



Evaluating Apache Cassandra as a Cloud Database

WHITE PAPER

By DataStax Corporation

March 2012

Contents

Introduction.....	3
Why Move to a Cloud Database?.....	3
The Cloud Promises Transparent Elasticity.....	3
The Cloud Promises Transparent Scalability.....	4
The Cloud Promises High Availability.....	5
The Cloud Promises Easy Data Distribution.....	5
The Cloud Promises Redundancy.....	6
The Cloud Promises Support for All Data Types.....	6
The Cloud Promises Easier Manageability.....	6
The Cloud Promises Lower Cost.....	6
Evaluating Apache Cassandra as a Cloud Database.....	7
What Is Apache Cassandra?.....	7
Why Cassandra?.....	7
Netflix – An Example of Succeeding in the Cloud with Cassandra.....	10
DataStax Enterprise – Certified Cassandra for Production Applications.....	11
Solving the Cloud Mixed-Workload Problem.....	11
Hadoop Analytics in the Cloud.....	13
Search With Solr in the Cloud.....	13
A Complete Big Data Platform for the Cloud.....	14
Visual Database Management.....	14
Enterprise Production Support and Services.....	17
Conclusion.....	17
About DataStax.....	18

Introduction

The amount of information that currently resides only in the cloud is small, but that's about to change. A recent study by IT industry analyst group IDC estimates that cloud computing accounts for less than 2 percent of IT spending today, but by 2015, nearly 20 percent of all information will be "touched" (stored or processed) in a cloud.¹ Moreover, IDC predicts that by that same year, as much as 10 percent of all data will be maintained in a cloud.²

Despite the growing movement toward cloud computing, some IT professionals remain standoffish toward the idea of porting a company's data onto a public cloud computing platform such as Amazon, Rackspace, and others. This position is understandable, given the current confusion over whether running a database in a cloud environment actually delivers tangible benefits – technical and otherwise –over keeping that same data on-premise.

Whether deciding to move a small or significant amount of data to a cloud database, today's IT decision-makers need to understand whether the solution they're considering is designed and/or implemented in a way that utilizes all the benefits and promises of cloud computing. This paper examines those key characteristics and discusses how Apache Cassandra™ stacks up from an evaluation perspective.

Why Move to a Cloud Database?

First, it should be understood that a cloud database is more than simply taking traditional relational database management system (RDBMS) software and running an instance of it on a cloud platform such as Amazon. Such a deployment in no way maximizes the capabilities of a cloud-computing environment.



¹ *Extracting Value from Chaos*, by John Gantz and David Reinsel, IDC, June 2011, <http://idcdocserv.com/1142>.

² *Ibid.*

But what constitutes a cloud-ready database? What features and functionalities must the database have to deliver on the potential that cloud computing offers? What follows is a discussion of some of the key promises of the cloud and the types of features a database should have to supply real benefits in a cloud environment.

The Cloud Promises Transparent Elasticity

The first promise of cloud computing worth noting is transparent elasticity. This equates to being able to add and subtract nodes (defined as actual physical machines or virtual machines) when the underlying application and business demands it.

This first attribute can be somewhat difficult for legacy RDBMSs, as well as some NoSQL solutions, as almost all are built on some master-slave architecture, which makes elastic expansion and contraction tricky and complex to manage.

Further, this capability should allow for online operations so that no downtime is experienced during expansion and contraction. In addition, a clustered database configuration should allow for some sort of easy load balancing so that an even distribution of the total workload is experienced.

The Cloud Promises Transparent Scalability

The primary reason for wanting elasticity in a cloud database is to scale properly to meet what businesses hope will be increasing customer traffic and accompanying workload. When talking about scale, there are normally two goals.

First, scaling out in the cloud carries with it an expectation that adding nodes will increase the performance of the database. The hope is that linear scalability will be experienced; for example, if two nodes are able to handle a throughput of 200,000 transactions, then four nodes should be able to manage 400,000.

