

IBM Informix OnLine Extended Edition Version 5.2 for UNIX and Linux

Highlights

- **Offers excellent reliability, creating a low total cost of ownership (TCO) at a reasonable price**
- **Provides fast response in low-to-medium transaction volume environments**
- **Features ease of administration**
- **Maintains data consistency at systems and transaction level without compromising ease of access and performance by large number of users**
- **Ensures high availability through multiple recovery mechanisms**
- **Supports industry standards such as Structured Query Language (SQL) 92, Open Database Connectivity (ODBC) and Java™ Database Connectivity (JDBC), which allows use by a wide variety of client development tools**

The requirements and priorities for information management solutions vary as widely as do the industries and institutions using relational database management systems (RDBMS). The prosperity of your business can be significantly affected by where you choose to store your business data.

For low-to-medium workloads, the IBM Informix® OnLine Extended Edition Version 5.2 database server offers superior online transaction processing support with the assurance of data integrity. This server also provides rich multimedia data management capabilities, with storage for a wide range of media such as documents, images and audio.

IBM Informix OnLine Extended Edition architecture

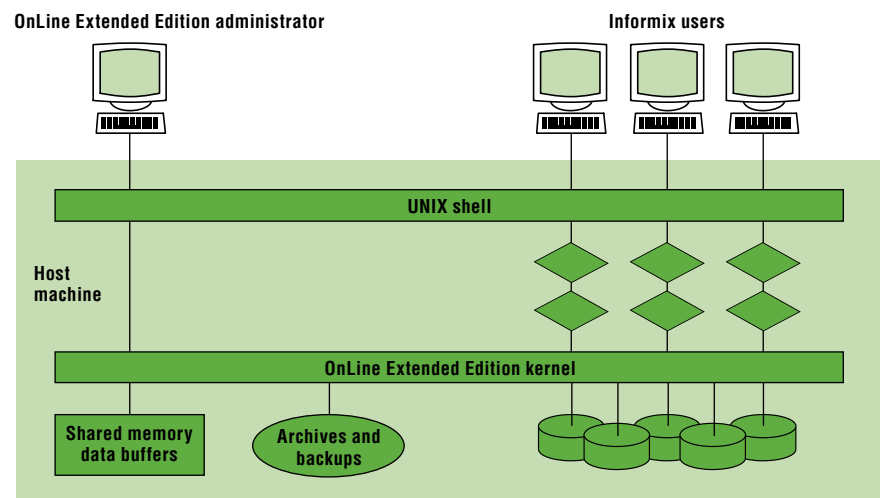


Figure 1. IBM Informix OnLine Extended Edition architecture: IBM Informix OnLine Extended Edition takes advantage of shared memory access and raw disk devices for data storage.

Store a wide range of media—documents, images, audio

Information management for demanding environments

With a proven track record of over 12 years, this reliable, embeddable database server offers the advantages of both the IBM Informix OnLine and IBM Informix-STAR, including:

- *High performance*
- *Data consistency*
- *Enhanced file-size support*
- *Industry-standard connectivity*
- *Distributed client/server*
- *Ease of administration*
- *Multimedia support*

Complementary products

Add-on products that extend functionality include:

- *IBM Informix Office Connect—Integrates seamlessly with Microsoft® Excel to provide an intuitive environment for developing and deploying Excel-based applications. It allows you to visualize Excel data effortlessly, without knowing SQL or the details of database structures. Additionally, it provides the capability to create a Web repository that allows sharing of spreadsheets in a safe and secure environment for workflow applications or simple precanned reports.*

- *IBM Informix 4GL—A fourth-generation programming language with facilities for producing database query and management using SQL, reports from a relational database or from other data sources, and form- and menu-based multiuser applications. It offers high performance in the production environment and integrates functionality to build complex applications. This easily portable language allows the optional use of a third-generation language and offers easy application maintenance.*

Pipe support

Informix OnLine Extended Edition also provides support for pipe connectivity. This allows the Rapid Application Development (RAD) tools—such as IBM Informix 4GL and IBM Informix SQL, Version 7.31 or later—to use pipes to communicate with Informix OnLine Extended Edition servers running on the same machine, instead of having to use a TCP/IP loopback connection with Informix-STAR or Informix Net. Customers who use Informix OnLine Extended Edition with the 7.31 tools therefore save the cost of purchasing Informix Net or Informix-STAR separately, and customers who are still using the Version 4.x tools can upgrade to Informix OnLine Extended Edition and take advantage of the many language enhancements in the Version 7.31 tools.

