International Journal of Computer Techniques – Volume 12 Issue 3, May - June - 2025

PARALLEL AND DISTRIBUTED DATABASE SYSTEM FOR BUSINESS INTELLIGENCE: AN EXPLORATIVE REVIEW

[1]Damingo, Levi A.S., [2]Elliot, Kizzy N., [3]Urang, Awajionyi S., [4]Diri, Grace O.

[2]Department of Computer Science and Informatics, and Cybersecurity, Federal University Otuoke, Nigeria.

[1,3,4]Department of Computer Science, Igantius Ajuru University of Education, Port Harcourt, Nigeria.

levi.damingo@iaue.edu.ng, elliotnk@fuotuoke.edu.ng, awajionyi@gmail.com, grace.diri@iaue.edu.ng

***Abstract*—In an ever-changing world of vast information collection and transmission, the value of meaningful, actionable business information for informed decision making cannot be overstated, and the need to make sense of BIG Data is always increasing. As a solution to the various issues that businesses face, business intelligence has developed as a veritable tool in data management for smart data analysis and decision making. This paper thus takes an exploratory look at the concepts of business intelligence in relation to database management in a parallel and distributed database environment, considering its benefits and strategies in light of the massive influx and growth in data available for informed business decision making to optimize business operation and promote the profitability of the business and efficiency of its operations.**

***Keywords*—Database management, Parallel and Distributed Database System, Business Intelligence, Exploratory review**

I. INTRODUCTION

Massive volumes of data ("big data") are being collected at an unprecedented scale and rate by modern industrial, government, and academic organizations. Many businesses keep track of client interactions, product sales, outcomes from online advertising initiatives, and other forms of data. Scientists are gaining access to huge amounts of data thanks to powerful telescopes in astronomy, particle accelerators in physics, and genome sequencers in biology [1]. The ability to execute timely and cost-effective analytical processing of such massive datasets in order to derive deep insights is now a critical success factor. These information can be used to drive automated ad placement processes, improve customer relationship management, and lead to key scientific advances [2].

Business intelligence (BI) refers to a collection of methodologies, frameworks, and instruments that turn unstructured data into information that is relevant to and useful for conducting business analyses [3]. BI encompasses business-centric processes and methodologies that can be used in a variety of high-impact applications, including e-commerce, market intelligence, healthcare, and security, in addition to the underlying data processing and analytical approaches. Big data analytics are a set of sophisticated and original data management, analysis, storage, and visualization tools that have been developed in response to the more recent explosion of data in order to support applications that are too large (from terabytes to Exabyte in size) and complex (from sensor to social media data) to be effectively supported by the earlier technologies. Big data analytics can provide firms with a competitive advantage and lead to commercial benefits such as more effective promotion and increased revenue [4].

A. Overview of Business Intelligence

Business intelligence (BI) is a technology-driven process for data analysis and information delivery that aids managers, employees, and executives in making wise business decisions. In order to make the analytics results available to business users for operational decision-making and strategic planning, organizations collect data from internal IT systems and external sources, prepare it for analysis, run queries against the data, and create data visualizations, BI dashboards, and reports [5]. Business intelligence combines business analytics, data mining, data visualization, data tools and infrastructure, and best practices to assist organizations in making decisions that are more data-driven. In actuality, you can tell if you have modern business intelligence when you can use your organization's data to drive change, eliminate inefficiencies, and respond quickly to supply or market changes. Modern business intelligence solutions place a high value on governed data on trusted platforms, flexible self-service analysis, empowered business users, and speed to insight.

Business intelligence (BI) systems provide historical, present, and predictive perspectives of corporate processes. Reporting, online analytical processing, analytics, data mining, business performance management, benchmarking, text mining, and predictive analytics are all common functions of business intelligence technology. Business intelligence is intended to aid in improved business decision-making. According to [6], business intelligence systems can also be referred to as a decision support system (DSS). Though the terms business intelligence and competitive intelligence are sometimes used interchangeably because they both support decision making, business intelligence analyzes mostly internal, structured data and business processes, whereas competitive intelligence gathers, analyzes, and disseminates information with a topical focus on company competitors. Competitive intelligence can be considered a subcategory of business intelligence.

