Eckerson Group | Report

Data Fabric

The Next Step in the Evolution of Data Architectures

JAY PISCIONERI

JANUARY 2023

RESEARCH SPONSORED BY

TIBCO™

THIS PUBLICATION MAY NOT BE REPRODUCED OR DISTRIBUTED WITHOUT ECKERSON GROUP'S PRIOR PERMISSION.

About the Author

Jay Piscioneri has over 25 years of experience in data technologies including data warehousing, business intelligence, data quality, and data governance. He's worked with organizations in a wide variety of industries helping them plan and build their next generation data platforms. As a leader of many initiatives, Jay brings extensive experience in the nuts-and-bolts of implementing data solutions, as well as the organizational challenges that go with adopting new priorities, processes, and tools.

About Eckerson Group

Eckerson Group is a global research and consulting firm that helps organizations get more value from data. Our experts think critically, write clearly, and present persuasively about data analytics. They specialize in data strategy, data architecture, self-service analytics, master data management, data governance, and data science. Organizations rely on us to demystify data and analytics and develop business-driven strategies that harness the power of data. **Learn what Eckerson Group can do for you!**

About This Report

Research for this report comes primarily from numerous briefings with software vendors. This report is sponsored by Nexla, InterSystems, and TIBCO, who have exclusive permission to syndicate its content.

Table of Contents

Executive Summary 4

Evolution of Data Architecture 6

Definition of Data Fabric 8

Potential Benefits and Pitfalls of Data Fabric 10

Characteristics of Data Fabric 11

Implementation Options 17

Conclusion: Data Fabric and Beyond 19

About Eckerson Group 20

About the Sponsor 21

Executive Summary

Over the years, organizations have tried many architectural approaches to deliver consistent, unified views of enterprise data to support analytics and other business needs. However, rapid changes brought on by e-commerce, social media, cloud computing, and an explosion in user demand have created new data challenges that previous approaches are not designed to handle.

In the 1990s, data warehousing emerged as the dominant architecture to consolidate enterprise data into a dimensionally-modeled, centralized repository for structured data to support business intelligence requirements. In the 2010s, data lakes emerged to centralize any type of data, not just structured, into a single repository adaptable enough for new uses in data science. But their extreme flexibility led to data quality issues and duplication of analytical effort.

More recently, the data fabric has emerged to provide a unified view of disparate and distributed data and to support any type of workload, from business intelligence to ad hoc analytics to data science. Data fabric uses metadata and artificial intelligence/machine learning (AI/ML) to automate data management functions such as onboarding new data sources and managing metadata. It also simplifies combining data of different types and from different locations to provide seamless access to enterprise data.

Key Takeaways

- *Data fabric is an architectural approach that uses metadata, machine learning, and automation to provide a unified view of enterprise data regardless of its format or location.*
- *Data fabric does not replace data warehouses and data lakes but rather encompasses them to support business intelligence, data science, and embedded analytics.*
- *Data fabric seeks to significantly reduce time to insight: the amount of time it takes to use data to get to an "aha" moment that informs business action.*
- *Data fabric incorporates AI-driven processes that automate many aspects of data management such as discovery, cataloging, preparation, validation, and monitoring to help data teams keep up with the never-ending demand for data.*
- *A data fabric is not a single product that you can buy. The options are to buy pre-integrated tools from one vendor or buy best-of-breed components from multiple vendors and integrate them yourself.*
- *Data fabric is best suited for organizations whose data landscape is growing rapidly, with a variety of data formats stored in multiple locations, and the need to democratize data access to meet many analytical demands.*

Recommendations

- *As with any new technology, companies considering a data fabric must carefully evaluate their needs, have realistic expectations, and experiment with small implementations before fully committing.*
- *Assess your organization's data management maturity to determine what areas of improvement will have the greatest business impact. Focus on data fabric components that address those areas.*
- *If your organization doesn't yet have a significant investment in data management tools—data cataloging, preparation, quality management—consider an integrated suite of components. This approach allows you to focus on delivering value faster, but means committing to a single vendor's offerings.*
- *If you do have some data management tools that are working well, then acquiring components that address your organization's needs and building automated integrations yourself is a better option.*
- *Implementing a data fabric requires a commitment of many different skills, such as data engineering, data science, software engineering, networking, security, and configuration management. So you need to determine if you have the required skills available or how you can build up your team.*

Evolution of Data Architecture

Data architecture has evolved significantly in recent years. But while the scope and complexity of data architecture has changed, its basic purpose has not. Data architecture translates business needs into system designs that deliver data throughout the enterprise to support current and future requirements.

Data architecture translates business needs into system designs that deliver data throughout the enterprise to support current and future requirements.

