WHITE PAPER

De-Risking Hybrid, Multi-Cloud Analytics

Stephen Swoyer

02.19 / EB10303 / VANTAGE / WHITE PAPER

teradata.

Table of Contents

A decade on, enterprises are still figuring out how to manage risk in the cloud. Risks that used to predominate in the imagination of the enterprise subscriber—namely, concerns about the availability, reliability, cost, and security of cloud services—are no longer perceived as, well, quite so risky.

This paper outlines a positive, constructive program an enterprise can use to achieve the benefits of cloud while at the same time pragmatically control the risks inherent in moving decision-support and analytic workloads to a cloud environment. It considers three related questions:

1. **What risks are unique to the cloud?**
2. **What unique risks must enterprises consider when they move decision-support and analytic workloads to the cloud?**
3. **What pragmatic steps can enterprises take to manage these risks?**

teradata.

Risk Specific to the Cloud

Like any technology, cloud computing comes with risks, the most notorious of which is the possibility of cloud provider lock-in. The increasing concern about cloud provider lock-in can plausibly be attributed to cloud's growing mission-criticality to and for business operations. As cloud becomes a more popular destination for business-essential workloads, it looms as a more obvious point of failure and source of risk.

The benefits of standardizing on a single cloud provider—e.g., Amazon Web Services (AWS), Microsoft Azure, or others—are now sometimes viewed as less compelling than the risk of being locked-in to a single environment. As a result, enterprises identify the risk of cloud provider lock-in—along with concerns about availability, reliability, and security—as among the riskiest aspects of cloud, according to multiple surveys. This concern is reflected in survey data where it is voiced by chief information officers, chief data officers, and other key decision-makers as well as in both academic and commercial research.

It is reflected in practice, too: increasingly, enterprises are pursuing multi-cloud strategies that distribute their storage and processing workloads among two or more cloud providers. In other words, companies are taking pragmatic steps to address the danger of cloud service provider lock-in.

Risk Pertaining to Decision Support and Analytics

Decision-support and analytic workloads were slow to move to the cloud. This had to do with the logistics of how queries, reports, dashboards, etc., are generated, distributed, and consumed in the enterprise. The upshot is that enterprises are now starting to move what are arguably their most important workloads to the cloud.

Data management workloads in general, and decision-support workloads in particular, are special—exceptional even. This is not to say that decision-support workloads cannot be moved to the cloud; it is to say that not all cloud environments are suitable for all types of decision-support workloads. Yes, cloud was an early destination for some kinds of big data analytics, but the characteristics of these workloads differ from those typical of decision support.

Decision making of one kind or another is a constant in the enterprise. The simultaneous use of queries, reports, dashboards, etc., by different decision makers in different roles places extraordinary demands on a decision-support infrastructure. These demands require an infrastructure capable of delivering extremely high levels of service availability. Practically speaking, this means that a decision-support infrastructure must provide granular workload management features if it is to support hundreds, thousands, or tens of thousands of users. Such an infrastructure—typically provided by an analytic database (data warehouse)—must be capable of prioritizing the allocation of compute and storage resources on a per-workload, per-user, per-class-of-user, and per-application basis.

In moving decision support and analytic workloads to the cloud, a different approach is called for: a hybrid strategy that balances the benefits of the cloud—e.g., ease-of-use, simplified configuration and management, reduced maintenance—with the flexibility to accommodate various kinds of on-premises deployments. Like it or not, on-premises analytic systems will continue to support vast amounts of current and historical operational data. But the use of cloud-like concepts and technologies will transform the on-premises context. In addition to conventional, non-virtualized on-premises deployments, enterprises will make much greater use of virtualization and automation technologies (or vendor-supported and/or managed services) to power on-premises private clouds.

What's needed is a strategy that spans public clouds and on-premises data centers: a hybrid, multi-cloud strategy that allows enterprises to run their decision-support and analytic workloads in the contexts in which it makes the most sense to run them. This is a strategy that controls different kinds of risk by distributing workloads, assets, and dependencies, and above all, lets cloud be cloud.

What would such a strategy look like? Is it even feasible at this stage of cloud adoption?

A Blueprint for Hybrid, Multi-Cloud Success

This section uses a commercial technology stack—specifically, that from Teradata—to explore and assess the challenge of scaling decision-support and analytic workloads in a hybrid, multi-cloud architecture.

Bill Graham, the legendary concert promoter, had this to say of the Grateful Dead:

"They're not the best at what they do, they're the only ones that do what they do."

For more than two decades, this was true of Teradata as well. No other vendor focused so fixedly on the very large analytic database market. Like the Grateful Dead, Teradata wasn't just the best at doing what it did, it was the only one that did what it did. Is this still the case in the Age of Cloud?