It is one of the resource-intensive applications in the in the recent past that consists of several data warehouses that were built by obtaining decision-support information from organization-wide databases. Through the execution of appropriate queries on the business processing and transactional databases, the data warehouses are updated frequently. The user-end interface of business intelligence known as online analytical processing (OLAP), is made to show end users multi-dimensional graphical reports. Data cubes are used in OLAP and are produced by multidimensional queries executed on a variety of data warehouses. An OLAP application also retrieves data from data warehouses, organizes it into extremely complex multidimensional data cubes, and displays it to users through customized GUI dashboards [7].

II. Literature Review

Traditional database systems are adapting to the new status quo, while large-scale data analytics solutions such as MapReduce and Dryad are gaining popularity for big data analytical workloads [8] [9]. Industry giants like Teradata, SAP, Oracle, and EMC/Greenplum have responded to this data volume growth by employing more powerful and parallel hardware in conjunction with sophisticated parallelization techniques in the underlying data management software. Twitter, LinkedIn, Facebook, Google, and others address the scalability challenge in their clusters by leveraging a combination of new technologies: key-value stores, columnar storage, and the MapReduce programming paradigm [10]. Finally, small and medium-sized businesses are gradually adopting new technology to meet their demands for finding, developing, and creating new strategic business possibilities. Big data analytics can provide firms with a competitive advantage and lead to commercial benefits such as more effective promotion and increased revenue. All the headings in the main body of your paper are numbered (automatically).

A. Benefit of Business Intelligence

Organizations always deal with a number of departments, and these departments also deal with a number of applications to meet the needs of their respective departments. This method according to [11], ensures that the data of an organization is always maintained separately in various sources. Business intelligence offers a method to manage this dispersed data and support strategic decision-making and a way to combine data from all accessible sources into a single, consistent repository, aiding not only business owners but also other users with the following:

i. Manage all type of risk well in advance.

ii. Understand the current position of an organization and its' position over a period of time.

iii. Fact based strategic, informed decision making at the right time.

iv. Find the problem areas and resolve them to improve the revenue

v. Identify all selling opportunities, i.e. cross and up for the business

vi. Understanding customer behaviour to serve them better and win more deals

vii. Understand customer segmentation, product analysis, cost, and operations improvements

B. Database Systems in Business intelligence

Database systems are essential for businesses because they are repositories for storing and processing data, and they play an important role in business information management. Understanding how networks operate will assist you in making the most of their services. Knowing the various types of data in a database and how to use it to better understand your business will allow you to create better reports and target your customers. Database management entails understanding how to implement and store data in a database, as well as how to maintain and clean it and retrieve it at any moment. When a company wants to make sense of its data, it must keep it all in a database. Database management helps an organization's data management approach. Database management, along with business intelligence (BI) solutions, is a critical component of a company's informational infrastructure, allowing informed decision-making and corporate development [12].

Databases are an essential component of corporate intelligence. Database systems are at the heart of contemporary enterprises, where they are used to process business transactions as well as to comprehend and manage the enterprise. They keep data that analysts, managers, and other stakeholders can use to make decisions. Databases provide data on the company's market share, sales patterns, and other information that can help the company make better business decisions. [13] asserts that in the case of relational databases, they are designed to keep data accurate and consistent, using SQL as a standardized language for accessing and modifying information on the database. They work well for structured data, much like organized file systems that allow for complex queries. However, real-world data are not always neatly structured.

This is where NoSQL databases come in. They are more flexible and suited for handling unstructured or messy data, such as social media posts and documents. Unlike relational databases, NoSQL comes in different types, including document stores and graph databases, each with its own strengths. These databases are particularly powerful in managing large-scale, diverse datasets and enabling real-time analysis, making them essential in the era of big data [13].