Generations of Data Architecture

Three generations of data architecture—data warehouses, data lakes, and data fabric—dominate how organizations manage and use their data (see figure 1). Data warehousing arrived on the scene first as the initial approach to creating a unified view of enterprise data. Then ecommerce and social media appeared, leading to an explosion in the volume and variety of data formats and generating data needs that warehouses were not designed for. Data lakes arose to pick up where data warehouses left off, providing much more flexibility to deal with new forms of data. But their flexibility led to many interpretations and conclusions, undermining a single version of the truth.

These shortcomings gave rise to the data fabric, whose objective is to combine the best of previous architectures as it deals with modern data challenges. Let's review each stage in this evolution.

Figure 1. Evolving Data Architectures

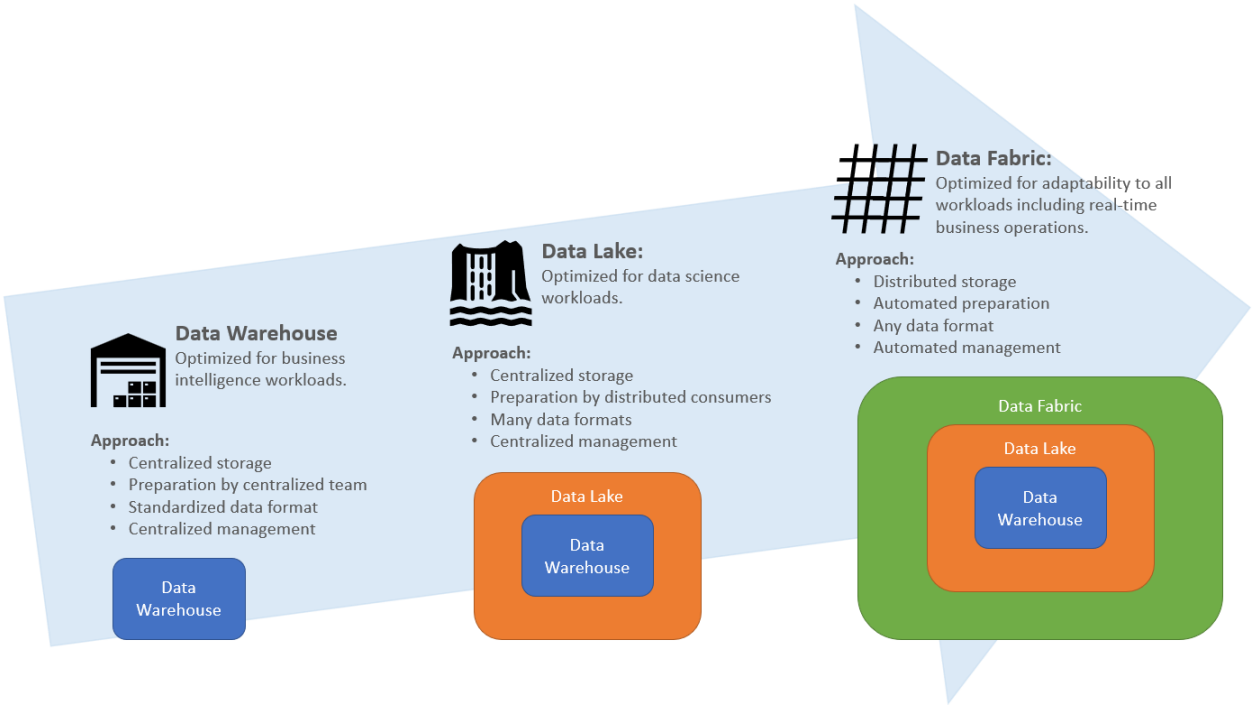

Data Warehouse

Data warehouse architecture was conceived in the 1980s to solve the same kinds of problems that we have today. A data warehouse is a database with centralized storage and uniform formats. Its purpose is to bring together data from different systems for a unified view of data to support management decisions. Many organizations still use data warehouses for this purpose. The primary consumers are analysts who produce business intelligence reports and dashboards.

A centralized data team integrates data in the data warehouse by transforming and structuring it to answer predefined categories of inquiry. This works well when analytic requirements are understood in advance, but it makes data warehouses unsuitable for ad hoc data exploration. Also, since a centralized team has to model and prepare data for consumption, business requests for data often take a long time to be fulfilled.

Data Lake

Data lakes arose to deal with the volume and variability of modern data and to support data science use cases, both of which data warehouses weren't designed to handle. A data lake brings data from distributed locations and stores it centrally, in the native format of the source system, without applying transformations. This provides the flexibility that data scientists need to explore multiple hypotheses and find unexpected insights. And it supports analytics functions integrated into applications.

The job of preparing data for particular uses is decentralized, meaning that it's up to the party who needs data to prepare it accordingly. Not only does this require data consumers to have a high degree of data engineering skill, it also leads to many interpretations of often undocumented source data and therefore conflicting conclusions.

Data Fabric

Neither data warehouses nor data lakes solve all enterprise data problems, but they each have a sweet spot in which they excel. Data fabric seeks to encompass both architectures, not replace them, to support business intelligence, data science, and integration with applications to enable embedded analytics.

