

What could you do if data management were no longer a limitation?

Any-Memory™

- NitroEDB™ integrates easily into server-side and client-side applications providing outstanding embedded data management performance and volume handling capabilities across the data management spectrum from extreme speed in-memory to extreme volume very large database on-disk deployments.

Benefits

- Enables unique product features
- Extreme product performance
- Unlimited product data volume
- Zero admin data management
- Minimizes hardware requirements
- Minimizes time-to-market

Features

- Extreme insertion rates
- Extreme in-memory speed
- Extreme VLDB data volumes
- Extreme SQL query speeds
- Linear CPU scalability
- Linear memory scalability
- Linear disk storage scalability
- Patent pending N-Tree indexing
- Simple memory vs. speed tuning
- Compressed in-memory indexes
- Embedded and/or C/S deployment
- API access to all features
- SQL, OO, and native APIs
- Automatic OOEI generator
- ODBC and JDBC support
- C, C++, Java and Delphi/Kylix APIs
- Multi-Platform availability
- Cross-Platform data compatibility
- GUI schema editor
- GUI data manager
- Expert integration support

NitroEDB, with its Any-Memory™ capability, is the only extreme performance relational data management product that supports the full spectrum of data management needs from extremely fast in-memory applications to huge-volume very large database (VLDB) on-disk server systems.



NitroEDB's speed, data volume handling, hardware requirements, and value characteristics are far superior to any of the commodity commercial or open-source data management products available – **often exceeding 100s of times the volume-performance characteristics!**

NitroEDB's capabilities and characteristics are the result of a unique combination of its patented extreme-performance N-Tree™ thread tuned indexing technology, highly tuned cache management techniques, SQL query optimization and storage management algorithms, and over one hundred staff-years of data management product development effort.

NitroEDB is easy to use and simple to deploy, and using its SQL, ODBC or JDBC interfaces developers will minimize the time-to-market of their products. Additionally, for those times when application tuning is required, NitroEDB provides developers with an automatic object oriented encapsulation interface (OOEI) generator, and a native API with access to all of its fundamental data management capabilities. NitroEDB can be used to augment existing enterprise data management capabilities by acting as a database accelerator cache, or it can be deployed as a fully embedded zero-administration stand-alone data management solution.

The United States DoD, DoE, other federal and state agencies, and commercial enterprises are using NitroEDB technology in mission critical applications.

NitroEDB History

Development of NitroEDB began as a system known as Sage (AdaSAGE/SageST) in 1983 at the United States Department of Energy's (DoE) Idaho National Engineering and Environmental Laboratory (INEEL) in Idaho Falls, Idaho. The INEEL, DOD, and DoE and their contractors used Sage as the foundation for numerous mission critical applications representing millions of lines of software. While in the government domain, over \$30 million and nearly 100 staff-years of effort were invested in the research and development of the technology.

Beginning in early 1998, the INEEL's Sage development team joined Nitro Data Systems (later becoming NitroSecurity) and created NitroEDB, the premier embedded database that so uniquely satisfies OEM requirements that it was awarded a patent in 2002 for its indexing technology. DoE, other federal and state agencies, and commercial enterprises are currently using NitroEDB technology in mission critical applications.

Advantages of NitroEDB

High insertion rates - In most implementations, NitroEDB will sustain insertion rates from 10 to 1,000 times faster than other database systems using similar database schema definitions and populations. For certain types of data, such as the very common "time series" data, insertion rates can remain at a constant high level rate in tables with populations increasing from thousands up to billions of records. It also supports insertion burst rates as high as five times faster than the standard rates. This insertion rate characteristic contributes to other very desirable benefits. It allows the developer to increase the number of indexed attributes resulting in improved query results while still being able to fulfill insertion performance requirements.

Large terabyte storage - NitroEDB is capable of storing many terabytes of data representing billions of records while sustaining excellent insertion and query rates. The physical distribution of information is user definable to suit the hardware environment.

Supports heavy multiple client request conditions - NitroEDB has been benchmarked to sustain acceptable insertion and query requests from many thousand concurrent users. One of the outstanding features of our patented N-Tree indexing technology is its parallel thread support.

Supports very fast "in-memory" hyper-speed requirements - The ability for the same database to be defined as totally "in memory" or a multi-terabyte Very Large Database (VLDB) system gives distinct advantages for distributed or multi-platform applications.

"Real Time" statistical and Meta data determinations - The ability to retrieve this type of data is designed into the patented indexing system. The capability of rapidly returning statistical information results in reduced storage requirements (where other database systems need to store separate meta data), the elimination of automatic background meta data gathering tasks, query path optimization based upon the actual data counts, and the retrieval of many "real time" statistical measurements.

Optimized SQL queries - The SQL processor for NitroEDB has been optimized to rapidly return a variety of query results as seen in the report from a standard Wisconsin benchmark test described later in this paper. In addition, those queries characteristic of the security, network traffic, and other similar arenas using such constructs as GROUP BY, COUNT, DISTINCT, SUM, MAX, MIN and AVG have been implemented using the "real time" statistical information retrieval built into the indexing system. It is not uncommon for customers to find that queries that previously took from tens of seconds up to many minutes will be performed in just a few seconds or even in sub-second times using NitroEDB.

Smaller storage requirements - NitroEDB requires less than one half the persistent disk storage of the most popular enterprise database systems

Numerous client/server configurations - NitroEDB supports multiple instances of three different types of database connection modes. These include being directly connected or linked, acting as a server, or accessing another server as a client. Connections may also be secured encrypted connections.

Circular cache - A definable circular cache allows NitroEDB to perform insertion and pruning/archival activities, almost doubling the efficiency of the separate operations.

Selective Indexing - This provides the ability to selectively include or exclude one or more record's (tuple's) field from its index. This can be used to provide "in-memory" speeds for VLDB tables of hundreds of millions or billions of records. Such indexes may be directly accessed from the API or through the use of SQL hints.

Triggers / Embedded methods / Remote procedures

Full SQL DDL and DML language support - NitroEDB supports ANSI standard SQL with many extensions. early 1998



What could you do if data management were no longer a limitation?

NitroSecurity, Inc.
230 Commerce Way
Suite 325
Portsmouth, NH 03801

PH: 603.766.8160
FX: 603.766.8169

Web: www.nitrosecurity.com
Email: info@nitrosecurity.com

