Simplify Your BI Architecture for Better Business Intelligence

Abstract

There is little doubt that data is the new competitive differentiator. But data alone is not enough; it has to be scrubbed and analyzed to generate real time, actionable insights for business success. This makes a robust business intelligence (BI) architecture a must-have for modern day enterprises. However, many organizations continue to manage an increasingly complex IT infrastructure due to the addition of independent applications and servers over time, and an expanding global user base.

Unifying and simplifying the BI architecture brings agility to IT infrastructure while ensuring high availability and energy savings. A streamlined and tightly integrated BI architecture also enables organizations to reduce the total cost of ownership (TCO) and take advantage of the next-generation information ecosystem.

How can the gaps in a traditional BI architecture be closed? Adopting a simplified approach, supported by modern tools, can help organizations facilitate real time and data-driven decision making for strategic business outcomes.

Why traditional BI architecture falls short in today's data-driven environment

Organizations are increasingly adopting modern analytics to connect disparate data sources, understand customer behavior, and predict product or services sales. But the reality is traditional BI architecture is incapable of processing and analyzing vast amounts of data from different sources including location, behavioral, and sensor data. The resulting architecture is cluttered and complex.

Mergers, acquisitions, and application integration along with incremental upgrades, result in a cluttered architecture.

While BI tools have evolved over the years, many organizations still follow a multi-tiered approach for data processing and transformation (see Figure 1). This includes source systems that provide transactional data to the enterprise data warehouse (EDW), either in a raw or semi-formatted state. The data is dropped into the staging area through processes that harmonize the data so that it can be consumed by the extract, transform, and load (ETL) layer.

Depending on the data processing strategy, the EDW layer processes data across multiple layers. This is followed by further processing in the reporting layer, before it is consumed by business users and corporate influencers.

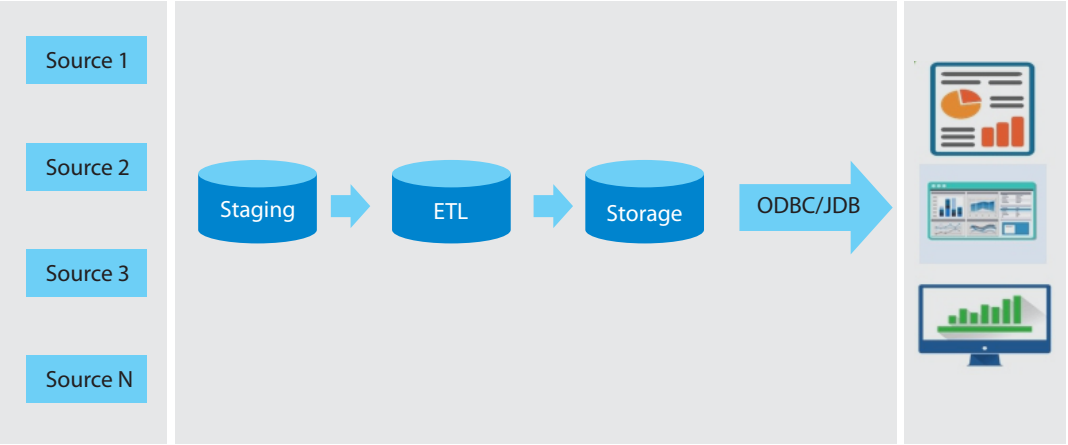

Figure1: The traditional distributed BI architecture

This distributed architecture of the data warehouse poses multiple issues. Firstly, maintenance of the interdependent components becomes a challenge, and needs the support of multi-skilled teams. Second, the architecture impacts data availability. In addition, traditional ETL tools work across multiple processing layers, requiring batch transformations to load the processed data into the EDW. A typical batch cycle for ETL takes nearly 24 hours. The time lag in getting the data to the next layer means that end users can utilize the data only

after it has loaded. Lastly, maintaining multiple layers incurs costs in terms of hardware and license fees, which increases the overall TCO.

A distributed architecture poses maintenance and data availability issues, and increases TCO.

