

NEWREC returns the number of the next available record in a data file.

Parameters:

FILE%

The number of the data file where you want to locate the next available data record.

DLOCK%

[MULTI] Specifies the type of data record lock operation requested for the data file to be accessed.

If another user holds an exclusive file lock on the data file specified by the FILE% parameter, LOKCOD returns a two, and the value returned by NEWREC has no meaning. If you have already opened a data file with either a shared or exclusive file lock, another program cannot get an exclusive file lock. Thus, once you open a data file, you should not run into a lock conflict when

calling NEWREC. LOKCOD should never equal one following a call to NEWREC.

Even though you do not normally need to test the LOKCOD function following a call to NEWREC, it can be an aid to debugging your program. See the SETLOK function for a description of lock codes.

Additional Comments:

NEWREC automatically reclaims previously deleted data records (see RETREC function) before extending the size of the data file. That is, if any deleted data records are available, NEWREC uses them first. If not, NEWREC increases the size of the data file to generate a new data record.

To determine the two high-order bytes of the next available data record number, call the DATVAL function immediately after NEWREC.

NEWREC must be called each time a new data record is required. If it is not, Access Manager does not track the actual size of the data file. The next time you attempt to use the RETREC, READAT, or WRDAT function, user error 52/CD might occur (indicating the data record number is beyond the end of the file).

The first available data record number NEWREC returns for a newly created data file is the first data record beginning on or after the 129th byte of the data file. For example, if the data file record length is greater than or equal to 128 bytes, the first available data record number is two. If the record length is less than 128, the first available record number ranges from three to thirty-three, depending on the record length. (For further information, refer to the OPNDAT function.)

Error Codes:

Table 3-20 lists the error codes for the NEWREC function.

Table 3-20. NEWREC error codes

Value	Code	Explanation
60	CL	Data file number (FILE%) is out of range.
69	DE	First byte of deleted data record not 0FFH.
177	KA	Lock code (DLOCK%) is out of range
183	KG	Bad parameter value.

Example:

FOR this example, assume the data file has been opened successfully.

```
dlock% = 2    REM Exclusive record lock request
drn% = newrec (File.No%, dlock%)
```

```
IF errcod <> 0 \
  THEN CALL Error.Handler (3)
IF lokcod <> 0 \
  THEN CALL Lock.Conflict (3)
```

NMNODS function Index file setup and maintenance function

Syntax: NMNODS (key%)

Explanation:

NMNODS returns a count of the number of used and usable records (that is, B-Tree nodes) in an index file. The number will not be more than two bytes in length.

Parameters:

KEY%

The number of the index file where the used and usable records are counted.

Additional Comments:

The count returned by the NMNODS function does not include the index file header record.

The NMNODS and GETDFS functions provide a way to compute how much disk space your Access Manager files occupy. (The example at the end of this section shows how to do this for an index file.)

Error Codes:

The NMNODS function causes these error codes:

Table 3-21. NMNODS error codes

Value	Code	Explanation
30	AN	Index file number (KEY%) is out of range.
46	BN	Index file is not open.
153	II	Bad parameter value.

Example:

This example shows how to compute the number of kilobytes occupied by an index file. Note that one is added to TOTAL.NODES to account for the index file header record.

```
Total.Nodes = nmnodes (Part.Key%)
IF Total.Nodes < 0 \
  THEN Total.Nodes = Total.Nodes + 65536.
kilobytes% = INT ( (Total.Nodes + 1.) * (No.Nodes.Sectors + 7) / 8)
```

NOKEYS function Index file setup and maintenance function

Syntax: NOKEYS (key%)

Explanation:

NOKEYS returns a count of the number of entries in an index file.

Parameters:

KEY%

The number of the index file where the entries are counted.

Additional Comments:

NOKEYS returns the two low-order bytes of the number of entries in the index file specified by KEY%. If your application program requires the two high-order bytes, call the DATVAL function immediately after NOKEYS.

Error Codes:

Table 3-22 lists the error codes for the NOKEYS function.

Table 3-22. NOKEYS error codes

Value	Code	Explanation
30	AN	Index file number (KEY%) is out of range.
46	BN	Index file is not open.
153	II	Bad parameter value.

Example:

Assume you have an inventory index file with more than 65,535 entries. This example might be used to count the number of entries or "parts" in that index file.

```
No.Of.Parts = nokeys (Part.Key%)  
IF No.Of.Parts < 0.0 \  
  THEN No.Of.Parts = No.Of.Parts + 65536.  
No.OF.Parts = No.Of.Parts + (65536. * datval)
```

NXTKEY function Index file search function

Syntax: NXTKEY (key%, file%, dlock%, idxval\$)

Explanation:

NXTKEY returns the data record number associated with the next entry in an index file. NXTKEY also places the key value it finds in the IDXVAL\$ parameter.

[SINGLE] NXTKEY is an efficient way to move through an index file sequentially when no updates are performed. You can interleave calls to NXTKEY for different KEY%, because separate position pointers are maintained for each index file.

[MULTI] NXTKEY is not normally used in the multiuser environment.

Parameters:

KEY%

The number of the index file where the search takes place.

FILE%

[MULTI] This parameter is ignored in a single-user environment. It is the number of the data file referenced by the KEY% parameter.

DLOCK%

[MULTI] This parameter is ignored in single-user environments. It specifies the type of data record lock requested for the data file referenced by the FILE% parameter. If the requested lock for the data record that NXTKEY found cannot be granted, LOKCOD returns a non-zero value. See the SETLOK function for a description of lock codes.

IDXVAL\$

This is an OUTPUT parameter. Access Manager places the key value found in the index file in IDXVAL\$. If no next index entry is found, NXTKEY returns a zero, and IDXVAL\$ is set to all blank spaces (which is not the same as a null string!...).

Before calling the NXTKEY function, IDXVAL\$ must be at least as long as the key length specified for the KEY% file. Normally, it is only necessary to initialize IDXVAL\$ once at the beginning of a program.

Additional Comments:

To determine the two high-order bytes of the associated data record number, call the DATVAL function immediately after NXTKEY.

Before the very first call to NXTKEY for a given KEY%, or after the index file associated with KEY% is updated, one of the other index search functions (except PRVKEY) must be called for the same KEY%. The other search functions establish the internal pointer used by NXTKEY and PRVKEY for sequential processing. Each time any of the search functions are called, including NXTKEY and PRVKEY, the internal pointer is appropriately updated. The internal pointers are not maintained when the index file is updated.

Error Codes:

Table 3-23 lists the NXTKEY function error codes.

Table 3-23. NXTKEY error codes

Value	Code	Explanation
30	AN	Index file number (KEY%) is out of range.
35	BC	IDXVAL\$ too short to contain key value.
46	BN	Index file is not open.
147	IC	Lock code (DLOCK%) is out of range.
153	II	Bad parameter value.

Example:

In the following example, SERKEY finds the first potential match for a customer name, and NXTKEY moves through the index file sequentially, to scan the remaining potential matches. This example is only for single-user environments, because NXTKEY is used.

The example listing assumes customer names are a maximum of sixteen bytes, and the names entered into the index file are truncated to nine bytes, and suffixed with a two-byte, unique sequence number by Access Manager. The data file record length is 64 bytes.

DEF space\$ (length)
 INTEGER length
 STRING space\$


```

REM .....1....+....2....+....3....+....4....+... (48 long)
space$ = LEFT$ ( \
"          " \
, length)
RETURN
FEND

```

```

dlock% = 0
IO.Buffer$ = space$ (48) + space$ (16)
Buffer.Ptr% = SADD (IO.Buffer$) + 2
entry$ = space$ (11)

```

```

INPUT "Enter Customer Last Name: "; Cust.Name$
Cust.Name$ = LEFT$ (Cust.Name$ + space$ (16), 16)

```

```

REM TARGET is padded with CHR$ (0), so Access Manager does not
REM pad with blank spaces. Note these last two bytes can
REM have any binary value, because they serve simply as
REM tie-breakers. If Access Manager were allowed to pad with
REM spaces (20H), some key values which matched in the
REM significant bytes (for example, the first nine) might
REM be skipped when SERKEY is called, because their last
REM two bytes were less than 2020H.

```

```

target$ = LEFT$ (Cust.Name$, 9) + CHR$ (0) + CHR$ (0)

```

```

drn% = serkey (Cust.Key%, file%, dlock%, target$, entry$)
WHILE LEFT$ (target$, 9) = LEFT$ (entry$, 9)
AND drn% <> 0
IF readat (file%, drn%, Buffer.Ptr%) <> 0 \
THEN PRINT "Read Error "; errcod :\
STOP

```

```

REM Assume Customer Name occupies the first sixteen bytes
REM of the data record.

```

```

IF LEFT$ (IO.Buffer$, 16) = Cust.Name$ \
THEN PRINT drn%; IO.Buffer$
drn% = nxtkey (Cust.Key%, file%, dlock%, entry$)
WEND

```

```

PRINT "Search for Customer Name ended."

```

Listing 3-2. NXTKEY function program code

OPNDAT function Data file setup and maintenance function

Syntax: OPNDAT (file%, dlock%, filename\$, reclen%)

Explanation:

OPNDAT opens a data file. A data file must be opened before it can be read, written, or erased.

At the conclusion of the function, OPNDAT returns the number of the opened data file.

Parameters:

FILE%

The number of the data file to be opened. The file number can be an integer between zero and NDATF%-1 (NDATF% is a parameter of the SETUP function. The maximum value of NDATF% depends on your operating system, and whether yours is a single-user or multiuser environment. Consult your Programmer's Guide.

Note that file numbers assigned by Access Manager are separate from those used with your application language. Assigning a particular number to an Access Manager data file does not preclude your using the same file number with an application language file.

If the FILE% parameter contains a value of -1, OPNDAT returns an integer value as follows:

- If an exact match for FILNAME\$ (see below) is found among the currently opened data files, OPNDAT returns the number assigned to that file.
- If no match exists for FILNAME\$, OPNDAT returns the first file number not in use.
- If no file number is available, OPNDAT causes user error 60/CL, indicating the file number is out of range.

[MULTI] Passing a FILE% parameter of -1 to OPNDAT allows different applications, or different users of the same application, to share the same data files without prior knowledge of file number assignments. If a file is already open and your application program does not use a FILE% parameter of -1, Access Manager returns user error 63/CO (indicating the file is already in use), or causes a lock conflict.

DLOCK%

[MULTI] Specifies the type of lock requested for the data file to be opened.

