



Scale your data lakehouse architecture with the power of Amdocs logical data model (aLDM)

Eran Katz, Data Architecture Lead, Amdocs
Yossi Hanein, Solution Architect, Amdocs

Understanding data lakehouse architecture

As AI transforms the telecommunications landscape, communications services providers (CSPs) face the emerging challenge of identifying a data architecture that goes beyond just supporting agility and scalability. To manage the increasing complexity and volume of data, they require a solution that guarantees reliability and usability. Data lakehouse architecture emerges as the ideal approach, combining the cost-effectiveness and flexibility of data lakes with the performance and structure of data warehouses.

Data lakehouse architecture is particularly suited for business intelligence, data science and GenAI applications, serving to facilitate efficient data management and ensure consistent quality across both structured and unstructured data. This empowers GenAI developers to build and refine models more effectively, supporting a wide array of use cases and development phases from industry taxonomy to conversational AI.

Successful adoption, however, requires more than just a new technical setup: It demands a shift in logical data architecture and data modeling practices. Without this shift, CSPs could face low adoption rates, high costs and a failure to achieve the intended business value. Indeed, this strategic transformation is crucial to avoid significant investments that do not yield full business benefits. A comprehensive approach that includes advanced data modeling is therefore crucial.

A key to leveraging the architecture's full potential lies in organizing the **bronze (raw data)** layer and integrating it into the **silver (curated data)** and **gold (analytical data)** layers for broader usability and integrity. A unified data model is also critical. Without it, the distinction between the silver and gold layers could lead to inconsistencies and complexity, impeding seamless data flow and resulting in discrepancies in analytical outputs – leading to hindered data-driven decision-making. For these reasons, establishing a cohesive data model that spans all layers is essential, ensuring consistency, reliability and efficiency.



How logical data models support data lakehouse architecture

This section outlines the principles and best practices for designing logical data models to effectively support a data lakehouse architecture.

General principles for data lakehouse logical data modeling

Data consistency:

Ensure consistent data representation across all layers using the same business language, particularly between the silver and gold layers. This will maintain data integrity and reliability, while supporting unified understanding and interpretation across operational and analytical contexts.

Modularity:

Design modular data models to enhance flexibility and scalability, enabling components to be independently added, removed, or modified without impacting the entire system.

Data lineage:

Track data origin and transformations with clear lineage information to ensure transparency, facilitate traceability and support compliance and auditing.

Separation of data layers:

Clearly differentiate between raw, integrated and analytical data within the logical data model to simplify management and governance by applying appropriate policies and controls to each data type.

Reusability:

Develop reusable entities and schemas to reduce redundancy, improve consistency and accelerate development timelines.

Telecom industry expertise:

Leverage deep industry knowledge to understand the unique data requirements and challenges in telecom. This expertise is crucial for creating a comprehensive and integrated data model that addresses the diverse and complex data needs of various lines of business, ensuring readiness for current and future use cases.

Domain-specific integration:

Source expertise in BSS, OSS and network systems to integrate data from network operations, customer interactions and billing systems, to build a cohesive data model that facilitates seamless data flow and enables comprehensive data consumption across data lakehouse architecture layers.



The critical role of aLDM in data lakehouse success

aLDM Overview

Amdocs' Logical Data Model (aLDM) is a well-established and widely adopted solution that was designed to address the unique challenges faced by CSPs in managing the increasing complexity and volume of data. It's a multi-layered data model certified by TM Forum, encompassing all telecom lines of business and products. Optimized for modern technologies and near real-time processing, the model supports over 60 operational and analytical domains, and is compatible with both Amdocs and non-Amdocs platforms.

aLDM delivers a host of beneficial business outcomes, including:

Faster time to market:

Speeds up data lakehouse implementations using prepackaged assets, enabling domain teams to quickly implement aLDM data products.

Enhanced governance and compliance:

Comes with extensive documentation and metadata, including detailed mapping to Amdocs and other BSS, OSS and network systems, facilitating easy data product production and consumption while enhancing data governance.

Easier adoption:

Includes predefined KPIs based on CSP business requirements and uses intuitive business language, making it easier for domain team and business user adoption.

Greater platform flexibility and choice:

Operates efficiently on-premise or in the cloud.