The Teradata approach suggests that it not only understands cloud, but understands what cloud is all about: cloud is a new frame or abstraction for IT. Teradata gets that over time, almost all IT resources—saving, of course, the most intractable of legacy workloads—will become cloudy. This means, specifically, that the technology attributes that make cloud "cloudy"—i.e., high-density hardware configurations, the virtualization and isolation of resources, the use of software to abstract complexity and automate critical IT operations—will become the norm. Adjectives like "public," "private," and "public-private," perhaps even the term "cloud" itself, will become less meaningful to most consumers, including many IT people. The Teradata strategy seems designed with such a future in mind.

Multi-Cloud? Check.

A multi-cloud architecture acts as a hedge against diverse kinds of risk. By distributing applications, services, and workloads across clouds, an enterprise can architect for resilience. By picking and choosing where workloads live and run, enterprises can more easily tune for performance. This combination permits an enterprise to target (and meet) aggressive service level agreements. More importantly, a multi-cloud architecture helps an enterprise protect against the danger of cloud provider lock-in.

Teradata IntelliCloud™, the company's as-a-service offering for analytics at scale, can be deployed as part of a multi-cloud architecture that spans AWS resources, Azure resources, and Teradata Cloud resources. (The last is a service that Teradata hosts on its own infrastructure in its own data centers.) In addition to its as-a-service options, the Teradata multi-cloud strategy also supports do-it-yourself deployment in customers' own Virtual Private Clouds (VPCs) and Virtual Networks (VNets) via the AWS and Azure Marketplaces. Teradata dubs this a "deploy anywhere, move anytime" model.

In practice, IntelliCloud seems to live up to this billing. For example, an enterprise can divide its IntelliCloud workloads between AWS and Azure and shift users and use cases between the different deployments as needed. Customers can likewise shift workloads to or from IntelliCloud instances and Teradata systems running in on-premises environments. "Shifting" a workload is not a turnkey event however. Customers must also move—at per-GB costs—the data associated with these workloads.

On the one hand, this portability helps mitigate the risk of cloud provider lock-in because it permits customers to shift workloads to and from Teradata solutions running in multiple scenarios. On the other hand, the actual process of moving decision-support workloads between environments cannot cheat space, time, and basic economics. If the workloads being relocated need to run against data that is "local" to their own cloud context, that data must move along with them. The upshot is that if relocating a workload involves moving a large volume of data, the rapidity at which it can be relocated and brought online in a new context is determined by how quickly the data it needs can be moved. This reality applies to any analytic workload running in any context and is not specific to the Teradata hybrid architecture.

teradata.

A case can be made that licensing, pricing, and deployment options from Teradata (see "Let Cloud Be Cloud," page 6) make this process easier—or at least more practicable. For example, Teradata provides several features that promote or simplify the hybrid, multi-cloud experience. Teradata QueryGrid™, which the company offers across all its deployment options, allows designers to build and expose virtual views of data that not only span different contexts (physical systems, public clouds, etc.) but geographical regions, too. Designers can use QueryGrid to build virtual or composite data views for different use cases or to suit a wide variety of requirements. A subscriber can likewise use QueryGrid to automate the preparation and movement of data—again, on a per-use case or per-role basis.

QueryGrid is designed to be smart about how it moves data, such as making use of available compute resources—e.g., different types of analytic engines including MPP SQL, Graph, Machine Learning, etc.—to transform or engineer data in the context in which it lives. This is critical. Not only is it faster to move a (smaller) derived data set between contexts, it is cheaper, too: most cloud platforms are biased in favor of data ingestion, not egress. The per-GB cost of moving data out of a cloud service comprises a practical barrier to data movement at a large scale. It also comprises a practical barrier to moving from one cloud provider to another. Technologies such as QueryGrid make the Teradata hybrid, multi-cloud strategy feasible from the perspectives of both IT and different kinds of business analyst and data science users.

Teradata Vantage, the company's flagship analytics platform software, is a critical factor as well. The last two decades have produced a profusion of database engines. There's no such thing as parity between MPP engines, however; some are faster than others. Some are faster for certain kinds of workloads, and some are more efficient at scaling up (better per-instance performance) or scaling out (better parallel performance across all instances).

Cloud doesn't change this. In most cases, to say an analytic database is "built" for the cloud is simply to say that it permits some degree of elasticity: the ability to add or subtract compute or storage resources independently of one another. Building and maturing a relational database management system is hard. It's impossible to cheat space and time in this regard, too: it takes a decade or more for an analytic database to fully mature; building and maturing a multi-node relational database management system takes even longer. Thanks to its pedigree, Teradata Vantage is designed to deal with the kinds of problems—i.e., unreliable service and de facto distribution—that are characteristic of cloud at this stage of its evolution.

teradata.

