

Top 10 Differences: Infobright Enterprise Edition and Infobright Community Edition

Infobright develops two versions of our innovative analytic database technology: Infobright® Enterprise Edition (IEE) which is offered under a commercial license, and Infobright Community Edition (ICE), an open source, freely available product under a GPLv2 license. Both are designed to provide excellent performance for query and reporting with minimal administrative effort.

IEE includes additional features designed to support the operational requirements of most enterprises, including enhanced performance, insert/update/delete capabilities, expanded loader options and other features. In addition, the IEE subscription includes a committed level of support from Infobright's service organization, product training and services, product warranty and indemnification and other advantages.

Let's examine the Top 10 differences in more detail:

1. Significant Query Performance Differences

While both IEE and ICE provide significantly better query performance than row-oriented databases such as Oracle, SQL Server and MySQL, IEE query performance can be as much as 50 - 500% faster than that of ICE.

The major difference between IEE and ICE is in the performance characteristics of the core engine. The workload manager in IEE utilizes all available resources including CPUs and memory to optimize the order of data access, and rapidly decompress and read data in parallel with query execution. This results in query performance increases of 50 - 500%. The workload management system also provides the ability for multiple concurrent queries to be balanced across all system resources. Introduced in IEE release 3.5, these new workload management features included parallel scanning and execution of individual queries, and automatic query throttling for load balancing data loads and query execution. These additions further enhance IEE performance.

With ICE, an individual query is limited to a single core to process the query. Other queries will share the resource, or use other available processing cores on a shared basis.

Key Advantage – Complex queries, particularly those that use filters and require access to a lot of data, benefit significantly from IEE's multi-threaded and multi-core capability.

2. DML Support

IEE supports the SQL Data Manipulation Language allowing the ability to INSERT, UPDATE, DELETE data in an IEE database.

IEE fully supports the SQL DML language commands. This includes INSERT, UPDATE, and DELETE along with ALTER TABLE. Getting data into ICE is done via the Infobright loader using the LOAD DATA INFILE command. Any changes to the database requires a reload of all the data.

Key Advantage – For organizations that require the ability to modify information in the database, IEE's support of DML will save significant work for database administrators. DML provides another alternative (although slower) to loading data into IEE as well as the ability to more easily manage the data previously loaded. Without DML support, ICE requires a reload of the data even if only a small portion of the data has changed.

3. Faster Data Loading and More Options

IEE offers expanded loader options for faster load speed or MySQL compatibility.

IEE includes a parallel, multi-threaded version of the Infobright high-speed loader that can achieve load speeds of over 300GB per hour and supports both ASCII and binary load files. IEE users also have the option of using the native MySQL loader, which provides additional flexibility for loading tables where embedded functions, transformations, and simulated auto-increment may be required.

New! Distributed Load Processor (DLP) available as add-on with IEE 4.0

DLP is designed for companies needing near real-time loading of very large volumes of data. With DLP, load speeds of over 2TB per hour into a single database table can be achieved. DLP delivers load speeds that scale linearly, as it compresses data and builds the Knowledge Nodes using multiple remote servers, then transfers the compressed data to the IEE database. Load performance gains result because CPU-intensive work is distributed across multiple machines, and query performance is enhanced by offloading this work from the database server.

ICE includes a single-threaded version of the Infobright loader, achieving load speed of up to 55GB per hour for ASCII files. ICE does not include the option of using the MySQL loader.

Key Advantage – Organizations requiring faster or more frequent loading of large volumes of data will benefit from IEE's faster loading options. For compatibility with MySQL, the option of using the native MySQL loader is also an advantage with IEE.

4. Replication and High Availability

IEE has support for statement-based replication, providing high availability, greater user concurrency and the ability to backup the database without shutting it down.

MySQL statement-based replication is supported with IEE. It doesn't require you to buy or install any additional software. You physically set up a slave server or multiple slaves and configure Infobright's replication (which is MySQL style) on both servers appropriately to begin replication.

Replication is also a good option for high availability configurations. It can also be used to address read-intensive scaling requirements. ICE does not support replication.

Key Advantage – IEE provides enterprise-class features such as replication and high availability that are required in production environments.

5. Concurrent Query While Load / DML Support

IEE executes queries while data loading, or DML statements are executing.

IEE supports the ability to query the database while data loads (or insert/update/deletes) are occurring, while maintaining consistency of the data and database. ICE does not have this feature. This means that while users are loading data into ICE, database queries wait until the load of data is completed and committed to the database.

Key Advantage – Many organizations need to query the database while loading new data, as they require continuous business operations or frequent data loads. Query While Load capabilities provides uninterrupted database accessibility to users.

