

Data Management Solutions
October 2001



© Copyright IBM Corporation 2001

IBM Corporation
Silicon Valley Laboratory
555 Bailey Avenue
San Jose, CA 95141
U.S.A.

Printed in the United States of America
10-01
All Rights Reserved

AIX, C-ISAM, IBM, the IBM logo, Informix and S/390 are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries or both.

Linux is a registered trademark of Linus Torvalds.

Microsoft and Windows NT are registered trademarks of Microsoft Corporation in the United States, other countries or both.

Sun and Solaris are trademarks of Sun Microsystems, Inc. in the United States, other countries or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Other company, product or service names may be trademarks or service marks of others.

References in this publication to IBM products or services do not imply that IBM intends to make them available in all countries in which IBM operates. Offerings are subject to change, extension or withdrawal without notice.



Printed in the United States on recycled paper containing 10% recovered post-consumer fiber.



GC27-1491-00



Low-overhead record management system

IBM Informix® C-ISAM® is a library of C-language functions that efficiently manages indexed sequential access method (ISAM) files. C-ISAM bypasses the overhead of a relational database management system (RDBMS) and provides fast and efficient access to records directly from applications. C-ISAM has been the industry-standard ISAM for UNIX® in the U.S. and Europe for more than 15 years, and Informix C-ISAM complies with the X/Open standard.

Technical overview

Organizations can use Informix C-ISAM to take advantage of the larger disks and increased capabilities of newer operating systems. Informix C-ISAM Version 7.25 supports 2GB files on systems that provide for files larger than 2GB. Before Version 7.25, corruption could occur if data, index, audit or transaction log files were larger than 2GB.

The Informix C-ISAM for ISAM file management is a highly reliable solution for quick data access for application development and low-level ISAM access without the RDBMS overhead. C-ISAM offers a wide variety of benefits, including:

- Quick data retrieval
- X/Open standards compliance
- Flexible indexing options
- Support for data integrity
- ISAM system maintenance
- File compatibility.

Quick data retrieval

With Informix C-ISAM, sifting through records to get to specific data is unnecessary. The C-ISAM B+ tree-index architecture makes data retrieval fast and easy. It uses index entries as keys that point to records, which allow specific pieces of data to be found without examining extra records. In addition, C-ISAM uses techniques to compress the keys for efficient index storage and processing. The reduced key size means fast response and outstanding performance for the end user.

Product availability

IBM Informix C-ISAM is available on the following platforms:

- Compaq/Digital UNIX
- Fujitsu-Siemens Reliant UNIX
- Hewlett-Packard HP-UX
- IBM AIX®
- Linux® for IBM S/390®
- Linux for Intel
- SCO OpenServer
- SCO UnixWare
- Sun™ Solaris™
- Microsoft® Windows NT®
- Other UNIX systems

Specifications

The following C-ISAM disk storage requirements are approximations:

Programs	355 kilobytes
Libraries	230 kilobytes
Include Files	18 kilobytes
Demo Files	29 kilobytes
Miscellaneous	79 kilobytes
Total	711 kilobytes

For more information

Please contact your IBM marketing representative or an IBM Business Partner, or call 1-800 IBM CALL within the U.S. Also, visit our Web site at ibm.com/software/data/informix



Index organization

C-ISAM indexes are organized in B+ trees that contain keys. Keys, the individual entries in an ISAM index file, are values from the data record that point to a record or multiple records, identified by record numbers. A key can be an employee number, for example. B+ trees also contain pointers that are arranged in a hierarchy. These pointers indicate another level in the index tree or a data record.

To find a record, the appropriate Informix C-ISAM function is executed and the key value for the search is supplied. The Informix C-ISAM function rapidly performs the search through the index. If it finds a match on the key value, it uses the pointer to read the data record. Informix C-ISAM then returns the data record to the program that requested the data.

