

ZL NoSQL DB™: Built for Analytics

ZL NoSQL DB™ is an excellent NoSQL database for large scale data persistence that efficiently addresses both storage and data access. for today's world of enterprise data.

BACKGROUND: THE NEW ANALYTICS PARADIGM

Today's analytics challenges are defined by vastly diverse data sources, rapidly evolving data types, and enormous scale. The traditional RDBMS approach, while proficient at certain processing tasks, is typically not efficient or flexible enough when used alone for today's world of enterprise data.

As digital footprints continue to grow, businesses are practically drowning in data proliferation. Previously, the most pressing challenge for the enterprise was to manage data for storage, operational, and compliance needs. However, these reactive requirements perhaps understated the most important asset: "People Data." Most corporate communications with co-workers, external customers, partners, and suppliers occur over email, documents, instant messaging, and social media. With this data being predominantly unstructured or multi-structured in nature, cutting-edge enterprises are now looking for ways to strategically mine this people-driven content to get a true 360° view of customers and internal personnel.

Challenges Facing Analytics Initiatives Today

The technology for traditional RDBMS was simply built to solve a different class of problems: specifically large-volume

transaction processing and various OLTP applications. Traditional data warehouses, which build large repositories of data generated by OLTP applications, extract-transform-load (ETL) the data for classic business intelligence: sales, inventory, and product reports. These reports are generated from "Business Data," or data generated by business applications like POS and ERP systems. The emergent trend in forward-thinking enterprises has been to supplement this business data with analysis of people data as well as basic "Machine Data" generated by system and event logs. The limitations of traditional RDBMS in terms of scale, total cost of ownership, and performance issues with large queries have led to a new breed of NoSQL databases.

Today, the enterprise wants to analyze people, business, and machine data together in a more cohesive, contextual way. This requires a data store to scale to the massive volumes of unstructured data like email, IM chats, and social media content while being able to process and analyze data streams in-line with a built-in text and graph computing engine. This provides the ability to perform Social Network Analysis (SNA) and graph searches. While all these technologies are available individually in different types of databases, the best solution thus far would involve stitching systems together, resulting in degradation of functionality

and tedious movement of data across components.

Introducing ZL NoSQL DB™

To analyze and derive insights from People Data, Business Data, and Machine Data together in one seamless system, ZL Technologies has engineered ZL NoSQL DB™, a new class of NoSQL analytical persistence engine.

ZL NoSQL DB™ is an excellent NoSQL database for large scale data persistence that efficiently addresses both storage and data access. ZL NoSQL DB™ is a core component of ZL Enterprise Analytics™ applications, as it provides the infrastructure for new use cases requiring the mining and correlation of people data, business data, and machine data.

It allows analytics that are flexible, fast, scalable, and deeply integrative: enabling a high-level analytics perspective that looks at the entire living enterprise. The core premise of ZL NoSQL DB™ is the adaptability to simultaneously handle billions of rows and thousands of columns: all of which can be stored and processed efficiently.

Data Type Support

ZL NoSQL DB™ supports multiple data types that are well suited for analytics and data processing. Out-of-box support includes multi-valued columns (string Array and Map) in addition to standard String, Integer, Long, Boolean, Float, Double, and Date data types. To significantly improve storage efficiency, two additional data types – Dictionary String and Dictionary – are also supported.

Hybrid Columnar Store

A pure columnar database – which stores tables as collections of compressed columns – possess certain deficiencies in range selection. Such data stores would need to decompress entire sets of columns before allowing users to filter down to a range of rows that might be of interest. While a columnar store provides the best compression efficiency, ZL NoSQL DB™ comes with a highly efficient storage and access mechanism which provides the right balance between pageable and columnar store models.

Dynamic Compression

An adaptive compression algorithm within ZL NoSQL DB™ lets the data decide the compression algorithm employed based on both the data type and the data distribution of that column within that page. Compression decisions are based on the data spread at the time of writing to the page. Compression and paging techniques provide storage and access efficiencies.

Data Processing

Text searching capabilities are built into the ZL NoSQL DB™ as an index component, which gives the enterprise the power to search both unstructured and structured content in context. Result sets can be filtered and sorted on multiple columns, which can then be used to build histograms in a single pass.