Data fabric accommodates both structured and unstructured data. Rather than creating a centralized storage location, it emphasizes leaving data in place and connecting to it as needed for both analytics and real-time business operations. Another key difference is the degree of automation that data fabric employs to reduce the workload of managing data. For example, data preparation is still decentralized, but non-technical data consumers use low-code/no-code tools to apply complex transformations that put data in the form they need for analysis.

Definition of Data Fabric

Data fabric is an architectural approach that uses metadata, machine learning, and automation to weave together data of any format in any location and make it easy to find and use by people or systems. It unifies the separate functions of data management—integration, preparation, cataloging, security, and discovery—into a cohesive process through intelligent automation.

Data fabric uses metadata, machine learning, and automation to weave together data of any format in any location and make it easy to find and use by people or systems.

While the functions of data management are not new, the degree of automation and integration prescribed by data fabric are. Data fabric encompasses these data management functions:

- **Integration.** Connecting to data sources to make data available in its native format.
- **Preparation.** Cleaning, normalizing, and reorganizing (i.e., transforming) data to suit a particular purpose.
- **Cataloging.** Creating, viewing, and managing metadata about data assets.
- **Access management.** Providing appropriate access based on what data is and who needs it.
- **Discovery.** Making data searchable by name, definition, concept, or related terms.

Data fabric automates functions such as onboarding new data sources, managing metadata, and combining data from different sources. It creates an abstraction layer that eliminates the complexity of working with distributed data and disparate formats. It provides seamless access to enterprise data, which helps data teams keep up with the exploding volume of user requests. Data fabric covers data at rest and data in flight. Most implementations focus on data at rest given the prevalence of analytics use cases.

What a Data Fabric Is Not

The data fabric incorporates familiar data management functions. However, it is distinct from components such as a data virtualization tool, data catalog, logical data warehouse, data lakehouse, and data mesh. It is also distinct from the discipline of DataOps. Let's consider the similarities and differences for each.

- **A data virtualization tool.** A data virtualization tool bridges data across distributed locations and varied formats without requiring data to be copied from the source to central location. As we'll explain later, data virtualization is one method of data abstraction, which is a critical function in data fabric. Abstraction makes data easier for consumers to find and use by removing the complexity of distributed sources and varied formats. Data fabric also includes the ability to create physical abstractions of source data through data pipelines.

- **A data catalog.** A data catalog stores and manages metadata. While data catalogs are evolving to include more functionality, their sweet spot has been storing business and technical metadata data for reference. Data cataloging is an integral part of data fabric, but data fabric actively uses all kinds of metadata, not just for reference but also to drive automation that reduces workloads for data teams and data analysts. Data fabric catalogs data in streams, APIs, files and other data systems that aren't necessarily covered by traditional catalogs.

- **A logical data warehouse.** A logical data warehouse is similar to data virtualization because it makes data from different sources and locations addressable for analytics. Data fabric is broader because, unlike a logical data warehouse, it also handles transactional applications and real-time streaming data.

- **Data lakehouse.** The data lakehouse is a form of data abstraction that handles data of many formats like a data lake, but also provides structured data sets that make data easier to use for data analysts. While it can be part of a data fabric, it lacks the automation functionality of data fabric.

- **Data mesh.** Data fabric and data mesh both seek to accelerate and expand support for all data use cases. However, data mesh emphasizes organizational changes, while data fabric focuses on the technology needed to design and implement data solutions at scale. Data mesh's decentralized approach calls for the teams whose systems create data to build data products that meet data consumer needs. Data fabric can provide the architecture for automating production and consumption of data products.

- **DataOps.** DataOps has many of the same objectives as data fabric. However, DataOps includes management of code, configuration of environments, builds, and deployments that data fabric does not address. Data fabric can contribute its copious active metadata to enhance DataOps' monitoring, validation, and operational functions. (For more information on DataOps, see the Eckerson Group report **DataOps Deep Dive: Different Approaches to the DataOps Platform**.)

Many of the components above are part of a data fabric. However, they do not by themselves define a data fabric. Integration between components and automation that reduces data management workloads completes the data fabric picture.

Potential Benefits and Pitfalls of Data Fabric

Benefits

Data fabric's objective is to increase the value organizations derive from their data in light of increasing data volumes, sources, formats, and uses. Business teams need to make fast data-driven decisions and data teams need new ways to keep up with the demand for data. The benefits of data fabric address these challenges.

- **Faster time to insight.** Time to insight describes a measure of data management effectiveness: the amount of time it takes to use data to get to an "aha" moment that informs business action. Data fabric seeks to significantly reduce time to insight. In addition to intelligent automation and enhanced discoverability, a data fabric hides the complexities of enterprise data so that consumers don't have to know where data is stored or what format it's in. Removing those details lets data consumers focus on conducting their analyses rather than wrangling the data they need.
- **Reduced data management workload.** Data fabric uses AI and machine learning to automate many data management tasks, such as cataloging new data sources and assisting data preparation. Automation can eliminate time consuming manual steps that make it difficult, if not impossible, for data teams to keep up with demand for data.
- **More effective data discovery.** A fully populated and up-to-date data catalog, combined with advanced search capabilities, enables data consumers to know what data is available to them. It's like having Google for enterprise data. Keyword search functions along with rich documentation about data assets enables business analysts, data scientists, and data teams to find data resources and evaluate their value to a given use case.

