

# Innovation in Cloud Analytics

---

An ENTERPRISE MANAGEMENT ASSOCIATES® (EMA™) White Paper  
Prepared for Teradata

February 2013



*IT & DATA MANAGEMENT RESEARCH,  
INDUSTRY ANALYSIS & CONSULTING*

# Innovation in Cloud Analytics

## Table of Contents

- Cloud Computing Overview ..... 1
  - Defining Cloud ..... 2
    - Service Models of Cloud Computing..... 3
    - Deployment Models of Cloud Computing..... 4
  - Cloud Value Proposition ..... 4
  - Cloud Adoption Statistics..... 4
  - Challenges and Opportunities..... 4
- Teradata Cloud Computing Solutions ..... 5
- Teradata ADW Private Cloud..... 5
- Teradata Virtual Machine Edition ..... 6
- Teradata Cloud Data Warehouse as a Service (DWaaS) ..... 6
- Teradata Express for VMware Player..... 6
- Teradata Express for Amazon EC2..... 7
- Teradata Aster for Amazon EC2 ..... 7
- Teradata Aster Express for VMware Player..... 8
- Teradata Applications in the Cloud ..... 8
- Hybrid Cloud Strategy ..... 8
- EMA Perspective..... 9



## Cloud Computing Overview

For years IT and business users have searched for a technology that delivers greater affordability and agility. Cloud computing delivers on both fronts by tapping into a new technology model that addresses capital and operational costs while delivering platforms and solutions designed to meet the needs of a more sophisticated analytics and business intelligence consumer. Innovative solution providers are delivering Cloud solutions in many variations.

Recent changes in enterprise data management strategies have created new opportunities for platforms such as Cloud and SaaS solutions to find a significant home within data management ecosystems of many companies. Four drivers are at the heart of this strategic shift and are creating new opportunities for IT departments to serve their users in a powerful and agile fashion:

1. **Maturing user community:** Over the past decade, demands of simple reporting have morphed into highly complex analytics. Consumers of business intelligence have changed dramatically from a few highly technical users within a company to a democratized culture where more and more employees are gaining access to information and driving the business with it.
2. **New technology:** Technology advancements are driving the adoption of Hybrid Data Ecosystems. Analytic platforms were the earliest to exist alongside the EDW, providing an additional platform for analysis. Cloud computing can match new workloads and data on the best possible platforms.
3. **Economics:** The overall cost of analytics plays a critical role in the adoption of a Hybrid Data Ecosystem. Lower capital costs associated with Big Data, Cloud and analytic platforms have contributed to enterprise adoption.
4. **Valuable data types:** Until recently, the enterprise has primarily focused on structured information best leveraged by SQL, and stored in relational databases. The task of analyzing high volume, high velocity and multi-structured information has often been too complex or expensive for most companies. New technologies have made it possible to incorporate new and highly valuable data (social, machine, sensor data) into analytic processes providing greater insights and predictive outcomes.

Many companies are adopting a Hybrid Data Ecosystem to leverage the best possible platform or combination of platforms to execute analytic and business intelligence workloads. Traditionally, the enterprise data warehouses have been at the center of the ecosystem, but recent acceptance of new and innovative computing platforms has opened the door for companies to adopt purpose-built solutions such as Cloud into their data management strategy. Recent EMA research, "[Big Data Comes of Age](#)," published in November 2012, shows that 67% of respondents are now utilizing between two and five different solution platforms to execute on the demands of business intelligence and analytics within their company.

---

Changes in enterprise data management strategies have created new opportunities for platforms such as Cloud and SaaS.