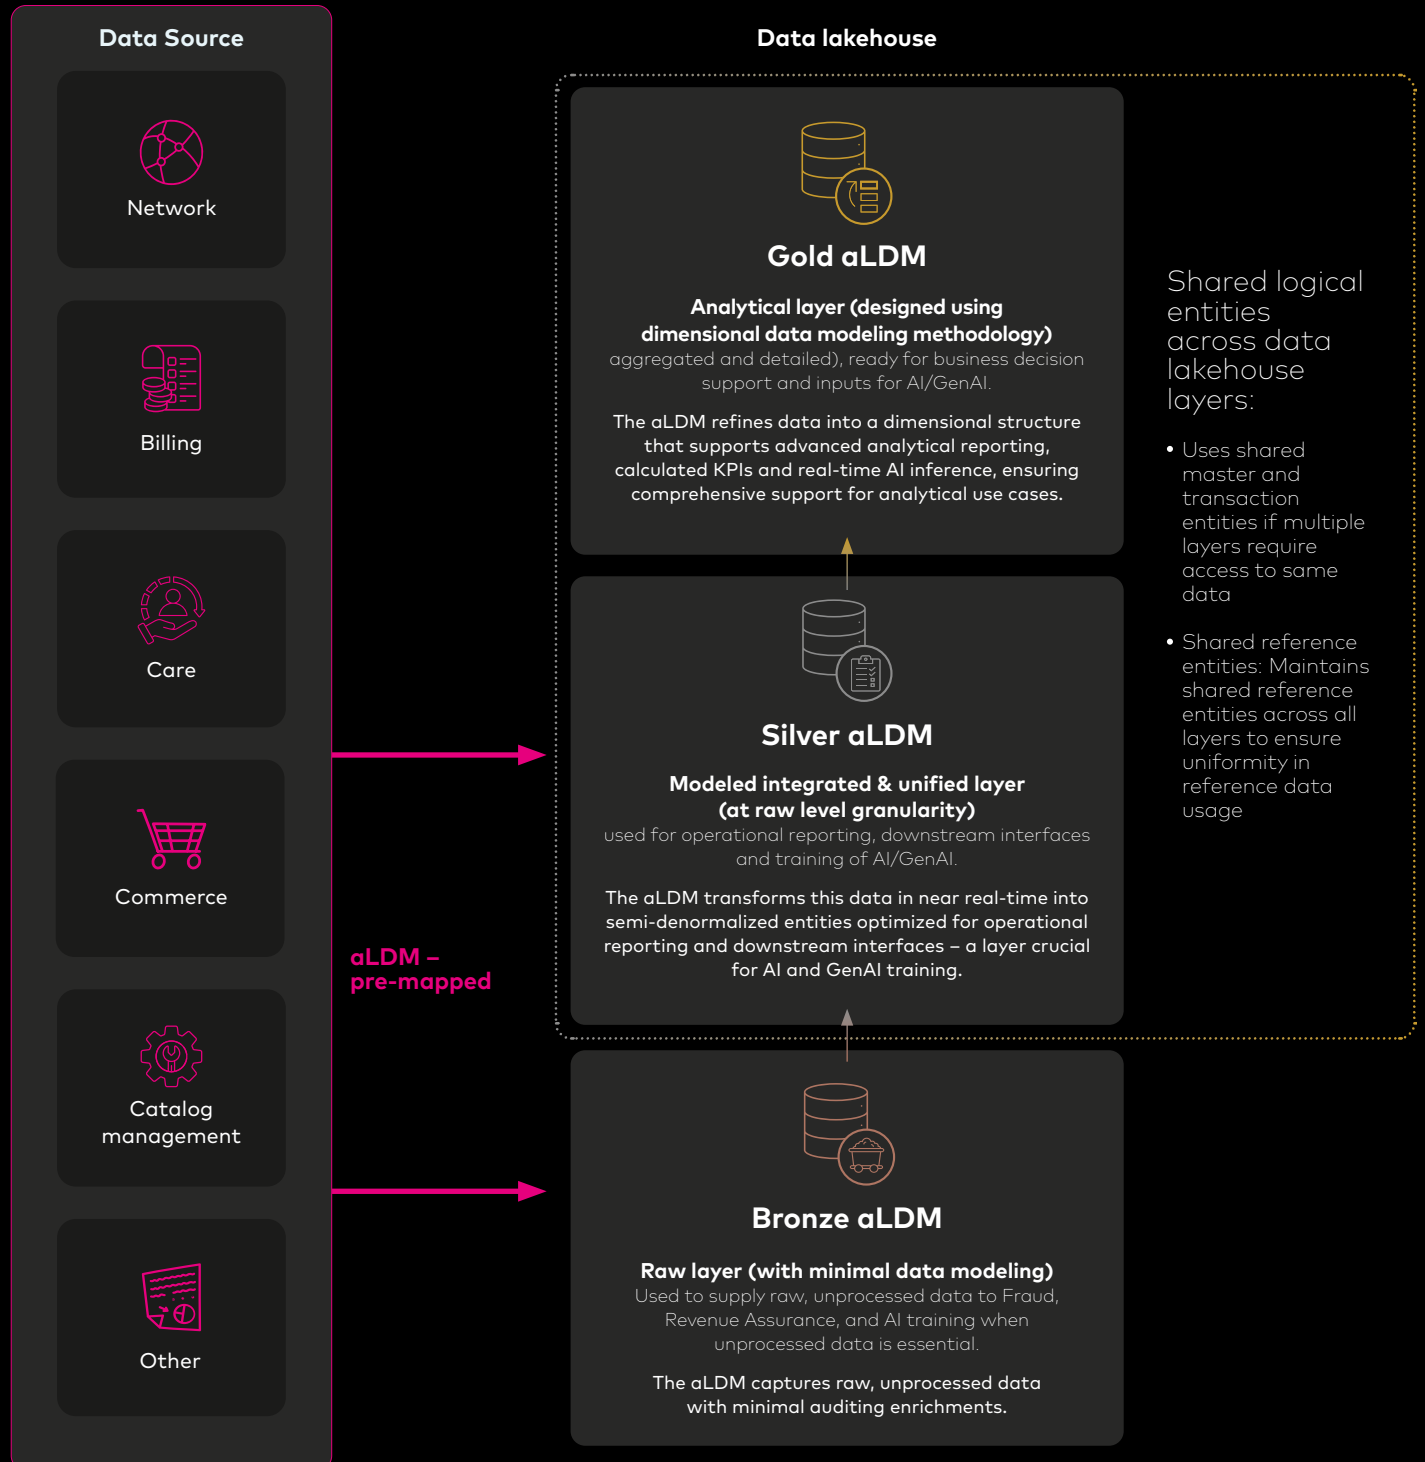
Near-real-time reporting:

Supports modern data warehouse architectures, enabling near-real-time operational reporting.



Layer-specific logical data modeling principles

aLDM facilitates data lakehouse implementation by organizing data into distinct **Bronze**, **Silver** and **Gold** layers, each designed for a specific purpose:



aLDM roles on each layer



aLDM Bronze layer



Raw entity representation:

Captures BSS, OSS and network-ingested data in its raw form, reflecting the source data structure without business transformations.



Raw training data for AI/GenAI:

Supplies raw, unprocessed data for AI and Gen AI training when unprocessed data is essential.



aLDM Silver layer



Semi-denormalized entities:

Created using a balanced data model approach that includes certain normalization for flexibility but retains limited denormalized structures to support operational reporting and downstream interfaces.



Operational focus:

Ensures the model supports near-real-time data processing and transformations to meet the needs of operational reporting and downstream interfaces.



Transformation and unifying rules:

Includes only transformations and unifying rules that do not require business rules implementation.



Training data for AI/GenAI:

Supplies high-quality, integrated data necessary for training AI and GenAI models.



aLDM Gold layer



Analytical reporting:

Supports analytical reporting and self-service capabilities, aligning with the need for denormalized, easy-to-query data structures.



Dimensional structure:

The use of star schemas and Snowflake schemas fits well with the data lakehouse principle of denormalized entities optimized for analytical performance.



Detailed and aggregated entities:

Includes both detailed and aggregated entities to support both granular and summary data for different analytical needs.



AI/GenAI support:

Provides high-quality, refined datasets, such as customer profiles, which AI/GenAI engines use in real time for inference and decision-making.



Calculated KPIs and business rules:

Implements predefined business rules and KPIs, aligning with the need for a dimensional data model that supports analytical business use cases.

aLDM solution benefits

Data quality and consistency:

By avoiding data duplication through effective data governance and using a reference model, aLDM ensures high-quality, consistent data, leading to more accurate and reliable customer interactions.

Efficiency and timeliness:

Reduced implementation costs and pre-mapped data sources speed up development of data products and services, resulting in faster response times to business users' requests and quicker innovation.

Immunity to frequent changes:

As business needs evolve, aLDM's adaptability ensures businesses remain responsive to customer demands, thereby enhancing customer satisfaction by incorporating many anticipated changes and simplifying the implementation of new ones.

Aligning business requirements and logical data models

To fully realize the potential of their data and AI investments, CSPs adopting a data lakehouse architecture need to ensure they implement a well-designed logical data model. By following best practices and principles, they can achieve consistent data representation and transformation throughout the system, enabling efficient data management and comprehensive analytics. Such an approach not only supports scalability and data integrity but also aligns with business requirements.

The Amdocs Logical Data Model (aLDM) underpins the data lakehouse strategy by adhering to principles that ensure comprehensive and flexible data structures across various layers. aLDM's structured methodology for integrating and processing data seamlessly aligns with the data lakehouse's need for a unified data model that supports both operational and analytical workloads, leveraging its robust data management capabilities across diverse business domains.





Amdocs helps those who build the future to make it amazing. With our market-leading portfolio of software products and services, we unlock our customers' innovative potential, empowering them to provide next-generation communication and media experiences for both the individual end user and enterprise customers. Our approximately 29,000 employees around the globe are here to accelerate service providers' migration to the cloud, enable them to differentiate in the 5G era, and digitalize and automate their operations.

Listed on the NASDAQ Global Select Market, Amdocs had revenue of \$4.89 billion in fiscal 2023.

For more information, visit Amdocs at www.amdocs.com