Set DLOCK% to three to request a shared data file lock. If granted, this lock stops any other user from gaining an exclusive file lock. Any number of users can hold shared file locks simultaneously.

Set DLOCK% to four to request an exclusive file lock. The lock is only granted if no other locks (file or record) are held on the data file by other programs.

If you hold a shared lock on a data file, you can only change to an exclusive

file or record lock if no other users hold shared locks at the time. Use the SETLOK function to attempt such a change.

DLOCK% values of one or two should not be used with the OPNDAT function, because they refer to individual record locks.

FILNAME\$

A character string to indicate the name of the data file you want to open. Access Manager always converts the string value to upper-case letters. The data filename can include a password. A password is designated by a file specification (which includes a semicolon (";")) followed by the password.

RECLEN%

A value to specify the length of the individual records in the data file. Records must be a minimum of four bytes in length.

Additional Comments:

You cannot open a corrupted data file with the OPNDAT function, user error 53/CE results. See the OPNDAT function for an explanation of opening corrupted data files.

Access Manager will assign a password to a data file created with the OPNDAT function if a password is part of the FILNAME\$ parameter (for example, B:CUSTOMER.DAT;SECRET where the password is SECRET), and the drive on which the file is created has been set for automatic XFCB (Extended File Control Block) creation. When both conditions are satisfied, the files created by OPNDAT are protected at the highest level; namely, READ protected (the password is required, even to read the file). See your operating system documentation to determine if the SET command is available, to enable automatic XFCB creation.

[MULTI] Even if automatic XFCB creation is not enabled in a multiuser environment, or if the file was not created with a password, the Access Manager background server will maintain temporary passwords, as long as the data file is open. However, the password will not be permanently recorded with the data file.

The first 128 bytes of a data file are reserved for status information. No other data can be stored in this area. Therefore, the first data record available to an application program use is the one beginning at or beyond the 129th byte of the data file. The first available data record number in a file can be determined by the following expression:

$$\text{INT} ((128 + \text{reclen}\% - 1) / \text{reclen}\%) + 1$$

where INT truncates its argument. For example, if the record length is 32 bytes, the first record available for data is number 5. Figure 3-2 shows how the beginning of a data file is organized when the record length is 32 bytes.

```
+-----+ +-----...  
| Reserved for data file || Available records -->
```

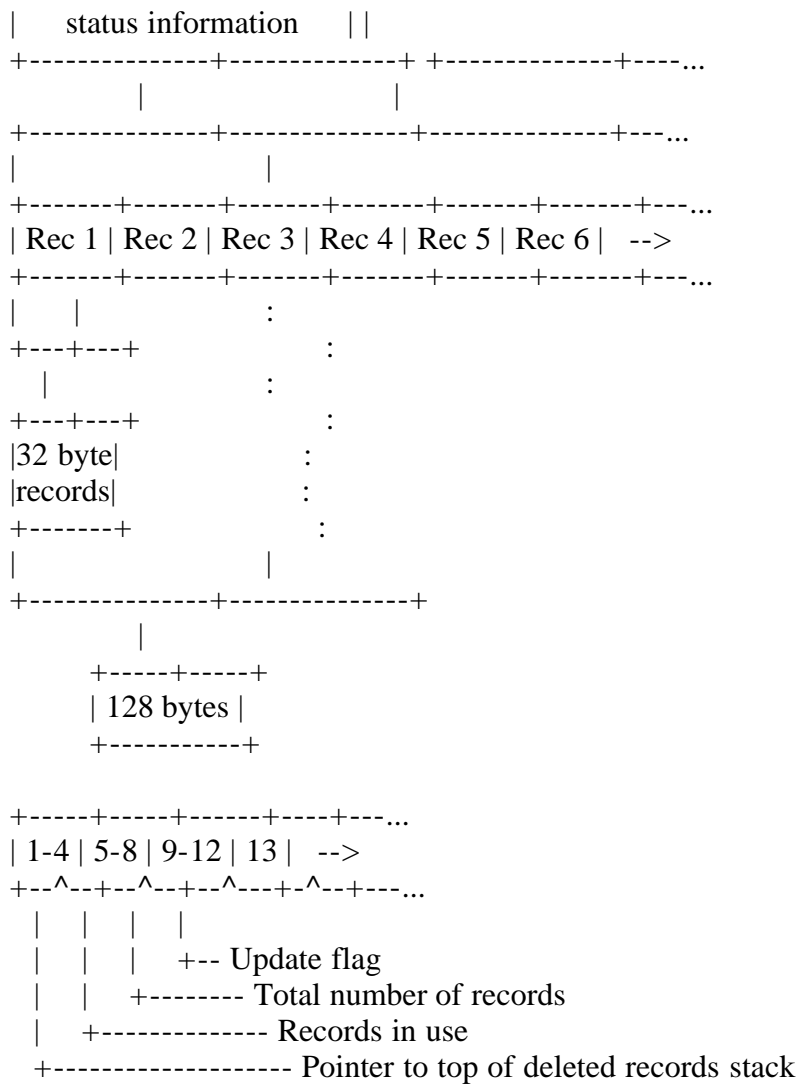


Figure 3-2. Sample data file layout for 32-byte records

Note that the NEWREC function automatically determines the number of the first usable data record; your application program need not compute this number.

Error Codes:

OPNDAT causes the following user errors:

Table 3-24. OPNDAT error codes

Value	Code	Explanation
53	CE	Data file is corrupt; no header record.
54	CD	No more directory space.
60	CL	Data file number (FILE%) is out of range.
61	CM	Data filename (FILNAME\$) incorrectly formed
63	CO	Data file (FILE%) is already in use.
70	DF	The data file is corrupt.
119	GG	Incorrect or missing password.
177	KA	Lock code (DLOCK%) is out of range.

183 KG Bad parameter value.

Example:

```
File.No% = -1      REM Automatic file number assignment
dlock% = 3        REM Shared file lock
Record.Len% = 100
File.Name$ = "E:CUSTOMER.DAT"
File.No% = opndat (File.No%, dlock%, File.Name$, Record.Len%)
IF errcod <> 0 \
  THEN CALL Error.Handler (1)
IF lokcod <> 0 \
  THEN CALL Lock.Conflict (1)
```

Note that the OPNDAT function assigns FILE.NO% a value. If a user error occurs when opening the data file, the program calls an error handling subroutine with a parameter indicating where the error occurred. The error handling routine can then call the ERRCOD function to get the actual error code value. Likewise, if Access Manager refuses the requested data file lock, control transfers to the lock conflict handler.

OPNIDX function Index file setup and maintenance function

Syntax: OPNIDX (key%, idxname\$, keylen%, keytyp%, dupkey%)

Explanation:

OPNIDX opens or creates an index file and, optionally, assigns the file number. When a file number is assigned, all subsequent references to the index file are made using this number.

Parameters:

KEY%

The number of the index file to be opened or created. The number must range from zero to NKEYS%-1 (see the SETUP function for an explanation of NKEYS%).

If KEY% passes -1 to OPNIDX, an integer value is returned as follows:

- If an exact match for the IDXNAME\$ parameter (see below) is found among the currently open index files, OPNIDX returns the index file number.
- If no match is found, OPNIDX returns the first unused index file number.
- If no index file number is available, user error 30/AN is returned,

indicating KEY% is out of range.

Index file numbers are independent of data file numbers assigned by Access Manager or the application language.

[MULTI] The ability to automatically assign KEY% parameters means application programs running simultaneously can share index files without any concern about how KEY% numbers are assigned.

IDXNAME\$

A character string to indicate the name of the index file you want to open or create. Access Manager converts the string to upper-case letters. The index filename might include a password. A password is designated by a file specification (which includes a semicolon (";")) followed by the password.

KEYLEN%

A value specifying the maximum length (in bytes) of the key values in this index file. The value must be at least one, and not greater than forty-eight. Integer keys must be at least two bytes in length.

KEYTYP%

A value indicating whether the keys in this index file are stored in alphanumericly increasing order or numeric order. Zero indicates alphanumeric keys; one indicates signed, integer keys. See the description of the ADDKEY function for additional information on coding integer key values.

All key values are passed to and from Access Manager as string-valued quantities, even if the key is designated as numeric. The KEYTYP% parameter affects only the order in which keys are stored.

Numeric key values are integers (not necessarily restricted to two bytes) stored with the least-significant byte first and the most-significant, including the sign of the integer, last. Negative integers are stored in two's complement form.

DUPKEY%

A value indicating whether or not duplicate key values should be handled by Access Manager. If DUPKEY% is one and KEYTYP% is zero (alphanumeric key values), Access Manager automatically assigns sequence numbers to the last two bytes of each key value. For example, if KEYLEN% is ten, bytes nine and ten of each key value are filled (or overwritten) by Access Manager with a sequence number guaranteed to be unique. These sequence numbers are in binary form, and usually do not represent valid ASCII symbols.

DUPKEY% has no effect if KEYTYP% is other than zero.

Note: Please read the DELKEY function description for important information concerning the adverse effect of automatic duplicate keys on the time required to delete them.

Additional Comments:

A corrupted index file cannot be opened using the OPNIDX function. Refer to the OPRIDX function for an explanation of opening corrupted index file.

Access Manager assigns a password to an index file created by OPNIDX if a password is part of the IDXNAME\$ parameter (for example, CUSTOMER.DAT;SCRAMBLE where the password is SCRAMBLE), and the drive on which the file is created is set for automatic XFCB (Extended File Control Block) creation. When both conditions are satisfied, the files created by OPNIDX are protected at the highest level; namely, READ protected (the password is required, even to read the file). See your operating system documentation to determine if the SET command is available to enable automatic XFCB creation.

[MULTI] Even if automatic XFCB creation is not enabled in a multiuser environment, or if the file was not created with a password, the Access Manager background server will maintain temporary passwords, as long as the index file is open. However, the password must be permanently recorded with the index file.

Based on the index file record length (refer to the NNSEC% parameter under the SETUP function), Access Manager automatically determines the maximum number of key values stored in each index file record. As the number of key values per index file record increases, the number of levels in the index structure decreases. The number of levels in the index structure determines the maximum number of disk accesses required to locate a key value. Of course, if an index file record is in an Access Manager I/O buffer, no disk access is required. Also, if the index file record length is greater than the physical sector size of the disk, each logical access can require two or more disk accesses.

