



Business intelligence Using Microsoft SQL Server

White Paper – Business intelligence Series

Scalable Systems Inc.

Contents

Executive Summary 2

Overview 2

OLAP 3

 ROLAP 3

 MOLAP 3

 HOLAP 3

 OLAP Analysis 3

Microsoft Business Intelligence Technologies 3

Integrate Enterprise Data Sources 4

 Optimize Database Performance 4

Building and Managing Sophisticated BI Solutions 5

 Developer Productivity 5

 Implement Best Practice Solutions 5

Increase Reporting Flexibility 5

Use a Single, Unified Tool 6

Monitor Data Warehouse Resources 6

 Scalable Analytics 6

 Scalable Reporting 7

Conclusion 7

Reference 7

Executive Summary

Business Intelligence systems allow enterprises to improve business decisions and achieve better business performance by using analytical data about customers, suppliers, and internal business operations etc. Business Intelligence (BI) is an integral part of Implementing Data Warehouses (DW) for decision supports. DW & BI involve:

- Extraction, Transformation and Loading from many sources such as Customer Relationship Management (CRM), Supply Chain Management (SCM), and Enterprise Resource Planning (ERP) systems, and other applications.
- Standardizing metadata information in repositories such as data warehouses and data marts.
- Data analysis using analytical tools that allow processing of large amounts of business data for decision support.

Data growth in enterprises offers various analysis patterns to help organizations grow by managing business critical data and taking the right decisions at the right time. Many business questions or situations are analyzed in order to achieve the business goals of an enterprise with the help of “Business Intelligence”. The latter provides extensive analysis patterns on the enterprise data that help reduce operating costs and provide for increased profitability, productivity, sales, and service, thereby enhancing decision making capabilities within a short time. Business Intelligence Models are based on multi-dimensional analysis and key performance indicators (KPI) of an enterprise.

Overview

Enterprises are realizing that they can only succeed in this competitive market by proactively identifying business trends and opportunities and by responding to new customer demands. To meet these challenges, enterprises need to gain actionable insight into the business so that they can make intelligent, informed decisions and contribute to business growth. Enterprises might have the following routine questions:

BI for Finance:

What are the actual net income, expenses, and profit for this month, quarter, and year?

BI for Accounts:

What is the total sales amount for this quarter?

BI for Purchase:

Who are the potential vendors to be contacted for various products purchase?

BI for Production:

What and how many products are manufactured in each production unit monthly, quarterly and yearly?

BI for Sales:

what is our Products sale in each area today, weekly, monthly, quarterly and yearly?

BI for Quality:

what is the number and name of defective products in various units weekly, monthly, quarterly and yearly?

BI for Service:

What is the level of our customer satisfaction and how we can improve?

These BI questions are related to why, what, how, when, and OLAP reports are the answers by means of dashboards, reporting, score cards, balance score cards that are helpful in executive decisions.

What is OLAP?

OLAP is an acronym for Online Analytical Processing which helps enterprises to take advantage of historical data. OLAP cubes provide the insight into data and help the enterprise management to take decisions in an efficient manner. Technically, OLAP cube allows for the analysis of data across multiple dimensions by providing a multidimensional view of aggregated, grouped data. OLAP applications include sales and customer analysis, budgeting, marketing analysis, production analysis, profitability analysis and forecasting etc.

ROLAP

Relational Online Analytical Process (ROLAP) provides multidimensional analysis of data stored in a Relational database (RDBMS).

MOLAP

Multidimensional OLAP (MOLAP) provides the analysis of data stored in a multi-dimensional data cube.

HOLAP

Hybrid OLAP (HOLAP) is a combination of both ROLAP and MOLAP which can provide multidimensional analysis of data stored in a multidimensional database and in a relational database (RDBMS) simultaneously.

OLAP Analysis

OLAP Analysis helps enterprises store data in a CUBE in a multi-dimensional format for analysis on a business subject area with drill-up and drill-down functionalities. It also provides for slicing and dicing of the information necessary for taking the right decision for business productivity.

Microsoft Business Intelligence Technologies

Microsoft SQL Server provides Business Intelligence (BI) offering that is based on a scalable data platform for data integration, data warehousing, analysis, and reporting, involving powerful and intuitive tools that end users can use to access and analyze business information. Microsoft SQL Server platform enables enterprises to:

- Unify storage and access for all data across the enterprise.
- Build and manage sophisticated BI solutions.

The specific technologies of SQL Server that form the basis of this powerful BI offering are described in the following table.