Scale and speed are of the utmost concern to modern businesses today, and for good reason. A brief look at the following facts³ from some of today's premier software and services providers demonstrates how performance impacts their business:

- Amazon says every 100 millisecond (ms) delay has the potential to cost the company 1 percent of sales (in 2009, that translated to US\$245 million)
- Microsoft discovered that an additional 500ms of delay on its webpage loads resulted in the loss of about 1.2 percent in revenue per user
- Mozilla shaved 2.2 seconds of load time off its landing pages and correspondingly increased download conversions by 15.4 percent, which translated into an additional 60 million downloads each year

³ For individual references for each statistic, go to <http://blog.equinix.com/2011/03/optimizing-internet-application-performance/>.

- Shopzilla reduced its page load time by five seconds and saw an increase of 25 percent in page views and a 7 to 12 percent increase in revenue
- Experts estimate that just a 10ms latency could result in a reduction in revenue for U.S. brokerage houses of 10 percent

Second, scaling out also equates to being able to tackle “big data”⁴ while maintaining response time service level agreements (SLAs). This means that no matter the data volume size, end user requests should be serviced every bit as fast as if the data volume size was much less.

The Cloud Promises High Availability

Another key promise of cloud computing is increased uptime. Like performance and scalability, availability is critical to successful businesses. The cost of database downtime can vary widely depending on the industry; however, average downtime costs can range anywhere from about US\$90,000 per hour in the media industry to about US\$6.48 million per hour for large online brokerages.⁵

A database implemented in the cloud should have features that piggyback off of a cloud provider’s infrastructure where high availability is concerned. Additionally, that database’s degree of availability often is affected by the next two cloud computing promises: easy data distribution and redundancy.

The Cloud Promises Easy Data Distribution

Cloud providers typically promise the ability to distribute compute resources and data across different geographies or “zones” that the cloud provider makes available. Where the underlying database of a cloud application is concerned, this usually equates to a couple of desirable features.

First is the ability to read and write from any node that makes up the cloud database. Again, master-slave architectures usually have a difficult time meeting this requirement, especially where writes are concerned. Reads may be served up across various slave servers, but writes are a different matter.

The multi-geographic data distribution also supplies the benefits of high availability if a particular region goes down, and simplified disaster recovery so the data is protected if a particular physical disaster befalls one of the cloud provider’s data centers. Of course, this assumes the database actually contains redundant copies of data spread across different data centers of the cloud provider, which brings us to the next point.

⁴ For more information on big data management, see *Big Data: Beyond the Hype, Why Big Data Matters to You*, DataStax, October 2011, <http://www.datastax.com/wp-content/uploads/2011/10/WP-DataStax-BigData.pdf>.

⁵ “How Much Does Downtime Really Cost?,” by Henry Martinez, InfoManagement Direct, August 6, 2009, http://www.information-management.com/infodirect/2009_133/downtime_cost-10015855-1.html.

The Cloud Promises Redundancy

One of the main draws of the cloud is redundancy for both compute resources and data stored in the cloud. Where the database is concerned, redundancy means keeping more than one copy of the underlying data, so if the primary copy is destroyed, another copy is available for use.

How such redundancy is managed can be very important. As stated in the prior section, multi-data center support in different geographies is important if a large-scale disaster occurs at a particular site.

However, machine outages within a single data center also can pose high availability issues (either a complete outage- or response time-related problems) for a database. Good redundancy support can mitigate such problems, though, through the use of data center rack awareness capabilities. This equates to the database being smart enough to store multiple copies of data on different physical racks, so if one rack experiences a machine failure, the other rack can continue to serve data.

Of course, the database can both support multi-rack awareness at single data centers and provide the ability to store redundant copies of data at other, remote data centers.

The Cloud Promises Support for All Data Types

Cloud computing, in general, attempts to welcome all data types and formats – structured, semi-structured, and unstructured. However, doing so is more easily said than done.

The database being implemented in the cloud is really the key determiner of whether the cloud provider can easily keep this promise. This is where today's NoSQL databases shine over legacy RDBMSs because they typically offer a flexible and dynamic schema that accepts all key data formats more straightforwardly.

