



The Power of an Enterprise Data Layer

The Power of an Enterprise Data Layer

20-49-35: That's the respective breakdown, in percentage terms, of planned hybrid cloud adoption among large enterprise, midsize companies, and small businesses, according to the research firm [Techaisle](#).

Clearly we are heading into the [hybrid cloud age](#), but many companies are struggling to even define what hybrid cloud means for them, let alone adeptly and affordably deploy their applications in hybrid cloud environments.

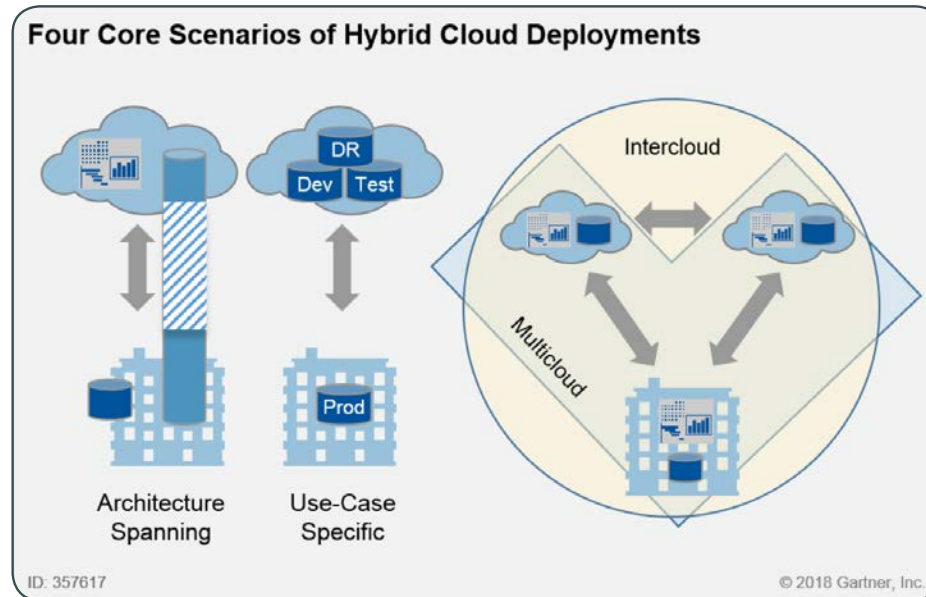
The answer to this challenge lies in the enterprise data layer. Hybrid cloud deployments provide enterprises with flexibility, global accessibility, and the capacity for rapid application development. But they also create data silos, in which cloud system data becomes isolated.

Breaking down these silos without draining IT resources presents a major challenge, the answer to which is an enterprise data layer that allows enterprises to seamlessly weave together their disparate data into a powerful whole that is far greater than the sum of its parts.

But before discussing exactly what an enterprise data layer is and does, let's discuss hybrid cloud and why it's so important.

A Closer Look at Hybrid Cloud

Gartner [defines](#) four different cloud deployment models that IT organizations are undertaking, all of which have some form of hybrid cloud at their core:



- > **Architecture Spanning** represents the case where the business is expanding its IT capacity with a cloud provider. This strategy may make sense for a variety of reasons that include the necessity to burst to the cloud for additional resources or to co-locate data where it is needed for performance or compliance factors. Note that with this form of hybrid cloud, the applications may live on infrastructure hosted by a cloud service provider and the data may reside within the company's self-hosted data centers.

A Closer Look at Hybrid Cloud (Cont.)

- › **Use Case Specific** describes when enterprises have versions of the same application deployed in different locations based on the stage of development. For example, a company may prefer that development and test occur in the cloud for flexibility where the production instance will run on specially provisioned in-house hardware. This case hinges on the fact that the application itself is not integrating with native cloud services.
- › **Multi-Cloud** is when a company runs different modern applications either all in the cloud or in an architecture spanning on-premises resources and also with those offered by multiple cloud providers. This model caters to a “developer-first” mindset, where the application developers are not siloed to a particular cloud and have the freedom to choose the technology that fits the task, whether it is offered by a cloud service provider or through custom, internal services.
- › **Inter Cloud** defines the evolution of the multi-data center, multi-residency design. This deployment features cross-cloud service provider traffic and boasts the ability to avoid vendor lock-in while taking advantage of the best of breed services. One example referenced by Gartner is: “Microsoft's PowerBI might connect to a Salesforce database residing outside of the Azure cloud infrastructure”.

