

Redefining modern data stacks: What's hot, what's not

Abstract

Data – while it's the most critical asset for any business to propel its growth, leveraging its true analytical value without a modern solution can be daunting. A few data-driven enterprises have already embraced modern data stack trends in AI/ML and BI, while others don't want to miss out on riding the modern data wave and be laggards.

Staying competitive in today's fast-paced business landscape requires a modern data stack that can keep up with your data-driven needs. Discover the essentials of a modern data stack and the crucial things to consider before its strategic implementation in our upcoming fireside chat. Join experts from Forrester and Gathr.ai as they provide insights on the latest trends and technologies that make up a state-of-the-art data stack, as well as real-world pragmatic approaches for its implementation.

How do you define a modern data stack, what does it offer to data consumers, data engineering teams, and dataops/devops teams?

The modern data stack is designed for elasticity and scalability. The platform decouples compute, storage, and the state of the data making it possible to support a wider variety of domains and use cases. Organizations are no longer constrained to monolith and bespoke warehouse implementations. Instead, the platform supports centralized, decentralized, and federated workloads. In addition, modern data engineering practices are supported. Data engineers are able to build data products to compose solutions quickly, have interoperability across systems and use cases, and portable across platforms and environments.

What is the buzz/fuss all about? What is so 'modern' about a modern data stack? What are some challenges?

The technical skill bar is lowered. Data engineers to data consumers have no-code/low-code tools to create, deploy, manage and share data pipelines, data sets, connectors, and ML models. However, the modern data stack still supports the use of opensource and code for more sophisticated use cases and solutions. Self-service is supported by data catalogs to find, access, publish, and managed data products and assets. And data governance is baked into the platform for wider democratization of data and insight.

Why is the modern data stack gaining popularity and when is it the right time for a business to implement a data stack?

More organizations are decentralizing their data development and analytic efforts. A. modern data stack is better prepared with scalable infrastructure for an enterprise foundation, and zones for business unit needs. The data estate handles analytic and operational workload while taking advantage of a data fabric designed for multiple batch and real-time data processing and speeds. Data governance is better supported by relinquishing top down policies and controls in favor of federated governance to deploy global policies but iterate and use local or use case specific policies. Cost is another benefit. With modern data stacks running in the cloud, organizations can do more with less, scaling to meet more data producers and use cases.

Is there anything you need to watch out for while setting up MDS?

Lift and shift is a common mistake during migration. For those deploying a new cloud environment, lift and shift increases cost due to poor configuration and optimization of workloads. Thus, cost benefits are unrealized. In addition, more attention is needed to orchestrate and coordinate data consistency between data sources. This is especially true as data is more virtualized to create data sets or satisfy queries and requests. Data flow between data sources, centralized cloud databases, and cache is designed for the consumption scenario, not a centralized data warehouse. In addition, data governance is built into the platform and not an afterthought. Data governance is overwhelmed due to the distributed nature of data and wide selection of use cases and domains supported if applied post migration. Iterate the MDS deployment by use cases and domains to avoid these pitfalls.

Is there a piecemeal approach to MDS for traditional data engineering teams? Does MDS have to be a big bang approach?

MDS by nature is built for iteration. MDS expects continuous evolution of existing use cases and the addition of new ones. Ideally, enterprise data management teams should build a foundation and core of enterprise data, data connectivity, and data policies as a launchpad. Data engineers, data scientists and analysts, data stewards, and data consumers build on top of the foundation. This ensures new data, new data policies, and connections and integration with other sources and systems is purposeful and contextual. Ultimately, this leads to stronger and tangible data monetization for the enterprise.

DEVELOPMENT, USE, AND CITATION GUIDELINES

The following summary of Forrester's guidelines applies to the use of the above post-engagement Q&A copy:

- Forrester Citation Policy states that all Q&A copy and usage, plus format, is subject to review prior to publication.
- Q&A copy must be company, product, and brand agnostic.
- When published, the Q&A must be client branded. Use of the Forrester logo is subject to additional guidelines and approval.
- Q&A copy contains a minimum of three (3) to a maximum of (5) five questions related to the original engagement content (i.e., webinar/video/speech content), with a two- to four-sentence answer per question.
- Q&A copy must link clearly to the analyst engagement that it goes along with and clearly identify the analyst as a guest speaker at that engagement.

- It must be clear that the questions are coming from the client (or from audience members if a post-event document) and that Forrester is answering them.
- The analyst works with the client to identify relevant questions related to the scope of the engagement, which are then reviewed, edited, and approved by the client.
- The analyst develops responses to the questions based on the engagement, their research coverage, and market knowledge.
- The analyst submits the completed copy for editing, citation review, and approval by Forrester.
- The client offers its audience the final deliverable in web (e.g., blog or landing page) or PDF format (e.g., post-event email offer or download). Alternative formats require citations approval.
- Q&A copy can be cited in other materials specifically related to the engagement. It must not be cited in materials unrelated to the engagement (e.g., quoted in a separate white paper or other presentations).



Scan and
start free 14-day trial



Scan to
schedule a demo