The maximum key values per index file record is computed as the largest integer less than or equal to

$$\left((nnsec\% * 128) - 10 \right) / (keylen\% + 4)$$

Further, the maximum number must be at least four, or user error 39/BG occurs. Access Manager restricts the maximum number of key values per index file record. These restrictions are stated in your Programmer's Guide.

Once you choose values for NNSEC% and KEYLEN%, you can compute MAXKV (the maximum number of key values per index file record). For a computed MAXKV, the minimum possible index file size (in bytes) for a given number of index entries is computed in Pascal/MT+ as

```
nodes = <entries / maxkv>    [where <X> is the smallest
                             integer greater than or equal to X]
index_nodes = nodes
WHILE index_nodes > 1
    index_nodes = <index_nodes / (maxkv + 1)>
    nodes = nodes + index_nodes
WEND
index_file_size = (nodes + 1) * nnsec% * 128
```

To compute the largest possible index file size, replace MAXKV by MAXKV/2 in the preceding algorithm. The minimum size computation assumes completely full nodes, while the largest size computation assumes half-full nodes. The

ordinary B-Tree structure ensures at least half-full nodes. Access Manager performs local node rotations that help maintain the smallest index file size by avoiding unnecessary node splitting.

For example, assume NNSEC% is four and KEYLEN% is ten. Access Manager computes MAXKV to be thirty-four. The minimum bytes required to index 10,000 key values under these circumstances is 156,672 (153 KB). The maximum number of bytes required is 320,512 (313 KB). Actual experience with Access Manager indicates for random insertions the index file records are approximately three-quarters full. This corresponds to an estimated file size of 210,432 bytes (206 KB).

Because user errors are automatically sent to the console in this example listing (ERROPT%, the second parameter of INTUSR, equals 0), there are no tests of the ERRCOD function.

```
REM -----  
REM  AM80 External Declarations  
REM -----
```

```
%INCLUDE am80extr.bas
```

```
REM -----  
REM  Parameter Setup  
REM -----
```

```
yes% = 1  
no% = 0  
progid% = -1  
Error.Trap% = no%  
timeout% = 0  
bufs% = 10  
keys% = 2  
nsec% = 4  
dfile% = 1
```

```
REM -----  
REM  Initialize AM80  
REM -----
```

```
progid% = intusr (progid%, Error.Trap%, timeout%)  
dummy% = setup (bufs%, keys%, nsec%, dfile%)
```

```
REM -----  
REM  Open Index Files  
REM -----
```

```
Cust.Idx$ = "B:CUST.IDX"  
Part.Idx$ = "C:PART.IDX"  
Cust.Type% = 0      REM Alphanumeric Key  
Part.Type% = 1     REM Integer Key  
Cust.Len% = 11  
Part.Len% = 4  
Cust.Dup% = yes%  
Part.Dup% = no%    REM No duplicate part numbers
```


Cust.Key% = opnidx (-1, Cust.Idx\$, Cust.Len%, Cust.Type%, Cust.Dup%)
Part.Key% = opnidx (-1, Part.Idx\$, Part.Len%, Part.Type%, Part.Dup%)

REM -----
REM Open Data File
REM -----

Inv.Dat\$ = "D:INVENTORY.DAT"
Inv.Len% = 192 REM Record length
S.File% = 3 REM Shared file lock
Inv.File% = opndat (-1, S.File%, Inv.Dat\$, Inv.Len%)

Listing 3-3. OPNIDX function program code

Error Codes:

Table 3-25 lists the OPNIDX function error codes.

Table 3-25. OPNIDX error codes

Value	Code	Explanation
23	AG	Index file is corrupt; no header record.
24	AH	No more directory space.
30	AN	Index file number (KEY%) is out of range.
31	AO	Illegal index filename.
33	AQ	Index file number (KEY%) is already in use.
34	BB	Key length (KEYLEN%) exceeds 48 bytes.
39	BG	Key length (KEYLEN%) too long for number of assigned disk sectors (NNSEC%%).
40	BH	Index file is corrupted.
89	EI	Incorrect or missing password.
53	II	Bad parameter value.

OPRDAT function Data file setup and maintenance function

Syntax: OPRDAT (file%, dlock%, filename\$, reclen%)

Explanation:

OPRDAT is used in place of the OPNDAT function to open a corrupted data file. A data file is corrupted when it has been updated but not subsequently closed by a SAVDAT or CLSDAT function. OPRDAT performs exactly like the OPNDAT function, except that it does not check to see if the data file is corrupted. OPRDAT can be used to open and reconstruct a corrupted data file.

At the conclusion of the function, the number of the opened data file is

returned.

Note: A corrupted data file opened with OPRDAT and then subsequently saved (SAVDAT) or closed (CLSDAT) will no longer appear corrupted to Access Manager, even if it still is corrupted. Therefore, use OPRDAT with due caution.

Parameters:

Parameters for OPRDAT are exactly the same as for OPNDAT. Please refer to OPNDAT for explanation of these parameters.

Error Codes:

OPRDAT causes the following user errors:

Table 3-26. OPRDAT error codes

Value	Code	Explanation
53	CE	No header record in data file.
54	CD	No more directory space.
60	CL	Data file number (FILE%) is out of range.
61	CM	Data filename (FILENAME\$) incorrectly formed.
63	CO	Data file (FILE%) number already in use.
119	GG	Incorrect or missing password.
177	KA	Lock code (DLOCK%) is out of range.
183	KG	Bad parameter value.

Example:

This example demonstrates how to create a data file that is guaranteed to be new. The old data file is erased first, even if it is corrupted.

```
File.No% = -1
dlock% = 4          REM Exclusive file lock
Record.Len% = 100
File.Name$ = "E:DUMMY.DAT"

File.No% = oprdat (File.No%, dlock%, File.Name$, Record.Len%)
IF errcod <> 0 \
AND errcod <> 53 \
  THEN CALL Error.Handler (1)
IF lokcod <> 0 \
  THEN CALL Lock.Conflict (1)

REM If no errors, erase old file and create new, empty file.

dummy% = eradat (File.No%, dlock%)
```

```
File.No% = opndat (-1, dlock%, File.Name$, Record.Len%)
IF errcod <> 0 \
  THEN CALL Error.Handler (2)
IF lokcod <> 0 \
  THEN CALL Lock.Conflict (2)
```

OPRIDX function Index file setup and maintenance function

Syntax: OPRIDX (key%, idxname\$, keylen%, keytyp%, dupkey%)

Explanation:

OPRIDX opens or creates an index file that cannot be opened or created using the OPNIDX function.

Each time OPNIDX opens an index file, it checks the integrity of the file. Integrity is lost if the file has been updated but not subsequently closed during the SAVIDX or CLSIDX functions. If OPNIDX discovers a loss of integrity in the index file, the file is not opened; instead, user error 40/BH is issued.

OPRIDX behaves exactly the same as the OPNIDX function, except it does not check the integrity of the index file.

Parameters:

Parameters for OPRIDX are exactly the same as for the OPNIDX function. Please refer to the OPNIDX function for a discussion of these parameters.

Additional Comments:

OPRIDX is typically used with the ERAIDX function. Any other use of OPRIDX can cause unexpected error conditions, including Access Manager internal consistency errors.

An index file opened with OPRIDX and subsequently closed during SAVIDX or CLSIDX might appear uncorrupted to Access Manager when, in fact, it is corrupted. If an index must be opened with OPRIDX, replace or rebuild it immediately.

Error Codes:

OPRIDX causes the error codes listed in Table 3-27.

Table 3-27. OPRIDX error codes

Value	Code	Explanation
23	AG	No header record in index file.
24	AH	No more directory space.
30	AN	Index file number (KEY%) is out of range.
31	AO	Illegal index filename.
33	AQ	Index file number (KEY%) is already in use.
34	BB	Key length (KEYLEN%) exceeds 48 bytes.
39	BG	Key length (KEYLEN%) is too long for the number of assigned disk sectors (NNSEC%).
153	II	Bad parameter value.

Example:

This example demonstrates how to create an index file guaranteed to be new and empty (possibly for temporary use). Note that OPRIDX should be used in conjunction with the ERAIDX and OPNIDX functions.

```
key% = -1
Key.Name$ = "DUMMY.IDX"
Key.Len% = 10
Key.Typ% = 0
Key.Dup% = 0

key% = opridx (key%, Key.Name$, Key.Len%, Key.Typ%, Key.Dup%)
IF errcod <> 0 \
AND errcod <> 23 \
  THEN CALL Error.Handler (1)

dummy% = eraidx (key%)
key% = opnidx (-1, Key.Name$, Key.Len%, Key.Typ%, Key.Dup%)
IF errcod <> 0 \
  THEN CALL Error.Handler (2)
```

PRVKEY function Index file search function

Syntax: PRVKEY (key%, file%, dlock%, idxval\$)

Explanation:

PRVKEY returns the data record number associated with the preceding key value in an index file. PRVKEY also places the value of the key it finds in the IDXVAL\$ parameter.

PRVKEY is an efficient way to move sequentially backward through an index file in a single-user environment, when no index file updates are performed. You

can interleave calls to PRVKEY for different KEY%s, because separate position pointers are maintained for each index file.

[MULTI] PRVKEY is not normally used in a multiuser environment.

Parameters:

KEY%

The number of the index file where the search takes place.

FILE%

[MULTI] This parameter is ignored in a single-user environment. It is the number of the data file referenced by the KEY% parameter.

DLOCK%

[MULTI] This parameter is ignored in single-user environments. It contains a code to specify the type of data record lock requested for the data file referenced by the FILE% parameter. If the requested lock for the data record PRVKEY found cannot be granted, LOKCOD returns a non-zero value. See the SETLOK function for a list of acceptable lock codes.

IDXVAL\$

This is an OUTPUT parameter. Access Manager places the key value found in the index file in IDXVAL\$ at the conclusion of the function. If no previous index entry is found, PRVKEY returns a zero, and IDXVAL\$ is set to all blank spaces (which is not the same as a null string!...).

Before calling the PRVKEY function, IDXVAL\$ must be at least as long as the key length specified for the KEY% file. Normally, it is only necessary to initialize IDXVAL\$ once at the beginning of a program.

Additional Comments:

To determine the two high-order bytes of the associated data record number, call the DATVAL function immediately after PRVKEY.

Before the very first call to PRVKEY for a given KEY% or after the index file associated with KEY% is updated, one of the other index search functions (except NXTKEY) must be called for the same KEY%. The other search functions establish the internal pointer used by PRVKEY and NXTKEY for sequential processing. Each time any of the search functions are called (including NXTKEY and PRVKEY), the internal pointer is appropriately updated. However, the internal pointers are not maintained when the index file is updated.