The addition of pipe connectivity in Informix OnLine Extended Edition offers Informix OnLine customers the opportunity to upgrade their legacy Informix 4GL Version 4.x applications to the benefits of the latest Informix 4GL Version 7.x:

- *GLS—A single version of the product is globalized and uses the latest internationalization techniques. A single binary supports different language settings, unlike Version 4.x, which required one binary for English and one for Asian Language Services (ALS).*
- *Server-friendly features—Support for NCHAR and NVARCHAR. (The previous release supports LongIDs.)*
- *Compatibility with Client Software Development Kit (CSDK) and the latest Informix servers*
- *Certification on the latest platforms*

Informix OnLine Extended Edition advantages

The IBM Informix OnLine Extended Edition Version 5.2 is a new product name for two mature, trusted products: IBM Informix OnLine and IBM Informix-STAR. The Informix OnLine Extended Edition establishes efficient storage for fast data access, buffers data in memory to minimize disk access, utilizes multiprocessor features to enable simultaneous access by different processors and determines the most efficient search strategy.

Industry-standard connectivity

The Informix Online Extended Edition supports the ODBC and JDBC industry standards. The Informix ODBC Driver is the Informix implementation of the ODBC 3.0 level 1+ standard and allows any ODBC-compliant application to connect to any Informix server. Informix JDBC Driver achieves optimal performance and robustness because it is a pure Java implementation and uses a native interface to connect to the database server.

Greater file-size support

High performance is essential for maintaining maximum throughput. To enhance performance, Informix OnLine Extended Edition provides support for greater than 2-GB chunk offsets and provides pipe connectivity. (Informix OnLine does not support file operations at greater than 2-GB offsets.) Today, almost all 32-bit platforms have support for files greater than 2 GB; hard disk drives are in the GB range and are continuing to increase rapidly. This limitation does not allow you to utilize the disk space beyond 2 GB unless it is partitioned. The Informix OnLine Extended Edition feature allows you to utilize the entire capacity of the hard disk drive without partitioning it into smaller logical devices and employing logical volume managers.

High performance

High-performance levels, which are essential for maintaining maximum throughput, are maintained through:

- *Raw disk management*
- *Shared memory, internal tables and buffer pool*
- *Multiprocessor features*
- *Cost-based optimization*

Raw disk management

Informix OnLine Extended Edition supports storage on both regular “cooked” operating system files, as well as on raw devices. For fast direct memory access (DMA), Informix OnLine Extended Edition achieves contiguous disk-space storage through raw disk devices. It uses raw disks for storage because the native UNIX[®] file system does not guarantee contiguous disk-space allocation. Raw disk-space allocation, conversely, makes it possible for Informix OnLine Extended Edition to create its own data storage system. Data storage on raw disk allows it to perform DMA. DMA writes data directly from memory to disk, bypassing intermediate UNIX buffering mechanisms that are required when data storage is performed through a standard UNIX file system. This process makes it faster to commit transactions to disk.

Shared memory, internal tables and buffer pool

All the important data is maintained in shared memory. Instead of being read from disk, the data is read from shared memory, which makes access faster. The shared memory consists of internal tables and a buffer pool. Internal tables monitor and manage high-demand resources such as buffers, locks and open data tables. The buffer pool comprises fixed-length pages that hold data during processing. When requested data is in the buffer pool, no disk I/O is required to access it. To read requested data that is not in the buffer pool, Informix OnLine Extended Edition copies the data from disk into a free buffer. The data is then available for processing.

Multiprocessor features

The Informix OnLine Extended Edition parallel-process sorting package streamlines processing by making multiple sorted runs in memory, and then merging the returned data into one result. The spin-and-test procedure eliminates unnecessary overhead from context switching. When two user processes need access to the same resource, this procedure allows the waiting user process to remain in the processor while retesting for access to the resource.