Pitfalls

Data fabric covers the full gamut of enterprise data management functions. The impact can be transformative if the benefits above are realized. However, with many people, processes, and technologies affected, the risks are big as well.

The impact of data fabric can be transformative. However, with many people, processes, and technologies affected, the risks are big as well.

- **Big scope, big impact, big risk.** Data fabric affects many data management functions, so it would be easy to take on too much at once and take too long to deliver improvements. Also, the promise of a utopian data future means stakeholder expectations could run wild if not managed, resulting in disappointment and wavering commitment to the end result.

- **Data fabric is not a product.** There is no data-fabric-in-a-box you can buy. Data fabric comprises a set of tools, some of which you may already have, and integrations between them. Be prepared for multiple tool evaluations and feature comparisons, as well as planning and configuring integrations.
- **Many vendors, many approaches.** Data management software companies are rushing to align their offerings with the data fabric concept. Some companies offer a suite of tools with the most mature components reflecting their historical area of expertise. Others focus on one product that plays a limited role in a data fabric. This complex product landscape makes the data team's job of designing solutions much harder. However, the product market is maturing. Expect to see more aspects of data fabric functionality integrated into product suites.

Characteristics of Data Fabric

There is no one way to build a data fabric. But for a data fabric to achieve its objectives of making data easy to find, use, and manage, it should consist of the following core components (see figure 2):