Accelerate business intelligence by simplifying the traditional data warehouse

To enable data analytics to keep pace with the speed of business in the digital age, companies will have to transform and simplify the traditional data warehouse, using replication technologies and in-memory data processing. The key steps involved in this are seen in Figure 2

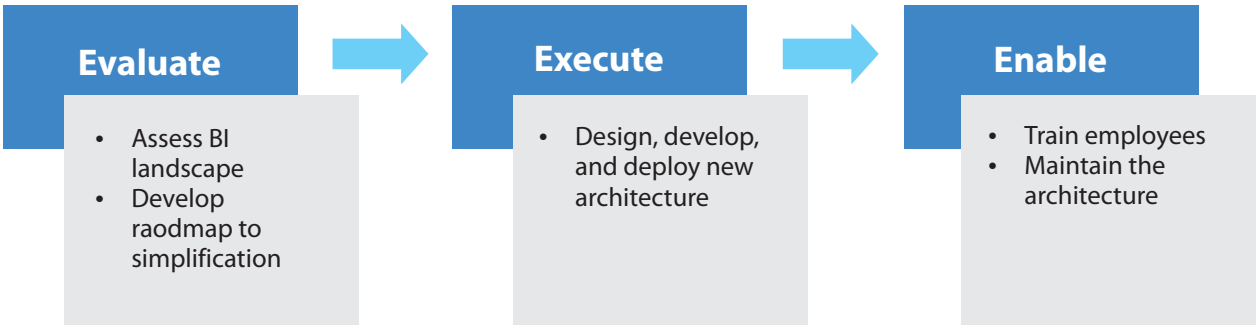

Figure2: Steps to simplify the data warehouse landscape

Phase 1 (Evaluate): Assess the current landscape with respect to feasibility, value, and priority of transformation. Then create a roadmap for steady transition to the simplified data warehouse, based on the findings of the evaluation.

Phase 2 (Execute): Following the roadmap creation, steps can be taken to design, develop, and deploy the real-time data warehouse.

Phase 3 (Enable): Change management and end user training is one of the most important tasks in this phase. It also involves knowledge transition to the managed services team for maintaining the data warehouse in real time.

The simplified data warehouse architecture is built on three layers as shown in Figure 3.

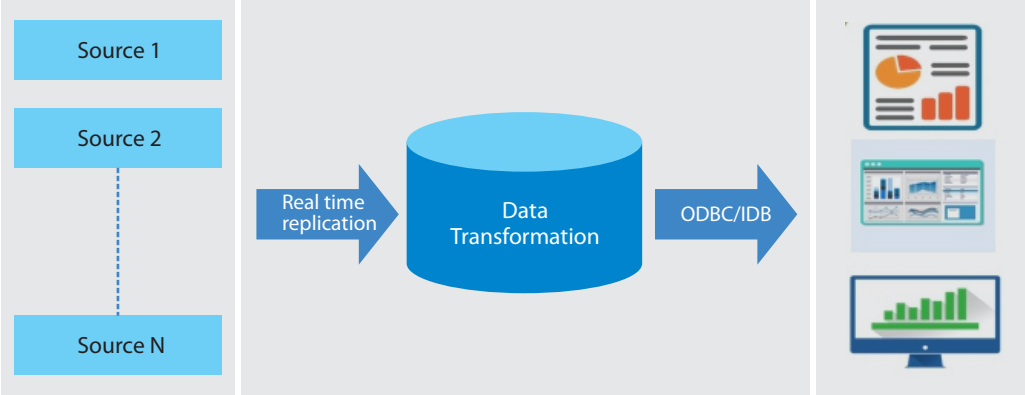

Figure3: Architecture of the simplified data warehouse

Source: As in the case of a traditional EDW, the source layer makes data available for consumption by the transformation layer.

Transformation: This layer accesses data from the source and transforms it to make it available for reporting requirements. Source data can be replicated in the database using trigger based change data capture techniques to ensure that real-time data is always available. Some technologies also allow access to source data as needed, eliminating data storage requirements. Business rules built into transformation logic are deployed in-memory to ensure faster responses to queries put forth by the reporting layer.

Reporting: Since complex calculations are now included in the transformation layer, the reporting layer is lighter and faster.

