

INFORMIX-SE

Version 7.2

For UNIX® and Linux

INFORMIX-SE is an SQL-based database engine for small- to medium-range applications. It is the ideal solution for businesses that want the power of SQL without complex database administration requirements. With over a decade of proven reliability, INFORMIX-SE has a strong track record among UNIX database servers.

Organizations don't always have the MIS resources or extensive database expertise required to administer complex database systems. However, they still need powerful information management systems—and for them INFORMIX-SE is the low-maintenance, high-reliability database solution.

INFORMIX-SE Technical Overview

INFORMIX-SE provides excellent performance, data consistency, client/server capabilities, and standards adherence with only minimal need for database administrator (DBA) expertise, making it the ideal multiuser database system for small corporations or self-contained departments.

Performance

Database performance is essential for maintaining maximum throughput. INFORMIX-SE maintains excellent performance levels through cost-based optimization and flexible indexing options.

Cost-Based Optimizer with Data Distributions

INFORMIX-SE's cost-based optimizer will automatically determine the fastest way to retrieve data from a database table based on detailed information about the distribution of that data within the table's columns. The optimizer can collect and calculate statistics about specific data distributions during non-peak processing hours to avoid any potential bottlenecks. It can then pick a return path that has the least impact on system resources.

Flexible Indexing Options

INFORMIX-SE provides a variety of indexing options to achieve the best possible performance. Indexes can include

from one to eight fields from the record, and each field can be a different datatype.

INFORMIX-SE supports different B+ tree indexing options that include unique and clustered indexes. In a unique index, every record must have a different key value—preventing duplicate entries. A clustered index organizes the table in the physical order of the index, improving performance.

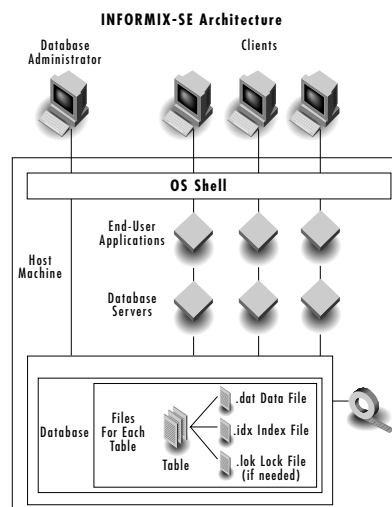


Figure 1: INFORMIX-SE uses the native OS file system for data storage and retrieval.

Data Consistency

INFORMIX-SE maintains data consistency via audit trails and transaction logging, and by establishing and enforcing locking procedures, isolation levels, and business rules.

Audit Trails and Transaction Logs

In the event of a system failure, audit trail facilities can be used to restore a table and all of the completed transactions from a backup copy of the table and the audit trail. Similarly, the entire database can be restored from a backup of the database and the transaction log file.

When an operation is unable to complete itself, the partially completed transaction must be removed from the database to maintain data consistency. To remove any partially completed transaction, INFORMIX-SE maintains a historical record of all transactions in the transaction log and automatically uses these transaction records as a reference to restore the database to the state prior to the transaction.

Locking and Process Isolation

Other important features for maintaining data consistency are locking procedures and process isolation. These security measures prevent other users from changing data that is currently being read or modified.

Locks

The database server prevents errors by imposing a system of locks. A lock is a claim or reservation that a program can place on a piece of data. The database server guarantees that, as long as the data is locked, no other database server process can modify it. When another program requests the data being modified, the database server either makes the program wait or turns it back with an error.

The throughput of transactions for a specific table can be affected by the locking strategy used for the table. Applications that use strategies of exclusive access to data might find that other database server processes are spending time waiting for access to the data. For this reason INFORMIX-SE provides several locking levels. The database server can place a lock on a single row, table, or database. Row-level locking is the default applied when the table is created. Table- and database-level locking are specified in the user's application.

Isolation Levels

The isolation level is the degree to which your read operation is isolated from concurrent actions of other database server processes—what modifications other processes can make to the records you are reading, and what records you can read while other processes are reading or modifying them.

INFORMIX-SE supports different levels of readonly isolation control. For example, a SQL SELECT request does not lock any records in the database, allowing full access to the data. This is true except when executing a SELECT FOR UPDATE request. While this type of request allows you to continue to access records, it does not allow other SELECT FOR UPDATE transactions to lock records already involved in update transactions.

Business Rules

Business rules enforce data consistency at the column level. They specify permissible data values, column defaults, and column-to-column relationships.

Business rules must be enforced by all applications that access or manipulate data. However, INFORMIX-SE does not depend on the application to implement these business rules. Instead, INFORMIX-SE enforces these rules independent of the user application. This centralization of responsibilities removes the burden from the user applications and guarantees adherence to business rules.

