

A Theoretical Study of Data Warehouse Architecture Models for Decision Support Systems

Vaivaw Kumar Singh

Research Scholar, Faculty of Business Management, Sarala Birla University, Ranchi, Jharkhand, India

DOI: <https://doi.org/10.51584/IJRIAS.2026.11080191>

Received: 09 September 2026; Accepted: 14 September 2026; Published: 25 September 2026

ABSTRACT

These days, businesses run on data like tons of it. They collect information from all sorts of sources: enterprise apps, transaction systems, customer relationship management platforms, financial software, websites, mobile apps, and even outside data providers. But having all this data lying around doesn't mean people will make better decisions. What organizations really need are solid ways to bring all this data together, store it, process it, analyze it, and finally present it so it actually makes sense.

That's where data warehouses come in. They've become a key piece of tech for Decision Support Systems (DSS) as they give managers integrated, historical, and analytical info to help them decide what comes next.

This theoretical study dives into the main types of data warehouse architectures and looks at how good they are for DSS. It covers the classic three-tier architecture, enterprise data warehouse setups, dimensional models, data marts, and hybrids. The researchers use a qualitative, theoretical approach, mainly by reviewing and putting together established sources.

Here's what they find: each architecture brings something different to the table. Enterprise-focused setups offer tight integration, good consistency, and strong governance. Dimensional models are all about easy analysis and fast queries. Data marts are flexible and you can get them running quickly, but if you build them separately, you risk duplicated or inconsistent data. The three-tier model keeps things organized and there's a clear split between where the data comes from, where it sits and gets integrated, how it gets analyzed, and what users actually see. The study also suggests that hybrid approaches can combine the best of both worlds like central oversight with business-focused analytics.

There are five conceptual diagrams in the study that map out these architectures and show how decision support actually works in practice. In the end, the researchers say that there's no one-size-fits-all answer for data warehouse architecture. The right choice really depends on a company's size, how complex its data is, what kind of analysis it needs, scalability, governance, and decision-making demands.

Keywords: Data warehouse, Decision Support Systems, data warehouse architecture, OLAP, dimensional modeling, data marts, enterprise data warehouse, business intelligence, decision making.

INTRODUCTION

Information technology has completely changed how organizations handle data. These days, companies gather an endless stream of information from sales transactions and ERP systems, to CRM tools, HR software, websites, phone apps, sensors, social media, and outside info providers. All this data opens the door to huge possibilities, but it also brings its own headaches: messy data, quality issues, confusing integration, storage struggles, and challenges in making that information truly useful and accessible for analysis.

Operational databases, the traditional backbone, were built mainly for day-to-day processing. Think sales, payments, new customer sign-ups, inventory moves, loan approvals all these systems are made to keep business humming, fast and reliable. But here's the thing: managers and decision makers need something different. They

want to spot trends over time, compare numbers across years, do multi-layered analyses, predict what's coming, and pull together results across departments. Transactional systems just aren't built for

complex, heavy-duty analysis.

That's where the idea of the data warehouse comes in. Inmon's work in the '90s spelled out why organizations needed a separate, integrated store for decision-making: a data warehouse. Instead of scattered, siloed, constantly changing data, a warehouse offers subject-focused, consistent, time-stamped info that's stable enough to analyze from a bunch of angles.

But the data warehouse isn't just a giant storage bin. It's a tool for turning scattered data into reliable insights that actually help managers make smart choices. Kimball and Ross pushed the idea of shaping data around business processes on dimensional models that make complicated data easy to understand and use.

Architecture matters and the design of a warehouse includes all the key pieces: the original data sources, extraction and transformation stages, holding areas, the main warehouse, smaller data marts, OLAP servers for analysis, repositories for metadata, dashboards, and tools for decision support. How you connect and organize these parts affects everything including data quality, speed of queries, how well the system scales up, how easy it is to maintain, how secure and well-governed the data is, and the usefulness of the system overall.

There isn't just one approach. Inmon's style puts everything into one central, well-integrated system. Kimball's model is all about organizing by business processes and using dimensional structures. Some companies go with independent data marts for different departments, while others use a layered approach, each tier with its own job. These days, a lot of organizations mix and match, choosing what fits best.