A simplified BI architecture for a high performance organization

The simplified data warehouse architecture eliminates multiple layers of data processing as well as the need for additional hardware and complex code. This makes it high available, easier to maintain, and reduces TCO. The elimination of intermediate layers also frees up resources and facilitates quick debugging through a centralized logic processing unit.

The combination of replication and in-memory database takes out the EDW layer completely, resulting in a reduced data footprint. In addition, since data is retrieved directly from memory, data processing is faster compared to storage-based databases. So, users no longer have to wait till a loading cycle is completed. Data is processed and transformed real-time facilitating faster query responses and decisions-making.

Take the case of a leading toy manufacturer. During the course of a massive global ERP rollout, the company deployed a simplified BI architecture, realizing significant benefits. The company was expanding into new markets and was faced with huge data explosion and global demand for immediate information. This resulted in the need for new code development for inclusion of market specific business logic into the EDW. The final BI landscape was not only complex, involving multiple technologies such as SQL server, Informatica, SAP Data Services, SAP HANA, and SAP Business Objects, but also slow due to data processing across multiple layers. With an average information delivery time of 24 hours, business users were accessing day-old data. Additionally, debugging took anywhere from several hours to weeks.

To resolve these issues, the company decided to simplify its architecture and moved data from the source using a real-time replication technology to the data transformation and storage layer, thereby eliminating all other layers. This approach helped them drive reporting as well as data driven decision-making in real-time, for superior business outcomes.

Transform data into a strategic differentiator

Digital, data-driven enterprises will need to revamp traditional BI to cater to the growing need for real-time visual analytics, as well as operational reporting. Moreover, with BI and analytics moving to the cloud, organizations must explore a new approach to architecting the data warehouse. According to a study by BARC Research and Eckerson Group, 78% of organizations are planning to increase the use of cloud for BI and data management in 2017[1].

Organizations also adopt cloud-based models for advanced and predictive analytics, operational planning and forecasting, as well as strategic planning and simulation. Moving forward, the need for a next-gen information ecosystem will compel enterprises to transform their BI architecture to realize better ROI from BI and analytics investments.

References

[1] BARC, BI and Data Management in the Cloud: Issues and Trends, January 2017, accessed May 2017, https://s3.amazonaws.com/eckerson/content_assets/assets/000/000/115/original/BARC_Research_Study_BI_and_Data_Management_in_the_Cloud_EN.pdf?1487101351

About The Authors

Vinayak Gole

Solution Architect, SAP Analytics Center of Excellence,

Enterprise Application Services

Mitesh Vijayvargiy

Technical Architect, SAP Analytics Center of Excellence,

Enterprise Application Services

About TCS' Consulting and Enterprise Solutions Unit

TCS Consulting and Enterprise Solutions professionals drive transformations which enable our clients to innovate business models and take advantage of emerging opportunities and market dynamics.

With a robust portfolio which includes business and IT consulting services and enterprise-wide software application development, implementations, and managed services, our consultants offer deep domain expertise and decades of experience to help our clients achieve their business goals.

Our proven methodology combined with our knowledge of next generation digital technology and our strategic partnerships with the world's leading enterprise software companies ensure business growth, operational efficiency, and competitive advantage for our clients across industry verticals.

Contact

Visit the Consulting and Enterprise Solutions page on www.tcs.com

Email: global.consulting@tcs.com

Blog: #Enterprise Insights

Subscribe to TCS White Papers

TCS.com RSS: http://www.tcs.com/rss_feeds/Pages/feed.aspx?f=w
Feedburner: http://feeds2.feedburner.com/tcswhitepapers

About Tata Consultancy Services Ltd (TCS)

Tata Consultancy Services is an IT services, consulting and business solutions organization that delivers real results to global business, ensuring a level of certainty no other firm can match. TCS offers a consulting-led, integrated portfolio of IT and IT-enabled, infrastructure, engineering and assurance services. This is delivered through its unique Global Network Delivery Model™, recognized as the benchmark of excellence in software development. A part of the Tata Group, India's largest industrial conglomerate, TCS has a global footprint and is listed on the National Stock Exchange and Bombay Stock Exchange in India.

For more information, visit us at www.tcs.com

TCS Design Services I M I 07 I 17