Figure 2. The Core Components of Data Fabric

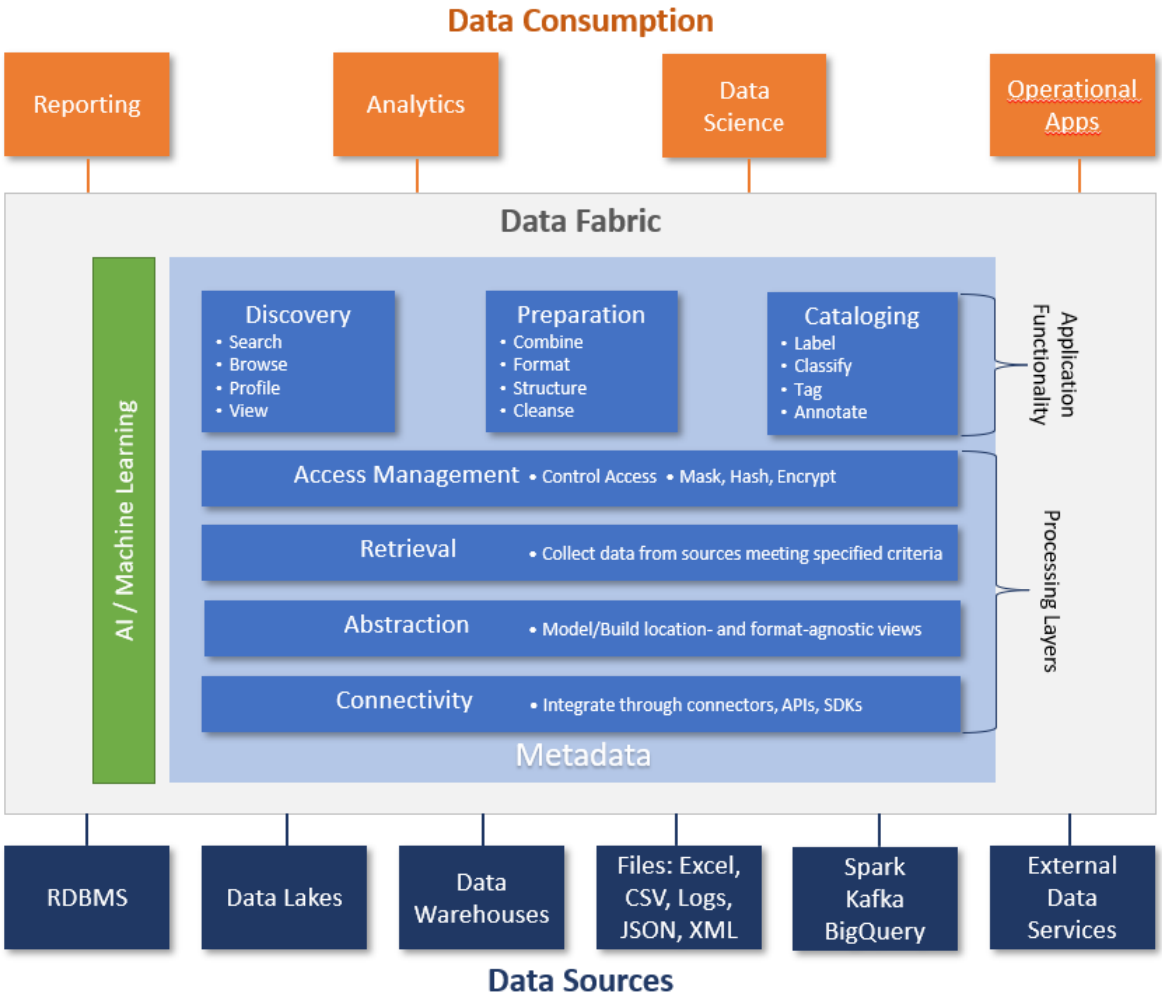

Metadata

Metadata is data that describes data. It serves as the foundation of the data fabric. Data fabric uses metadata to create an abstraction layer that hides the complexity of working with distributed data and disparate formats. It also uses metadata to automate functions such as data cataloging and preparation. Data fabric uses and generates metadata that spans four broad categories.

- **Technical metadata.** Technical metadata includes the characteristics of data that systems need to work with it, such as format, type, length, and location.
- **Business metadata.** Business metadata uses business language to provide context for the data that appears in applications, reports, and dashboards. It includes elements such as terms, definitions, and classifications.
- **Operational metadata.** Operational metadata provides information on how data is used and what happens when it's used. It includes information from a variety of sources such as execution logs, rule engines, error logs, and audit registers, with dates and times of events.
- **Social metadata.** Social metadata captures the enrichment that comes from user collaboration. It includes elements such as tags, ratings, and comments.

Metadata serves as the foundation of the data fabric. It's used to hide the complexity of working with distributed data and disparate formats, and to automate functions such as data cataloging and preparation.

Application Functionality

The interactive functions of discovery, preparation, and cataloging enable business users, power users, and data teams to find and work with data.

- **Discovery.** Discovery is the data fabric function that makes data easy to find. Without effective discovery, the data needed for a particular purpose is lost among terabytes or petabytes of other data. It enables users to browse through or search for available data sets, and preview data from different sources to better understand whether it helps them in their analysis. Browsing allows users to navigate categories of data, such as business subject or data source, often presented as a hierarchy. Search functions vary from basic key word search to AI-assisted search with natural language processing (NLP).

- **Preparation.** Preparation is the data fabric function that enables power users to access data from one or more sources without having to deal with different formats or distributed locations. For example, data scientists use preparation functions to gather and organize the data they need for their statistical models, while data teams use preparation functions to conform and aggregate data for reporting.
- **Cataloging.** Cataloging is the data fabric function that enables users to create, view, and manage metadata. It's the interface that displays information about data resources during discovery and allows users to enrich the metadata with tags, comments, ratings, and certifications. Cataloging functions also capture how data is used, such as how often tables are accessed, what queries they appear in, and who is using them. This helps consumers evaluate the trustworthiness of data resources available to them.

Processing Layers

The operational work of a data fabric involves providing connections to data sources, managing workflows, abstracting data to create a unified view, and ensuring access to data is appropriately controlled.

- **Connectivity.** The connectivity layer automates access to a data source through an application programming interface (API) that enables communication between data sources and the data fabric. Data fabric products include many source-specific connectors for different platforms, such as **Snowflake**, **Salesforce**, **AWS S3**, and **MongoDB**, to name a few.
- **Abstraction.** The abstraction layer hides the disparate formats and distributed locations of data sources from consumers making it easier for them to use enterprise data. In a data fabric, consumers don't interact directly with data sources. They work with abstracted versions of source data—we'll call them abstracted data objects or ADOs—so they don't have to worry about where the data is or what platform it's coming from. ADOs can be virtual, meaning that they reference data stored in the sources. They also can be physical, meaning that data is copied to another location in the fabric, usually for faster retrieval.
- **Retrieval.** The retrieval layer consists of a query engine that collects data from source systems that meet specified criteria for a given purpose. Every time a user runs a report, a data scientist prepares data for use in a model, or a data pipeline process runs to update a physical ADO, the retrieval layer determines where the data is and how to optimize its delivery.
- **Access management.** The access management layer determines who can use certain data based on the role of the consumer and the classification of data being requested. At a minimum, the security policies of the data source govern who can access its data. Data teams can define additional levels of access control, such as masking, obfuscation, and encryption.

Artificial Intelligence/Machine Learning (AI/ML)

AI and ML help deliver on the data fabric promise by intelligently automating data management workloads that data teams are hard pressed to keep up with manually. Using the fabric's metadata, AI and ML improve scalability by automating many aspects of the processing and application functions. The examples of automation in table 1 reduce the workload of data teams and users.

AI and ML help deliver on the data fabric promise by intelligently automating data management workloads that data teams are hard pressed to handle manually.