On the other hand, business intelligence is the use of data analysis to enhance enterprise management and routine business operations such as intelligent supply chain management. Business intelligence systems can help organizations in gaining a deeper understanding of the variables affecting their operations, such as metrics for sales, production, and internal operations, and they can aid organizations in making more informed business choices. Data mining, data farming, and data warehouses are examples of business intelligence apps and technologies that can assist businesses in storing and analyzing data. Large quantities of data are stored in data warehouses. Short response periods are crucial for online decision support in On-Line Analytical Processing (OLAP), where this data is

frequently used. The query performance is one of a data warehouse server's most crucial needs [14].

The market for data storage has grown incredibly in recent years. Despite the sophistication and maturity of traditional database technologies, the growing size of corporate databases and the emergence of a brand-new global Internet database raise the possibility that new computing models will soon be necessary to completely support a number of essential data management duties. Given their inherent ability for concurrent computation and data access, the use of parallel algorithms and architectures in particular offers great potential. In that they enable users to evaluate the development of an organization in terms of a number of important data attributes or dimensions, data warehouses can be compared to decision support systems. These attributes are typically extracted from a variety of operational sources, whether they be relational or not, before being cleaned up and normalized and then loaded into a relational database. Users can drill down or roll up on hierarchies, split and dice specific characteristics, or perform various statistical operations like ranking and forecasting by taking advantage of multi-dimensional views of the underlying data warehouse. This strategy is known as "On-Line Analytical Processing." (OLAP).

C. *Parallel Database Systems for Business Intelligence*

In a parallel database machine environment, instead of a single backend, there are multiple backend machines connected either to a single host or to multiple hosts. The backends control multiple disk units containing the database. The entire database may be replicated on several disks or the files may be partitioned over the disks. Data partitioning can be done by placing new records on successive disks in round-robin fashion, by hashing on some attribute, or by range partitioning, which means placing records by designating disks according to a range of values for a certain attribute. When a query is processed, since the required data may reside on different disks, the query is decomposed into subqueries that are then processed using the appropriate partition of the database [14].

In very big databases with sophisticated and time-consuming querying and processing requirements, high-performance parallel database systems are displacing traditional systems. Because they frequently require the analysis of several relations, relational queries are well suited to parallel execution. We can parallelize the data retrieval required for Data Mining as well as the data processing required to apply the various algorithms used to accomplish the Mining activities. According to [15], these processors can be programmed to;

i. share memory,

ii. share disks, or

iii. share nothing.

Parallel processing can be done with a shared address space, which allows for efficient communication. However, sharing nothing is the most scalable approach because it eliminates overheads. The inherently parallel nature of most Mining operations in Data Mining allows us to use processors that only need to communicate periodically, resulting in efficiency in both speed and scalability.

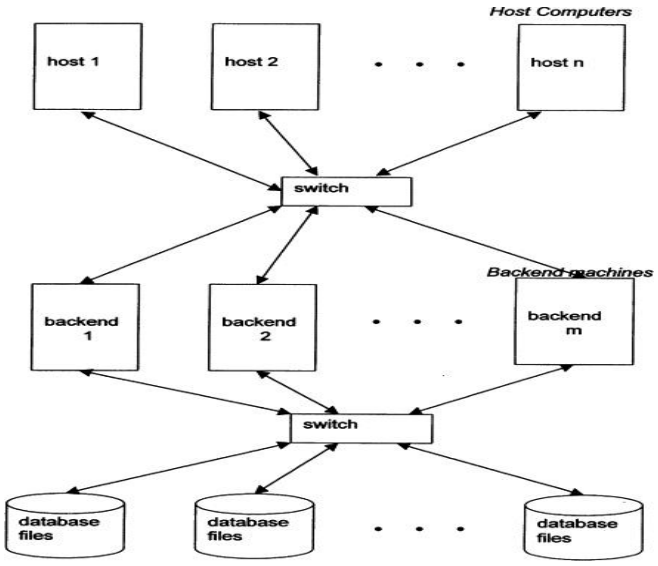

Figure 1: Parallel Database architecture [16].