The Cloud Promises Easier Manageability

All cloud providers tout their ease-of-use, but the reality is that managing a distributed and geographically dispersed database is, normally, anything but easy. Web-based interfaces supplied by the cloud provider often can help with management duties. But for specific monitoring and management of the cloud database's particular feature set, the database vendor should provide a tool or set of tools that helps carry out routine administrative operations.

The Cloud Promises Lower Cost

Many IT professionals who begin looking at the cloud as an alternative to typical on-premise architectures assume they will experience cheaper software costs in a cloud implementation. However, they often have a rude awakening when they fail to experience cost benefits from managing data in the cloud. The cold reality is that some traditional RDBMS providers are every bit as expensive in a cloud implementation as they are in a standard on-premise implementation.

When looking to implement a database in the cloud, IT professionals should seek a cost structure that is friendly to scaling out horizontally, regardless of machine size or the data volume being managed. Otherwise, there is risk of unpleasant cost increases when the underlying business becomes very successful and more nodes are needed to manage ballooning data volumes and increased concurrent users.

Evaluating Apache Cassandra as a Cloud Database

So what is Apache Cassandra and how does it stack up against the criteria for cloud databases previously discussed? Following is an overview of Cassandra, including a description of the key technology differentiators that make it a stand-out cloud database. Also discussed is how DataStax Enterprise – a smart data platform powered by Cassandra – provides the best possible cloud database option for those needing to manage both real-time and analytic data in a cloud environment.

What Is Apache Cassandra?

Apache Cassandra is a highly scalable and high-performance distributed database management system that excels at being a real-time data store (i.e., the “system of record”) for online/transactional applications that need extremely fast read and write operations. Cassandra can manage the distribution of data across multiple data centers and offers incremental scalability with no single point of failure.

Cassandra was originally incubated at Facebook and is based upon Google’s BigTable and Amazon’s Dynamo software. The end result is an extremely scalable and fault-tolerant data infrastructure that solves both small and big data problems, handles write-intensive user traffic, delivers sub-millisecond caching layer reads, and supports demanding workloads involving petabytes of data.

Why Cassandra?

Key technical differentiators that make Cassandra a winning choice in a cloud computing environment include the following:

- A built-for-scale architecture that can handle petabytes of information and thousands of concurrent users/operations per second as easily as it can manage much smaller amounts of data and user traffic
- Peer-to-peer design that offers no single point of failure for any database process or function; every node is the same, so there is no concept of a master node or anything similar
- Online capacity additions that deliver linear performance gains for both read and write operations

- Read/write anywhere capabilities that equate to a true network-independent method of storing and accessing data
- Guaranteed data safety that ensures no loss of data, no matter what node is written to in a cluster
- Tunable data consistency that allows Cassandra to offer the data durability and protection like an RDBMS, but with the flexible choice of relaxing data consistency when application use cases allow
- Flexible/dynamic schema design that accommodates all formats of big data applications, including structured, semi-structured, and unstructured data; data is represented in Cassandra via column families that are dynamic in nature and accommodate all modifications online
- Simplified replication that provides data redundancy and is capable of being multi-data center and cloud in nature
- Data compression that reduces the footprint of raw big data by over 80 percent in some use cases
- A SQL-like language (CQL – Cassandra Query Language) that lessens the learning curve for developers and administrators coming from the RDBMS world
- Support for key developer languages (e.g., Java, Python) and operating systems
- No requirement for any special equipment; Cassandra runs on commodity hardware
- Very easy installations in cloud environments including Amazon Machine Images (AMIs) that enable a user to be up and running with a multiple-node cluster in minutes

Cassandra is built with the assumption that failures can and will occur in a data center or cloud infrastructure. Therefore, data redundancy to protect against hardware failure and other data loss scenarios is built into and managed transparently by Cassandra. Furthermore, this capability can be configured so that big data applications can use a single large database distributed across multiple, geographically dispersed data centers, between different physical racks in a data center, and between public cloud providers and on-premise managed data centers.