Component	Description
SQL Server Database Engine	A scalable, high-performance data storage engine for extremely large volumes of data making it an ideal choice for consolidating business data from across the enterprise into a central data warehouse for analysis and reporting
SQL Server Integration Services	A comprehensive platform for extract, transform, and load (ETL) operations that enables the population and synchronization of your data warehouse with data from the disparate data sources that are used by your business applications throughout the organization
SQL Server Analysis Services	Provides an analytical engine for Online Analytical Processing (OLAP) solutions, including business measure aggregation over multiple dimensions and key performance indicators (KPIs), and for data mining solutions that use specialized algorithms to identify patterns, trends, and associations in business data
SQL Server Reporting Services	An extensive reporting solution that makes it easy to create, publish, and distribute detailed business reports both within and outside the enterprise

Integrate Enterprise Data Sources

SQL Server data source views enable integration of data and access to data that originates from heterogeneous enterprise-wide data stores such as SQL Server, Oracle, DB2, and Teradata. These views also provide an OLAP store of enterprise scale. SQL Server Integration Services support extraction of data from diverse data sources of existing business applications. Therefore, unifying the data in all of your enterprise data sources is easy to accomplish regardless of whether you want to build an abstraction layer through a data source view or use an ETL process to synchronize a dedicated data warehouse for analysis and reporting.

Additionally, through support for Web services and the Microsoft .NET Framework, SQL Server 2008 supports interoperability with multiple platforms, applications, and programming languages so that investment in new and existing systems can be maximized by integrating and connecting the disparate data sources. Support for existing and emerging open standards such as HTTP, XML, SOAP, XQuery, and XSD, further facilitates communication across the extended enterprise systems.

Optimize Database Performance

SQL Server 2008 includes a high-performance relational database engine that enables one to build highly effective data warehouse solutions. Innovations such as query optimizations for star schemas and tools to help one tune indexes and data structures make SQL Server a natural choice for a heavily queried data warehouse. With the introduction of the MERGE Transact-SQL statement in SQL Server 2008 developers can more effectively handle common data warehousing scenarios such as checking whether a row exists and then executing inserts or updates. Additionally, an extension to the GROUP BY clause enables users to define multiple grouping in the same query. Grouping Sets produce a single result set that is equivalent to a UNION ALL of differently grouped rows, which makes aggregation querying and reporting easier and faster.

SQL Server 2008 supports partitioned tables to help one optimize the performance and management of large tables, and with new support for partitioned table parallelism, one can significantly optimize a data warehouse by using partitioned tables.

Building and Managing Sophisticated BI Solutions

As organizations demand ever increasingly complex analytics, the need to be able to build and deliver effective BI solutions quickly while reducing the management overhead of the BI infrastructure has become a major consideration. SQL Server 2008 includes innovative tools that increase developer productivity and manageability, which enables faster capitalization on new analysis and reporting capabilities even as it reduces administrative overhead.

Developer Productivity

SQL Server 2008 simplifies the development of business intelligence solutions. BI developers benefit from easy-to-use utilities and tools that increase control and automate routine, time-consuming tasks. They can thus use the productivity features of SQL Server 2008 to create effective analysis and reporting solutions more quickly than ever so that your organization can take advantage of them sooner.

Implement Best Practice Solutions

Enabling developers to build solutions more quickly is effective only if those solutions are optimally designed. To help ensure the best possible performance and correct functionality, SQL Server 2008 includes the following development environment features that promote best practices and help developers create effective analysis solutions:

- A consistent development environment for all BI solutions, including Analysis Services, OLAP, and data mining applications.
- Built-in support for the full development lifecycle, including design, build, debug, and deploy operations; and support for team-based development through integrated support for source control.
- A number of intuitive designers and wizards that make it easy to create Analysis Services solutions quickly.
- An attribute relationship designer that includes built-in validations to help create optimal dimension designs.
- A dimension editor that has been slimmed down to provide better productivity and to ensure that the presence of parent-child relationships is automatically detected.
- A cube designer that has been streamlined and improved to provide better detection and classification of attributes along with identification of member properties.
- Aggregation of individual partitions which enables you to optimize measures from different periods or areas.
- The aggregation designer has a new algorithm to help create initial aggregations. The aggregation designer is optimized to work with usage driven aggregations. You can now look at the aggregations that have been created and add to those aggregations or remove them. Intelligent support is provided to help in merging existing and new aggregation designs.

Additionally, SQL Server 2008 provides AMO warnings to alert developers when their design breaks one of over 40 best practices. These warnings are integrated into real-time designer checks and provide a non-intrusive way for developers to detect potential problems with their design.

Increase Reporting Flexibility

Reporting is a significant element of any BI solution, and business users are demanding

increasingly complex reports. SQL Server Reporting Services provides the following features to make it easy to build reporting solutions:

- A Visual Studio-based report development interface in Business Intelligence Development Studio that developers can use to build, debug, and deploy reports.
- A business-focused report development tool named Report Builder that business users can use to create and deploy reports.
- A wide range of data display structures, including tables, matrices, lists, and charts.

Use a Single, Unified Tool

SQL Server 2008 provides DBAs with SQL Server Management Studio, a single, unified management tool that provides integrated management of Analysis Services, Reporting Services, Integration Services, and multiple versions of SQL Server from the same interface for increased DBA productivity, flexibility, and manageability across extended SQL Server implementations.