D. *Distributed Database System for Business Intelligence*

A distributed database is a collection of many networked databases that are physically scattered over multiple places. All of the databases that are linked communicate with one another via a network. A distributed database management system controls the distributed database in such a way that it seems to users to be a single database. It is not a disjointed file system. The primary purpose of a distributed database is to maximize performance by distributing data across multiple networks or locations. It also aids in the efficient, dependable, and transparent use of IT resources. It also ensures fault tolerance. When all other components fail, this database system allows access to resources. This will be accomplished with the assistance of data replication [17].

According to [18], a distributed database management system (DDBMS) is a network-distributed collection of numerous, logically interconnected databases. They provide a mechanism for making data distribution visible to users. The degree of distribution in distributed database management systems is significant. Transparency simplifies database access for the using client, but it creates a more complex working environment and increases the load on the software to do translation, location, and distribution within itself. Low-transparency systems, on the other hand, require a well-educated and skilled client and user group because they must specify translation, location, and distribution. This necessitates that the client software and users understand the complexities and actively construct the transaction distribution needed.

In contrast to centralized databases, where all data is maintained at a single location and it is assumed that individual transaction processing is basically sequential. Distributed databases process information requests using a client/server architecture and can actively construct the transaction distribution that is needed. The two types of distributed database systems are homogeneous distributed database systems and heterogeneous distributed database systems. The former consists of the same database program

on all computers. Oracle is used as a DBMS system, whereas at least one of the databases in a mixed distributed database system is not an Oracle database. It is common for distributed databases to be confused with mirrored databases. It may appear to be the same to some extent, but the distinction is founded on its purity. The system handles a single copy of all data and supporting database objects in a pure (non-replicated) distributed database. Distributed database applications typically use distributed transactions to access both local and remote data and change the global state [19].

Figure 2: Distributed Database System Architecture (Source: [17])

The data in a distributed database is stored in storage devices placed in separate physical locations or geographical areas. They are not linked to a single CPU, but the database is managed by a centralized DBMS. To gain access to data from a distributed database. The replication and duplication methods that are employed to keep the database current. Creating and maintaining a distributed database is a more difficult undertaking than managing a centralized database.

E. Importance of Databases to Business Intelligence

In the early and middle 1990s, as data warehousing and business intelligence (BI) gained popularity, some technologists objected to the idea of copying data into a different database where reports and analytics would then be performed. They believed that storage was still a somewhat valuable resource and that replicating data would be expensive. In addition, every extraction, transformation, and loading (ETL) operation that copied data from one or more source systems into a data warehouse was a potential source for data anomalies and mistakes. Not minding that the quality of the original-form data housed in many applications was itself highly suspect. Databases can be linked to the Web via middleware software, allowing users or clients to access company data via a Web browser interface. Such software could include an application server, a proprietary software program, or common gateway interface (CGI) scripts. In a client/server arrangement, the DBMS may sit on a special specialized machine known as a database server. Web interfaces are simple to use and need little or no changes to the internal database [20].

The following are some of the importance of databases to the effective functioning of a business intelligence system:

i. Keep track of fundamental transactions of the organization,

ii. Provide data that will assist the organization in running its operations more efficiently.

iii. Assist managers and staff in making more informed decisions.

There is a substantial amount of work that applies parallel processing techniques to relational database systems. The goal is to improve performance by utilizing either inter-query or intra-query parallelism. Much research in the topic has focused on query processing in parallel and distributed databases [21]. Although standardized relational algebra operators used in databases lend themselves well to parallelization, concerns have been expressed about join processing and communication overheads associated with both joins and other operation needs to communicate data.

Parallel Hash Join algorithms [22] and Placement Dependency algorithms [23] are two examples of usual ideas for reducing some of the cost involved with join operations. Another key difficulty in parallel and distributed database architectures which was also discussed. Partitioning data for independent processing of fragments can often cut response time dramatically while increasing total effort due to the need to interchange and integrate data and partial results. [24]. This might cause significant speed degradation for complex heavy join queries. Data warehouses are precisely the type of situation where node partitioning can improve response time while also causing significant extra overheads.