Error Codes:

Table 3-28 lists the PRVKEY function error codes.

Table 3-28. PRVKEY error codes

Value	Code	Explanation
30	AN	Index file number (KEY%) is out of range.
35	BC	IDXVAL\$ is too short to contain key value.
46	BN	Index file is not open.
147	IC	Lock code (DLOCK%) is out of range.
153	II	Bad parameter value.

READAT function Data file update function

Syntax: READAT (file%, drn%, buffer%)

Explanation:

READAT reads a specified data record into a buffer area in memory. If the read operation is successful, READAT returns a zero at the conclusion of the function; otherwise, a non-zero user error code results.

Parameters:

FILE%

The number of the data file from which the data record is read.

DRN%

The number of the data record to be read from the file specified by the FILE% parameter.

BUFFER%

The address of a buffer area in memory. When the data record is read from the data file, it is placed in this buffer area for subsequent processing.

Additional Comments:

Your application program must ensure that sufficient space is allocated at the location specified by the BUFFER% parameter to accommodate the data records. In other words, make sure your buffer area is large enough to contain your data records.

Your application program must also parse the data record into the desired variables for use. This is easily accomplished in languages permitting based variables and/or data structures.

If the data record number (DRN%) exceeds two bytes, call the SETDAT function immediately before READAT to set the two high-order bytes of the data record

number.

Error Codes:

READAT returns zero if the read operation is successful; otherwise, a non-zero value is returned. User errors that might result are listed in Table 3-29.

Table 3-29. READAT error codes

Value	Code	Explanation
52	CD	Data record number (DRN%) is zero or beyond the logical end-of-file.
53	CE	Attempt to read past physical end-of-file, or read unwritten data.
60	CL	Data file number (FILE%) is out of range.
183	KG	Bad parameter value.

Example:

```
Read.Buffer$ = space$ (Record.Length%)  
Adr.Buffer% = SADD (Read.Buffer$) + 2
```

```
IF readat (File.No%, drn%, Adr.Buffer%) <> 0 \  
  THEN CALL Error.Handler
```

```
FOR fld% = 1 TO No.Fld%  
  field$ (fld%) = MID$ (Read.Buffer$, Field.Beg% (fld%), \  
    Field.Len% (fld%) )  
NEXT fld%
```

RETREC function Data file update function

Syntax: RETREC (file%, dlock%, drn%)

Explanation:

RETREC returns data records to the pool of available records for subsequent reuse with the NEWREC function. If the RETREC function is successful, zero is returned at its conclusion; otherwise, a non-zero user error code results.

Parameters:

FILE%

The number of the data file where data records are reclaimed.

DLOCK%

[MULTI] This parameter is ignored in a single-user environment. It contains a code specifying the type of data record lock requested for the data file to be accessed. See the SETLOK function for a list of acceptable lock codes.

If Access Manager grants the requested data record lock, RETREC places the returned data record number at the top of a logical stack of deleted records. If a lock conflict occurs, no action is taken and LOKCOD returns the appropriate non-zero value.

Note that, unless RETREC is called with a DLOCK% parameter of zero, you must test the return value of LOKCOD immediately after calling RETREC. If LOKCOD returns a non-zero value, the data record is not deleted and your application program must take appropriate action. Further, if RETREC is successful, the data record lock held by the calling program is automatically released.

RETREC automatically releases the data record lock after deleting the data record. This ensures that no other program gets the just deleted record from NEWREC before the deleting program releases its data record lock. In short, the deletion and release-lock operations work as an automatic operation. Therefore, it is unnecessary to call FRELOK after RETREC.

DRN%

The data record number of the record to be returned to the pool.

Additional Comments:

If the data record number requires more than two bytes, call the SETDAT function immediately after RETREC.

The stack of deleted records is implemented by a linked list that ensures fast operation, because the NEWREC and RETREC functions manipulate only the topmost record in the stack. After a successful call to RETREC, the record in the data file with record number DRN% has two fields written over the previous contents. These fields are defined in the following table:

Table 3-30. Contents of deleted data record

Byte position	Contents
-----	-----
0	0FFh (255 decimal). This byte serves as a flag for deleted data records. Your program should reserve this byte, and ensures it never otherwise contains 0FFH
1-3	A link to the next available data record.
4+	Undefined.

The 0FFh value in the first byte of deleted data records serves two purposes.

First, Access Manager checks the first byte of every deleted record it is about to reuse, to see if it contains 0FFH. If it does not, user error 69/DE results. Second, your application program can use the first byte as a deleted record flag. Therefore, when reading through a data file without accessing the index files, you can disregard data records with 0FFh in the first byte. This latter use assumes you either reserved the first byte for the delete flag, or your valid data can never contain 0FFh in the first byte.

Error Codes:

 Table 3-31 lists the error codes caused by the RETREC function.

Table 3-31. RETREC error codes

Value	Code	Explanation
52	CD	Data record number (DRN%) is zero, or beyond the logical end-of-file.
60	CL	Data file number (FILE%) is out of range.
69	DE	First byte of record is not 0FFH.
177	KA	Lock code (DLOCK%) is out of range.
183	KG	Bad parameter value.

Example:

 In this example, DRN2% contains the two high-order bytes, and DRN% the two low-order bytes of the data record number to be returned to the pool of available data records. Note the use of SETDAT to initialize DRN2%. The arguments of the exception processing functions indicate the location of the exception.

```
dlock% = 2          REM Exclusive record lock request code
CALL setdat (drn2%)
```

```
IF retrec (File.No%, dlock%, drn%) <> 0 \
  THEN CALL Error.Handler (4)
IF lokcod <> 0 \
  THEN CALL Lock.Conflict (4)
```

SAVDAT function Data file setup and maintenance function

 Syntax: SAVDAT (file%)

Explanation:

 SAVDAT forces data file updates to the disk, while keeping the data file open

for further use. If the save operation is successful, a zero is returned at the conclusion of the function; otherwise, a non-zero user error results.

SAVDAT has the same effect on a data file as issuing a call to CLSDAT followed by a call to OPNDAT, except that, if there are no updates, SAVDAT simply returns without performing any actions.

Parameters:

FILE%

The number of the saved data file.

Additional Comments:

The main use of the SAVDAT function is to save data file updates at critical points in your application program. Hardware or software failures cannot corrupt a saved data file, unless additional updates have been performed, subsequent to the save.

If you make changes to a data file, your application program must call the SAVDAT or CLSDAT function to force the updates to the disk file. If you do not do this, the integrity of the data file can be disrupted, because the Access Manager header record could be incorrect. Further, the operating system might be holding data in internal buffers. SAVDAT and CLSDAT force the updated header record to the disk, and flush the operating system buffers. Calling SAVDAT after each new record is added to a data file will degrade system performance. You must balance the desire for security with the need for quick response time.

Error Codes:

Table 3-32 lists the user errors for the SAVDAT function.

Table 3-32. SAVDAT error codes

Value	Code	Explanation
60	CL	Data file number (FILE%) is out of range.
73	DI	Could not reopen data file during save.
177	KA	Lock code (DLOCK%) is out of range.
183	KG	Bad parameter value.

Example:

```
IF savdat (File.No%) = 0 \
  THEN PRINT "File save was successful."
```

SAVIDX function Index file setup and maintenance function

Syntax: SAVIDX (key%)

Explanation:

SAVIDX forces index file updates to the disk, while keeping the file open for further use. If the save operation is successful, zero is returned at the conclusion of the function; otherwise, a non-zero user error code results.

SAVIDX has the same effect on an index file as issuing a call to CLSIDX followed by a call to OPNIDX, except that, if there are no updates, SAVIDX simply returns without performing any actions.

Parameters:

KEY%

The number of the saved index file.

Additional Comments:

The main use of the SAVIDX function is to save index file updates at critical points in your application program. Hardware or software failures cannot corrupt a saved index file, unless additional updates have been performed since SAVIDX was used.

If you make changes to a index file, your application program must call the SAVIDX or CLSIDX function to force the updates to the disk file. If this is not done, the integrity of the index file is not ensured, because some updated nodes might still be in an I/O buffer and/or the header record might be incorrect. Calling SAVIDX after each update of an index file will degrade system performance. You must balance the need for security with that for speedy updates.

Error Codes:

The SAVIDX function causes the following user errors:

Table 3-33. SAVIDX error codes

Value	Code	Explanation
-----	----	-----
21	AE	Disk or directory full.

30 AN Index file number (KEY%) is out of range.
43 BK Could not reopen index file during save.
46 BN Index file (KEY%) is not open.
153 II Bad parameter value.

SERKEY function Index file search function

Syntax: SERKEY (key%, file%, dlock%, keyval\$, idxval\$)

Explanation:

SERKEY returns the data record number assigned to the first entry (in key-sequential order) that is equal to or greater than a specific key value. SERKEY also places the key value it finds in the IDXVAL\$ parameter.

SERKEY can be used to locate an entry in an index file when only the beginning portion of the key value is known.

Parameters:

KEY%

The number of the index file where the search takes place.

FILE%

[MULTI] This parameter is ignored in a single-user environment. It is the number of the data file referenced by the KEY% parameter.

DLOCK%

[MULTI] This parameter is ignored in single-user environments. It contains the code specifying the type of data record lock requested for the data file referenced by the FILE% parameter. If the requested lock for the data record that SERKEY found cannot be granted, LOKCOD returns a non-zero value. See the SETLOK function for a list of acceptable lock codes.

KEYVAL\$

The value of the key record being used as a reference.

IDXVAL\$

This is an OUTPUT parameter. Access Manager places the key value found in the index file in IDXVAL\$. If no index entry is found with a key value equal to or greater than KEYVAL\$, SERKEY returns a zero, and IDXVAL\$ is set to all blank spaces (which is not the same as a null string!...). Also, see "Additional Comments" below.

Additional Comments:

If you must determine the two high-order bytes of the data record number found by SERKEY, call the DATVAL function immediately after SERKEY.