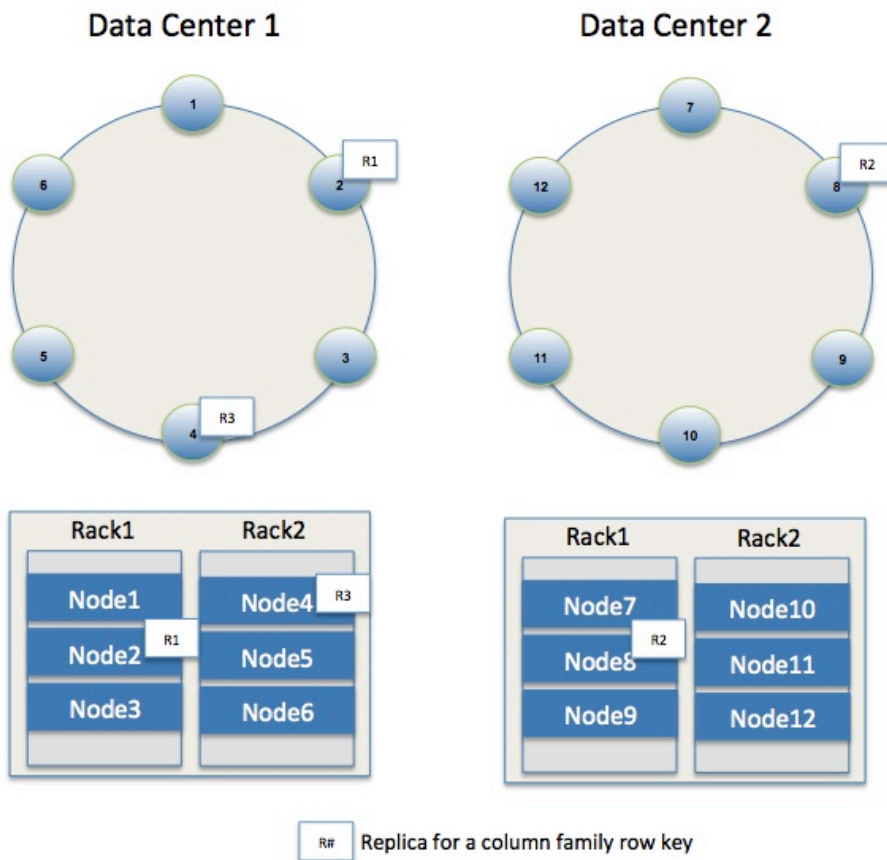
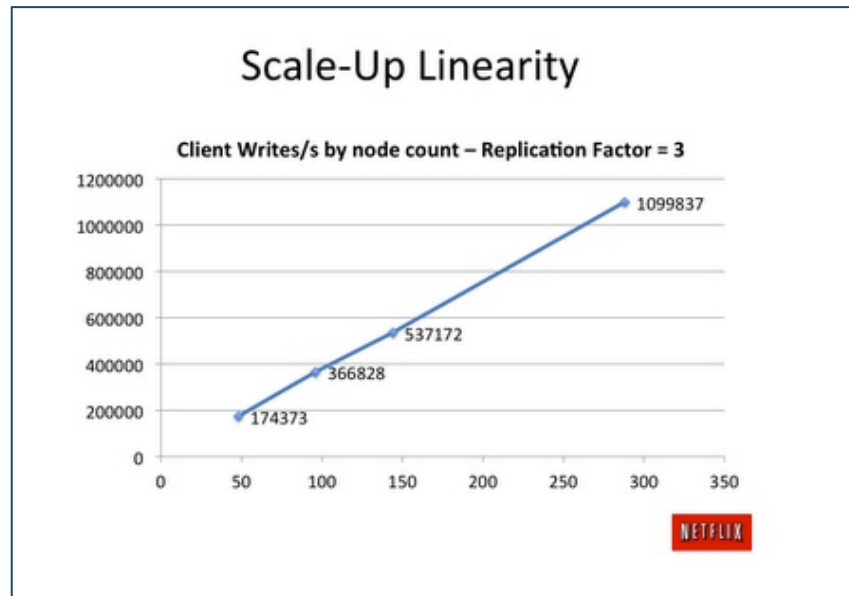


Figure 1: Cassandra multi-data center capabilities

These and other capabilities make Cassandra and DataStax Enterprise the smart choice for modern businesses with big data management needs that have outgrown traditional RDBMS software.