Cost-based optimization

The cost-based optimizer considers all possible query plans and compiles an estimate of the number of records to be examined and number of disk pages to be accessed for each plan. The optimizer then assigns a cost for each alternative and chooses the least costly plan to perform the data access.

The optimizer can also be instructed to return information about the chosen query plan to the application development tool prior to executing the DMA. This allows the database administrator (DBA) to evaluate the effect of the DMA request on current system resources and decide whether to proceed with the query or wait for a later time to execute the query.

High availability—when and where you need it

Informix OnLine Extended Edition's multiple recovery mechanisms allow recovery from hardware or system failures—with a minimal amount of application downtime.

Mirroring

Units of Informix OnLine Extended Edition disk space are paired with equal-sized mirror units. The mirror maintains a complete copy of the primary data that can transparently take the place of the primary disk should a hardware failure occur.

Data restore

A complete archive of the databases and a record of all transactions—both stored on tape—provides recovery from a hardware failure. An archive can be created while the databases are in use. If the host machine does not have the tape facilities required for an archive, any machine on the network with tape facilities can be used.

Fast recovery

Any time Informix OnLine Extended Edition is brought up after being stopped unexpectedly—as by a power failure—it automatically initiates a fast recovery. All committed transactions are preserved and all uncommitted transactions are rolled back to bring the system online in a consistent state without data loss.

Data consistency

Data consistency protects data integrity with built-in mechanisms— at the system and transaction levels— without posing unnecessary barriers to information access and smooth performance. Informix OnLine Extended Edition maintains data consistency using:

- *Transaction logging—To remove any partially completed transaction, a historical record of all transactions is maintained in the logical logs. This record is automatically used to restore the database to the state prior to the failed transaction.*

- *Internal consistency checking—Data-level checks can detect data inconsistencies that might be caused by hardware or operating system errors. If inconsistencies are detected, messages are written to the message log. To identify the cause of the inconsistency, Informix OnLine Extended Edition administrators can instruct users to set consistency-checking environment variables. These variables generate diagnostic output—for example, contents of shared memory when the inconsistency occurred.*

Locking and process isolation

Important features that prevent other users from changing data that is currently being read or modified include:

- *Locks—Four levels of locks are designed to prevent errors. 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.*

Informix OnLine Extended Edition also prevents deadlocks, a situation where two users have each locked data that the other user needs. For example, user A has locked one row and won't release it until user A can access the locked row that user B has. Informix OnLine Extended Edition detects a deadlock immediately and returns an error message to the second program to prevent a standstill.

The throughput of transactions for a specific table can be affected by the table's locking strategy. 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 OnLine Extended Edition provides several locking levels. The database server can place a lock on a single row, page, table or database.

Row- and page-level locking are specified when the table is created or altered; table- and database-level locking are specified in the user's application.

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

Isolation levels are in effect only for reads; they are not used for statements that insert, update or delete. Informix OnLine Extended Edition has the following four isolation levels:

- **Dirty read**, the simplest isolation level, provides no isolation at all. When a program accesses a row using a dirty read, it places no locks and respects none. For dirty reads, users are able to read all the data—committed or uncommitted.

- **Committed read** guarantees that Informix OnLine Extended Edition will read only rows that have been committed to the database. Before retrieving a row, the database server tests to see if an updating process has placed a lock on the row. If no locks have been placed, the database server accesses the row. The committed read isolation level ensures that the database server does not read uncommitted data because locks are placed on the rows that have been updated but not committed.

- **Cursor stability** places a lock on the latest-row read. Only one row is locked at a time; each time a new row is read, the previous lock is released. This isolation level ensures that the current record will not change while the program examines it.

- **Repeatable read** guarantees that the results obtained during a read operation will be identical if the read is repeated later in the same transaction. Not only does a repeatable read place shared locks on the records that were read, it also prohibits other users from adding records to the database or modifying records if they would have satisfied the criteria of the read request had they been in the database earlier.

Business rules

Business rules enforce data consistency at the column level. They specify possible data values, column defaults and column-to-column relationships. Business rules must be enforced by the applications that access or manipulate data. However, Informix OnLine Extended Edition enforces these rules independent of the user application.