There are three additional notes to consider, concerning the IDXVAL\$ parameter:

- 1) Access Manager never changes the address or length of IDXVAL\$ (which is a string-valued parameter). If IDXVAL\$ is not initially set to a string value that is long enough to contain the string-valued index entries assigned to IDXVAL\$, Access Manager cannot make the assignment (this results in user error 35/BC). At the beginning of an application program, you must set up one or more string variables that can be used in subsequent calls to SERKEY and to the other Access Manager functions that return key values from the index. This can be accomplished by assigning sufficiently long, blank strings to these variables.
- 2) If IDXVAL\$ is longer than the KEYLEN% associated with KEY%, IDXVAL\$ is padded on the right with blank spaces. (See OPNIDX function description.)
- 3) KEYVAL\$ and IDXVAL\$ cannot be the same variable, because IDXVAL\$ might be initialized to all blank spaces before the actual search of the index file.

As mentioned under "Duplicate Key Value" in the ADDKEY function, situations can arise where key values must be modified to accommodate duplicate key values. When this is done, you can no longer use GETKEY to locate such an index entry, because the sequence number modifies the original value. SERKEY can find the first candidate for a match with the specified key value. Subsequent candidates are found using the NXTKEY and/or AFTKEY function. KEYVAL\$ must be set up to avoid automatic padding by Access Manager.

Error Codes:

Table 3-34 lists the SERKEY function user errors.

Table 3-34. SERKEY error codes

Value	Code	Explanation
30	AN	Index file number (KEY%) is out of range.
46	BN	Index file is not open.
147	IC	Lock code (DLOCK%) is out of range.
153	II	Bad parameter value.

Example:

This example shows how to use SERKEY to find the first key value in a set of duplicates. Assume the function SPACE\$ has been defined and returns a blank

space string equal to its argument value. Also, assume KEY.LEN% is the key length, including the two-byte sequence number automatically set by the ADDKEY function.

```
idxval$ = space$ (Key.Len%)
```

```
INPUT "Enter the target key value: "; Key.Value$
```

```
Key.Value$ = LEFT$ (Key.Value$ + space$ (Key.Len%), Key.Len% - 2)
```

```
Key.Value$ = Key.Value$ + CHR$ (0)
```

```
drn% = serkey (key%, file%, dlock%, Key.Value$, idxval$)
```

SETDAT function Data file setup and maintenance function

Syntax: SETDAT (drn%)

Explanation:

SETDAT passes the two high-order bytes of the four-byte data record number when needed. If you do not call SETDAT, the two high-order bytes are set to zero.

Use SETDAT only when the input data record number parameter (DRN%) exceeds the two-byte capacity of an ordinary integer variable (65,535).

Parameters:

DRN%

The two high-order bytes of the data record number.

Additional Comments:

SETDAT does not return a value; it sets the two high-order bytes for data record numbers used as input parameters in other Access Manager functions (such as RETREC).

SETDAT is always used with another Access Manager function, and is called immediately before that function.

Error Codes:

SETDAT does not cause user errors.

Example:

CALL setdat (Seg.No%)
Ret.Code% = addkey (key%, File.No% (1), Xclsv.Lock%, keyval\$, drn%)

IF errcod <> 0 \
 THEN CALL Error.Handler
IF lokcod <> 0 \
 THEN CALL Lock.Conflict
IF Ret.Code <> 1 \
 THEN CALL Addkey.Problem

SETLOK function Data locking function

Syntax: SETLOK (file%, dlock%, drn%)

Explanation:

SETLOK sets a lock on a data file or data record. If the lock operation is successful, zero is returned at the conclusion of the function; otherwise, a one or two is returned. This function has no effect in a single-user environment.

Parameters:

FILE%

The number of the data file for which the lock is requested.

DLOCK%

The code to specify the type of lock operation requested. Table 3-35 lists the lock requests that can be requested with this function.

DRN%

The number of the data record for which the lock is requested.

Additional Comments:

If the requested lock is granted, LOKCOD is set to zero, and the calling program (see PROGID% parameter under the INTUSR function) holds the lock.

If the requested lock cannot be granted, LOKCOD is set to one or two. One indicates a record lock conflict; two indicates a conflict with an exclusive file lock.

Acceptable lock request codes are shown in Table 3-35.

Table3-35. Data file/data record lock requests

DLOCK%

value Lock request

- 0 Locks are ignored, but LOKCOD returns zero.
- 1 Attempt a shared record lock on record number DRN%.
- 2 Attempt an exclusive record lock on DRN%.
- 3 Attempt a shared file lock on FILE%.
- 4 Attempt an exclusive file lock on FILE%.

The ignore lock request (DLOCK% = 0) always returns a successful LOKCOD result. Therefore, unless your program relies on locking protocols outside Access Manager, use caution when passing DLOCK% with zero. This lock request is most commonly used when your program must search an index file, but you do not want to access the data record associated with the returned key value.

Use a shared data record lock (DLOCK% = 1) when you want to review the record, but have no intention (at least not yet) of updating it. The shared lock blocks other users from gaining an exclusive lock, but still allows them to access the record for review.

Use an exclusive data record lock (DLOCK% = 2) when you want to update a data record. This prevents two users from updating the same record at the same time.

A user holding a shared record lock can try to change to an exclusive lock; but an attempt to change to an exclusive lock on a data record that already carries shared locks by other users fails with a LOKCOD of one.

A shared file lock (DLOCK% = 3) ensures that the user cannot be subsequently blocked from the file by another user's exclusive file lock.

Use an exclusive data file lock (DLOCK% = 4) when it is imperative that no other users have access to the records in a data file. An exclusive file lock is granted only if no other users have shared locks on the file, and there are no active data record locks.

When one user holds an exclusive file lock, all other non-zero lock requests (that is, DLOCK% = 1, 2, 3, or 4) by other users are denied with a LOKCOD of two.

Index file functions can set data record locks at the time the index file operation takes place. Therefore, calling SETLOK when a data record is located by the index file functions is not normally necessary.

Error Codes:

SETLOK function error codes are listed in Table 3-36.

Table 3-36. SETLOK error codes

Value	Code	Explanation
-------	------	-------------

52 CD Data record number (DRN%) is zero, or beyond
the logical end-of-file.
60 CL Data file number (FILE%) is out of range.
177 KA Lock code (DLOCK%) is out of range.
183 KG Bad parameter value.

Example:

For this example, assume a shared record lock is held on the data record (DRN%). The purpose is to upgrade the shared lock to an exclusive record lock.

```
IF setlok (File.No%, 2, drn%) <> 0 \  
  THEN PRINT "Exclusive lock failed."
```

```
IF errcod <> 0 \  
  THEN CALL Error.Handler
```

SETUP function System initialization and maintenance function

Syntax: SETUP (nbufs%, nkeys%, nnsec%, ndatf%)

Explanation:

SETUP only applies in single-user environments. Calls to SETUP are ignored in a multiuser environment, because the background server automatically calls the function.

SETUP prepares the special Access Manager buffer area, and specifies the basic characteristics of the index files. You must complete the SETUP function before opening or using any index and/or data files in your program.

Parameters:

NBUFS%

A value specifying the number of index file I/O buffers to be used. There must be at least three of these buffers.

As the number of buffers increases, the time to access a key value decreases, and the memory space required increases. All index files share the same buffer space, thereby minimizing the memory required in your program. Even if several index files are used simultaneously, it is not necessary to use more than three buffers.

NKEYS%

A value specifying the maximum number of index files the program will use

simultaneously. The value must be at least one.

NNSEC%

Determines or specifies the length of the records in an index file. Specifically, NNSEC% is the number of 128-byte disk sectors in each index file record. Each record corresponds to a B-Tree node. The more sectors per record, the more key values stored per node. The more key values stored per node, the fewer accesses required to find a key value.

NNSEC% must be at least one. There is no upper limit for this parameter, but because there are limits on the number of key values stored per node, you should consider the following practical limitations.

- Access Manager forces all index files open at the same time, to have the same record length.
- To maintain compatibility with application software from other vendors, Digital Research suggests a value of four as an informal standard for the NNSEC% parameter. Compatibility is particularly important for multiuser environments, because different software can operate simultaneously.

NDATF%

A value specifying the maximum number of data files that are open at one time. It is not necessary to use the Access Manager data file functions. You can use routines coded in the application language, or in whatever form is appropriate for your programs.

Additional Comments:

See the "Access Manager Design Constraints" section of your Programmer's Guide for additional information on the maximum values permitted for SETUP parameters.

Although unlikely, you might want to change the SETUP parameters during the course of your application program. If you do, be certain to close all index and data files before subsequent calls to SETUP.

Error Codes:

The SETUP function returns a zero if the parameter values fall within their legal ranges and the buffer area is large enough. Error codes that can result are listed in Table 3-37.

Table 3-37. SETUP error codes

Value	Code	Explanation
-------	------	-------------

209	MA	Illegal parameter value, or buffer area too small.
-----	----	--

Example:

See the INTUSER function description.

UPDPTR function Index file update function

Syntax: UPDPTR (key%, file%, dlock%, keyval\$, drn%)

Explanation:

UPDPTR changes the data record number assigned to an existing key value. At the conclusion of the function, a value is returned to indicate the success or failure of the update operation (see "Additional Comments" below).

Without UPDPTR, you would first have to delete KEYVAL\$ from the index file, and then reinsert it with its new data record number (DRN%). UPDPTR provides a more efficient method for doing this.

Parameters:

KEY%

The number of the index file where the data record number is changed.

FILE%

[MULTI] This parameter is ignored in a single-user environment. It contains the number of the data file referenced by the KEY% parameter.

DLOCK%

[MULTI] This parameter is ignored in single-user environments. It contains the code specifying the type of data record lock requested for the data file referenced by the FILE% parameter. See the SETLOK function for a list of acceptable lock codes.

To ignore locking protocols in a multiuser environments, assign DLOCK% a value of zero. This procedure, however, is not recommended.

KEYVAL\$

The key value for the record where the data record number is changed.

If sequence numbers are assigned for duplicate keys by Access Manager (see ADDKEY function), KEYVAL\$ must include the proper sequence number.

DRN%

The data record number assigned to the key value contained in KEYVAL\$. Specifically, this is the new number assigned to the key value.

Additional Comments:

UPDPTR returns one of the values listed in Table 3-38.

Table 3-38. UPDPTR function values

Value Meaning

- 0 KEYVAL\$ was not found in the index.
- 1 The data record number was successfully changed.
- 4 The requested DLOCK% was not granted for the specified data record (DRN%). The change is not made, nor is the index file searched for KEYVAL\$.

Error Codes:

Table 3-39 lists the user errors for the UPDPTR function.

Table 3-39. UPDPTR error codes

Value Code Explanation

- 30 AN Index file number (KEY%) is out of range.
- 46 BN Index file is not open.
- 147 IC Lock code (DLOCK%) is out of range.
- 153 II Bad parameter value.