Netflix – An Example of Succeeding in the Cloud with Cassandra

With more than 25 million members worldwide, Netflix, Inc. (Nasdaq: NFLX) is the world's leading Internet subscription service for enjoying movies and TV shows. Netflix allows its members to



instantly watch unlimited movies and TV episodes streaming over the Internet to computers and TVs.

Figure 2: Performance results from Netflix's benchmark of Cassandra in the cloud

Cassandra and DataStax are a key part of Netflix's database infrastructure, with everything being hosted in the cloud. Netflix gave a presentation at the 2011 High Performance Transaction System workshop that demonstrated both the ease of use and linear performance capabilities of using Cassandra in the cloud. The following is an excerpt from a Netflix blog post summarizing the presentation:

" The automated tooling that Netflix has developed lets us quickly deploy large scale Cassandra clusters, in this case a few clicks on a web page and about an hour to go from nothing to a very large Cassandra cluster consisting of 288 medium sized instances, with 96 instances in each of three EC2 availability zones in the US-East region. Using an additional 60 instances as clients running the stress program we ran a

*workload of 1.1 million client writes per second. Data was automatically replicated across all three zones making a total of 3.3 million writes per second across the cluster.*⁶

The linear performance capabilities are illustrated well in the Netflix benchmark, delivering a very impressive 1.1 million writes per second. The ease with which Cassandra nodes can be configured and implemented in the cloud is also clear.

DataStax Enterprise – Certified Cassandra for Production Applications

Cassandra is a top open source project for the Apache foundation and enjoys strong community support and developer involvement. New community releases and patches are produced very quickly, with the understanding that community builds are not put through any enterprise-styled quality assurance process, and often contain a mixture of enhancements plus bug fixes.

By contrast, DataStax Enterprise only contains selected Cassandra releases chosen by the expert staff and committers at DataStax. Each selected release is then put through a rigorous certification process designed by DataStax engineers and QA staff to ensure it is stable and ready for enterprise production systems. Any found issues are immediately fixed and applied to the DataStax Enterprise server.

In addition, DataStax provides enterprises with predictable, certified quarterly service pack updates as well as other software benefits such as emergency hot fixes (for production outages) and bug escalation privileges that prioritize customers' issues over community-submitted bugs.

DataStax Enterprise – Real-Time, Analytics, and Search in the Cloud

DataStax is the leading provider of enterprise NoSQL software products and services based on Apache Cassandra. Through its offerings, DataStax supports businesses that need a progressive data management system that can serve as a real-time datastore for critical production applications, and delivers built-in analytic and search capabilities for analyzing and searching that data once it is in Cassandra.

DataStax Enterprise inherits Cassandra's entire, powerful feature set for servicing modern real-time applications, and uses it to merge in a fault-tolerant, analytics, and enterprise search platform that provides Hadoop MapReduce, Hive, and Pig support for analytics and uses Apache Solr for fast enterprise search.

Solving the Cloud Mixed-Workload Problem

A primary benefit that DataStax Enterprise provides to enterprises needing smart big data management capabilities is its ability to service real-time, analytic, and enterprise search data

⁶ "Benchmarking Cassandra Scalability on AWS – Over a million writes per second," by Adrian Cockcroft and Denis Sheahan, *The Netflix "Tech" Blog*, November 2, 2011, <http://techblog.netflix.com/2011/11/benchmarking-cassandra-scalability-on.html>.

operations in the same database cluster without any of the loads impacting the other. The key to making this possible is the underlying architecture of Cassandra.

Hadoop Analytics in the Cloud

Built into DataStax Enterprise is an enhanced Hadoop distribution that utilizes Cassandra for many of its core services. DataStax Enterprise provides integrated Hadoop MapReduce, Hive, Pig, Mahout, and job/task tracking capabilities, replacing Hadoop's HDFS storage layer with Cassandra (CassandraFS). The end product is a single integrated solution that provides increased reliability, simpler deployment, and lower total cost of ownership (TCO) than a traditional Hadoop solution. DataStax Enterprise also is fully compatible with existing HDFS, Hadoop, and Hive tools and utilities.