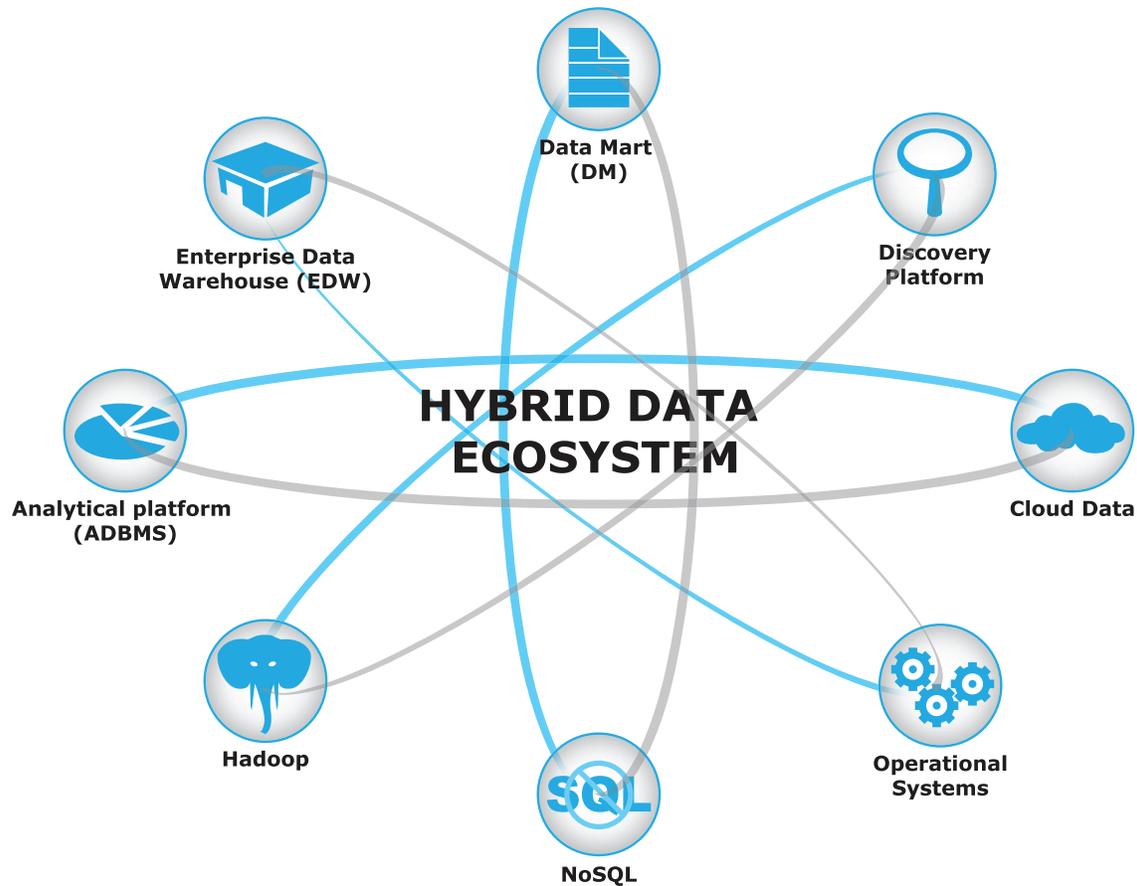
---

---

Many companies are adopting a Hybrid Data Ecosystem to leverage the best possible platform or combination of platforms to execute analytic and business intelligence workloads.

---

# Innovation in Cloud Analytics



## Defining Cloud

Cloud Computing is defined by the U.S. National Institute of Standards and Technology (NIST) as:

*“A model for enabling convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction.”*

Supporting the technology are five characteristics generally present in the Cloud environment; however, as more vendors enter the market this definitional framework is being stretched and modified to meet the needs of the consumers and innovation of the vendors.

1. **On-demand self-service:** A consumer can unilaterally provision computing capabilities, such as server time and network storage, as needed automatically without requiring human interaction with each service’s provider.
2. **Broad network access:** Capabilities are available over the network and accessed through standard mechanisms that promote use by heterogeneous thin or thick client platforms (e.g., mobile phones, laptops and PDAs).