Example:

```
CALL setdat (drn2%)
IF updptr (key%, 0, 0, Key.Value$, drn%) <> 1 \
  THEN PRINT "Index file pointer not updated."
IF errcod <> 0 \
  THEN CALL Error.Handler
```

WRTDAT function Data file update function

Syntax: WRTDAT (file%, drn%, buffer%)

Explanation:

WRTDAT writes a data record into a data file. The data record to be written is

taken from a buffer area in your computer's memory. If the write operation is successful, a zero is returned at the conclusion of the function; otherwise, a non-zero user error code results.

Parameters:

FILE%

The number of the data file into which the data record is written.

DRN%

The relative number of the data record to be written into the file specified by the FILE% parameter.

BUFFER%

The address of a buffer area from which the data record is written into the data file.

Additional Comments:

Your application program must ensure that sufficient space is allocated at the location specified by the BUFFER% parameter to accommodate the data records. In other words, make sure your buffer area is large enough to contain your data records.

If the data record number (DRN%) exceeds two bytes, call the SETDAT function immediately before WRTDAT, to set the two high-order bytes of the data record number.

Error Codes:

If WRTDAT returns a non-zero value, a user error occurred during the write operation. Listed in Table 3-40 are the error codes for the WRTDAT function.

Table 3-40. WRTDAT error codes

Value	Code	Explanation
51	CC	Disk or directory is full.
52	CD	Attempt to write record zero, or write past logical end-of-file.
60	CL	Data file number (FILE%) is out of range.
183	KG	Bad parameter value.

Example:

Write.Buffer\$ = " "

```
FOR fld% = 1 TO No.Flds%
  Write.Buffer$ = Write.Buffer$ + field$ (fld%)
NEXT fld%
```

```
Adr.Buffer% = SADD (Write.Buffer$) + 2
IF wrtdat (File.No%, drn%, Adr.Buffer%) <> 0 \
  THEN CALL Error.Handler (locale%)
```

Section 4: Access Manager error codes

4.1 Error types

Access Manager functions generate two types of errors: internal consistency errors and user errors.

4.1.1 Internal consistency errors

Access Manager generates internal consistency errors when a B-Tree index file or a data file does not satisfy the internal consistency checks the Access Manager functions perform when a file is used. Such errors do not occur, unless your application is processing corrupted files, or a pointer variable is not properly initialized. If an Access Manager internal consistency error does occur, the console displays an error message of the following form:

Access Manager Internal Error ...XY...

where "XY" is replaced by one of the codes listed in Table 4-1.

If, after reviewing your application program, you cannot find any obvious cause for the internal error, and cannot successfully recreate the file, please contact Digital Research. Provide as much information about the error as possible, including the two character error code that is displayed.

4.1.2 Error codes

User error codes occur when Access Manager finds avoidable problems; such as no more space on a disk, or an illegal value for the KEY% parameter. You determine how Access Manager handle user errors by the value passed in the ERROPT% parameter via the INTUSR function. You have two options:

- 1) If the INTUSR function is called with a non-zero value in the ERROPT% parameter, user errors can be trapped by testing the ERRCOD function for a non-zero value after calls to Access Manager functions. When a user error occurs, ERRCOD returns one of the values listed in Table 4-1.
- 2) If the INTUSR function is called with a zero value in the ERROPT%

parameter, the console displays a user error message of the following form:

USER ERROR ...XY... CHECK ACCESS MANAGER MANUAL

where "XY" is replaced by one of the codes listed in Table 4-1. Control then returns to the operating system.

In Table 4-1, the symbol "*" marks errors that send a complete error message to the console. The symbol "^" marks errors that are trapped by Access Manager only if your operating system supports extended BDOS error returns.

Table 4-1. Access Manager user error codes

Value	Code	Explanation
21	AE	Cannot write an index file record. Check to see if the disk or its directory is full.
23	AG	Cannot read index file record. You might be attempting to use a newly created index file that was never closed properly.
24	AH	There is no more directory space on the disk. This occurs while trying to create a new index file.
25	AI	Access Manager could not find the name of the index file in the disk directory while attempting to close the file. This might occur if overlays destroy the Access Manager data areas.
30	AN	KEY% parameter value is out of range. The value must satisfy the following equation: $0 \leq \text{KEY\%} \leq \text{NKEYS\%}$
31	AO	The IDXNAME\$ parameter value in an OPNIDX function contains unacceptable information. Check to see if the parameter contains a null value.
33	BA	You are attempting to reuse an index file number (KEY% parameter value) already assigned to an open index file.
34	BB	The KEYLEN% parameter value in an OPNIDX function exceeds the maximum allowable value of 48.
35	BC	The IDXVAL\$ parameter has been initialized with an improper value. Check to make sure the parameter value is at least as long as the key length.
36	BD	While using the ADDKEY function, an attempt was made to assign zero as the data record number.
39	BG	The key length (KEYLEN%) is too long for the number of disk sectors assigned (NNSEC%). The number of sectors must be sufficient to accommodate at least four key values.

- 40 BH When trying to open an index file with the OPNIDX function, Access Manager found the file to be corrupted. Index files become corrupted when they are not closed (CLSIDX or SAVIDX functions) following updates. The index file must be rebuilt.
- 43 BK Access Manager was unable to reopen an index file while performing the SAVIDX function. This error should not occur under normal circumstances. If it persists, contact Digital Research.
- 44 BL You have attempted to close an index file that is not currently open.
- 46 BN You have made an attempt to use an index file that is not currently open.
- 51 CC Access Manager cannot write the data record into the data file. The disk or directory might be full.
- 52 CD While using the READAT, WRDAT, or RETREC functions, you have attempted to use a data record number of zero, or a data record beyond the logical end of the data file. If applicable, check the way you have used the NEWREC function.
- 53 CE Access Manager cannot read the data record your program has requested. You might be attempting to read a newly created data file that was not properly closed, read past the physical end of the data file, or read unwritten data.
- 54 CF There is no more directory space on the disk where you are attempting to create a data file.
- 55 CG Access Manager cannot close the data file specified in a CLSDAT function parameter. This might occur if an overlay destroys the Access Manager data areas.
- 60 CL The data file number specified in a FILE% parameter is out of range. The file number must satisfy this equation:
- $$0 \leq \text{FILE\%} < \text{NDATF\%}$$
- 61 CM The FILNAME\$ parameter value in an OPNDAT or OPRDAT function contains unacceptable information. Most likely cause is a null filename.
- 63 CO You have attempted to use a data file number (FILE% parameter) already assigned to another data file.
- 69 DE The NEWREC function attempted to reuse a data record in which the first byte did not contain OFFH. (The RETREC function sets this byte.) You might be attempting to use a data file that has not been properly closed.

- 70 DF The data file you have attempted to open with the OPNDAT function is corrupted. The file can be corrupted if updates are made but not posted with the CLSDAT or SAVDAT functions. The data file can be opened using the OPRDAT function.
- 73 DI Access Manager cannot reopen a data file during a SAVDAT function. See discussion of user error 43/BK.
- 74 DJ You have attempted to close a data file (FILE% parameter) that is not currently open.
- 76 DL You have attempted to use a data file (FILE% parameter) that is not currently open.
- 83^ EC A physical error occurred on a disk during input/output to an index file.
- 84^ ED You have attempted to update an index file on a disk with Read-Only status.
- 85^ EE You have attempted to update an index file with Read-Only status.
- 86^ EF An index filename references a nonexistent drive, or the File Control Block is no longer active.
- 87 EG Index file opened by another process in a multitasking environment.
- 88 EH FCB checksum error on index file close.
- 89 EI Incorrect or missing password for an index file.
- 90^ EJ OPNIDX attempts to create a second copy of file, instead of simply opening existing file. Should not occur. Report to Digital Research.
- 91^ EK An illegal character (?) in index filename.
- 92 EL Operating system open file limit exceeded during index file open.
- 93 EM Space in system lck list exhausted during index file operation.
- 113^ GA Physical error on disk during data file I/O.
- 114^ GB Attempt to update data file on Read-Only disk.
- 115^ GC Attempt to update Read-Only data file.
- 116^ GD Data filename references nonexistent drive, or File Control Block is no longer active.

- 117 GE Data file opened by another process in a multitasking environment.
- 118 GF FCB checksum error on data file close.
- 119 GG Incorrect or missing password for a data file.
- 120^ GH OPNDAT attempts to create a second copy of file, instead of simply opening existing file. Should not occur. Report to Digital Research.
- 121^ GI An illegal character (?) in data filename.
- 122 GJ Operating system open file limit exceeded during data file open.
- 123 GK Space in system lock list exhausted during data file operation.
- 145 IA You have attempted to erase an index file (KEY% parameter) that is not currently open.
- 146 IB The filename of the index file you are attempting to erase cannot be found in the disk directory.
- 147 IC The type of lock you have requested in a DLOCK% parameter is unacceptable.
- 153 II You have passed an unacceptable value in a function parameter. This code usually indicates a problem with the interface between the application program and Access Manager.
- 154 IJ Usually indicates a bad language interface and/or failure to call the INTUSR function in your program.
- 155 IK The number of B-Tree nodes in an index file exceeds 65,535.
- 161 JA* The value passed in the PROGID% parameter is out of range.
- 162 JB* Access Manager cannot open the queue for a shared, multiuser module. Make sure a Resident System Process has been included as part of the GENSYS procedure.
- 163 JC* Access Manager cannot open the queue for a user, because the PROGID% parameter value is out of range.
- 175 JO You have attempted to erase a data file (FILE% parameter) which is not currently open.
- 176 K@ Access Manager cannot find the name of the data file you are attempting to erase in the disk directory.
- 177 KA You have passed an illegal value in the DLOCK% parameter.

- 183 KG Same as error 153/II above.
- 184 KH Same as error 154/IJ above.
- 209 MA While using the SETUP function, you have passed a bad parameter value, or the buffer area is too small.

Section 5: RECREATE utility program

Access Manager provides a utility program for recreating index and data files when necessary. Normally, the only time you need to recreate files is when they are vulnerable and the hardware loses power, or the application program terminates abnormally.

The source code for the RECREATE program is provided in each language supported by Access Manager. Your Programmer's Guide contains instructions and examples for using the RECREATE program with each of these languages.

To simplify the recreation process, the RECREATE program uses a parameter file to describe the characteristics of the index and data files.