Table 1. Examples of AI / Machine Learning in Data Fabric Application Functionality

Function	AI/Machine Learning Process
Discovery	**NLP search.** NLP search enables users to enter a natural language sentence such as "How many customers are likely to churn in the next three months?" to find data resources that can be used to answer the question.
Preparation	**Program code generation.** Preparation is assisted by AI and machine learning processes that automatically generate program code with appropriate joins, filters, sorting, grouping, and subtotaling. AI/ML also improve data operations through self-regulating validation and error detection.
Cataloging	**Automated cataloging.** AI/ML functions automate the generation of technical, business, and operational metadata when a new data source is onboarded, and capture operational metadata about how data is used.

Data Abstraction

It's hard to overstate the importance of abstraction in data fabric. This is the function that delivers data to users, which is the point, after all. So let's define the term. Abstraction is the process of removing details that are not relevant in a certain context in order to emphasize other details that are relevant to a given analysis. Data fabric uses abstraction in a variety of ways. The first step is to remove the details of source data format and location, thus emphasizing the content of the data that the source provides (see figure 3).

Abstraction is the process of removing some details to focus on others that are most relevant to a given analysis.

Figure 3. Data Fabric Creates Abstracted Data Objects

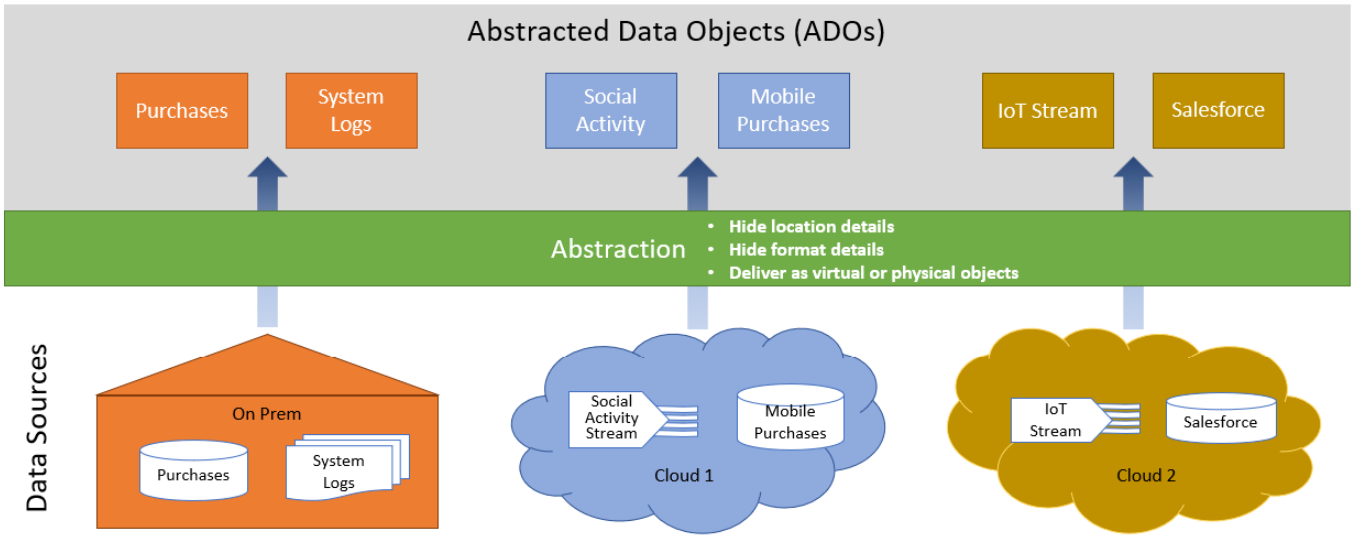

Data Sources to Abstracted Data Objects

The abstraction layer delivers data in either physical or virtual form. Data pipelines create physical abstracted data objects (ADOs), while data virtualization techniques create virtual ADOs. Virtualization is preferable because it reduces cost and risk by not making copies of data. However, sometimes performance issues or security constraints require data to be copied and moved. So, a data fabric needs to do both.

ADOs and Derived Data Views

Abstracted data supports a wide range of analytic uses. ADOs that mirror source objects create a logical data lake that supports the data science need for flexible data exploration and mash-up. In addition, power users can access ADOs through the preparation function to create derived data views that are optimized for specific analytic uses. These derived data views can support multiple scenarios. For example, they can serve as conformed data sets for business intelligence, or serve as prepared data inputs for predictive models (see figure 4).