# Innovation in Cloud Analytics

- 3. Resource pooling.** The provider's computing resources are pooled to serve multiple consumers with different physical and virtual resources dynamically assigned and reassigned according to consumer demand. The customer generally has no control or knowledge over the exact location of provided resources but may be able to specify location at a higher level of abstraction (e.g., country, state, or datacenter). Examples of resources include storage, processing, memory, network bandwidth and virtual machines.
- 4. Rapid elasticity:** Capabilities can be rapidly and elastically provisioned, in some cases automatically, to quickly scale out and rapidly released to quickly scale in. To consumers, the capabilities available for provisioning often appear to be unlimited and can be purchased in any quantity at any time.
- 5. Measured service:** Cloud systems automatically control and optimize resource use by leveraging a metering capability at some level of abstraction appropriate to the type of service (e.g., storage, processing, bandwidth and active user accounts). Resource usage can be monitored, controlled and reported providing transparency for both providers and consumers.

In addition to these characteristics of Cloud Computing it's important to understand there are several service models that can be employed to utilize Cloud Computing platforms, In addition each has a unique value proposition and can serve different needs of end users.

## Service Models of Cloud Computing

- 1. Software as a Service (SaaS):** Provides applications running on a Cloud infrastructure. Applications are accessible from various client devices through a thin client interface such as a Web browser (e.g., Web-based email). Consumers do not manage or control the underlying Cloud infrastructure including network, servers, operating systems, storage, or even individual application capabilities, with the possible exception of limited user-specific application configuration settings.
- 2. Platform as a Service (PaaS):** Deployment onto Cloud infrastructure for consumer-created or acquired applications using programming languages and tools supported by a provider. Consumers do not manage or control the underlying Cloud infrastructure including network, servers, operating systems, or storage, but have control over deployed applications and possibly application hosting environment configurations.
- 3. Infrastructure as a Service (IaaS):** Provisioning of processing, storage, networks and other fundamental computing resources where the consumer is able to deploy and run arbitrary software, which can include operating systems and applications. The consumer does not manage or control underlying Cloud infrastructure, but has control over operating systems, storage, deployed applications and possibly limited control of select networking components (e.g., host firewalls).

Respondents to EMA research have ranked SaaS service models as the most adopted technology (66%), far eclipsing PaaS and IaaS models.

The last important component to understand is the manner in which Cloud technologies can be deployed. There are three major deployment models that most solutions leverage.

# Innovation in Cloud Analytics

## Deployment Models of Cloud Computing

1. **Private Cloud:** Cloud infrastructure is operated solely for an organization. It may be managed by the organization or a third party and may exist on-premise or off-premise.
2. **Public Cloud:** Cloud infrastructure is made available to the general public or a large industry group and is owned by an organization selling Cloud services.
3. **Hybrid Cloud:** Cloud infrastructure is a composition of two or more Clouds (private, community, or public) that remain unique entities but are bound together by standardized or proprietary technology enabling data and application portability (e.g., Cloud bursting for load balancing between Clouds).

EMA Research indicates that in early adoption cycles, many companies opt to utilize public Cloud deployments. Hybrid Models, however, are ranked highest in customers planning or researching projects.

## Cloud Value Proposition

76% of respondents to EMA's Cloud research, [\*Cloud Business Intelligence and Data Management as a Service: A Global Survey on Adoption, Challenges and Outlook\*](#), indicate that Cloud computing has delivered real, measurable cost savings to their organizations. These companies are realizing significant upfront capital cost reductions as nearly 40% of respondents report 30% or greater cost savings compared to traditionally deployed projects. Ongoing operational costs are affected as well, as 23% of the respondents report 30% or greater savings on these costs. Financial impact is important, as it remains a top driver of Cloud adoption, but should not overshadow the improvement in IT service quality or increase in flexibility and agility that often comes from Cloud solutions.