Search

The power of the text engine provides additional context and insight when an enterprise mines unstructured or people data. The embedded search capabilities in the database provide a 360° view for unparalleled text mining. This is not possible with existing technologies in the market without stitching together multiple products -- with exorbitantly high cost, and at the expense of efficiency.

Security and Manageability

Security and manageability are provided by the built-in component called BigDB and ZL Vault: core components of the ZL UA platform. BigDB functionality provides centralized cataloging, which carries information on all tables, indexes, and associated partitions. BigDB's integration with ZL Vault supports a myriad of application needs and promotes greater manageability.

Global Enterprise Dictionaries

ZL NoSQL DB™ supports the notion of Global Enterprise Dictionaries, thus allowing the platform to truly be the master search engine for enterprise data. Global Enterprise Dictionaries can be used across various NoSQL tables and can be built from available data sets, such as the Wikipedia corpus. It can be used across enterprise applications and grow with the company. The core infrastructure requirements remain intact, even though dictionaries may vary depending on the industry or application.

Dynamic Tables

To support near real-time data processing and data availability for search, ZL NoSQL DB™ provides a Dynamic Table feature that can create an implicit index which can be generated many times: either fully or partially. It can then be merged with the table's main index allowing the user to search seamlessly across multiple generations of the data in the table or

KEY FEATURES

ZL BIG DATA API

Make the most of existing data and resources

DYNAMIC TABLES

Seamlessly search across multiple generations of the same table

GRAPH ENGINE

Built-in graph engine makes it easy to explore relationships across diverse data types

TEXT ENGINE

Built-in text engine navigates the human content of enterprise unstructured data

MULTIPLE INDEX SUPPORT

Dynamic tables allow multiple index generation for real-time searches

VAST DATA TYPE SUPPORT

Native support for string array, Map, string, integer, long, boolean, float, double, date, and more

GLOBAL DICTIONARIES

Easily re-purpose dictionaries created from available data sets

OS AGNOSTICISM

Use whatever operating system works best for the enterprise

KEY BENEFITS

ARCHITECTURE

Engineered directly into the ZL UA information governance platform

SPEED

High-speed data streaming for faster insight

NO MACHINE AFFINITY

Processing power of existing infrastructure is maximized

HIGH AVAILABILITY

Less downtime for deeper insights into big data

EASE OF USE

Built to help streamline IT management and analytics initiatives

SECURE CONTROL

Consolidated control with log tracking and other features for enterprise security

LEADING SUPPORT

Enterprise-class software and processing with support services that can meet the demands of business, 24/7

query processed data in near-real time fashion.

ZL NoSQL DB™ Use Cases

The ZL UA platform inherently meshes with ZL NoSQL DB™, which enables several example use cases that tap into the global repository of enterprise content.

LOG ANALYTICS

ZL UA generates real time transactional logs that are distributed across many servers, which are streamed in real time and merged in a central location using ZL NoSQL DB™. IT administrators get the power to search across many log files from a single user interface that can proactively help in identifying issues before they happen.

COMMUNICATIONS ANALYTICS

In a large organization, the corpus of unstructured data generated from human-tohuman communication can easily total several hundred GB... or more. For all this data to be analyzed in concert, inline processing can be leveraged with ZL NoSQL DB™ to derive meaningful patterns and insight from emails, files, IMs, social media, and more.

QUERY ACCELERATION

Database query latency grows problematic with scale. However, ZL NoSQL DB™ has a built-in graph that minimizes lag and has the capacity to store very large-scale graphs, and can filter based on granular criteria. The result is functional graphs with powerful cross-data business insights that show relationships and their strength.

NETWORK ANALYSIS

The modern enterprise has the potential to generate billions of data points and events every day about employees. By

doing Social Network Analysis (SNA) on emails and other forms of human-generated content, the business can better understand who knows whom, who knows what, where there are bottlenecks, and how information flows through the organization.

Conclusion

ZL NoSQL DB™ is a robust, large-scale database that is both storage and access efficient, built for the demands of scale and data diversity. With ZL NoSQL DB™ at the core of the ZL Enterprise Analytics™ product, the enterprise can address an entirely new class of massively integrative analytics use cases by joining People Data, Business Data, and Machine data under the same roof.