

Membase Announcements Hadoop World 2010



Membase – Company Background



- ズ Database software company
- ¥ Founded: January 2009
- 🛛 Proven Leadership
 - <u>Management team</u> combines 50+ years experience from Informix, VMware, IBM, Intel, Netscape, Hotmail, Yahoo!
 - <u>Founders</u> have previously founded companies worth over \$250 million at exit (Kiva, Akimbi, Fifth Generation Systems)
- **X** Customer and Partner Advisory Board
 - Cadir Lee (Zynga), Bobby Johnson (Facebook), Mike Olson (Cloudera), Alan Kasindorf (Six Apart)
- X Strong balance sheet with top-tier investors
 - \$15 million funding (Series A and B)
 - Accel Partners, Mayfield Fund, North Bridge Venture Partners, Zynga

Hadoop World 2010 Announcements



Monday 11 October 2010

X Membase is generally available

- Membase Server Enterprise Edition
- Membase Server Community Edition
- Membase source code Apache 2.0 license
- X NorthScale, Inc. is now Membase, Inc.





Membase and Cloudera announce partnership, integration and joint customers

- Membase Server: low-latency access to data
- Cloudera Distribution for Hadoop: batch data analytics
- Bidirectional data integration via Flume and Sqoop
- AOL Advertising and ShareThis joint customers

Membase Generally Available

Where to get Membase



membase.org



membase.com



Source code. Community binary. Developer community. Certified product. Membase Network. Services and support.



Membase Enterprise Edition - Pricing

	Membase Basic	Membase Standard	Membase Premium	Memcached Basic	Memcached Standard	Memcached Premium	Membase Free
Pricing and Licensing							
Price	\$999	\$1,999	\$2,999	\$499	\$999	\$1,499	\$0
Licensing	1 Node	1 Node	1 Node	1 Node	1 Node	1 Node	1 Node, Max 2 Node Cluster
Licensing Restrictions	None	None	None	None	None	None	Max 2 Node Cluster for Production; Unlimited for Eval
Support Entitlements							
Hours	10 x 5	10 x 5	24 x 7	10 x 5	10 x 5	24 x 7	n/a
Hours of Operation	7am - 5pm PT	7am - 5pm PT	n/a	7am - 5pm PT	7am - 5pm PT	n/a	n/a
Support Channel	Forums, Email, Web	Forums, Email, Web, Phone	Forums, Email, Web, Phone	Forums, Email, Web	Forums, Email, Web, Phone	Forums, Email, Web, Phone	Forums Only
Number of Cases	5	Unlimited	Unlimited	3	Unlimited	Unlimited	n/a
Response Time - P1	1 Business Day	5 hour (within business hours)	2 hours	1 Business Day	5 hour (within business hours)	2 hours	n/a
Response Time - P2	1 Business Day	1 Business Day	5 hour	1 Business Day	1 Business Day	5 hour	n/a
Product Update Entitlements							
Major Releases (1.0)	x	×	X	x	x	x	x
Minor Releases (1.x)	x	x	x	x	x	x	x
Maintenance Releases (1.0.x)	x	x	x	x	x	x	x
Hotfixes	x	x	х	X	x	X	n/a
Membase Network	x	x	x	x	x	x	n/a
Update Notification Service	x	x	x	х	х	x	n/a
Technical Alerts	x	x	x	x	x	x	n/a

Membase-Cloudera Partnership

"AOL serves more than 5 billion impressions per day from our ad serving platforms, and any incremental improvement in processing time translates to huge benefits in our ability to more effectively serve the ads to needed meet our contractual commitments. Traditional databases like MySQL lack the scalability required to support our goal of five milliseconds per read/write. Creating user profiles with **Hadoop**, then serving them from **Membase**, reduces profile read and write access to under a millisecond, leaving the bulk of the processing time budget for improved targeting and customization."

