

DataStax

Page 16

WHY APIS ARE  
KEY TO BUILDING ON  
A DATABASE

**database**  
TRENDS AND APPLICATIONS



# MODERNIZING YOUR ENTERPRISE DATA ARCHITECTURE

**database**  
TRENDS AND APPLICATIONS

Best Practices Series





# DATA-TO-VALUE: DESIGNING A MODERN ENTERPRISE DATA ARCHITECTURE THAT DELIVERS MORE FOR THE BUSINESS

## Best Practices Series

Data-driven enterprises do not exist on data on their own; they require an advanced, responsive data architecture to gain ground within their markets. “It’s what organizations are able to accomplish with data that counts, whether that’s improved profitability, increased innovation, or better customer experiences,” observed the authors of a recent study of 311 data executives, “Turning Data Into Unmatched Business Value,” published by Harvard Business Review Analytic Services. The study looked closely at the best practices of leaders in data-driven architecture who have established high-producing data architectures.

Such best practices can be critical, as companies that don’t adhere to an effective data architecture that supports initiatives such as artificial intelligence potentially leave up to \$460 billion in incremental profit on the table, a report out of Infosys Knowledge Institute estimates. “Despite heavy investments,

data and AI are failing to deliver business value,” the study’s authors state.

In both the Harvard Business Review and Infosys studies, authors identified the

*Despite heavy investments, data and AI are failing to deliver business value.*

common characteristics of organizations demonstrating extraordinary value:

**Data leaders have a clear enterprise data strategy:** This is perhaps the most important piece of the data value proposition, the Harvard Business Review authors state. Just about all leaders, 97%, have established such a comprehensive strategy for internal data, versus 59% of all other respondents. In addition, 99% have a similar strategy for incorporating external sources of data, versus 53% of all others. “Clarity around data management

and governance and a clear direction about what an organization wants to accomplish with data-driven insight are both critical. You can collect and clean

and distribute all this data, but if you don’t have any idea what you want to do with it, you’re going to fail. You have to understand both the architecture and the task at hand.”

- **Data leaders share their data across enterprises—and beyond:** Developing a sharing ecosystem of partners and customers “generate[s] greater value from their data,” the Infosys report states. “Data increases in value when treated like currency and circulated

Successful *data-driven* enterprises aren't *built* within single *departments*—  
they require the *involvement* of the *entire enterprise*.

through hub-and-spoke data management models. Companies that refresh data with low latency generate more profit, revenue, and subjective measures of value.”

- **Data leaders build cross-functional teams:** Successful data-driven enterprises aren't built within single departments—they require the involvement of the entire enterprise. That would include “data wranglers, enterprise architects, data scientists, AI experts, and data-savvy businesspeople who can identify the problems that data might solve,” according to the Harvard Business Review authors.
- **Data leaders have made the move from data compliance to data trust:** “Companies highly satisfied with their AI (currently only 21%) have consistently trustworthy, ethical, and responsible data practices,” according to Infosys. “These prerequisites tackle challenges of data verification and bias, build trust, and enable practitioners to use deep learning and other advanced algorithms.”
- **Data leaders continually assess their data competencies:** More than half of data executives in the Harvard Business Review study (58%) report their organizations are capable of connecting datapoints across assets, devices, and services, compared with just 10% of all others. They connect these environments “to access and analyze data in real time; and to automate data-driven insights.” Three-quarters of leaders say “their organization's overall performance and success are strongly linked to each of these capabilities.”

- **Data leaders develop real-time capabilities:** A majority of data leaders (58%) also say their organizations have mature, real-time data access and analytics capabilities versus just 6% of the rest, according to the HBR study. At least two-thirds of leaders (67%) say they have mature automation of data-driven insights using machine learning built into their workflows, compared to only 5% of the rest. “The fastest-growing area in terms of investment is to get to as real-time as possible. That is, in part, a reaction to faster business cycles and the drive toward faster decision cycles,” the HBR researchers state. Nearly a third of leaders (32%) say they have a broad ability to access data across systems and locations versus just 10% of other respondents.
- **Data leaders have a multi-cloud strategy:** Overall, the HBR study shows that 83% of respondents turn to cloud services that allow their organizations to deliver new capabilities. In addition, 68% currently have or are pursuing a multi-cloud strategy with that number likely to increase to 71% in 2 years' time. “Leaders are even more likely to be embracing a multi-cloud approach, with nine out of 10 either having or pursuing that today.”

A number of business drivers are fueling multi-cloud adoption rates, the study also shows. These include greater operational efficiency (cited as a top-three driver by 42% of all respondents), performance optimization (35%), cost optimization (33%), increased agility (31%), risk mitigation (30%), and access to best-in-breed solutions (24%).

