

Infobright Community Edition (ICE) 3.1.1

The Infobright Community Edition (ICE) Open Source data warehouse is designed to run high-performance ad hoc, complex queries on very large data sets without the need of manual tuning, data partitioning, or index creation.

ICE achieves query speed and performance by creating a “Knowledge Grid” of the data during load, and running queries through the associated Infobright Optimizer. When data is loaded, it is compressed to approx 10:1 and stored in “data packs.” The Knowledge Grid automatically creates a highly compact set of metadata, designed on rough set theory, which stores information about the relationship between data packs and statistical information about the contents. When a query is initiated, the Infobright query Optimizer uses the Knowledge Grid to determine which data packs need to be decompressed. The Knowledge Grid eliminates the need for specialized partitioning of data and indices.

FEATURES

SQL Support	Infobright Community Edition (ICE) has full ANSI SQL-92 with some SQL-99 extensions including full support for VIEWS and stored procedures
Flexible Schema Support	ICE supports all schema designs
Industry Standard Interfaces	ICE supports standard database interfaces, including ODBC, JDBC and native connections
Supported APIs	Supported APIs include C, C++, C#, Borland Delphi (via dbExpress), Eiffel, SmallTalk, Java (with a native Java driver implementation), Lisp, Perl, PHP, Python, Ruby, REALbasic, FreeBasic, and Tcl
Concurrent Users	ICE supports up to 500 database users with up to 24 concurrent queries (assumed to be complex analytic queries)
Operating Systems	Windows XP (32-bit), Red Hat Enterprise Linux 5 (64-bit), Red Hat Enterprise Linux 5 Advanced Server (64-bit), Debian ‘Lenny’ (64-bit), CentOS 5 (64-bit), Fedora 9 (32-bit), Ubuntu 8.04 (32-bit)
Processor Support	ICE supports 32-bit and 64-bit Intel and AMD x86-based servers
High Performance Data Loader	A custom ICE loader is included with ICE for high performing data loads, with parallelization across multiple tables

BENEFITS

Small Data Footprint	On average, ICE compresses data at a ratio of 10:1 but can achieve much higher compression levels. Clients have reported ratios in the 30-40:1 range. Variability of compression ratios depends on the type of data loaded into ICE
High Scalability	ICE excels with large data volumes and can scale up to 30 TB (compressed down to less than 3TB of disk space) on a single server implementation
Column Based Advantage	When a traditional database retrieves a row of data, it must read the entire row off disk. ICE identifies and reads the relevant columns off disk
Off-the-shelf Hardware Support	ICE supports industry standard Intel and AMD x86 servers
Simplified Administration	MySQL Administration tools are included to easily manage ICE
Low Maintenance	ICE automatically maintains all Knowledge Grid structures across every column in the database
Load and Go	ICE is load-and-go using your existing schema and does not require materialized views, data partitioning, or the implementation of indices
Supported Load Formats	The ICE loader can be used to load text files
BI Flexibility	ICE is integrated with MySQL, through which virtually all BI tools are supported including Cognos™, Business Objects™, SAS™, Pentaho™, JasperSoft™, MicroStrategy™, etc.