Integrity constraints

The IBM Informix implementation of American National Standards Institute (ANSI) SQL-compliant integrity constraints ensures that information is not improperly deleted and that inserted data meets column specifications. Informix OnLine Extended Edition provides two types of integrity constraints:

- **Referential integrity** allows users to define and enforce relationships between columns. For example, information about an entry in a master table will not be 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** enforces acceptable data values for particular columns. This check allows the DBA to specify a range of permissible values. Default values allow you to specify a default of any value that is compatible with the column data type.

Stored procedures

SQL commands and program statements can be named and stored in the database. These stored procedures maintain common, optimized application routines in the database rather than in the application program.

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.

Security

Two levels of access privileges 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 OnLine Extended Edition supports *alter*, *insert* and *delete* security at the table level, while enforcing *select* and *update* security at either the table or column level. Separate privilege statements are used to grant and revoke the appropriate access level to users. No separate database login is required because Informix OnLine Extended Edition applies security at the user's login level.

Stored procedures provide an additional security mechanism by establishing their own permissions, different from the permissions for the data that they access. 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.

Distributed client/server capability

Informix OnLine Extended Edition users can transparently manipulate multiple databases at different locations. A two-phase commit protocol keeps all transaction interactions consistent with multiple database servers.

With the addition of Informix-STAR, Informix OnLine Extended Edition can join, view and modify multiple databases at different locations as if they were one common database. With Informix OnLine Extended Edition's two-phase commit protocol, data consistency is guaranteed at a global level.

In transactions that include modifications of multiple servers at different locations, one of the servers assumes the role of coordinator for the global transaction. The coordinator routes the transaction work and tracks the progress of the transaction for all Informix OnLine Extended Edition database servers that participate in the transaction.

The two-phase commit protocol occurs in two stages. First is the precommit phase, in which the coordinator directs each participating database server to prepare to commit the transaction. The coordinator then waits for each participating Informix OnLine Extended Edition server to respond, indicating whether or not the participant can commit the transaction. If every participant indicates it can commit the modification, the coordinator makes the decision to commit the transaction. If any participant indicates it is unable to commit the modifications, the coordinator makes the decision to roll back the transaction.

After the decision is made, the second phase of the protocol begins. During the post-decision phase, the coordinator directs each database server to either commit the changes or roll back the transaction. If the coordinator directs the participants to commit the transaction, it waits to receive acknowledgments.

Software specifications

System capacity	IBM Informix OnLine Extended Edition Version 5.2 systems per computer (dependent on available system resources)	255
	Maximum number of accessible remote sites	Machine specific
Table capacity	Maximum number of rows per table	4 278 190 080
	Maximum number of data pages per table	16 777 216
	Maximum number of data bytes per table (excludes binary large object [BLOB] data)	33 822 867 456
	Maximum length of a row	32 767 bytes
	Maximum amount of BLOB data	2631 bytes
	Maximum number of columns per table	32 767
	Maximum number of indexes per table	100
	Maximum number of columns per index	16
Field capacity	Maximum character field size	32 511 bytes
Access capacity (per Informix OnLine Extended Edition system)	Maximum number of defined databases	100
	Maximum number of defined tables	16 777 216
	Maximum number of active users	1000
	Maximum number of open tables	32 000
	Maximum size of BLOB table	int
	Maximum size of user table	32K
	Maximum size of logical log partition	1 000 000 pages
Product availability	IBM Informix OnLine Extended Edition Version 5.2 is available on the following platforms and operating systems: <ul style="list-style-type: none">• Sun Solaris• IBM AIX®• Hewlett Packard HP-UX• SCO Open Server• Linux for Intel®	



For more information

Please contact your IBM Marketing Partner, or call 1 800 IBM CALL (1 800 426-2255) within the U.S. Also, go to our Web site at **ibm.com/software/data/informix**

© Copyright IBM Corporation 2003

IBM United States
Software Group
Route 100
Somers , NY 10589
U.S.A.

Produced in the United States of America
01-03
All Rights Reserved

IBM, the IBM logo, AIX, DB2 and Informix are trademarks of International Business Machines Corporation in the United States or other countries, or both.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Microsoft, is a trademark of Microsoft Corporation in the United States, other countries, or both.

Intel is a trademark of Intel Corporation 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 and service names may be trademarks or service marks of others.

References in this publication to IBM products and services do not imply that IBM intends to make them available in all countries in which IBM operates.