Figure 4. Derived Data Views

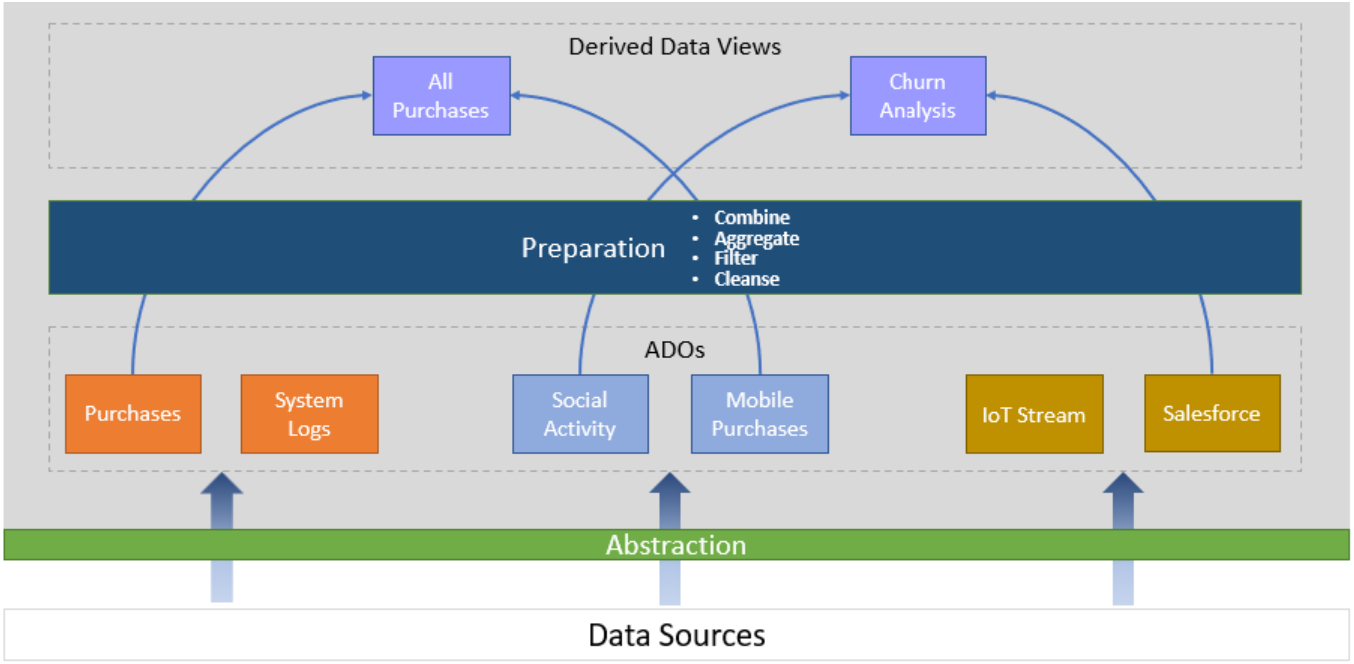

Data Fabric Approaches

There are many approaches to designing a data fabric. The challenges and opportunities your organization faces determine what aspects of data fabric should be emphasized. Here are three approaches, each of which focus on different business problems:

> **Empowering business people.** Empowering business people with access to data so that they can answer questions and make informed decisions is an important objective of data fabric. If this is your primary concern, then the ease of discovery and preparation functions are critical. For example, natural language processing (NLP), a form of artificial intelligence, makes finding data easier. It takes regular sentences or phrases and intelligently parses the words to find relevant data resources.

> **All-purpose analytics.** Supporting data science, embedded analytics, and streamlining data management are broader, more advanced objectives of data fabric. Data scientists need the unaltered data of ADOs (see figure 3 above) to organize and refine data for AI/ML models, and for real-time analytics embedded in applications. Companies with large, complex data environments have a greater need for automated data management functions to keep up with cataloging, data quality, and compliance.

> **Operational data management.** Up to now, we've been discussing how a data fabric supports analytics use cases. But data fabric can also play a role in the transactional world of operations. Managing real-time integration between applications to bridge data silos and reaching to the edge of the data landscape where IoT devices live are examples of transactional use cases that require the most advanced approach to data fabric.

Implementation Options

As described in our "Characteristics" section, a data fabric comprises components such as data catalogs, transformation tools, and query engines. To implement a data fabric, you need many of these components. The question is whether to buy most or all of them from a single "suite" vendor that has designed them to work together in a data fabric. Alternatively, you can buy best-of-breed tools from multiple vendors and integrate them yourself into a data fabric.

The question is whether to buy data fabric components from a single vendor that has designed them to work together or buy best-of-breed tools from multiple vendors and integrate them yourself into a data fabric.

The following criteria can help you choose between a suite and best-of-breed approach. Which option to choose depends on your organization's needs, what your present data architecture consists of, and the resources available.

- **Your organization's challenges.** First, you need to understand the data issues that affect your organization's ability to fulfill its objectives. What are your data strengths, weaknesses, opportunities, and threats? Does your organization have unique challenges that require custom solutions? Know what problems you need to solve and in what priority order.
- **Your enterprise data architecture.** You may have already invested in some of the components of data fabric, such as one or more data catalogs, data integration tools, or data virtualization tools. How deeply embedded are these tools in your current environment? How well are they working? The elements of your data architecture that you want to retain versus those you need to replace will inform your decision.
- **Available products.** The market of data fabric products is evolving quickly. Some companies offer a suite of integrated products that address many data fabric functions. Others focus on one or two core fabric capabilities with APIs to integrate with other products. There are no pure, all-encompassing data fabric products on the market. Most products are from data management incumbents. Each company's data fabric offerings emphasize the functionality that they have developed expertise in.
- **Development resources.** Building a data fabric from multiple best-of-breed components requires a significant commitment of development effort. It also requires many different skills such as data engineering, data science, software engineering, networking, security, and configuration management. Before choosing the best-of-breed approach, determine whether you have the required skills available or how you can build up your team.