Picking the right setup really is a strategic decision. A small company with just a handful of simple reporting needs doesn't need the same complex architecture as a global giant dealing with hundreds of different systems. Companies focused on tight data control want centralized integration. Others that need speed and flexibility lean toward data marts.

This paper explores the main theories behind these architecture models, specifically for Decision Support Systems (DSS).

Problem Statement

Even though data warehouses and business intelligence tools are everywhere, organizations still struggle to pick the architecture that works best for them. Each model has different strengths whether it's integration, performance, flexibility, speed, governance, or ability to grow.

This study looks at the real issue: there isn't a simple, clear framework that connects how you design a data warehouse to how effective your Decision Support System is.

Objectives of the Study

Here's what this study aims to do:

- Lay out the core theories behind data warehouse architecture.
- Break down the main architecture models.
- Analyze how architecture choices impact Decision Support Systems.
- Weigh the pros and cons of each model.
- Propose a theoretical framework that helps organizations choose the right architecture for effective DSS.

LITERATURE REVIEW

Concept of Data Warehouse

Think of a data warehouse as the nerve center of a company's data which is one place where everything comes together, ready for deep dives and big-picture analysis. Inmon (1992) pointed out that companies can't just rely on their everyday operational systems for decision-making. They need a separate setup focused on analytics.

There are four main things that make a data warehouse stand out:

- **Subject orientation:** Everything gets grouped by major themes like customers, products, sales, finance, employees, that sort of thing.
- **Integration:** Data from various systems is merged and cleaned up so it fits together neatly.
- **Time variance:** You don't just see what's happening now and you can look back at years of history.
- **Non-volatility:** This isn't where constant changes happen. Data gets loaded in batches for analysis, not modified every time someone makes a sale. (Inmon, 1996)

These features really separate data warehouses from regular transactional databases.

Decision Support Systems

Decision Support Systems, or DSS, are tools that help managers make better choices. They pull together information, run analyses, and give you different options to consider like handling anything from straightforward decisions to the really tricky ones.


Data warehouses play a huge role here. They basically power DSS by giving managers access to reliable, organized info about what's happened in the past, how different parts of the company stack up, trends, customer patterns, key metrics, and future forecasts.

It looks something like this:

Operational Data → Data Warehouse → Analytics → Information → Decision Support → Decisions

So, the quality of decision support really depends on how solid and well-designed your underlying data setup is.

Diagram 1: Data Warehouse as the Foundation of Decision Support



Here's how it works: Raw data comes in from various operations. Through ETL or ELT processes, that data is extracted, cleaned, and merged, then loaded into the warehouse. From there, OLAP and analytical tools dig into it. The results include reports, dashboards, forecasts, and more and give managers what they need to guide the business.

Evolution of Data Warehouse Architecture

Data warehouses have come a long way. At first, companies used big, centralized repositories. Over time, things moved toward smaller, department-focused data marts and smarter designs that made it easier to analyze data quickly. These days, you see multi-layered systems, stronger metadata management, distributed setups, cloud infrastructure, and advanced analytics coming into play.

We can break this evolution into four main phases:

- Centralized enterprise warehouses
- Departmental or dimensional data marts
- Integrated, multi-layered architectures
- Hybrid and distributed analytical environments

This whole progression mirrors how organizations and their data needs have grown more complex, and how decision-making keeps demanding more flexible, powerful solutions.

MAJOR DATA WAREHOUSE ARCHITECTURE MODELS

Three-Tier Data Warehouse Architecture

The three-tier architecture crops up all the time when people talk about data warehouses. It splits the system into three main parts: where the data lives and gets cleaned up, where people actually analyze it, and where the results show up for users.

At the bottom, you've got your raw source data, ETL processing, staging areas, and the main warehouse storage. The middle's where the heavy lifting happens that is OLAP servers run calculations and let analysts slice and dice the data. Up top, you find all the stuff users see: dashboards, reports, visual tools, decision support systems, you name it.

This sort of separation does a few things. It keeps the heavy analytical work away from the operational systems so we don't slow down core business transactions just because someone's running a sales report. End users get to query and explore data without messing with complicated source databases.