INFORMIX-SE supports integrity constraints, stored procedures, and triggers to enforce business rules.

Integrity Constraints

Informix's implementation of ANSI SQL-compliant integrity constraints ensures that information is not improperly deleted and that inserted data meets column specifications. INFORMIX-SE provides two types of integrity constraints: referential integrity and entity integrity.

Referential Integrity

Referential integrity allows users to define and enforce relationships between columns. For example, INFORMIX-SE's referential integrity guarantees that information about an entry in a master table is not deleted when the corresponding information still exists in the detail table. This prevents users from deleting a customer record if an order still exists for that customer in an order table.

Entity Integrity

Entity integrity enforces acceptable data values for particular columns. This check allows the DBA to specify a range of permissible values. Default values allow users to specify a default of any value that is compatible with the column datatype.

Stored Procedures

Stored procedures are SQL commands and program statements that are stored in the database as a named procedure to ensure consistent implementation of commonly used operations. Stored procedures maintain common, optimized application routines in the database rather than in the application program.

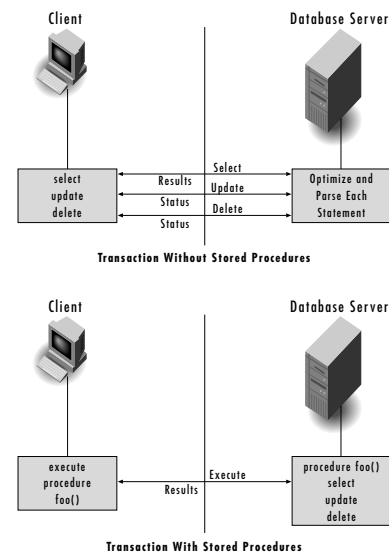


Figure 2: Stored procedures reduce the amount of communication between the client and the database server.

Stored procedures reduce the amount of network traffic for database operations because stored procedures can handle multiple tasks—such as insert, update, and delete—with a single command. After processing the request, the stored procedure returns only the results of the request rather than the numerous result sets for the individual queries.

Triggers

A trigger is an alternate method for invoking a stored procedure. Rather than the application program calling a stored procedure to enforce a business rule, a trigger can be defined, which will cause the database server to automatically execute a stored procedure (or SQL statement) when any attempt is made to INSERT, DELETE, or UPDATE a field in the table. Since triggers are stored in the system catalogs, the need for application programs to maintain redundant code has been removed. Likewise, consistent integrity constraint enforcement across all transactions is guaranteed since triggers cannot be bypassed.

Security

INFORMIX-SE provides two levels of access privileges to ensure database security. Database privileges control

access to the database and the privileges for creating tables and indexes in the database. Table privileges specify the operations that a user is allowed to perform against a specified table.

INFORMIX-SE supports alter, insert, and delete security at the table level while enforcing select and update security at the column level. Separate privilege statements are used to grant and revoke the appropriate access level to users. No separate database login is required since INFORMIX-SE applies security at the user's login level.

Stored procedures provide an additional security mechanism by establishing their own permissions, different from the data permissions. The owner of a stored procedure grants users the right to execute the stored procedure—which allows the user to perform all the SQL operations in the procedure—but restricts other access against the database. By using stored procedures to prohibit users from performing operations against the database, except through authorized stored procedures, DBAs can elevate security to the procedure level.

Client/Server Ready

With INFORMIX-SE, client/server connectivity is automatically built in—ensuring that any Informix application using connectivity libraries can transparently view and modify databases in either client/server or local mode.

INFORMIX-SE System Administration

INFORMIX-SE does not require UNIX kernel modifications for shared memory or the use of character-special devices—raw devices. Since INFORMIX-SE takes advantage of the native UNIX operating system and file manager, backing up databases and tables is as simple as backing up any other UNIX file—maintenance requires no special skills and no complicated utilities must be learned.

INFORMIX-SE supplies easy-to-use utilities that allow users to migrate INFORMIX-SE data in ASCII format to another INFORMIX-SE database or to Informix Dynamic Server™,

Informix's high-performance, parallel processing database server for the UNIX operating system.

Another important aspect of database administration is index verification. Should INFORMIX-SE ever indicate that an index on a table has been corrupted by a system failure, the bcheck utility can quickly verify the integrity of the index file and rebuild those indexes that have been damaged by the failure.