- **Data leaders measure, measure, and measure some more:** A majority of data executives overall (80%) say it's important to “measure and report on the business value or business outcomes of their data and analytics investments.” At the same time, the study finds, only 26% say they are extremely effective at measuring and reporting. “Connecting the dots between the capabilities that enable insight and the resulting business value is key,” the report's authors state. Among data leaders, 93% measure their data's business value, and just over half (58%) “indicate that they are extremely effective at measuring and reporting the business impact of data analytics and investments.”

What kinds of metrics are involved? Data leaders have been able to measure significant increases in operational efficiency over the past year (46% say so versus 20% of others), but also saw gains in more strategic metrics of performance, the HBR authors state. “Nearly twice as many leaders as the rest have seen both revenues and profitability increase significantly.”

As both research studies show, enterprise data leaders have designed and built data architectures that connect to and leverage people and resources from across the enterprise. Key to these efforts are open and accessible data and insights that can be deployed in a variety of ways as business needs demand—from artificial intelligence to customer relationship management. ■

—Joe McKendrick

# Why APIs Are Key to Building on a Database

# DataStax

By Ed Anuff, Chief Product Officer at DataStax

An application isn't worth much until it's put into production. For developers, getting to this point quickly means easy access to data they need to build with, without having to worry about the details of spinning up, managing, and maintaining databases.

APIs have become the de facto standard for connecting applications to databases, but it wasn't always that way. Here we'll discuss what has changed in the software world to elevate the importance of exposing databases as APIs, and we'll discuss Stargate, the open-source project that simplifies how developers can work with these APIs. (The latest version of Stargate, which includes improvements in scalability and flexibility, [was released recently](#)).

## HOW WE BUILD SOFTWARE THAT USES DATABASES

Database administrators (DBAs) used to be the ones in charge of queries, as it required a database expert to design the way that developers interacted with data. But until recently, developers were expected to know a lot about interacting with databases, even if they weren't experts at SQL, queries, or data models.

In the network world, however, we have new concepts of identity, permissions, and security that are completely separate from the query language. The notion of a driver simply communicating on a binary protocol doesn't work well in the cloud. HTTP is the foundational application layer protocol for the cloud, but most HTTP-based APIs are slow. A low-latency option, like gRPC, is critical for real-time communications in distributed systems.

## HOW SOFTWARE TALKS TO SOFTWARE

The standard way that clients or app servers used to talk to databases involved drivers that ran within the data center. Now everything is running in the cloud, but developers write in a wide range of languages. JavaScript, Python, or any of a host of different frameworks might be used, so the means by which software accesses the data needs to be abstracted in ways that are different from classic database drivers—using APIs (JSON, gRPC, GraphQL, or Document) that developers are comfortable with.

The modern way that software talks to software is an API; it's the abstraction layer that hides the complexities of the database.

## THE NATURE OF DATA

Data used to be much more uniform; it fit neatly into rows and columns in database tables. But the dynamics of data have changed. Data moves from in-memory representations back

to the database in a seamless way without much intervening software. And we're dealing with new types of data formats that are much more robust than the data primitives that people used to deal with—or the half-dozen or so data types that SQL could handle.

## A GATEWAY TO SIMPLICITY, SCALABILITY, AND EXTENSIBILITY

A [data API gateway](#) is a piece of software infrastructure that provides access to data via APIs of various styles including REST, gRPC, and others. The gateway abstracts the details of storing and retrieving data using one or more persistent stores. This enables application developers to focus on writing business services that access data via easy-to-use APIs instead of having to learn the intricacies of a database query language.

*The modern way that software talks to software is an API; it's the abstraction layer that hides the complexities of the database.*

[Stargate](#) is an open-source data API gateway that sits between an app and the databases it needs to query. It was first introduced in September 2020, and [Apache Cassandra](#) was chosen as the first database in part because it solves the world's hardest scalability and availability challenges.

Stargate offers developers the power of Cassandra by presenting REST, Document, and GraphQL as simple APIs. It also offers a set of gRPC libraries for doing CQL over gRPC as an easier, lightweight, and more cloud-friendly alternative to native drivers for CQL without sacrificing performance. With the [recent release](#) of Stargate v2, the engineering team at DataStax made some significant updates that make it easier for developers to use and the community to contribute to the project.

Stargate is offered via [DataStax Astra DB](#), the database-as-a-service built on Cassandra. Together, these essential pieces of data infrastructure are key to shortening that critical path to getting an application to production.

Learn more about [DataStax](#) or give [Astra DB](#) a try.

DATASTAX

<https://www.datastax.com>