Another benefit of using Hadoop in DataStax Enterprise is that it eliminates the complexity and single points of failure of the typical Hadoop HDFS layer. From an operational standpoint, there is no need to set up a Hadoop name node, secondary name node, Zookeeper, and so on.

Instead, DataStax Enterprise provides a single layer in which every node is a peer of the others and automatically knows its position in the cluster. On startup, all DataStax Enterprise nodes automatically start a Hadoop task tracker, and one of the nodes is elected to be the job tracker. If the job tracker node fails, the job tracker is automatically restarted on a different node. DataStax Enterprise utilizes full data locality awareness for Hadoop task assignment.

Search With Solr in the Cloud

DataStax Enterprise includes strong enterprise search support via Lucene and Apache Solr. Coming from the Apache Lucene project, Solr is the most popular open source enterprise search platform in use today.

Solr's primary features include robust full-text search, hit highlighting, faceted search, rich document (e.g., PDF, Microsoft Word) handling, and geospatial search.

By integrating Solr into the DataStax Enterprise big data platform, DataStax extends Solr's capabilities and overcomes a number of shortcomings that native Solr has such as:

- Lack of data durability (community Solr has no write-ahead log, so data can be lost if a node crashes). No chance of data loss exists with Solr in DataStax Enterprise
- Solr's write bottleneck, as all writes go through a single master. But with DataStax Enterprise, users can read and write to any Solr node in the cluster
- Replication and sharding of Solr, which is a manual process and requires careful planning for scaling and failover. DataStax Enterprise, however, supplies automatic sharding and no single point of failure
- Manual re-indexing of data. Indexes can be automatically rebuilt in DataStax Enterprise
- Writes to indexes in community Solr cannot span multiple data centers; there is only a single master that replicates via rsync. But, in DataStax Enterprise, multiple writes to

search indexes in different data centers are merged together (i.e., writes can occur anywhere)

- Solr indexes in DataStax Enterprise can be dropped/recreated/rebuilt on the fly (versus how things are done in native Solr)

In essence, in the same way that DataStax Enterprise takes Hadoop and delivers a fault-tolerant, no single point of failure, and dynamically scalable Hadoop/analytics system, it automatically does the same thing for Solr and enterprise search operations. Using Cassandra as the underlying foundation, DataStax Enterprise allows search data to be written to any participating search node in a DataStax Enterprise cluster. New search nodes can be added online to increase both fault tolerance and performance, with gains being near linear in nature.

Those currently using Solr will be at home with DataStax Enterprise, as the solution is 100 percent Solr compatible, with all Solr utilities, APIs, and so on, included.

A Complete Big Data Platform for the Cloud

A key benefit of DataStax Enterprise is the tight feedback loop it has between real-time applications and the analytics and search operations that naturally follow. Traditionally, users would be forced to move data between systems via complex ETL processes, or perform both functions on the same system with the risk of one impacting the other. In big data environments, this process can be time-consuming and burdensome.

With DataStax Enterprise, real-time, analytic, and search big data operations take place in the same distributed system, but users have the ability to dedicate certain nodes solely for analytics or search so their workloads don't slow down real-time processing. Users simply define one or more replica groups, and configure the role of each – one or more Cassandra, Hadoop, or HDFS (i.e., HDFS without job/task tracker), and search/Solr nodes. Writes are instantly replicated between all nodes.

With DataStax Enterprise, users truly have the best of all worlds for big data management. They have all the power of Cassandra serving their highest-volume and high-velocity, real-time applications; the power of Hadoop, Hive, and Pig working directly against the same data for analytics; and Solr for enterprise search in the same distributed database. The result is smart workload isolation for big data applications that is much simpler to manage and more reliable than any alternative.

