

Calpont InfiniDB[®] Enterprise

InfiniDB Enterprise Key Features

- 100% Columnar architecture
- Multi-threaded, distributed database operations
- Scale-out MPP
- UDFs for In-database Analytics
- Integrated map reduction operations
- Shared Nothing data cache
- Real-time compression
- Automatic partitioning
- Partition drop
- High speed data loader
- Query snapshot read
- Performance diagnostic utilities
- Transactional Insert/Update/Delete
- Linux and Windows platforms

“InfiniDB delivers important benefits for our research and analytics processes. Compared to our legacy MySQL based data warehouse, we can process data an order of magnitude faster and use a lot less computing resources.”

Andy Lucas
CTO, Cognitive Match

The Database Conundrum - Exploding Data and Deep Analytics Latency

In today's competitive business climate, organizations have become more analytical in the way they look at business operations, customer interactions, and expansion into new markets. Executives push for increased visibility and accountability to key business metrics, driving greater operational integration of processes and data systems. Data capture is exploding as e-commerce, blogs and forums, social networking, mobile media, and online customer interactions have become mainstream. As a result, line of business managers, analysts and IT departments have matured in their use of Business Intelligence (BI), analytics, and data warehouse systems.

BI and analytics have become the most strategic weapons for managing business. The promise of near real-time analytics is quickly becoming a necessity for many organizations, but is often compromised by the enormous volume of data needed for analysis at a deep, granular level. Traditional database and data warehouse technologies are unable to keep pace with both data volume and rapid decision making. This has led to the emergence of new, analytic-specific database technologies.

Calpont InfiniDB Enterprise - The Next Generation Data Platform for Deep Analytics

Fast and Scalable

InfiniDB Enterprise is a high performance, scalable database built from the ground up for analytics, BI, data warehousing and read intensive applications. Designed for data professionals who need to augment or replace their traditional RDBMS technologies, InfiniDB is a column-oriented, massive parallel processing (MPP) technology that processes data queries extremely fast and scales linearly with any type of storage hardware, growing as analytics needs and data grow.

The Columnar Advantage

InfiniDB's column-oriented architecture overcomes query limitations that exist in traditional row-based RDBMS systems by storing, managing, and querying data based on columns rather than rows. Only the necessary columns in a query are accessed, reducing I/O activities by circumventing unneeded rows. This makes InfiniDB ideal for BI and analytics because query processing and results response against large and very large data occur in much less time than standard OLTP and OLAP databases. InfiniDB also has high query response predictability, providing dependability in meeting service-level requirements imposed by the business.

InfiniDB Enterprise Advantages

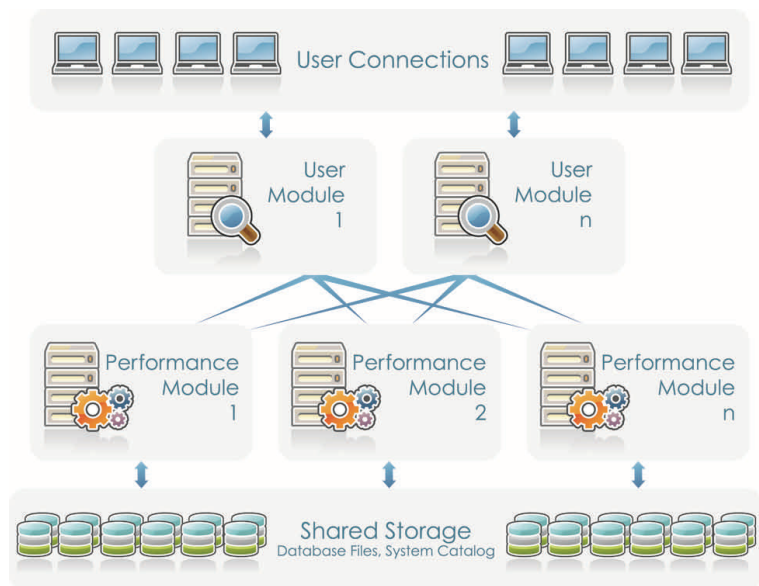
- Query performance 10x to 1000x over row-oriented technologies
- High speed data loading—millions of rows per second
- Very large data scaling on commodity hardware
- Self learning and managing
- No indexes, materialized views or summary tables required
- High concurrency
- High availability
- Affordable, simple licensing model—per software node

InfiniDB delivers fantastic price versus performance for our MySQL data intensive application. With our recent FCC partnership, we are faced with a significant data collection task in analysing end user broadband experience. InfiniDB takes away the scalability concerns we had with our previous MySQL database, whilst at the same time requiring little or no redevelopment effort."

**Sam Crawford,
Founder, SamKnows**

Scale-out MPP Architecture

InfiniDB scales out in MPP fashion across multiple hardware machines to achieve linear increases in overall performance. Adding inexpensive commodity hardware to an InfiniDB configuration allows the database to increase processing power, cutting response times in half with each successive addition of a new node. The result is the ability to adjust for growing data volumes, increased user activity, or the desire to meet better overall performance goals. What's more, scale out occurs dynamically, without the need to take InfiniDB offline as nodes are added.



Distributed Performance

InfiniDB uniquely distributes and processes queries using a two-tiered design. The User Module handles concurrency scaling, breaking down SQL requests and distributing the various parts to one or more Performance Modules. The Performance Module is responsible for storing, retrieving, and managing data, as well as processing block requests for query operations. The User Module assembles all the query results from the various participating Performance Modules to form the complete query result set that is returned to the user. User Modules and Performance Modules can be configured in any combination to meet user, data, and performance requirements.

In-Database Analytics

InfiniDB enables creation of fully parallel and distributed user defined functions (UDFs) that run as an integrated operation within the storage engine. This enables implementation of custom analytic functionality that leverages the full benefits of InfiniDB's integrated map reduction capability, distributing the custom functionality across all available cores within the distributed layer.

Calpont InfiniDB® Enterprise

The High Performance, Scalable Data Platform that enables Deep Analytics

To learn more, please visit www.calpont.com/infinidb or call 214-618-9500

To download a free trial, please visit www.calpont.com/tryinfinidb