Business intelligence is becoming increasingly important as it connects businesses to vast databases, enabling them to extract valuable insights. It enhances database systems, improving operational efficiency and decision-making. Artificial intelligence plays a key role by providing businesses with crucial market trends and consumer data, helping them stay ahead of competitors. Business intelligence also fosters a data-driven culture, allowing firms to predict market trends and uncover new opportunities. Ultimately, these tools empower businesses to make informed decisions and drive innovation [13].

F. Business Intelligence and Data Warehousing

A data warehouse collects and stores data from multiple sources. Housing or storing data in a digital "warehouse" is analogous to storing papers or photographs in the cloud. Having a place to keep your data makes it easier to use and delivers more insights on a wider scale. You can report the most recent and integrated data by importing historical data or timely data streams. A data warehouse can be used for analytical and commercial reporting reasons. To fully utilize all of your data, you must develop an integrated data strategy as well as an authentic audience story. As a result, integrating data warehouse and business intelligence technologies will increase organizational decision making and productivity [25].

Business intelligence (BI) applications frequently use data from a data warehouse. However, not all data warehouses are utilized for business intelligence, and not all business intelligence applications necessitate the usage of a data warehouse. Data warehousing and business intelligence are phrases used to describe the process of storing all of an organization's data in internal or external databases from diverse sources, with the goal of analyzing it and providing actionable insights utilizing online BI tools [26] [6]. When implemented correctly, data warehousing and business intelligence can serve as an organization's information backbone, assisting in the alignment of every line of business to support a fully data-driven operation.

Data warehouses are the backbone of data storage in business intelligence. Complex inquiries and data comparisons are used in business intelligence to inform everything from daily choices to organizational-wide shifts in focus. Business intelligence is made of three major activities to accomplish this: data wrangling, data storage, and data analysis. Data wrangling is typically aided by extract, transform, and load (ETL) technologies, which we'll go over in depth below, while data analysis is performed using business intelligence tools. Data warehouses, which facilitate data storage via OLAP, are the glue that holds this process together. They combine, summarize, and change information to make it easier to examine [27].

G. Business Intelligence Architecture

The term "business intelligence architecture" refers to the standards and policies for organizing data using computer-based techniques and tools to build business intelligence systems used for online data visualization, reporting, and analysis. Data warehousing is a component of business intelligence (BI) infrastructure. A central repository system, namely a data warehouse, must organize, store, clean, and extract data, which is regarded the fundamental component of business intelligence [26]. As organizations begin to implement business intelligence, one critical duty is to ensure that they adhere to a proper business intelligence (BI) architecture design during the implementation process in order to assure the success of their business intelligence (BI) investment.

A business intelligence (BI) architecture is a framework that deals with the various components of BI, such as data, people, processes, technology, and management, and how these components interact to maintain the system's seamless operation. [28]. A good business intelligence (BI) architecture must include information such as the data categories that must be collected, the methods to be used to evaluate data, and the manner in which specific information must be presented. Inconsistencies between the various components may occur if the underlying architecture is not adequately planned.

Inconsistencies that occur among the multiple components if the underlying architecture is not adequately planned may lead to difficulties such as difficulty to share information among the components, inability to meet business requirements, and poor business performance. In the worst-case scenario, a poor business intelligence (BI) design may result in the delivery of incorrect information to the wrong person at the wrong time. Even if the systems are functioning despite poor architecture, organizations will be unable to maximize the value they should have received from their business intelligence (BI) investments. [29]. According to [30], integrating and simplifying the business intelligence (BI) architecture increases information and technology (IT) infrastructure agility while ensuring high availability and energy savings. A streamlined and closely linked BI architecture also allows enterprises to lower total cost of ownership (TCO) and capitalize on the next-generation information ecosystem.

Business intelligence architecture is made up of several components and levels, each of which serves a specific purpose, as mentioned below:

1. Collection of data: The first step is related to the collection of relevant data from various external and internal sources which can be databases, ERP- or CRM systems, flat files, or APIs, just to name a few.