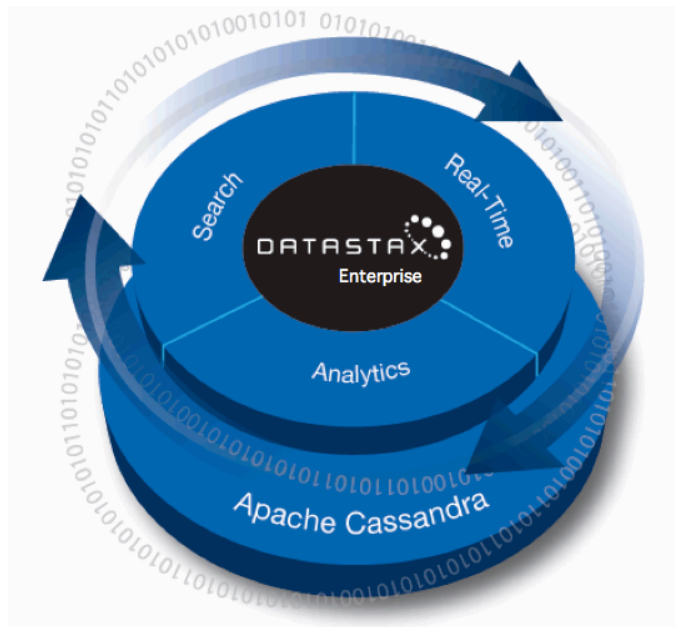


Figure 3: DataStax Enterprise – real-time and analytic, and search data in one cloud database

Visual Database Management

DataStax Enterprise includes a visual, browser-based management solution named OpsCenter Enterprise to manage and monitor cloud database deployments. OpsCenter Enterprise allows a developer or administrator to manage and monitor the health of cloud databases from a centralized web console.

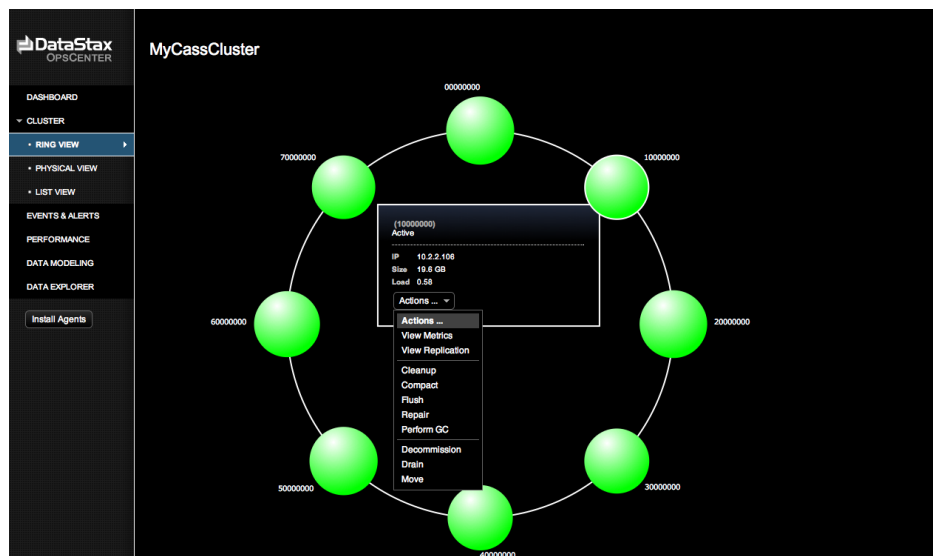


Figure 4: OpsCenter Enterprise database cluster ring view

OpsCenter Enterprise uses an agent-based architecture to monitor and carry out tasks on each node in a DataStax Enterprise cluster. Through a graphical and intuitive point-and-click interface, a user can understand the state of a cluster, which nodes are up and down, and what type of performance users are experiencing. Key events are reported into a centralized dashboard displayed along with other vital statistics.