- **Budget.** Either approach will cost money. Licensing expenses, staff augmentation, infrastructure upgrades, and cloud storage and compute costs are all likely to be involved in implementing a data fabric. Also plan for increases in spending to account for growth in data storage, processing, and software licensing.

If your organization doesn't yet have a significant investment in data management tools—data cataloging, preparation, quality management—an integrated suite of tools is a better implementation option. This approach allows you to focus on delivering value faster. But it means committing to a single vendor's offerings.

If you do have investments in some tools that are adopted and working well, or if your organization has unique must-have requirements, then choosing individual components and building automated integrations is the way to go. This approach enhances your existing architecture and can meet specific requirements. However, you need enough development resources and the right skills to make it work.

In either case, you should plan to implement your data fabric incrementally focusing first on the functions that will solve your organization's most pressing problems.

Conclusion: Data Fabric and Beyond

Since the introduction of data warehouse architecture decades ago, we've been pursuing the same elusive objectives: having one version of the truth, taking data-driven decisions and actions, and democratizing data. Data fabric is the latest approach to achieving these objectives in today's vast and complex data landscape. It addresses the reality that manual methods of data management cannot scale to meet modern data demands.

Data fabric uses pervasive metadata and AI/ML automation to reduce data management workloads and help teams keep up with never-ending demand. Through advanced discoverability, it enables consumers to find the data they need and understand its value with rich documentation in a fully populated data catalog. Data fabric can reduce time to insight by removing the complications of disparate and distributed enterprise data. This allows users to focus on conducting their analyses rather than wrangling data.

Data fabric will continue to evolve as preceding data architectures did. Products will continue to mature and close gaps in functionality and integration. The scope of data fabric will likely expand to encompass other aspects of data management, such as DataOps, which includes managing program code, configuration of data environments, automated testing, and data quality monitoring. We can also expect

to see a blending of the organizational innovations of data mesh—domain-oriented data ownership and data as a product—with the architectural patterns of data fabric. To realize such opportunities, enterprise data teams should learn the fundamentals of the data fabric and determine how it applies to their environments.

Recommendations

- As with any new technology, companies considering a data fabric must carefully evaluate their needs, have realistic expectations, and experiment with small implementations before fully committing.
- Assess your organization's data management maturity to determine what areas of improvement will have the greatest business impact. Focus on data fabric components that address those areas.
- If your organization doesn't yet have a significant investment in data management tools—data cataloging, preparation, quality management—consider an integrated suite of components. This approach allows you to focus on delivering value faster, but means committing to a single vendor's offerings.
- If you do have some data management tools that are working well, then acquiring components that address your organization's needs and building automated integrations yourself is a better option.
- Implementing a data fabric requires a commitment of many different skills, such as data engineering, data science, software engineering, networking, security, and configuration management. So you need to determine if you have the required skills available or how you can build up your team.

About Eckerson Group

Wayne Eckerson, a globally-known author, speaker, and consultant, formed **Eckerson Group** to help organizations get more value from data and analytics. His goal is to provide organizations with expert guidance during every step of their data and analytics journey.

Eckerson Group helps organizations in three ways:

- **Our thought leaders** publish practical, compelling content that keeps data analytics leaders abreast of the latest trends, techniques, and tools in the field.
- **Our consultants** listen carefully, think deeply, and craft tailored solutions that translate business requirements into compelling strategies and solutions.
- **Our advisors** provide one-on-one coaching and mentoring to data leaders and help software vendors develop go-to-market strategies.

Eckerson Group is a global research and consulting firm that focuses solely on data and analytics. Our experts specialize in data governance, self-service analytics, data architecture, data science, data management, and business intelligence.

Our clients say we are hard-working, insightful, and humble. It all stems from our love of data and our desire to help organizations turn insights into action. We are a family of continuous learners, interpreting the world of data and analytics for you.

Get more value from your data. Put an expert on your side. **Learn what Eckerson Group can do for you!**

Eckerson Group

RESEARCH CONSULTING ADVISORY

GET • MORE • VALUE • FROM • YOUR • DATA

About the Sponsor

TIBCO Software unlocks the potential of real-time data for making faster, smarter decisions. Our Connected Intelligence Platform seamlessly connects any application or data source; intelligently unifies data for greater access, trust, and control; and confidently predicts outcomes in real time and at scale. Learn how solutions to our customers' most critical business challenges are made possible by TIBCO at www.tibco.com.