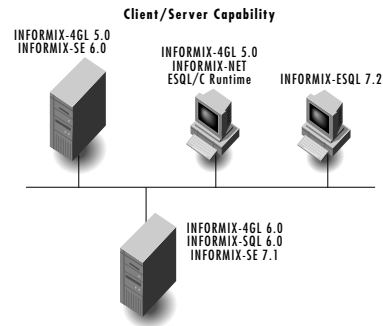


Figure 3: Informix tools and servers are built with connectivity libraries and require no additional networking products—although Informix 5.0 tools in a remote client/server mode require INFORMIX-NET to establish connection.

Standards Compliance

INFORMIX-SE meets the requirements for ANSI SQL-92 Entry Level with a few exceptions. (Notable exceptions include no support for delimited identifiers, or transactions that can be serialized.) It also supports numerous Intermediate and Full SQL-92 features, and some SQL3 features and functions.

INFORMIX-SE passes a subset of the National Institute of Standards and Technology (NIST) Federal Information Processing Standard (FIPS), number 127-2 test suite for Entry Level SQL, database sizing, and the FIPS Flagger.

NIST certifies conformance with ANSI and FIPS standards. The NIST SQL test suite is currently the only test suite available for validating ANSI SQL compliance.

Database Migration

If your information management requirements change in the future, INFORMIX-SE applications for UNIX can be easily migrated to Informix Dynamic Server.

Informix Dynamic Server is Informix's powerful, multithreaded database server designed to exploit the capabilities of both symmetric multiprocessor and uniprocessor architectures to deliver breakthrough database scalability, manageability, and performance. Informix Dynamic Server provides superior transaction processing, high availability, data integrity, mainframe-caliber administration, and multimedia capabilities—all within a single, client/server-ready package.

X/Open Native Language Support

Informix's implementation of native language support (NLS) is based upon the X/Open® XPG3 specification. INFORMIX-SE 7.2 NLS capability supports single-byte, eight-bit NLS platforms. This allows INFORMIX-SE 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 NLS also provides worldwide support of database applications, so applications can be migrated to multiple languages while maintaining the same functionality.

Note: These specifications were compiled on a Sun™ SPARCstation™ with SunOS™ 4.1.1. These specifications may be different for other machines and operating system configurations, and should be used only as guidelines. Dependent on disk space, system configuration and/or the operating system configuration.

INFORMIX-SE Version 7.2 Specifications	
Table Capacity	
Maximum number of rows per table	1 billion or prior disk space limitation
Maximum length of a row	32,767 bytes
Maximum number of columns per table	32,767
Maximum number of indexes per table	unlimited
Maximum number of columns per index	8
Maximum size of any index key	120 bytes
Field Capacity	
Maximum character field size	32,511 bytes
Maximum size of a quoted string	256 bytes
Access Capacity	
Maximum number of defined databases per INFORMIX-SE system	unlimited
Maximum number of defined tables per INFORMIX-SE system	unlimited
Maximum number of active users per INFORMIX-SE system	system dependent
Maximum number of open tables per INFORMIX-SE system	UNIX open files limit
Maximum number of tables in a join	UNIX open files limit
Maximum number of locks per INFORMIX-SE system, database, user	system dependent
Maximum number of open databases per user	1
Resource Consumption	
Memory Requirements Fixed memory for INFORMIX-SE database server process (re-entrant)	1.2 megabytes
User process virtual address space requirements (average varies by application)	180 kilobytes
Disk Requirements	
Disk space required for INFORMIX-SE	6 megabytes

About Informix

Informix Corporation, the world's database technology leader, provides innovative database products that enable the world's major corporations to attain competitive advantage. Informix, based in Menlo Park, California, is widely recognized as the technology leader for corporate computing environments ranging from small workgroups to very large parallel processing applications.

Informix's database servers, application development tools, superior customer service, and strong partnerships enable the company to be at the forefront of major information technology solution areas including data warehousing, OLTP, and Web/content management.

For more information about Informix products and services, please contact the sales office nearest you, or visit us on the Web at www.informix.com.

INFORMIX REGIONAL SALES OFFICES	
• Asia/Pacific	65 298 1716
• Canada (Toronto)	1 416 730 9009
• Europe/Middle East/Africa	44 181 818 1000
• Federal	1 703 847 2900
• Japan	81 3 5562 4500
• Latin America	1 305 265 7545
• North America	1 800 331 1763
	1 650 926 6300



4100 Bohannon Drive
 Menlo Park, California 94025
 1 650 926 6300
 World Wide Web: www.informix.com

©1998 Informix Corporation. All rights reserved. The following are trademarks of Informix Corporation or its affiliates: Informix®, Informix® Dynamic Server™, Universal Data Option™, and DataBlade®. The Informix logo is registered with the United States Patent and Trademark Office.