2. Data integration: At this stage, the data collected is integrated into a centralized system, often with the help of ETL processes. Here the data is also cleaned and prepared for analysis.

3. Storage of data: This is where data warehousing comes into the picture. A warehouse is a place in which structured data is stored. It makes it available for querying and analysis.

4. Data analysis: After the information is processed, stored, and cleaned it is ready to be analyzed. With the help of the right tools, the data is visualized and used for strategic decision-making.

5. Distribution of data: The data, now in the form of graphs and charts, is distributed in different formats. This can be online reporting, dashboarding, or embedding solutions.

6. Reaction based on insights: The final stage of the architecture process is to extract actionable insights from the data and use them to make improved decisions to ensure company growth.

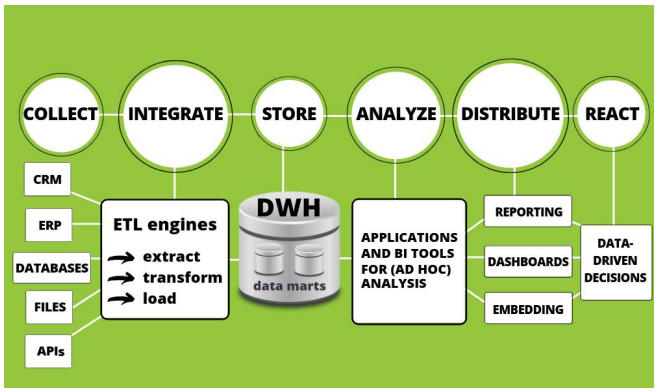

Figure 3: Business intelligence Architecture (Source: [26])

III. CONCLUSION

Business intelligence employs parallel and distributed databases that are driven from a central point, and it is generic enough to be used in a variety of situations, including information and data warehouses. Parallel programming standardizes codification and reduces

execution time for determined queries. The system allows for a consolidated view of the information in distributed databases, which is accomplished through the publication of the primary queries in the portal. It is important to note that for effective business decision making the need for robust database and data warehousing architecture facilitate the efficiency of business analytics especially in this era of massive data generation and application. With upsurge of Big structured and unstructured data in business domain has made it necessary for the implementation of business intelligence application for enhanced business decision making process.

REFERENCES

[1] Cohen, J., Dolan, B., Dunlap, M-, Hellerstein, J.M & Welton, C. (2009). MAD Skills: New Analysis Practices for Big Data. Procrrdings of the VLDB Endowment, 2(2), 1481–1492.

[2] Frankel, F & Reid, R. (2008). Big Data: Distilling Meaning from Data. Nature, 30(4), 455-460.

[3] Chen, H., Chiang, R.H & Storey, V.C. (2012). Business Intelligence and Analytics: From Big Data to Big Impact. MIS Quarterly: Management Information Systems, 36(4), 1165-1188.

[4] Herodotou, H. (2016). Business Intelligence and Analytics: Big Systems for Big Data. Cyprus University of Technology.

[5] Stedman, C. (2023). Business Intelligence. Retrieved from Tech Target:

[6] Cebotarean, E. (2011). Business Intelligence. Journal of Knowledge Management, Economics and Information Technology, 1(2), 2-10.

[7] Glaser, J & Stone, J. (2008). Effective use of Business Intelligence. Healthcare Financial Management, 62(2), 68-72.

[8] Dean, J & Ghemawat, S. (2008). MapReduce: Simplified Data Processing on Large Clusters. Communications of the ACM, 51(1), 107–113.

[9] Isard, M., Budiu, M., Yu, Y., Birrell, A & Fetterly, D . (2007). Dryad: Distributed DataParallel Programs from Sequential Building Blocks. ACM SIGOPS Operating Systems Review, 41 (3), 59–72.