Back in 1997, Chaudhuri and Dayal pointed out that data integration, how the warehouse stores info, OLAP tools, and analysis apps all matter in data warehousing.

Advantages

- Each part of the system has a clear job.
- Analytics run faster because they're not fighting for resources.
- Operational systems don't get bogged down by extra work.
- Easier to fix or upgrade parts of the system.
- Plays well with lots of analytical tools.

- Works for complicated decision support environments.

Limitations

- More moving parts means more complexity.
- You need skilled people to build and maintain it.
- You have to plan carefully so each layer talks smoothly to the others.

Enterprise Data Warehouse Architecture

With enterprise data warehouse (EDW) architecture, everything's about keeping all business data in one central spot. This idea often connected to Inmon's top-down approach focuses on collecting data from the whole organization, cleaning it up, and making it available from a single source.

The big selling point is that everyone works from the same version of the truth. The whole organization benefits from clear definitions and solid data governance, sticking to approved standards and tracking everything with proper metadata.

Key characteristics:

The data's stored centrally.

Data from across the business gets integrated.

Strong focus on governance.

Keeps historical data for long-term analysis.

Uses consistent definitions across departments.

Centralized management of metadata.

This is a solid choice when you have to keep things consistent organization-wide. Imagine making sure "customer" means the same thing for sales, finance, marketing—everyone. With a centralized warehouse, it's possible.

Dimensional Architecture

Kimball's take on data warehousing is all about dimensional modeling, which skips complex, heavily normalized designs. Instead, it sets up data in facts and dimensions.


Here, a fact table holds measurable events like sales amounts or number of units sold. Dimension tables add context, such as "customer," "product," "location," or "time." It's all about making analysis easier and more intuitive.

Kimball and Ross (2013) said you've got to understand the business process, set the detail level, pick out dimensions, and identify which facts matter for your model.

Star Schema

A star schema keeps things simple with a central fact table, directly linked to several dimension tables.

Diagram 2: Dimensional Star Schema



Let's take sales data: the fact table holds things like sales amount, quantity, discounts, and profit. The dimension tables will be break down by time (day, month, quarter, year), customer (who bought what), product details, and location. This setup lets managers analyze performance from lots of angles quickly.

Data-Mart Architecture

Data marts are smaller, specialized versions of data warehouses. They're built for single departments, business functions, specific user groups, or even narrow topics. Think:

- Sales data mart.
- Finance data mart.
- Marketing data mart.
- HR data mart.
- Customer analytics data mart.

Some data marts take their data from the main enterprise warehouse which is a dependent data mart. Others grab data right from operational sources or outside feeds, which makes them independent.

- Dependent Data Mart

Dependent marts pull their data from a central warehouse, keeping things consistent across departments.

- Independent Data Mart

Independent marts don't rely on the main warehouse, which means they're faster to deploy but can get messy with data integration issues, mismatched definitions, you name it.

The big plus with data marts is agility. Teams don't have to wait for IT as they can build what they need, when they need it. But go too far, and you risk data silos, conflicting definitions, and duplicate information. For example, sales and finance might define "revenue" in two different ways if their marts grow in isolation.

Hybrid Architecture

Hybrid architecture mixes pieces from all the other models. You might have a central enterprise warehouse to keep the data clean and standardized, but still build dimensional data marts on the side for specific departments that need tailored analytics.


We do this because needs vary. Some teams need company-wide consistency; others need fast, flexible analysis. Hybrid models let you have both.

For example:

- The main warehouse provides master data.
- Dimensional marts give analysts simple, user-friendly access to the pieces they need.
- OLAP tools power up detailed, slicing-and-dicing analysis.
- DSS systems help with complex decisions.

Hybrid architecture matters because in reality, organizations rarely stick to just one model. By mixing and matching, we get the structure, flexibility, and performance we need to handle all kinds of business requirements without getting boxed in.

Diagram 3: Comparative Architecture Models



The diagram lays out how each main architecture model focuses on something different. Enterprise architecture is all about integration and governance. Dimensional architecture puts the spotlight on analytical performance and usability. Data marts give departments the flexibility they need. Three-tier architecture organizes everything into clear layers, while hybrid architecture mixes these strategies to tackle more complicated business needs.