Monitor Data Warehouse Resources

SQL Server 2008 includes performance data collection and warehousing, which enables monitoring and reporting of resources across your data services solution.

Scalable Analytics

The premise of Online Analytical Processing (OLAP) is that instant access to accurate information enables end users to answer even the most complex questions at the speed of thought. Thus, the aim to continuously excel in providing even faster query times and data refresh rates is a priority during the development process of any SQL Server Analysis Services release, an aim that naturally has been driving the release of SQL Server 2008 Analysis Services.

SQL Server 2008 includes Analysis Services that enable you to drive broader analysis with enhanced capabilities, including complex computations and aggregations. Analysis Service provides enterprise-scale performance through:

- A flexible caching model. With Analysis Services, you can control how data and aggregations are cached to optimize query performance while maintaining an acceptable level of latency between the cache and its underlying data store.
- Declarative attribute relationships. In an Analysis Services dimension, you can explicitly declare relationships between attributes in a hierarchy. This enables Analysis Services to pre-generate aggregations when a cube or dimension is processed, which improves runtime query performance.
- Block computation. Block computation eliminates unnecessary aggregation calculations (for example, when the values to be aggregated are NULL) and provides a significant improvement in analysis cube performance, which enables users to increase the depth of their hierarchies and complexity of computations.
- Write-back to MOLAP. Analysis Services 2008 removes the requirement to query ROLAP partitions when performing write-backs, which results in huge performance gains.
- Scale-out Analysis Services. A single read-only copy of an Analysis Services database can be shared between many Analysis Servers through a virtual IP address. This creates a highly scalable deployment option for an Analysis Services solution.

- Execution plan persistence. SQL Server 2008 provides functionality to lock down query plans so that, to the maximum extent possible for correctness, the query plans survive server restart, server upgrade, and production deployments. This ensures consistent optimal performance of queries against SQL Server data.

Scalable Reporting

For many organizations, getting the right information to the right people at the right time is a significant challenge. SQL Server 2008 provides a high-performance reporting engine for processing and formatting reports along with a complete set of tools for creating, managing, and viewing reports. An extensible architecture and open interfaces enable easy integration of reporting solutions in diverse IT environments.

You can generate reports from multiple diverse data sources, including SQL Server, DB2, and Oracle, without first building a centralized data warehouse. You can deliver reports throughout the organization both internally and externally through the simple deployment and configuration capabilities that are provided by Reporting Services. This enables users to easily create and share reports of any size or complexity. You can also deliver reports to customers and suppliers easily by posting reports over the Internet.

Reporting Service provides support and the ability to control server behavior with memory management, infrastructure consolidation, and straightforward configuration through a centralized store and an API for all configuration settings.

Conclusion

In today's struggling economy with competition controlling and shrinking the space for business, it is important to make sound business decisions based on complete data. With proper Business Intelligent implementation, businesses can make decisions and feel comfortable that they are provided with the proper tools and data needed to believe in their decisions. Without the correct Business Intelligence solution even well planned and well executed data warehouse architectures can fail. Honestly it is only when one can access this information rapidly and easily in order to take appropriate action or make the right decision on what is actionable that business success is achieved.

Reference

1. <http://msdn.microsoft.com>
2. <http://technet.microsoft.com>

About Scalable Systems:.

Scalable Systems is a global software consulting, development and IT outsourcing company providing both offshore and onshore software solutions and integration services to business enterprises around the globe. Scalable Systems has proven expertise in encompassing low cost, but high quality and reliable software solutions and services in areas like Data Management, Business Intelligence, Content Management and Application Development.

Scalable Systems Inc

525 Milltown Road, Suite 303
North Brunswick, NJ 08902, USA
Tel : (732) 993 4320
Fax: (732) 909 2732
Email: info@scalable-systems.com
Web: www.scalable-systems.com

Copyright © 2008 Scalable Systems. All Rights Reserved.

While every attempt has been made to ensure that the information in this document is accurate and complete, some typographical errors or technical inaccuracies may exist. Scalable Systems does not accept responsibility for any kind of loss resulting from the use of information contained in this document. The information contained in this document is subject to change without notice.

Scalable Systems logos, and trademarks or registered trademarks of Scalable Systems or its subsidiaries in the United States and other countries.

Other names and brands may be claimed as the property of others. Information regarding third party products is provided solely for educational purposes.

Scalable Systems is not responsible for the performance or support of third party products and does not make any representations or warranties whatsoever regarding quality, reliability, functionality, or compatibility of these devices or products

While every attempt has been made to ensure that the information in this document is accurate and complete, some typographical errors or technical inaccuracies may exist. Scalable Systems does not accept responsibility for any kind of loss resulting from the use of information contained in this document.

This page shows the publication date. The information contained in this document is subject to change without notice.

This text contains proprietary information, which is protected by copyright. All rights are reserved. No part of this document may be photocopied, reproduced, stored in a retrieval system, transmitted in any form or by any means, or translated into another language without the prior written consent of Scalable Systems Inc.

This edition published August 2008