All of the above types of hybrid cloud deployments come with their own array of caveats. It is important to design for the model that best fits the needs of the enterprise, where the balance of SLAs, cost, and technology offerings should drive decisions.

What Does a Data Management Strategy for Hybrid Cloud Architecture Look Like?

Having a robust, [hybrid cloud data management strategy](#) involves both physical architecture and application architecture. In order to understand why having a data management strategy is mission-critical, it's important to recognize the three main goals that cloud computing promises to fulfill:

1. Effortless scale
2. Geo-distribution
3. Accelerated time to value

To be able to fulfill these promises, enterprises must calculate the implications their architectural decisions have on each of these goals and shape their data strategy accordingly.

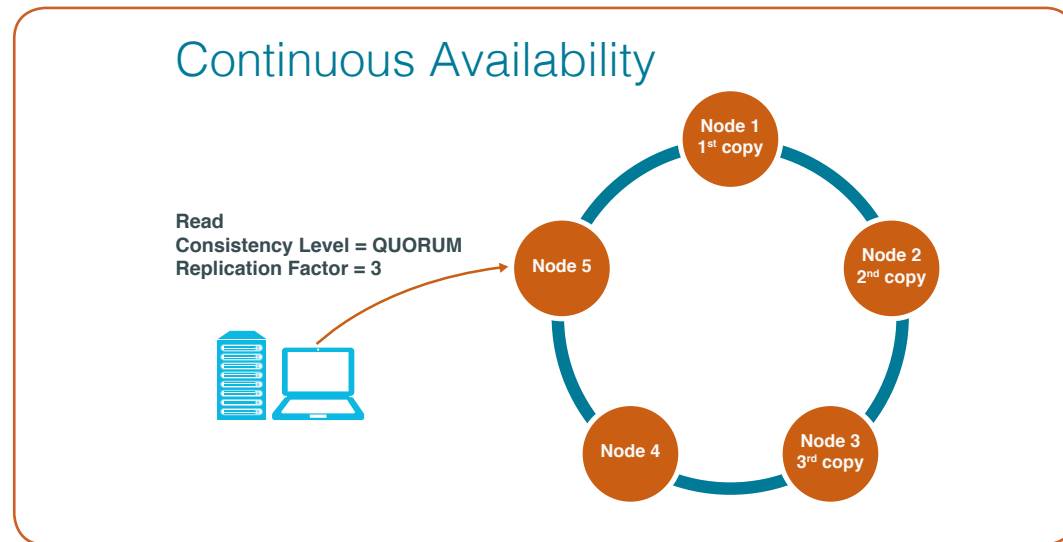
The advantage of the cloud lies in the ability to scale up and down to keep up with peak season demands and enjoy the cost savings associated with variable resource utilization. Users expect applications to scale with the demands of our connected economy. Viral trends impose unprecedented scalability demands on IT systems. The solutions to these problems are distributed computing and horizontal scalability, but to harness them, application and data architects must face a paradigm shift from the monoliths of the past.

Unfortunately, most [legacy applications](#) are not designed with geographic distribution in mind, and many enterprise data strategies do not take this into account. To accomplish the goal of geographically distributed applications, architects must incorporate things like replication, data locality, network failure, machine failure, and many other

What Does a Data Management Strategy for Hybrid Cloud Architecture Look Like?

(Cont.)

considerations into their data management strategies. Planning for continuous availability means ensuring that the system automatically knows how to respond when disruptions to your network or hardware occurs.



Then, there is the issue of attaining data availability seamlessly and without any constraints. New customer cloud applications need to rely in part on master data or historical data that already exists in the enterprise. Often times, new applications will need to pull data from disparate legacy systems and database management systems that are systems of record. The newer applications must be resilient and cannot depend on brittle technology that does not scale, does not support hybrid cloud, and does not support the rapid pace of modern data flows from multiple sources.

To truly achieve these goals of scalability, geographical distribution, and accelerated time to value, the architectural pattern architects should incorporate into their data management strategy is the use of an enterprise data layer.

What is An Enterprise Data Layer?