RESEARCH METHODOLOGY

Research Design

This study uses a qualitative, theory-driven approach. There's no original survey data here, and we aren't diving into statistical hypothesis tests. Instead, the work digs into well-established concepts, pulling together what's already known about data warehouse architecture and decision support systems (DSS).

At its core, it's an exploratory, descriptive, and comparative study.

Research Approach

Everything here centers on reviewing the literature. The backbone comes from foundational work on data warehousing, dimensional modeling, OLAP, database systems, and DSS. The analysis draws especially on Inmon's ideas about enterprise data warehousing, Kimball and Ross's approach to dimensional modeling, and the research of Chaudhuri and Dayal on data warehousing and OLAP.

The study drills down into several main points:

- Data integration
- Data quality
- Query performance
- Scalability
- Implementation complexity
- Flexibility
- Governance
- How well they fit with DSS

Comparative Analytical Framework

Each architecture model gets evaluated based on those criteria. The aim isn't to hand out empirical scores, but to highlight each model's theoretical pros and cons.

FINDINGS

Finding 1: Architecture Shapes the Analytical Landscape

Here's what stands out: the way you build your data architecture shapes how your organization turns raw data into actionable insight.

When the architecture's thought-out, everything connects. Data flows from the source, gets integrated, finds its spot in storage, gets processed, and finally reaches the analytics layer all feeding decision-making.

But things fall apart if you miss a step. Weak integration creates messy, unreliable results. Poor storage drags down performance. A flimsy analytics layer means people can't find the insights they need.

So, if you want your decision support systems to work, you need a solid data warehouse architecture at the core.

Finding 2: Enterprise Architecture Brings True Integration

Enterprise data warehouses pull everything into one place. That means you don't have different departments stuck in their own corners; instead, sales data sits right next to customer info, finance, product details, and even geographic data.

Managers love this when they need to see the big picture and not just what's happening in one area, but how it all connects.

You end up with more reliable, consistent information across the organization. That organizational consistency is the real strength of enterprise architecture.

Finding 3: Dimensional Architecture Makes Analysis Simple

With dimensional architecture, your business users can actually understand the data models. Instead of fighting with complicated database tables, they see data as customers, products, locations, times, and sales.

This model is tailor-made for dashboards, BI tools, or OLAP. People can add things up, compare, drill down for details, or slice and dice the data however they want (Kimball & Ross, 2013).

It just works better for everyday business analysis.

Finding 4: Independent Data Marts: Speed at a Cost

Independent data marts can get results fast with no waiting around for central IT. But there's a catch: you might end up with conflicting numbers.

Say sales reports ₹100 million in revenue, but finance only sees ₹95 million. Different versions of the truth confuse everyone.

That's why it's smarter to manage data marts with shared definitions, clear metadata, and good integration. That way, people compare apples to apples.

Finding 5: Three-Tier Architecture Keeps Things in Order

The three-tier setup keeps things clear: you have a layer for pulling in and storing data, another for analytics, and a final one for users.

This keeps your analytical systems organized and stops operational chaos from reaching the end users. Plus, it lets you swap out or upgrade analytical tools without users needing to learn all the technical details.


Finding 6: Hybrid Architecture Delivers Flexibility

Hybrid architectures blend centralized control with the freedom for departments to do their own analysis.

You keep a trusted enterprise warehouse as your "source of truth," but you also let groups spin up dimensional data marts for their unique needs.

This is the sweet spot if you've got lots of departments, different data sources, piles of historical information, strict governance rules, or changing analytical demands. It's flexible enough to grow with your decision support needs.

Diagram 4: Proposed Hybrid Architecture for Decision Support



The hybrid architecture suggested combines different architectural approaches. Data from various heterogeneous sources is first extracted, transformed and loaded (ETL/ELT), and subjected to data quality processes. The enterprise data warehouse offers an integrated and governed information source. Department-level data marts are derived from the centralized warehouse. OLAP, business intelligence and advanced analysis tools leverage on this to offer information in the form of dashboards forecasts trend analysis and what-if analysis. Finally, decision makers use the derived information for strategic, tactical and operational decisions.