Empowering? Check.

There is a misconception that technical debt becomes somebody else's problem once a workload shifts to the cloud. This is categorically not the case.

Whenever subscribers use techniques—e.g., user-defined functions, stored procedures, scripts, etc.—to work around the limitations, quirks, or bugs of a service, they will incur technical debt. An added complication is that all cloud services are constantly evolving. Customizations designed to work around a "feature," bug, or limitation in one iteration of a service might break with a new iteration.

This is as true in the era of cloud and distributed services architecture as it was in the days of the monolithic software stack, but even more so because platform change is accelerated. The pace at which new features, parameters, and runtime options are added or removed increases with cloud. Take application programming interface (API) versioning, for example. It is not uncommon for a cloud provider to introduce new versions of core APIs when it makes significant changes to its service. Owing to the scope of these changes, it might be impossible for the provider to maintain backward-compatibility with deprecated APIs, causing software and services that depend on them to break.

The Teradata MPP database engine is highly mature. Many of the other components of Vantage are likewise mature. The APIs, parameters, features, and behaviors of these components of the Teradata hybrid, multi-cloud stack are evolving, to be sure, but are more stable and established than those of its competitors. On the other hand, IntelliCloud itself—the company's as-a-service offering—is evolving rapidly. In just the last year, for example, Teradata introduced support for fully elastic operations so IntelliCloud compute and storage resources can be added or subtracted, started, paused, or turned off, independently of one another. The feature set of IntelliCloud that enables self-service-like ease-of-use and automates common tasks is evolving along with those of other cloud services.

There's also an undeniable sense in which some customers really are "locked in" to Teradata. On the one hand, they have requirements they believe they cannot cost-effectively address using any other mix of products, services, or open-source projects. On the other, they're dealing with data processing problems at such a massive scale they feel they cannot cost-effectively move to another platform. As a director of database solutions with one Teradata cloud customer puts it: "Our data will still be [in a Teradata solution] in 10 years. Teradata just does what nobody else can do. That won't change."

This is one kind of lock-in. Call it the "benign" variant—if only because in the final analysis, a customer feels okay with it. "Benign" lock-in happens all the time. It is a familiar story: a customer has demanding or esoteric requirements, a very limited number of technologies or providers address those requirements, the customer chooses from among one of these vendors, and the resulting relationship between customer and vendor grows as the vendor works to ensure that its future roadmap addresses the needs and priorities of its customers. The result is a mutually-beneficial relationship.

Let Cloud Be Cloud? Check.

In terms of its feature set, capabilities, and pricing model, Teradata IntelliCloud is an "analytics platform" delivered in an as-a-service model. There is considerable merit to the strategy.

In the first place, IntelliCloud is not just a repackaging of venerable software capabilities from Teradata. It provides a browser-based management console intended to simplify the process of monitoring and controlling Teradata instances in the cloud. IntelliCloud also offers several complementary technologies—such as AppCenter, Data Mover, and QueryGrid, to name a few—as optional services. These are essential enabling technologies for subscribers with existing Teradata investments and/or subscribers pursuing hybrid, multi-cloud strategies.

Secondly, IntelliCloud lets cloud be cloud. IntelliCloud is a fully elastic data-warehouse-as-a-service offering. "Elasticity" in the context of cloud means subscribers can add new compute instances or provision additional storage volumes to yield extra capacity as needed.

This is typically referred to as scale up/down (from smaller to larger individual compute instances) or scale out/in (from fewer to greater numbers of parallel processing clusters). Elasticity is not simply a function of scale up/down or scale out/in. However, in a practical sense, elasticity also permits subscribers to shut down unneeded resources as desired. This makes it possible to manage (and optimize) cloud spending.

That being said, elasticity of this kind is not necessarily a critical requirement for traditional decision-support and analytic workloads. These workloads often have predictable connection and concurrency characteristics that allow their capacity requirements to grow at predictable rates.

Some analytic use cases do require IntelliCloud's high levels of elasticity, however. To take one example: a data scientist tasked with building machine learning models will have demanding requirements with respect to the engineering, storage, and processing of the data she needs to train and score these models. The problem is that she cannot predict how much of which resources (compute, storage, network) she will need, nor, crucially, when or for how long she will need them.

Data science and other types of advanced analytics workloads are characterized by what is called "burstiness"—i.e., irregular spikes in demand for resources. An additional complication is that this inconsistency is asymmetrical: some workloads will require exponentially more compute than storage or network, while others will need exponentially more storage than compute or network. Burstiness is a feature, not a bug, of data science and other advanced analytics workloads.