6. Hadoop Connector

DLP, available with IEE, includes a connector for Hadoop.

The Distributed Load Processor, available only with IEE, also includes a Hadoop connector that provides a simple way to extract data from the Hadoop HDFS file system and load it into IEE at very high speeds.

Key Advantage – The Hadoop connector and DLP make it much easier to leverage Hadoop's large-scale distributed batch processing benefits with Infobright's fast ad-hoc analytic capabilities.

7. ALTER TABLE/RENAME TABLE Support

IEE's support of alter and rename table allows simplicity and flexibility in managing the database.

IEE support of ALTER TABLE allows you to modify the table structure by adding, changing, and deleting columns and attributes as well as renaming the table. ICE supports RENAME TABLE only.

Key Advantage – IEE reduces the administrative overhead to manage a frequently changing database model. If there is a need for frequent changes or updates of the database schema, the administrator will require both ALTER TABLE and RENAME TABLE support. IEE gives DBAs complete control over how to update or modify table definitions. The overall capability to modify the data is still available in ICE, but it requires the DBA to drop the table and all its data and recreate the table with the new schema.

8. Temporary Table Support

IEE supports the use of Temporary tables.

IEE fully supports temporary tables using the CREATE TABLE SQL Statement. A temporary table is visible only to the current connection, and is dropped automatically when the connection is closed. This means that two different connections can use the same temporary table name without conflicting with each other or with an existing non-temporary table of the same name. (The existing table is hidden until the temporary table is dropped.) ICE does not support temporary tables.

Key Advantage – The use of temporary tables is another way to save time and effort, particularly when dealing with several queries that return very large and similar result sets. When a developer finds themselves dealing with the same subset of table data several times in a single session, it's often faster and more economical to obtain it once and store it in a temporary table, rather than either saving the data in a programming structure (such as an array or hash) or repeating a complex join several times.

9. Enterprise-Level Support Options

IEE subscriptions provide a low, all-in-one price for end-user organizations. OEM pricing and program benefits are offered to qualifying software and technology vendors.

IEE offers three annual subscription options. Each plan includes a product license and updates, documentation, a defined level of support, availability of training and much more. Details can be found at <http://www.infobright.com/Services/Enterprise-Subscriptions-&-Services/>. For software companies and technology providers interested in becoming OEMs, custom pricing and partner program support is offered. Learn more at http://www.infobright.com/Customers_Partners/OEM_Partners/.

ICE does not include an option for contracted Infobright support, but ICE users can post questions and issues to the ICE forums at <http://www.infobright.org/Forums> and receive prompt replies.

Key Advantage – Enterprise customers receive a full subscription license and all the associated benefits and support commitments. Software and technology vendors can take advantage of custom pricing, OEM program benefits and a commercial license.

10. Commercial License Agreement, Indemnification and Warranty

IEE is licensed under a commercial license agreement while ICE is licensed under GPLv2. IEE also provides IP indemnification and a limited warranty.

ICE is licensed under the GPLv2 open source license. As such, it is generally advisable that companies not use ICE with commercial solutions that they are planning to distribute or sell. If they do, they may be required to release their own commercial code, both binary and source, under the GPLv2 open source license. However, if a company is building an open source solution licensed under the GPLv2 license, then it is acceptable for them to use and distribute ICE with their own open source solution, provided the entire bundle or stack is compatible with the GPLv2 open source license.

IEE on the other hand is licensed under a commercial license and therefore, a company is not at risk in distributing IEE with other commercial off the shelf software or hardware solutions, provided they have obtained the appropriate commercial license from Infobright. With IEE, enterprises are also provided with indemnification protection against intellectual property infringement.

About Infobright

Infobright's high-performance database is the preferred choice for applications and data marts that analyze large volumes of "machine-generated data" such as Web data, network logs, telecom records, stock tick data and sensor data. Easy to implement and with unmatched data compression, operational simplicity and low cost, Infobright is being used by enterprises, SaaS and software companies in online businesses, telecommunications, financial services and other industries to provide rapid access to critical business data. For more information, please visit www.infobright.com or join our open source community at www.infobright.org.

CONTACT US

Corporate Headquarters:

47 Colborne Street, Suite 403
 Toronto, Ontario M5E1P8 Canada
 Tel. 416 596 2483
 Toll Free 877 596 2483
 Fax 416 352 5674

Sales :

Toll Free 877 596 2483 x225
 416 596 2483 x225
info@infobright.com

European Office:

Block 4, Suite 204
 Europa Business Centre
 Harcourt Road
 Dublin 2 Ireland
 Tel: +353 (0)87 743 710

www.infobright.com

www.infobright.org

Twitter: Infobright

YouTube: infobrightdb