DISCUSSION

The findings shows that the way you set up your data warehouse has a big impact on how well your Decision Support Systems work. But different architectures fit different company needs.

If your organization values consistency, strong governance, and smooth integration across departments, an enterprise data warehouse makes sense. Centralizing all your data helps lock down common definitions and cuts down on scattered information.

If quick and easy analytics matter most, you'll want a dimensional approach. Using star schemas and dimensional models lets people dig into business data by things they already know like time, customer, product, or location (Kimball & Ross, 2013).

Data marts, on the other hand, bring more flexibility and can be set up fast. They're good for specific departments that need tailored analytic tools. The catch? Stand-alone data marts can create silos. Building data marts off a main integrated warehouse tends to keep things more consistent.

The three-tier architecture isn't really a rival to the Inmon or Kimball models but it works more as a framework. You can fit centralized, dimensional, or hybrid models inside it. Its main strength is separating data storage, analytics, and presentation into clear layers.

Hybrid architecture looks good on paper for big, complicated organizations. It mixes the benefits of different models: you get centralized oversight but departments still have the freedom to use their own dimensional models and analytical tools.

Table 1: Theoretical Comparison of Data Warehouse Architecture Models

Factor	Enterprise/Inmon	Dimensional/Kimball	Data Mart	Three-Tier	Hybrid
Enterprise integration	Very High	High	Low–Moderate	Very High	Very High
Query performance	High	Very High	High	High	Very High
Governance	Very High	High	Low–Moderate	Very High	Very High
Flexibility	Moderate	High	Very High	High	Very High
Implementation speed	Low–Moderate	High	Very High	Moderate	Moderate
Business-user usability	Moderate	Very High	Very High	High	Very High

Scalability	Very High	High	Moderate	Very High	Very High
DSS suitability	Very High	Very High	High for departmental DSS	Very High	Very High

Bottom line: there's no one-size-fits-all here. Enterprise warehouses excel at integration and control. Dimensional and data-mart setups are better for usability and speed. Hybrid models cover the widest range of needs.

Architecture and Decision Quality

A key takeaway here is that your data warehouse's architecture shapes decision quality, but indirectly.

Here's how it plays out:

Better Architecture → Tighter Data Integration → Higher Data Quality → Better Analytics → Smarter Decisions

Get your architecture wrong, and you risk:

- Duplicated data
- Confusing or inconsistent terms
- Gaps in your historical info
- Sluggish analytics
- Frustrating data access
- Eroded trust in reports

Build it right, and you get:

- More information at everyone's fingertips
- Consistent definitions
- More flexibility with analytics
- Strong historical reporting
- Quick reports
- Leadership that actually trusts the data

In short, architecture doesn't make decisions for you as it just gives you better, more accessible information to work with.

What to Think About When Picking an Architecture

Before choosing a structure, you'll want to weigh a few factors:

- Organizational Size

If you're running a big company with many divisions, you'll need strong integration and firm data governance.

- Data Complexity

The more varied and scattered your source systems are, the more robust your integration tools need to be.

- Analytical Needs

If your teams count on advanced OLAP, forecasting, or multi-dimensional analysis, your setup has to support that level of analysis.

- Scalability

Your architecture should be able to grow as your data, users, and analytics needs expand.

- Governance

If you're in a highly regulated field, you'll probably need centralized rules for data definitions, access, security, and quality.

- Implementation Speed


If time is short, the dimensional or data-mart routes roll out faster. Still, don't sacrifice long-term data governance for a quick launch.

PROPOSED THEORETICAL FRAMEWORK

After digging into the literature and comparing different studies, this research suggests that several information-management capabilities sit between architecture and how well decision support systems (DSS) work.

The framework sees architecture shaping DSS effectiveness through things like data integration, data quality, data accessibility, analytical performance, information consistency, and governance.

Diagram 5: Proposed Theoretical Framework



Organizational requirements form the basis of deciding which data warehouse architecture is most suitable. The chosen architecture then sets the scene for information capabilities like integration quality accessibility, performance, and governance. These are all factors in the enhancement of analytical capabilities like OLAP, business intelligence, forecasting, and what-if analysis. Higher quality analysis in turn leads to more effective Decision Support Systems and ultimately perhaps better decision-making in the organization.