IntelliCloud's ability to elastically scale up or down makes it a good fit for advanced analytics and other bursty workloads. IntelliCloud's elasticity permits subscribers to start or stop resources as needed, too. In most cases, for example, an enterprise subscriber will maximize its use of provisioned resources during the business workday, e.g., from 8:00 AM to 10:00 PM, but will use fewer resources during the overnight period. True, a substantial proportion of resources might be used to power critical after-hours workloads (ETL batch jobs, scheduled backups, routine data processing), but not all of them.

Another upshot is that IntelliCloud's elasticity enables an enterprise to adjust the capacity it uses on a granular basis, stopping resources if or when they aren't needed, and restarting and growing them if or when they are. Instead of paying for capacity on a 24/7 basis, subscribers can exploit IntelliCloud's elasticity to reduce their consumption of resources. The earlier on-premises model was characterized by capital-intensive (albeit much longer) IT buying cycles. As a result, enterprises tended to over-provision resources to meet future growth, which meant paying upfront for more compute and storage capacity than was initially required. In the IntelliCloud as-a-service model, subscribers may add or subtract capacity as and when needed—one of the signature cost-control benefits of the cloud consumption model.

One unique benefit to Teradata is its licensing model. The company offers subscription licensing for both cloud and on-premises instances of its software. A customer running Teradata software using Teradata IntelliFlex™ instances in an on-premises data center, for example, can shift some or all those software licenses to IntelliCloud.

To facilitate portability between cloud and on-premises contexts, Teradata developed a metric it calls "TCore." TCore allows software pricing to reflect the actual capabilities of CPU and I/O resources in different environments. This is important because CPU and I/O performance is typically constrained in the cloud. Instead of charging for capacity on a per-node basis, TCore charges are based on the underlying infrastructure's CPU rating and I/O bandwidth. For example, workloads running in an on-premises Teradata IntelliFlex configuration will typically have access to more much compute power and I/O capacity on a per-node basis than other IntelliCloud instance types. TCore adjusts for this asymmetry: if workloads require more virtual compute instances to achieve the same levels of performance as a physical node, that's okay because pricing is a function of system TCore value, not number of instances. TCore is a means of ensuring that subscribers pay the same amount for compute and I/O performance irrespective of how that capacity is resourced.

TCore is a pricing metric that's designed for the future, not just for cloud. Enterprise workloads will continue to shift to the cloud. Some will run in public cloud services hosted by large cloud providers, while others will run in enterprise private clouds hosted in on-premises data centers, in co-location facilities, or as managed services by regional hosting providers. Over time, these distinctions will collapse as cloud becomes more of a general metaphor—the de facto frame—for designing, managing, and thinking about IT architecture. The TCore pricing metric is ahead of this curve.

Pragmatic Strategies Control Risk

Traditionally, enterprises have been forward-thinking with respect to cloud. In the early stages of cloud adoption, and in the midst of substantial hype, firms made pragmatic use of cloud services to power low-risk, high-reward workloads and use cases (such as test/dev and prototyping) that they could not cost-effectively support and manage in on-premises contexts.

As cloud has evolved into a less risky option for a greater variety of workloads and use cases, organizations have adjusted their use of it accordingly. And as new risks have surfaced, such as the growing potential for cloud provider lock-in, companies have embraced hybrid and multi-cloud options as pragmatic strategies they can use to control risk.

It is critical that enterprises retain this pragmatic focus as they move decision-support and analytic workloads to the cloud. Cloud promises to be a force for tremendous value in combination with these workloads. It gives enterprises new flexibility with respect to how they acquire, configure, provision, access, manage, and yes, pay for the resources they depend on to power these workloads. The hybrid, multi-cloud technology stack from Teradata provides a compelling demonstration of what this flexibility looks like in practice.

About Teradata

With all the investments made in analytics, it's time to stop buying into partial solutions that overpromise and underdeliver. It's time to invest in answers. Only Teradata leverages all of the data, all of the time, so you can analyze anything, deploy anywhere, and deliver analytics that matter most to your business. And we do it on-premises, in the cloud, or anywhere in between. We call this pervasive data intelligence. It's the answer to the complexity, cost and inadequacy of today's analytics. And how we transform how businesses work and people live through the power of data. Get the answer at **Teradata.com.**

About the Author

Stephen Swoyer is a technology writer with 20 years of experience. His writing focuses on business intelligence, data warehousing, and analytics. Swoyer has an abiding interest in tech, but he's particularly intrigued by the thorny people and process problems technology vendors never, ever want to talk about.

17095 Via Del Campo, San Diego, CA 92127 Teradata.com

teradata.