[10] Thusoo, A., Sarma, J. S., Jain, N., Shao, Z., Chakka, P., Anthony, S., Liu, H., Wyckoff, P & Murthy, R. (2010). Hive: A Warehousing Solution over a Map-Reduce Framework. VLDB Endowment, 2(2), 1626–1629.

[11] Kumar, A. (2014). Business Intelligence – An Overview. Retrieved from Technology advise: https://www.developer.com/design/business-intelligence-an-overview

[12] Rangan, K. (2021). Database Management: Improve Data-Driven Decision Making. Retrieved from G2: https://www.g2.com/articles/database-management

[13] Hosen, M.S., Islam, R., Naeem, Z., Folorunso, E.O., Chu, T.S., Al-Mamun, M.A & Orunbon, N.O. (2024). Data-Driven Decision Making: Advanced Database Systems for Business Intelligence. Nanotechnology Perceptions, 20(3), 687-704.

[14] Kadry, S & Smaili, K. (2008). Massively Parallel Processing Distributed Database for Business Intelligence. Information Technology Journal, 7(1), 70-76.

[15] McClean, S. (2003). Data Mining and Knowledge Discovery. In H. Bidgoli, Encyclopedia of Physical Science and Technology (pp. 229-246). Elsevier.

[16] Ricardo, C. (2013). Database Machines. In H. Bidgoli, Encyclopedia of Information Systems (pp. 403-410). Elsevier

[17] Verma, G., Saxena, K., & Saxena, S. (2018). DISTRIBUTED SYSTEM: Learn to Design and Develop the Distributed Applications. BPB Publications.

[18] Rouse, M. (2012). Distributed Database Management System. Retrieved from Techopedia: https://www.techopedia.com/definition/14686/distributed-database-management-system-ddbms

[19] Patil, H.S., Mukhtar, Y.S., Pawa, S.U & Patil, D.W. (2011). Distributed Databse - An Relevance to business organization. Journal of Information and Operations Management, 2(1), 21-24.

[20] Bourgeois, D.T., Smith, J.L., Wang, S & Mortati, J. (2019). Information Systems for Business and Beyond. John Wiley & Sons.

[21] Chen, H & Chengwen, L. (2000). An efficient algorithm for processing distributed queries using partition dependency. . International Conference on Parallel and Distributed Systems, (pp. 339-346.).

[22] Liu, C & C. Hao. (2006a). A hash partition strategy for distributed query processing. 5th International Conference on Extending Database Technology: Advances in Database Technology (p. 37). Springer-Verlag.

[23] Liu, C & C. Hao. (2006b). A heuristic algorithm for partition strategy in distributed query processing. ACM Selected Areas in Cryptography SAC.

[24] Copeland, G.P., William, A., Boughter, E & Keller, T. (2008). Data placement in bubba. ACM SIGMOD Record, 17(1), 99-108.

[25] Tikait, P. (2023). BI/DW: What is Business Intelligence and Data Warehousing? Retrieved from Select Hub: https://www.selecthub.com/business-intelligence/business-intelligence-and-data-warehousing/

[26] Calzon, B. (2022). The Role Of Data Warehousing In Your Business Intelligence Architecture. Retrieved from Data Pine Blog: https://www.datapine.com/blog/data-warehousing-and-business-intelligence-architecture/

[27] David, M. (2021). How to Use Data Warehouses in Business Intelligence. Retrieved from Charito: Data Tutorial: https://chartio.com/learn/business-intelligence/how-to-use-data-warehouses-in-business-intelligence/

[28] Rob, P & Coronel, C. (2009). Database Systems: Design, Implememtation and Management. Course Technology.

[29] Rasmussen, N.H., Rasmussen, N., Chen, C.Y & Bansal, M. (2009). Business Dashboards: A visual catalog for design and deployment. John Wiley & Sons.

[30] Gole, V & Vijayvargiy, M. (2022). Simplify Your BI Architecture for Better Business Intelligence. Tata Consultancy Services. Retrieved from https://www.tcs.com/content/dam/global-tcs/en/pdfs/insights/whitepapers/simplify-bi-architecture-better-business-intelligence.pdf