CONCLUSION

Data warehouse architecture is at the heart of today's Decision Support Systems. This study makes it clear: a data warehouse isn't just about storing loads of data but it's about turning messy, scattered information into something integrated, reliable, easy to access, and truly useful for analysis.

The research breaks down five big approaches: enterprise data warehouse architecture, dimensional architecture, data-mart architecture, three-tier architecture, and hybrid architecture. Each style has its strong suits.

Enterprise architecture focuses on pulling everything together, keeping things consistent, and making sure governance is in place. Dimensional models make analysis and queries faster and easier. Data marts let you move quickly and adapt, though if teams build them in isolation, you end up with silos. The three-tier style separates the heavy lifting that is data management, analysis, and display so it all stays organized. Hybrid architecture mixes and matches, making it a good fit for organizations dealing with lots of complexity.

There's no one-size-fits-all answer here. The right setup depends on things like how big your organization is, how complicated your systems are, how much data you're dealing with, what kind of analysis you need, how much control matters to you, and what kinds of decisions you're making.

The main takeaway looks like this:

Data Warehouse Architecture → Information Capabilities → Analytical Capability → DSS Effectiveness → Decision Quality

In other words, the architecture sets everything in motion, but it's not the end goal. What really matters is giving decision makers information they can trust when they need it.

Looking ahead, research can test this framework in the real world by studying how different architectural features actually affect things like information quality, analytical performance, user satisfaction, how well DSS works, decision speed, and overall results for the organization. Researchers could use quantitative methods like SEM, dive into case studies, or compare organizations to see if the theory holds up.

REFERENCES

1. Chaudhuri, S., & Dayal, U. (1997). An overview of data warehousing and OLAP technology. *ACM SIGMOD Record*, 26(1), 65–74. <https://doi.org/10.1145/248603.248616>
2. Chowdary, V. G. R., Rao, K. S., Uppu, A., Kumar, V., Mukherjee, S., & Roy, R. (2026). The extent and impact of artificial intelligence adoption in India's logistics sector. *Lecture Notes in Electrical Engineering*, 134–146. https://doi.org/10.1007/978-3-032-25038-4_12
3. Inmon, W. H. (2005). *Building the data warehouse* (4th ed.). Wiley.
4. James, L. T., Kaushik, A., Singh, V. K., Kumaravel, S. K., Jasim, J. A., & Hasan, R. (2025). Neural networks-based fraud detection system for secure digital banking transactions. In *2025 International Conference on Recent Innovation in Science Engineering and Technology (ICRISET)* (pp. 1–6). IEEE. <https://doi.org/10.1109/ICRISET64803.2025.11252084>
5. Kimball, R., & Ross, M. (2013). *The data warehouse toolkit: The definitive guide to dimensional modeling* (3rd ed.). Wiley.
6. Kimball, R., Ross, M., Thornthwaite, W., Mundy, J., & Becker, B. (2008). *The data warehouse lifecycle toolkit* (2nd ed.). Wiley.

7. Palopoli, L., Terracina, G., Ursino, D., & Vocaturo, G. (2002). A novel three-level architecture for large data warehouses. *Journal of Systems Architecture*, 47(11), 937–958.
[https://doi.org/10.1016/S1383-7621\(02\)00056-5](https://doi.org/10.1016/S1383-7621(02)00056-5)
8. Ponniah, P. (2010). *Data warehousing fundamentals for IT professionals* (2nd ed.). Wiley.
<https://doi.org/10.1002/9780470604137>
9. Roy, R., Das, T., Pal, D., Singh, V. K., & Mukherjee, S. (2026). The dark side of ChatGPT usage and its impact on students: An empirical evidence. *Lecture Notes in Electrical Engineering*, 178–187.
https://doi.org/10.1007/978-3-032-25038-4_16
10. Srinivasa Rao, K., Chowdary, V. G. R., Uppu, A., Kumar Singh, V., Mukherjee, S., & Roy, R. (2026). Big data analytics and industrial performance: A strategic perspective. *Lecture Notes in Electrical Engineering*, 158–167. https://doi.org/10.1007/978-3-032-25038-4_14