An enterprise data layer, also known as an [operational data store](#), is a highly distributed, always available data store that supports master and operational data to provide real-time data availability and functionality to all consumers and endpoints. Key logical entities that span more than one kind of data system reside in this data store, providing an interoperability that would otherwise not be possible due to data silos.

In enterprises that rely on shared data, an enterprise data layer makes it easier to solve complex problems across diverse groups while reducing data silos by creating a way to serve data sets that span applications. This data layer can reside on-premises, in hybrid cloud, or in more than one cloud, and can transform legacy applications that have trouble with availability, real time, and scale, into modern cloud applications that are contextual, always on, real-time, distributed, and scalable (ie, “CARDS”).



Contextual

Contextual relevance by serving intelligent information to the user at the right time, through the right channel(s).



Always On

Zero downtime across many years.



Real Time

Anytime, anywhere access with instant responsiveness.



Distributed

Span multiple regional and global data centers, clouds, and service providers.



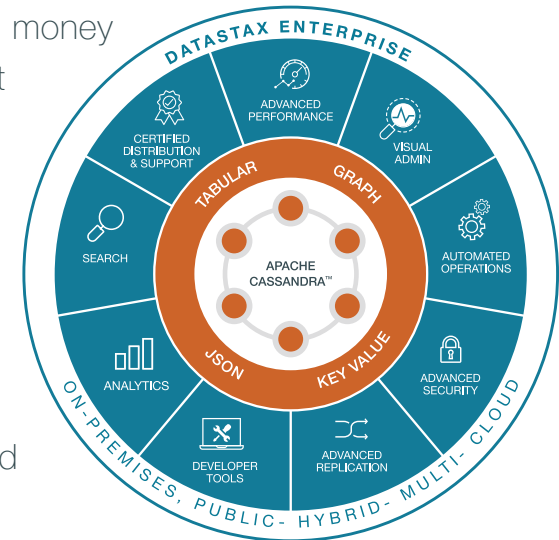
Scalable

Linear scalability. Add nodes and compute capacity with scale-out architecture and no downtime.

DataStax Enterprise Powers the Modern Enterprise Data Layer

Enterprises today are investing valuable time and money on building a robust database infrastructures that enable them to seamlessly deploy applications on-premises and in the cloud

[DataStax Enterprise \(DSE\)](#) addresses these challenges by providing an active everywhere distributed database architecture, built on Apache Cassandra™, that delivers a consistent data layer between on-premises data centers and public clouds.



An enterprise data layer powered by DSE enables enterprises to:

- > Make their enterprise data hybrid capable.
- > More rapidly develop and deploy applications with similar delivery approaches across lines of business and across teams, using existing infrastructures security architectures.
- > Build a set of reference architectures that support a common technical design across lines of businesses.
- > Pull multiple data domains together so insights can be generated across systems.
- > Transform systems that have trouble with availability, scale, and data visibility into cloud applications that are contextual, always on, real time, distributed, and scalable.

DSE powers a data layer that brings together mission-critical, user-facing applications and the various source or legacy systems that enterprises are still using. This allows enterprises to effortlessly accelerate their application development across hybrid and multi-cloud environments without incurring the additional costs of changing or removing their legacy systems.

To learn more, watch [this video](#).

About DataStax

DataStax delivers the always-on, active-everywhere distributed hybrid cloud database built on Apache Cassandra™. The foundation for full data autonomy and personalized, real-time applications at scale, DataStax Enterprise makes it easy for enterprises to exploit hybrid and multi-cloud environments via a seamless data layer that eliminates the issues that typically come with deploying applications across multiple on-premises data centers and/or multiple public clouds.

We help many of the world's leading brands across industries transform their businesses by eliminating data silos and powering modern, mission-critical applications. For more information, visit www.DataStax.com and follow us on [@DataStax](https://twitter.com/DataStax).

© 2017 DataStax, All Rights Reserved. DataStax, Titan, and TitanDB are registered trademarks of DataStax, Inc. and its subsidiaries in the United States and/or other countries.

Apache Cassandra, Apache, Tomcat, Lucene, Solr, Hadoop, Spark, TinkerPop, and Cassandra are trademarks of the Apache Software Foundation or its subsidiaries in Canada, the United States and/or other countries.