## Cloud Adoption Statistics

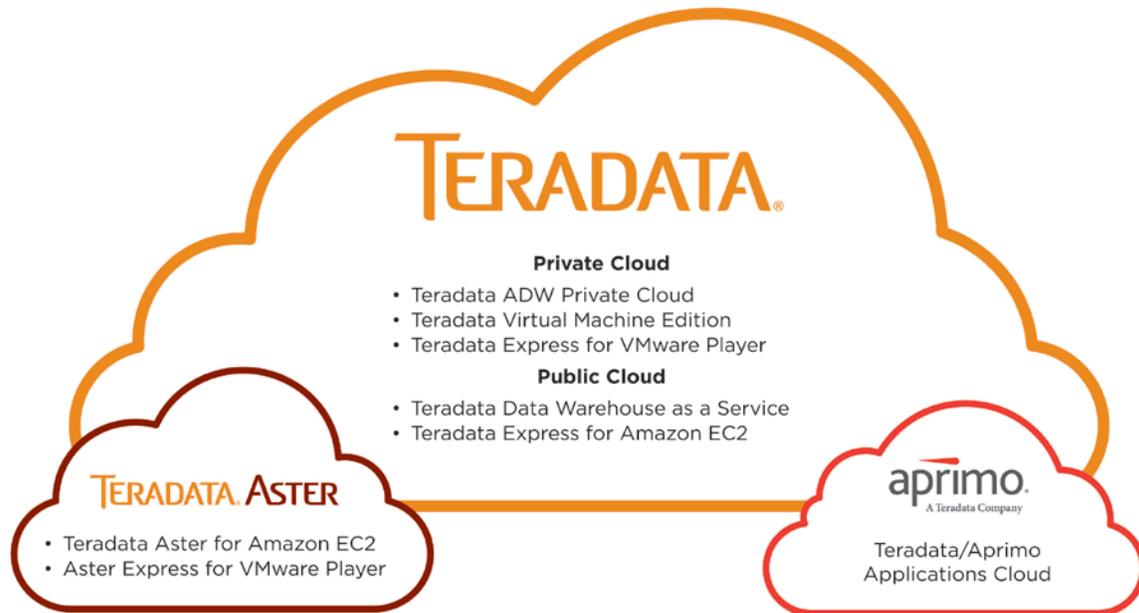
Already, 41% of EMA research respondents have implemented an Enterprise BI/DM system in a Cloud computing environment, and 28% have implemented departmental systems. Prototyping and short-lived projects are less popular with 18% and 13% respectively, but are an excellent example of Cloud flexibility and agility. However, there are significant differences in geographies: Cloud computing for prototyping is more popular in Europe with 34% vs. 11% in North America, while in North America there are significantly more enterprise BI/DM systems implemented in the Cloud compared to Europe (48% vs. 27%).

## Challenges and Opportunities

Any new technology comes with challenges and for public Cloud, data integration and security are both common concerns. Many solution providers are building data integration technology into the platform or leveraging market leading partners to deliver this feature set. Moving data to and from the Cloud is a strategic decision driven by use case and data volume. Often it makes better sense to leverage data within the originating system. With security and compliance in mind, many companies opt to leave highly regulated (HIPAA, Personally Identifiable Data) data behind the firewall and utilize data integration techniques that enable Cloud applications to leverage data from that location. To alleviate concerns over security, companies are working to achieve various levels of security certification including SAS70, SSAE 16, and ISO 27001, 27002. Early on, data integration and security slowed Cloud adoption, but as the market matures, both have become a more manageable concern for most companies.

## Teradata Cloud Computing Solutions

Teradata has designed a wide and flexible array of Cloud solutions to serve its customers. They have taken advantage of public and private deployment models along with leveraging investments in Aprimo, Teradata Aster and the Teradata Platform. Their Cloud solutions span the needs of developers, marketers and consumers of both analytic and data warehouse information.



## Teradata ADW Private Cloud

Elasticity and flexibility are key characteristics of Cloud computing architectures. The ability to grow and contract an analytic computing environment on demand goes well beyond traditional feature sets. The Teradata® Active Data Warehouse (ADW) Private Cloud delivers real-time, on-demand elasticity in a metered pay as you go model.

