



How well does it improve database performance?

On average, the performance improvements yielded by DBXten in a typical application are that it:

- decreases total database size by a factor of 25
- generates these indexes more than 60 times faster
- speeds up insertion of rows by a factor of 3
- improves data retrieval times for single and concurrent queries by a factor of 10 or more

Rather than store the data records in separate rows, DBXten stores groups of data rows as single database rows, using high levels of column compression in the process. To provide fast search and retrieval capabilities, DBXten then takes advantage of this storage mechanism to generate fast and compact indexes on the stored data. As a result, the DBXten database extension dramatically speeds up the ingestion and indexing of data sequences while reducing storage space requirements too. Since much more data can now reside in in-memory data caches, queries are also speeded up considerably.

In 2010 IBM provided us with access to a large System x configuration at their San Mateo Innovation Center to benchmark DBXten with the latest IBM Informix DBMS. The results of this testing are in: http://www.barrodale.com/docs/Impact_of_DBXten_on_IBM_Informix_Performance.pdf

Encouraged by these results, we then performed similar tests with Informix on a modest PC workstation at our offices. The results of these tests are reported in: http://www.barrodale.com/docs/Using_IBM_Informix_on_a_workstation_with_DBXten.pdf

We then repeated the PC testing using DBXten and PostgreSQL – a free and open source DBMS. The results of these tests are reported in: http://www.barrodale.com/docs/PostgreSQL_Performance_is_Enhanced_with_DBXten.pdf

The table below summarizes the results contained in each of these three White Papers, by indicating the savings factors due to DBXten, averaged over the various tests performed for each configuration. The test data employed was typical oceanographic model data obtained from a UK web site, and the tables ranged in size from 55 million rows to 1,447 million rows. The three White Papers referenced above contain extensive details of the computations performed.

DBMS with DBXten	Storage Space savings	Load Time savings	Query Time savings
IBM Informix (on System x)	25 x	3 x	10 x
IBM Informix (on Office PC)	15 x	2.5 x	10 x
PostgreSQL (on Office PC)	60 x	10 x	10 x

Table Description: Approximate improvement factors due to inclusion of DBXten – averaged over several test runs with different tables, various queries, and single and multiple concurrent users.