C-ISAM datatypes

The C-ISAM data storage method is machine-independent, eliminating any confusion about machine-dependent data representation. The C-ISAM datatypes and their C-language equivalents are listed below:

C-ISAM Data Type	C-Language Data Type
CHARTYPE	char
INTTYPE	int
LONGTYPE	long
FLOATTYPE*	float
DOUBLETTYPE*	double
DECIMALTYPE**	typedef struct decimal dec_t

*stored machine independently
**C structure (struct decimal)

X/Open standards compliance

C-ISAM is the worldwide industry-standard UNIX file access method. X/Open standards-compliant C-ISAM is designed to pass the X/Open XPG3 test suite for ISAM and includes variable-length record support, as well as additional status-variable support, as specified in XPG4.

X/Open global language support

Informix C-ISAM 7.25 uses Informix global language support (GLS) to support multiple languages. The Informix implementation of GLS is based upon the X/Open XPG4 specification. GLS supports single-byte and multibyte characters to allow Informix C-ISAM to collate character strings, print dates and accept currency input in the rules and formats required by the country where the products are being used. X/Open-compliant GLS also provides worldwide support of database applications, so applications can be migrated to multiple languages while maintaining the same functionality.

Flexible indexing options

Informix C-ISAM indexing capabilities allow for any number of ISAM file indexes. Indexes can include up to eight parts, and each part can be a different datatype.

With Informix C-ISAM, you can build indexes on multiple fields, a single field or parts of a field. And you can specify ascending or descending order for data in different parts of the same index. Informix C-ISAM also allows duplicate and unique key values. After indexes have been created, Informix C-ISAM automatically updates the indexes while the records are being changed, simplifying data management tasks.

Data integrity

C-ISAM provides several mechanisms to preserve data integrity. Transaction management as well as logging and recovery features allow you to recover your data in the event of a system or disk failure, and locking mechanisms manage concurrent use of the data. A bcheck utility ensures the integrity of indexes.



Transaction management

Informix C-ISAM has built-in routines for transaction management that provide additional security for managing your data. With these routines, you can treat multiple operations as a single transaction by defining where the transaction should begin and where it should end (commit point). If necessary, you can roll back to the beginning of the transaction to ensure data integrity and avoid partially completed transactions.

Logging and recovery

Logging facilities help ensure data integrity by recording each transaction in a log file. If the work in progress must be rolled back, the ISAM files can be restored from the log file. The log file also provides a recovery mechanism if the ISAM file becomes corrupted or destroyed.

Locking

Informix C-ISAM has locking options that ensure data integrity when multiple users concurrently access the same row. When a transaction begins, the Informix C-ISAM routines that modify records also lock them. The records remain locked until the transaction is either committed or rolled back. These field-level and record-level locking capabilities help prevent data inconsistency caused by interference from simultaneous multiple users.

bcheck

The bcheck program is an Informix C-ISAM utility program that checks and repairs ISAM index files to maintain data integrity. This program can be run whenever the integrity of an index may have been compromised—generally because of a system failure while updating the index.

ISAM system maintenance

Informix C-ISAM streamlines ISAM system maintenance, and the Informix C-ISAM built-in routine reduces the time required to write maintenance tasks. After linking Informix C-ISAM with applications, the built-in routines perform the following maintenance tasks:

- Create indexed file systems
- Add and delete indexes
- Add, delete and modify records
- Reorder data records
- Rename and erase ISAM files
- Lock records or files.

File compatibility

Because Informix C-ISAM is the base of IBM Informix Standard Engine (SE), all applications built with Informix development tools can be used against appropriately constructed C-ISAM files that are compatible with Informix SE. Compatibility at the ISAM level allows data to be transferred between Informix products to combine the strengths of the various features. For example, Informix C-ISAM fast-file access and high-level forms utilize Informix SQL and Informix SE reports for easy access to complex data.

Technical configuration

Operating system file organization

A C-ISAM file consists of two operating system files: one to hold the data records and another to hold the index information. The data file extension is .dat, and the index file extension is .idx. These two operating system files are always used together as a single, logical ISAM file.