

# TERADATA'S AMBITIONS SOAR WITH TERADATA CLOUD

// BY STEPHEN SWOYER

**A closer look at the company's latest push into the cloud.**

The cloud might seem like a turbulent atmosphere in which to run a data warehouse (DW).

That's what makes Teradata Corporation's latest push into the cloud so intriguing.

At its Partners conference in October 2013, Teradata unveiled the Teradata Cloud, a new cloud analytics service that consists of Teradata Data Warehouse and Teradata Aster Discovery platform, along with a Hadoop offering.

This isn't Teradata's first foray into cloud. Four years ago, it announced its Teradata Enterprise Analytics Cloud, which consisted of both public (external) and private (on-premises) cloud offerings for test and development, including a version of Teradata for Amazon EC2; last year, it introduced its Teradata Active Data Warehouse Private Cloud with self-service capabilities and burst elastic performance.

The Active Data Warehouse Private Cloud is based on Teradata's bread-and-butter data warehouse platform. Teradata Cloud, on the other hand, enlists Teradata Database and Teradata Aster Discovery Platform, along with its Hadoop offering, to address analytic discovery, advanced analytics, and a number of other use cases—in addition to bread-and-butter business intelligence (BI) and decision support.

"The cloud is going to mimic what people want [in their] on-premises [resources] today. If there's a value for the core Teradata Data Warehouse, if there's a value for Aster Discovery, if there's a value for Hadoop, a growing subset of companies are going to want that in a cloud delivery model," says Ed White, general manager for cloud solutions with Teradata.

Over the next five years, White suggests, as many as 50 percent of DW workloads could move to the cloud. "If you think about what we're doing with the cloud, we're going to set it up [for the customer], we're going to manage basic operations—we're taking all of those administration-related, behind-the-scenes, under-the-covers tasks and automating them."

## SCALING IN THE CLOUD

Most of Teradata's cloud offerings have had a massively parallel processing (MPP) component, and Teradata Cloud is no different. One reason for this, White maintains, is that MPP scales about as well in the cloud as it does on-premises.

"Whether you're accessing [a data warehouse] in a remote data center or on-premises, you still need the performance that we provide with our MPP optimizer and [workload] provisioner. Knowing that you're going to get linear scalability [because of MPP] is even better," White explains. What differentiates Teradata Cloud from on-premises offerings are "the easy and fast provisioning, the flexible scale-up, and the competitive pay-as-you-go pricing," he points out. "The core MPP architecture processing, the optimizer, is 100 percent the same as you get with [on-premises] Teradata [Warehouse]."

According to White and Teradata, MPP is actually a good fit for the cloud. This has to do both with the way in which MPP is typically scaled—i.e., by adding nodes—and with the economics of the cloud: if you add four extra nodes to a four-node MPP configuration, you'll effectively double its performance. For this reason, White says, it's possible to predictably scale an MPP data warehouse, regardless of whether it's on-premises or in the cloud. In the cloud, he notes, it's actually much easier to do so: instead of budgeting for, acquiring, installing, and configuring a physical system, you're simply spinning up analytics "as a service."

In addition, White points out, it's easier and considerably more cost-effective to *prototype* a data warehouse in the cloud; it's likewise easier to add or subtract production capacity as needed. For these and other reasons, he suggests that most of the traditional barriers to cloud adoption—e.g., concerns about security, questions about availability or service levels—have taken a lower priority and now customers are *demanding* cloud options.

“We’re already seeing a tremendous amount of what I call cloud mandate demands ... [from] customers. Some of it is driven by Cap Ex and Op Ex [concerns]—they want a cloud delivery model—but the other thing that cloud can be used for is before [a customer] make[s] a multi-million dollar commitment, some of them want to test it for a few months,” White says. He says that it’s easier, faster, and much, much cheaper to “try out” a data-warehouse-as-a-service (DWaaS) offering. If you buy an on-premises data warehouse and that system doesn’t live up to expectations, he observes, you’re nonetheless locked into it—until you can afford to purchase a replacement.

“How do you really figure out what performance you need?” he asks. “The first part [figuring out the performance] is going to be easier, because if you buy four nodes today and that doesn’t work, you can go to six nodes or eight nodes—so it isn’t as if you’re stuck with a configuration for a few months or years.”

## PRIVATE AND PUBLIC—TERADATA CLOUD HAS ASPECTS OF BOTH

The second part of White’s question—i.e., actually *getting* the desired level of performance in the cloud— isn’t as easy, however.

In an on-premises MPP configuration, DBAs will take great care to minimize skew effect. In most cases, a DBA focuses on tweaking data distribution across an MPP cluster precisely in order to minimize data skew. In an MPP cloud configuration using virtualized resources, skew could pose significant problems, chiefly because storage is virtualized and non-local.

Teradata Cloud addresses this and other potential performance concerns by using dedicated physical systems to power its MPP nodes, White says.

“We’re going to have our [MPP] production nodes [on] dedicated [physical hardware] and we’re going to virtualize all around that, so the BI servers ... [will be] on quasi-dedicated [hardware],” he explains, noting that Teradata Cloud, unlike most cloud services (including Salesforce.com), offers service-level agreements (SLAs). “We can put multiple customers in the same rack, so we have kind of a hardware multi-tenancy [by means of multiple dedicated systems in the same rack], but offer dedicated production systems.”

As mentioned earlier, Teradata is no newcomer to cloud; it launched its first cloud offering in 2008. Even without this experience, Teradata Cloud would rocket into the high end of the

cloud troposphere, thanks in part to a comparative dearth of DWaaS offerings, to say nothing of *MPP* DWaaS products. One obvious potential competitor is Amazon’s Redshift DWaaS, but Redshift lacks Teradata-like workload management capabilities, doesn’t have anything comparable to Teradata’s Aster Discovery Platform, and isn’t available with a dedicated (i.e., physical) cloud compute option. This last makes Redshift harder to predictably scale—especially for mission-critical applications or services, which require high levels of availability. The upshot is that Teradata Cloud breaks new ground.

“We’re addressing the key performance issues with cloud data warehousing, including linear scalability, performance, and privacy,” White points out, suggesting that Teradata Cloud’s use of dedicated cloud compute units comprises an additional privacy safeguard, too. ●

*Stephen Swoyer is a contributing editor for TDWI.*

## ABOUT OUR SPONSOR



THE BEST  
DECISION  
POSSIBLE™

### teradata.com

Teradata is the world's largest company focused on integrated data warehousing, big data analytics, and business applications. Our powerful solutions portfolio and database are the foundation on which we've built our leadership position in business intelligence and are designed to address any business or technology need for companies of all sizes.

Only Teradata gives you the ability to integrate your organization's data, optimize your business processes, and accelerate new insights like never before. The power unleashed from your data brings confidence to your organization and inspires leaders to think boldly and act decisively for the best decisions possible. Learn more at [teradata.com](http://teradata.com).

## ABOUT TDWI



### tdwi.org

TDWI, a division of 1105 Media, Inc., is dedicated to educating business and information technology professionals about the best practices, strategies, techniques, and tools required to successfully design, build, maintain, and enhance business intelligence and data warehousing solutions. TDWI offers a worldwide membership program, five major educational conferences, topical educational seminars, role-based training, on-site courses, certification, solution provider partnerships, an awards program for best practices, live Webinars, resourceful publications, an in-depth research program, and a comprehensive website, [tdwi.org](http://tdwi.org).