Figure 5: OpsCenter dashboard

Analytic operations also can be monitored and controlled from within OpsCenter Enterprise:

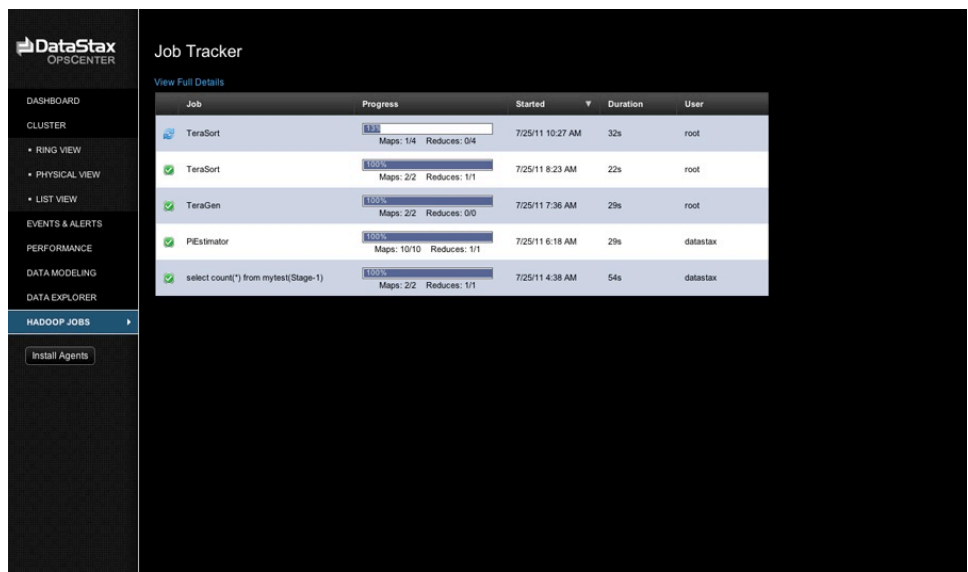


Figure 6: OpsCenter analytic operations monitoring

Enterprise Production Support and Services

Cloud implementations often require fast access to skilled expertise. DataStax Enterprise includes experienced production support and consultative services from Cassandra experts who know cloud environments. IT professionals can choose the right production support package for their business needs, including rapid response SLAs and consultative help.

Additionally, DataStax offers professional training on Cassandra and Hadoop, with classes offered in many major cities as well as on-site for corporations that need many staff members trained at once.

Conclusion

Moving to a cloud-based infrastructure necessitates choosing a database that is capable of fully utilizing all the benefits the cloud provides. The following table illustrates how Apache Cassandra meets all key attributes required of a cloud database, and how DataStax Enterprise, which is powered by Cassandra, delivers a unique, smart data platform ready for the cloud and capable of managing both real-time and analytic data in the same database cluster.

Cloud Database Requirement	Meet?	Additional Information
Transparent elasticity	Yes	Peer-to-peer architecture makes scaling out in an online fashion easy
Transparent scalability	Yes	Linear performance gains delivered through node additions; big data-capable
High availability	Yes	No single point of failure
Easy data distribution	Yes	Simple replication among one or more data centers, geographies, and the cloud; read/write anywhere design
Data redundancy	Yes	Data is easily replicated across many regions, with rack awareness also supported
Support for all data formats	Yes	Structured, semi-structured, and unstructured data is supported in data model
Simple manageability	Yes	Easy AMI installs for the cloud; web-based tools are also provided
Low cost	Yes	Apache Cassandra is free/open source; DataStax Enterprise is 80 to 90 percent less expensive than legacy RDBMS software

To find out more about DataStax Enterprise and obtain software, please visit www.datastax.com or email info@datastax.com.

About DataStax

DataStax offers products and services based on the popular open-source database, Apache Cassandra™ that solve today's most challenging big data problems. DataStax Enterprise combines the performance of Cassandra with analytics powered by Apache Hadoop and enterprise search with Apache Solr, creating a smartly integrated, big data platform. With DataStax Enterprise, real-time, analytic, and search workloads never conflict, giving you maximum performance with the added benefit of only managing a single database.

The company has over 140 customers, including leaders such as Netflix, Disney, Cisco, Rackspace and Constant Contact, and spans verticals including web, financial services, telecommunications, logistics and government. DataStax is backed by industry leading investors, including Lightspeed Venture Partners and Crosslink Capital and is based in San Mateo, CA.

For more information, visit www.datastax.com.