The Teradata ADW Private Cloud is ideal for customers across all industries. For example, retailers often experience spikes in computing needs throughout the year for sales driven programs. Sometimes, retailers will utilize the Public Cloud to augment their seasonal computing needs. This practice is often called “cloud bursting” where a company utilizes extra computing power from the cloud for short durations. This practice comes with risks and challenges that include data integration, security and extra management concerns. An ideal alternative is to have burst processing power – paid for with OPEX budgets – on the data warehouse platform the company is already using for analytics. The Teradata ADW Private Cloud does just that by allowing clients to provision extra computing power within their platform in an on-demand and metered model. This ability to “elastically burst” on the Teradata Active Data Warehouse platform is unique and highly valuable to companies who require flexibility within their analytic and data warehousing systems.

In addition, innovative companies interested in driving their business with insights from their business practitioners are adopting self-service analytics. Teradata has a self-provisioning feature made possible

# Innovation in Cloud Analytics

with Teradata Data Lab, which is a set of portlets where customers can create their own “sandbox” or analytic environments within the production data warehouse. Teradata Data Lab allows a customer to run proof-of-concept projects, special applications and testing with minimal setup effort. The creation of data labs empowers Cloud-based, self-service analytics, but also provides automatic management and governance of the environment.

## Teradata Virtual Machine Edition

Leveraging your existing server infrastructure by utilizing virtual machine technology creates an opportunity to add flexibility and scalability to your existing ecosystem while enjoying significant cost reductions. Teradata Virtual Machine Edition (VME) integrates with VMware ESXi private cloud environments and allows users to quickly deploy preconfigured instances of the Teradata Database without extensive support from IT Systems Administration or Teradata Customer Services personnel. Teradata Virtual Machine Edition delivers Cloud elasticity by allowing developers to quickly create a test and prototyping environment separate from the production systems that can be discarded as quickly as it was created. Application developers can choose from three templates sizes: small 80GB, medium 320GB or larger 2.4TB. Teradata VME will operate on VMware ESX 4.0 or VMware ESXi5.0 or higher. Teradata VME plugs directly into the VMware Management Console making it easy to control and manage each instance of the solution.

## Teradata Cloud Data Warehouse as a Service (DWaaS)

The Teradata DWaaS solution is a highly integrated, hosted and managed data warehousing platform that bundles Teradata hardware, software, managed services and customer services into a packaged solution. Teradata DWaaS is priced with a monthly payment model allowing companies to swap traditional capital costs for monthly operational costs. The platform has the flexibility to host Teradata Data Labs portlets to deliver self-service data warehouse environments, as well as hosting applications such as Teradata Integrated Marketing Management, Microstrategy, Microsoft BI and MDM, among others. The solution is delivered in three sizes depending on customer requirements: 1-4 TB, 5-9 TB and greater than 9TB. There are also modular service options for a hosted, managed platform or fully managed solution. Optionally, industry starter packages are available for retail and healthcare with predefined industry standard reporting/analytic templates and data models to get up and running quickly. Or customers can take advantage of Teradata Consulting Services for migrating an existing data warehouse or designing and implementing a custom solution. Teradata utilizes its proven data warehouse appliance to power its Cloud-based DWaaS solution.

## Teradata Express for VMware Player

Teradata Express for VMware Player is a fully functional, on-premises, private Cloud version of the Teradata Database and Teradata Studio Express delivered in a VMware® Player image with 64-bit SUSE Linux 10. The solution is designed to run on a 64-Bit compatible server or personal computer (Windows, Linux, MacOSx) and be easily installed and completely operational in less than five minutes. The solution comes in three sizes: 4 GB, 40GB and 1TB versions. Teradata Express provides an opportunity to work with the Teradata Database as a fully configured system for development, testing and learning database management. The available versions of Teradata in VMware are 13, 13.10 and 14.0. Teradata Express runs Teradata ViewPoint for system monitoring and management and also includes EZLoader Utility to enable fast and easy data loading to the platform.