5.1 Recreate parameter file

The parameters you place in this file tell the RECREATE program how to rebuild a particular index or data file. When you run the RECREATE program, it asks you to enter the name of the Recreate Parameter File to use. The filename is the only thing you are asked to enter.

You can modify the RECREATE program to automatically read the appropriate parameter file without any user intervention, or to be started automatically by the application program when error trapping detects corrupted files.

The Recreate Parameter File is the only place Access Manager explicitly represents the relationships between index and data files. There are four types of records in the Recreate parameter File:

- 1) recreate header record
- 2) data file records
- 3) index file records
- 4) key part records

Table 5-1 illustrates the contents of each Recreate Parameter File record type.

Table 5-1. Recreate Parameter File record contents

Parameter	Meaning
-----	-----

Header record

=====

No.Data.File% The number of data files to be recreated.

mnsec% The size of index file records in sectors.

Data file record i

=====

filename\$ The data filename.

reclen% The data record length.

No.Index.Files% The number of associated index files.

Beg.Rec% The number of first record to be scanned.

Index file record i-j

=====

idxname\$ The index filename.

keylen% The length of index file key values.

keytyp% The key type (alphanumeric or integer).

dupkey% The duplicate key flag.

No.Key.Parts% The number of data record fields,
concatenated to form a key value.

Blank.Key.To.Null\$ How to deal with blank key values.

Key part record i-j-k

=====

Beg.Byte% The starting byte of key part.

Key.Part.Length% The length of key part.

In Table 5-1, BEG.REC% in the Data File Record specifies the first data record used for actual data. If BEG.REC% is zero, the RECREATE program automatically computes its value as the first data record to begin on or after the 129th byte of the file.

In the Index File Record, NO.KEY.PARTS% specifies the number of data record fields to be concatenated to form a key value. This information is used directly in conjunction with the Key Part Record. BLANK.KEY.TO.NULL\$ determines whether or not to convert a key value equal to all blank spaces to a null string. A "Y" indicates conversion is necessary; a "N" means no conversion is required. Convert a blank space string to a null key value when a blank space field means there is missing data and there is no key value to be entered. Null key values are not added to the index by the ADDKEY function, but a blank space string is.

In the Key Part Record, BEG.BYTE% specifies the byte position of a key part, and KEY.PART.LENGTH% specifies its length. For example, if BEG.BYTE% is two and KEY.PART.LENGTH% is five, the data in the second through sixth bytes of the data record forms a key part. Multiple key parts can be combined to form a complete key value. Note that the first byte of the record is considered byte position number one.

When you run the RECREATE program, the order of the duplicate keys is determined by the order in which they are present in the data records. This might not be the same order as before the recreation, due to the reclamation of deleted records.

Figure 5-1 illustrates the required order for records in a Recreate Parameter File.

```
Header record
Data file record 1
  Index file record 1-1
    Key part record 1-1-1
    ...
    Key part record 1-1-r
  Index file record 1-s
    Key part record 1-s-1
    ...
    Key part record 1-s-r
Data file record w
  Index file record w-1
    Key part record w-1-1
    ...
    Key part record w-1-r
  Index file record w-s
    Key part record w-s-1
    ...
    Key part record w-s-r
```

Figure 5-1. Recreate Parameter File record ordering

Figure 5-2 illustrates the contents of an example Recreate Parameter File. In this example, the header record indicates two data file, and 512-byte index file records (NNSEC% equals four). CUSTOMER.DAT has a 100-byte record length, and three associated index files. The RECREATE program automatically determines the first record with actual data, because BEG.REC% is zero. Notice that the key parts for NAME.IDX and ZIPC.IDX are two bytes shorter than the key length declared in the respective index file records. This is because both of these index files are using the automatic suffixing of the key values to accommodate duplicates. If the key parts were two bytes longer, the suffix would replace the last two bytes anyway.

In Figure 5-2, INVENTORY.DAT has 192-byte records, and only one associated index file. However, the key values for INVENTORY.IDX are constructed from three fields of the INVENTORY.DAT records: bytes 15-19, 100-101, and 4-6. That is, the ten bytes comprising these three fields are combined to form one key value.

```
2,4
CUSTOMER.DAT,100,3,0
NAME.IDX,10,0,1,1,Y
22,8
NUM.IDX,4,0,0,1,N
2,4
ZIPC.IDX,11,0,1,1,Y
84,9
INVENTORY.DAT,192,1,0
INVENTORY.IDX,12,0,1,3,Y
15,5
100,2
```

Figure 5-2. Example Recreate Parameter File

5.2 Data file recreation

The RECREATE program checks the integrity of each data file specified in the Recreate Parameter File. If a data file is not corrupt, the RECREATE program does not modify the data file. If the data file is corrupt, the RECREATE program corrects the data file in place. It reads each record in the data file to determine whether or not the record is active or deleted. If deleted, the RECREATE program automatically links the record onto the stack of deleted records for the data file. Once the entire file is read, the header record is rebuilt and written to the front of the data file.

A data record is considered deleted if the first byte of the record contains 0FFh. Access Manager automatically writes 0FFh into the first byte of each record deleted with the RETREC function. Therefore, if your application program uses RETREC to return deleted data records and, either reserves the first byte of each record for the delete flag, or ensures that no valid data can cause 0FFh in the first byte, you can use the RECREATE program as distributed.

If your application program conflicts with the use of 0FFh as a delete flag in the first byte of the data records, you must modify the RECREATE program to recognize deleted records some other way.

5.3 Index file recreation

The Recreate Parameter File indicates which index files are related to each data file. If the data file corresponding to an index file is not modified by the RECREATE program, and the index file is neither corrupt nor empty, the index file is not recreated.

If the RECREATE program modifies the data file corresponding to an index file, or the data file is corrupt, the index file is erased and recreated from scratch.

To speed recreation of the index file, the RECREATE program buffers scanning of the associated data file to extract key values. That is, the RECREATE program processes a large number of data file records before making the corresponding index file entries. This reduces disk head movement significantly. Further, the RECREATE program sorts the buffered key values before adding them to the index file. This also enhances the speed of the recreate process.

5.4 Recreate messages

The RECREATE program generates two types of messages: those that report on the progress of the recreate process, and those that report on unexpected conditions. Conditions that terminate the recreate process are caused by illegal values in the parameter file, lack of disk or directory space, inability to secure exclusive data file locks in a multiuser environment, or system inconsistencies.

If a system inconsistency arises (such as an inability to close a file, or to reopen a file previously closed), be sure to check the integrity of your hardware and operating system. If these are operating correctly and there is no apparent cause for the RECREATE error message, you should go back to your most recent data base back-up files, instead of trying to rebuild from your current, but corrupted data.

Appendix A: Access Manager function index

Table A-1 is an alphabetical list of Access Manager functions, cross-referenced to the mnemonic function name. These functions are organized alphabetically by mnemonic function name, and discussed in detail in Section 3.

Table A-1. Access Manager functions

Function description	Name
-----	-----
ADD key to index file	ADDKEY
CHANGE data record number in index file	UPDPTR
CLOSE data file	CLSDAT
CLOSE index file	CLSIDX
COUNT entries in an index file	NOKEYS
COUNT nodes (used and available) in an index file	NMNODS
COUNT records in a data file	GETDFS
COUNT records used in a data file	GETDFU
DELETE key from index file	DELKEY
DETERMINE results of data lock/unlock request	LOKCOD
ERASE data file	ERADAT
ERASE index file	ERAIDX
FIND error code value	ERRCOD
FIND first entry in index file	FRSKEY
FIND two high-order bytes of data record number	DATVAL
FIND index entry equal to or greater than a specific key value	SERKEY
FIND index entry following a specific key value	AFTKEY
FIND index entry in index file	GETKEY
FIND index entry preceding a specific key value	BEFKEY
FIND last entry in index file	LASKEY
FIND next available data record number	NEWREC
FIND next entry in an index file	NXTKEY
FIND previous entry in an index file	PRVKEY
INITIALIZE Access Manager	INTUSR
INITIALIZE Access Manager for single-user environment	SETUP
LOCK data file or data record	SETLOK
OPEN corrupted data file	OPRDAT

OPEN corrupted index file	OPRIDX
OPEN data file	OPNDAT
OPEN index file	OPNIDX
READ data record from data file	READAT
RECLAIM available data file records	RETREC
RELEASE data file or data record lock	FRELOK
SAVE data file updates	SAVDAT
SAVE index file updates	SAVIDX
SET two high-order bytes of data record number	SETDAT
WRITE data record in data file	WRTDAT

Index

(To be done...)

EOF

(Retyped by Emmanuel ROCHE.)

Digital Research
Access Manager
Function Summary

First Edition: March 1983

Table of Contents

- Introduction
- Functions by category
- Function index
- Function syntax and description
- Function parameters
- Parameter use table
- Usr error codes
- Data file/record lock request codes
- Data file/record lock release request codes

Introduction

This "Function Summary" contains information extracted and summarized from your "Access Manager Reference Manual". Explanations of the functions, parameters, lock codes, and error codes found here are brief. If you require more detailed information, consult your "Access Manager Reference Manual".

This "Function Summary" does not contain information specific to any operating system or programming language. Consult the appropriate "Access Manager Programmer's Guide" for this information.

Functions by category

System initialization
and maintenance

+-----+

INTUSR	
SETUP	
ERRCOD	
LOKCOD	

V V

DATA FILE FUNCTIONS

INDEX FILE FUNCTIONS

=====

=====

Data file setup
and maintenance

Index file setup
and maintenance

CLSDAT
DATVAL
ERADAT
GETDFS
GETDFU
OPNDAT
OPRDAT
SAVDAT
SETDAT

CLSIDX
ERAIDX
NMNODS
NOKEYS
OPNIDX
OPRIDX
SAVIDX

Index file update

Data file update

ADDKEY
DELKEY
UPDPTR

NEWREC
READAT
RETREC
WRTDAT

Index file search

Data locking

AFTKEY
BEFKEY
FRSKEY
GETKEY
LASKEY
NXTKEY
PRVKEY
SERKEY

FRELOCK
SETLOCK

Function index

ADD key to index file
ADDKEY

CHANGE data record number in index file entry
UPDPTR

CLOSE data file
CLSDAT

CLOSE index file
CLSIDX

COUNT entries in an index file
NOKEYS

COUNTS nodes (used and available) in an index file
NMNODS

COUNT records in a data file
GETDFS

COUNT records used in a data file

GETDFU

DELETE key from index file

DELKEY

DETERMINE results of lock/unlock request

LOKCOD

ERASE data file

ERADAT

ERASE index file

ERAIDX

FIND error code value

ERRCOD

FIND first entry in index file

FRSKEY

FIND two high-order bytes of data record number

DATVAL

FIND index entry equal or greater than key

SERKEY

FIND index entry following a specific key value

AFTKEY

FIND index entry in index file

GETKEY

FIND index entry preceding a specific key value

BEFKEY

FIND last entry in index file

LASKEY

FIND next available data record number

NEWREC

FIND next entry in an index file

NXTKEY

FIND previous entry in an index file

PRVKEY

INITIALIZE Access Manager

INTUSR

INITILIZE Access Manager for single-user environment

SETUP

LOCK data file or data record

SETLOK

OPEN corrupted data file
OPRDAT

OPEN corrupted index file
OPRIDX

OPEN data file
OPNDAT

OPEN index file
OPNIDX

READ data record from data file
READAT

RECLAIM available data file records
RETREC

RELEASE data file or data record lock
FRELOCK

SAVE data file updates
SAVDAT

SAVE index file updates
SAVIDX

SET two high-order byte of data record number
SETDAT

WRITE data record in data file
WRTDAT

Function syntax and description

ADDKEY (KEY%, FILE%, DLOCK%, KEYVAL\$, DRN%)
Add key value and data record number to index file.