Teradata Studio Express is a powerful information discovery tool that queries data from Teradata Database Systems. Its features include Data Source Explorer for viewing database objects, such as schemas, tables, macros, stored procedures, user-defined functions, user-defined types, views and triggers, as well as Connection Profile Wizard to create connection profiles using the Teradata JDBC Driver. Also included is SQL Query Builder to visually build and edit SQL queries and Data Wizards to extract data from the Teradata Database to a file on the desktop as well as a loading wizard to load data from the desktop file system to the Teradata Database.

Teradata Express for VMware Player is a free solution and is ideal for developers who need to work in a private on-premises Cloud or desktop deployment model while retaining the option to move their work to the public Cloud or to Teradata environments. To download copies of Teradata Express Private go to: <http://downloads.teradata.com/download/database/teradata-express/vmware>.

## Teradata Express for Amazon EC2

Teradata Express for Amazon EC2 is a Public Cloud version of the Teradata Database and Teradata Studio Express products designed to leverage the Amazon EC2 platform. Teradata Express is fully functional and runs on Novell SLES 10 64-bit Linux hosted on Amazon's Elastic Compute Cloud (EC2). The EC2 platform is an industry leading service that provides a scalable IaaS computing environment primarily controlled by the user. Developers who use this service pay only for capacity they actually use and avoid the upfront costs often associated with traditional computing environments. Teradata offers a 1TB version of Teradata Express for Amazon EC2. Implementation takes only minutes and can be executed via the Amazon EC2 Dashboard. Developers can choose between Teradata versions 14 and 13.10.

Many developers utilize public Cloud environments to build out proof of concept projects before rolling them into production. The Public Cloud enables fast system provisioning, an isolated yet enterprise class platform for testing, and flexibility to experiment.

Teradata Express for Amazon EC2 is a free solution for non-production environments and ideal for developers who want to leverage the public Cloud and its economical pay as consumed model. To register for free access to Teradata Express for Amazon EC2 go to: <http://downloads.teradata.com/download/database/teradata-express/amazon-ec2>

## Teradata Aster for Amazon EC2

Big data analytics in the Cloud is now available through Aster Database Cloud Edition. Aster Database Cloud Edition leverages the elasticity and extreme scalability of Cloud computing while delivering massively parallel software that embeds MapReduce to perform analytics on multi-structured data. The Teradata Aster Discovery Platform can be deployed quickly and managed via the Amazon Web Services (AWS) Console. Teradata Aster has also partnered with leading infrastructure providers to offer a private cloud version of the solution. The Teradata Aster Discovery Platform enables applications to be fully embedded within the Aster Database engine for in-database analytics delivering deep analysis on massive data sets at high speed. The AWS deployment model allows clients to forgo the upfront cost often associated with analytic platform solutions as well as the ability to reduce costs by only paying for the system while it's in use. Other features of the Aster Database Cloud Edition include online fault

# Innovation in Cloud Analytics

tolerance and recovery, and automated parallel back up. Online Precision Scaling is also available – a one-click provisioning tool for increasing or decreasing capacity or performance across multi-functional tiers – giving the user fine level control of the environment and its costs.

## Teradata Aster Express for VMware Player

As previously detailed in this white paper, Teradata offers virtual images available for development, testing and learning with Teradata Express. In March 2012 Teradata introduced Aster Express for VMware player to help companies gain insights on multi-structured data. While not licensed for production use, Aster Express for VMware Player is a fully functional Aster cluster, a perfect choice for “taking a test drive” and exploring features and capabilities of MapReduce analytics on Big Data. The Aster Express cluster has two images – a Queen and Worker. Each image has been created for the VMware Player environment and runs on a 64-bit SLES Linux instance. A 64-bit capable CPU is also required to run Aster Express for VMware Player. To download copies of Aster Express for VMware Player go to: <http://developer.teradata.com/aster/articles/introducing-aster-express>.

## Teradata Applications in the Cloud

Teradata Applications include Cloud-based Integrated Marketing Management solutions for Marketing Operations, Campaign Management and Digital Messaging. The solutions are highly configurable with modules that add extensive feature sets empowering marketers to manage demand generation, multi-channel campaign management, digital marketing and performance monitoring. The Teradata Applications portfolio for the Cloud includes both SaaS and Private Cloud service models, and can be combined with Teradata Applications on-premise and integrated with other on-premise or third-party Private Cloud, Public Cloud or on-premise solutions. Users benefit from “Teradata At Your Service” (T@YS) a comprehensive customer service system. Teradata offers commercially available certifications for Cloud service offerings, which include SAS 70 Type 2, SSAE 16 SOC 2 or ISO 27001 compliance, and can offer SLAs for disaster recovery operations and system redundancy. Teradata Applications in the Cloud delivers a fully featured marketing suite that allows markets to execute on mission critical tactics while leveraging the convenience of Cloud Computing

## Hybrid Cloud Strategy

Hybrid Cloud Deployment is the most heavily researched deployment model by respondents to recent EMA Cloud research. Early Cloud adoption was generally focused on Public Cloud Deployment but as the industry has matured, enterprise adopters of Cloud technology are looking for strategies that reduce risk and still take advantage of the value found in Cloud Computing.

A hybrid Cloud use case is where a developer might utilize Teradata Express hosted on Amazon EC2 infrastructure for testing and development purposes, and then bring those results back into the ADW Private Cloud onsite. With a hybrid Cloud Computing model, developers become more effective and the company benefits from projects that bridge between private and public platforms depending on a company’s specific needs and workloads.

Hybrid models allow for Public Cloud applications to leverage data from behind the firewall without actually moving the data or exposing it to undo risk. Teradata’s line up enables developers to build applications within this model.

## EMA Perspective

Cloud has proven itself valuable to the enterprise. Early on, many vendors held true to working within the strict limits of Cloud definitions and guidelines. As the market has matured, users have become less interested in Cloud definitions and more focused on value and capabilities. Teradata clearly understands the value proposition of Cloud and builds or custom fits its solutions to deliver business results. As the industry evolves, EMA expects to see more solutions entering the market that borrow and blend Cloud features with an eye to value delivery. In the end, Cloud definitions are less important than saving capital and operations costs, improving IT service and adding agility to Business Intelligence (BI) environments.

---

Teradata clearly understands  
the value proposition  
of Cloud and builds or  
custom fits its solutions to  
deliver business results.

---

### About Enterprise Management Associates, Inc.

Founded in 1996, Enterprise Management Associates (EMA) is a leading industry analyst firm that provides deep insight across the full spectrum of IT and data management technologies. EMA analysts leverage a unique combination of practical experience, insight into industry best practices, and in-depth knowledge of current and planned vendor solutions to help its clients achieve their goals. Learn more about EMA research, analysis, and consulting services for enterprise line of business users, IT professionals and IT vendors at [www.enterprisemanagement.com](http://www.enterprisemanagement.com) or [blogs.enterprisemanagement.com](http://blogs.enterprisemanagement.com). You can also follow EMA on [Twitter](#) or [Facebook](#).

---

This report in whole or in part may not be duplicated, reproduced, stored in a retrieval system or retransmitted without prior written permission of Enterprise Management Associates, Inc. All opinions and estimates herein constitute our judgement as of this date and are subject to change without notice. Product names mentioned herein may be trademarks and/or registered trademarks of their respective companies. "EMA" and "Enterprise Management Associates" are trademarks of Enterprise Management Associates, Inc. in the United States and other countries.

©2013 Enterprise Management Associates, Inc. All Rights Reserved. EMA™, ENTERPRISE MANAGEMENT ASSOCIATES®, and the mobius symbol are registered trademarks or common-law trademarks of Enterprise Management Associates, Inc.

Teradata is a registered trademark of Teradata Corporation and/or its affiliates in the U.S. and worldwide.  
EB-7503> 0213

#### Corporate Headquarters:

1995 North 57th Court, Suite 120  
Boulder, CO 80301  
Phone: +1 303.543.9500  
Fax: +1 303.543.7687  
[www.enterprisemanagement.com](http://www.enterprisemanagement.com)  
2610.021313