AFTKEY (KEY%, FILE%, DLOCK%, KEYVAL\$, IDXVAL\$)
Find index entry following a specific key value.

BEFKEY (KEY%, FILE%, DLOCK%, KEYVAL\$, IDXVAL\$)
Find index entry preceding a specific key value.

CLSDAT (FILE%)
Close a data file.

CLSIDX (KEY%)
Close an index file.

DATVAL

Find two high-order bytes of a data record number.

DELKEY (KEY%, FILE%, DLOCK%, KEYVAL\$, DRN%)

Delete entry (key value) from an index file.

ERADAT (FILE%, DLOCK%)

Erase a data file.

ERAIDX (KEY%)

Erase an index file.

ERRCOD

Find the value of the error code.

FRELOCK (FILE%, DLOCK%, DRN%)

Release a data file or data record lock.

FRSKEY (KEY%, FILE%, DLOCK%, IDXVAL\$)

Find first entry in an index file.

GETDFS (FILE%)

Count the records in a data file.

GETDFU (FILE%)

Count the records used in a data file.

GETKEY (KEY%, FILE%, DLOCK%, KEYVAL\$)

Find specific entry (key value) in an index file.

INTUSR (PROGID%, ERROPT%, TIMEOUT%)

Initialize Access Manager.

LASKEY (KEY%, FILE%, DLOCK%, IDXVAL\$)

Find last entry in an index file.

LOKCOD

Determine results of a request to lock a data file or data record.

NEWREC (FILE%, DLOCK%)

Find next available data record number in a data file.

NMNODS (KEY%)

Count the number of used and available nodes in an index file.

NOKEYS (KEY%)

Count the entries (number of key values) in an index file.

NXTKEY (KEY%, FILE%, DLOCK%, IDXVAL\$)

Find next entry (key value) in an index file.

OPNDAT (FILE%, DLOCK%, FILNAME\$, RECLLEN%)

Open a data file.

OPNIDX (KEY%, IDXNAME\$, KEYLEN%, KEYTYP%, DUPKEY%)
Open an index file.

OPRDAT (FILE%, DLOCK%, FILNAME\$, RECLLEN%)
Open a corrupted data file.

OPRIDX (KEY%, IDXNAME\$, KEYLEN%, KEYTYP%, DUPKEY%)
Open a corrupted index file.

PRVKEY (KEY%, FILE%, DLOCK%, IDXVAL%)
Find previous entry (key value) in an index file.

READAT (FILE%, DRN%, BUFFER%)
Read a record from a data file.

RETREC (FILE%, DLOCK%, DRN%)
Delete data record and make available to NEWREC.

SAVDAT (FILE%)
Save updates made to a data file.

SAVIDX (FILE%)
Save updates made to an index file.

SERKEY (KEY%, FILE%, DLOCK%, KEYVAL\$, IDXVAL\$)
Find the index entry equal to or greater than a specific key value.

SETDAT (DRN%)
Set two high-order bytes of a data record number.

SETLOK (FILE%, DLOCK%, DRN%)
Request a lock on a data file or data record.

SETUP (NBUFS%, NKEYS%, NNSEC%, NDATF%)
Initialize Access Manager for a single-user environment.

UPDPTR (KEY%, FILE%, DLOCK%, KEYVAL\$, DRN%)
Change data record number in an index file entry.

WRDAT (FILE%, DRN%, BUFFER%)
Write a record in a data file.

Function parameters

BUFFER%
Pointer to data file I/O buffer area.

DLOCK%
Lock request code.

DRN%
Data record number.

DUPKEY%

Determines if Access Manager should handle duplicate keys in an index file.

ERROPT%

Determines if Access Manager should trap errors. Zero value disables trapping; nonzero enables.

FILE%

Data file number.

FILNAME\$

Data file name.

IDXNAME\$

Index file name.

IDXVAL\$

OUTPUT parameter. Contains key value returned from searching an index file.

KEY%

Index file number.

KEYLEN%

Number of bytes in an index file key.

KEYTYP%

Specifies alphanumeric or integer key values. Zero indicates alphanumeric; nonzero indicates signed, integer values.

NBUFS%

Number of buffers available for index files.

NDATF%

The maximum number of data files used in your application program.

NKEYS%

The maximum number of index files used in your application program.

NNSEC%

The number of 128-byte disk sectors in an index file record.

PROGID%

Unique identifier for each program calling the background server in a multiuser environment.

RECLN%

Data file record length.

TIMOUT%

Delay (in seconds) to determine if background server is operational.

Parameter use table

User error codes

- 21/AE Cannot write an index file record.
- 23/AG Cannot read an index file record.
- 24/AH No more directory space.
- 25/AI Index filename not found in directory during close.
- 30/AN KEY% parameter value out of range.
- 31/AO Invalid IDXNAME\$ parameter in OPNIDX function.
- 33/BA Index file number already assigned to open file.
- 34/BB KEYLEN% parameter in OPNIDX function greater than 48.
- 35/BC IDXVAL\$ parameter contains invalid data.
- 36/BD Data record number in ADDKEY function is zero.
- 39/BG KEYLEN% too long for number of assigned disk sectors.
- 40/BH Index file is corrupted.
- 43/BK Cannot reopen index file following SAVIDX function.
- 44/BL Attempt to close index file that is not open.
- 46/BN Attempt to use index file that is not open.
- 51/CC Cannot write data record into data file.
- 52/CD Data record number is zero or beyond end of file.
- 53/CE Cannot read requested data record.
- 54/CF No more directory space for data file.
- 55/CG Cannot close specified data file.
- 60/CL Data file number is out of range.
- 61/CM FILNAME\$ parameter contains invalid data.
- 63/CO Data file number is already assigned to a data file.
- 69/DE First byte of record in NEWREC function is not 0FFH.
- 70/DF Data file is corrupted.
- 73/DI Cannot reopen data file following SAVDAT function.
- 74/DJ Attempt to close data file that is not open.
- 76/DL Attempt to use data file that is not open.
- 83/EC Physical error on disk during index file I/O.
- 84/ED Attempt to update index file on Read-Only disk.
- 85/EE Attempt to update Read-Only index file.
- 86/EF Index file name references non-existent drive, or file control block is no longer valid.
- 87/EG Index file opened by another process in a multitasking environment.
- 88/EH FCB checksum error on index file close.
- 89/EI Incorrect or missing password for index file.
- 90/EJ OPNIDX attempts to create a second copy of file, instead of simply opening existing file. Should not occur. Report to Digital Research.
- 91/EK An illegal "?" in index file name.
- 92/EL Operating system open file limit exceeded during index file open.
- 93/EM Space in system lock list exhausted during index file operation.
- 113/GA Physical error on disk during data file I/O.
- 114/GB Attempt to update data file on Read-Only disk.
- 115/GC Attempt to update Read-Only data file.
- 116/GD Data file name references non-existing drive, or file control block is no longer valid.
- 117/GE Data file opened by another process in a multitasking environment.
- 118/GF FCB checksum error on data file close.
- 119/GG Incorrect or missing password on data file.
- 120/GH OPNDAT attempts to create a second copy of file, instead of simply opening existing file. Should not occur. Report to Digital Research.

- 121/GI An illegal "?" in data file name.
- 122/GJ Operating system open file limit exceeded during data file open.
- 123/GK Space in system lock list exhausted during data file operation.
- 145/IA Attempt to erase index file that is not open.
- 146/IB Cannot find index file name for ERAIDX function.
- 147/IC Invalid DLOCK% parameter value.
- 153/II Bad parameter value.
- 154/IJ Failure to call INTUSR function (bad language interface).
- 155/IK B-tree nodes in index file exceeds 65,535.
- 161/JA PROGID% parameter value out of range.
- 162/JB Cannot open queue for shared, multiuser module.
- 163/JC Cannot open queue (PROGID% out of range).
- 175/JO Attempt to erase data file that is not open.
- 176/K@ Cannot find data file name for ERADAT function.
- 177/KA DLOCK% parameter value is invalid.
- 183/KG Bad parameter value.
- 184/KH Failure to call INTUSR function (bad language interface).
- 209/MA Bad parameter value passed in SETUP function.

Data file/record lock request codes

- 0 Use this code to ignore the locking facilities. The LOKCOD function returns a zero.
- 1 Requests a shared record lock on the data record specified by the DRN% parameter.
- 2 Requests an exclusive record lock on the data record specified by the DRN% parameter.
- 3 Requests a shared file lock on the data file specified by the FILE% parameter.
- 4 Requests an exclusive file lock on the data file specified by the FILE% parameter.

Data file/record lock release request codes

- 0 No action. The LOKCOD function returns a zero.
- 1 Release the shared lock on the data record specified by the DRN% parameter.
- 2 Release the exclusive lock on the data record specified by the DRN% parameter.
- 3 Release the shared lock on the data file specified by the FILE% parameter.
- 4 Release the exclusive lock on the data file specified by the FILE%

parameter.

- 5 Release either the shared or exclusive record lock on the data record specified by the DRN% parameter.
- 6 Release all locks held by the calling program on the data file specified by the FILE% parameter.
- 7 Release all locks held by the calling program on all data files with a number greater than or equal to that specified by the FILE% parameter.

EOF