> Pero Subasic Chief Architect, AOL

Membase-Cloudera Partnership



- Joint development of bi-directional software integration between Membase and Hadoop
 - Membase NodeCode Module streaming interface to Cloudera Distribution for Hadoop via Flume interface
 - Sqoop-derived command line utility for bi-directional batch movement of data between Membase and Cloudera Distribution for Hadoop
- Joint marketing and sales of integrated distributed OLTP-OLAP solution
 - Membase the distributed OLTP solution
 - Cloudera the distributed OLAP solution
- X Cloudera to distribute integration components

Customer use case – Ad targeting





What is Membase?

Before: Application scales linearly, data hits wall





Users

Membase is a distributed database





In the data center

On the administrator console

Built-in Memcached Caching Layer





Fact: Membase development team has also contributed over half of the code to the Memcached project.

Membase is Simple, Fast, Elastic





Five minutes or less to a working cluster

- Downloads for Linux and Windows
- Start with a single node
- One button press joins nodes to a cluster
- **¤** Easy to develop against
 - Just SET and GET no schema required
 - Drop it in. 10,000+ existing applications *already* "speak membase" (via memcached)
 - Practically every language and application framework is supported, out of the box
- **Easy to manage**
 - One-click failover and cluster rebalancing
 - Graphical and programmatic interfaces
 - Configurable alerting

Membase is Simple, Fast, Elastic





X Predictable

- "Never keep an application waiting"
- Quasi-deterministic latency and throughput
- ズ Low latency
 - Built-in Memcached technology
 - Auto-migration of hot data to lowest latency storage technology (RAM, SSD, Disk)
 - Selectable write behavior asynchronous, synchronous (on replication, persistence)
- - Multi-threaded
 - Low lock contention
 - Asynchronous wherever possible
 - Automatic write de-duplication

Membase is Simple, Fast, Elastic



		Alanda - Companya - Mala	
NorthScale	O WEB CONSOLE		- agricul
MONITOR	MANAGE SERVERS		
Cluster Overview	Active Servers Pending Rebalance	Stop Rebala	nce
Data Buckets			
Alerts	Up east1.campus1.mycorp.net	90% C	omplete
Logs	up 💼 📓 east2.campus1.mycorp.net	90% 0	omplete
MANAGE	Up 💼 📦 east3.campus1.mycorp.net	90% C	omplete
Data Buckets			
Servers	Up 🔤 🗰 newservertoadd.net	ADDING SERVER 90% C	omplete
SETUP	Up 💼 📓 XXXXXX.net	REMOVING SERVER 90% C	omplete
Settings			

Zero-downtime elasticity

- Spread I/O and data across commodity servers (or VMs)
- Consistent performance with linear cost
- Dynamic rebalancing of a live cluster
- X All nodes are created equal
 - No special case nodes
 - Clone to grow
- **¤** Extensible
 - Filtered TAP interface provides hook points for external systems (e.g. full-text search, backup, warehouse)
 - Data bucket engine API for specialized container types
 - Membase NodeCode [FUTURE]

Proven at small, and extra large scale



- Leading cloud service (PAAS) provider
- ▼ Over 65,000 hosted applications
- ▼ Over 2,000 users to date
- Membase Server serving over 3,000 Heroku customers



- Social game leader FarmVille, Mafia Wars, Café World
- Over 230 million monthly users
- Membase Server is the 500,000 ops-per-second database behind FarmVille and Café World

After: Data layer scales like application logic layer Data layer now scales with linear cost and constant performance.



Scaling out flattens the cost and performance curves.

A practical path to NoSQL database adoption



Technical Details - Backup

vBucket mapping





Membase "write" data flow – application view User action results in the need to change the VALUE of KEY Application updates key's VALUE, 2 performs SET operation Membase (memcached) client hashes KEY, identifies KEY's master server 3 SET request sent over network to master server Membase replicates KEY-VALUE pair, caches it in memory and stores it to disk

22

Membase data flow – under the hood



23

Membase Architecture





Membase Architecture



Data buckets are secure membase "slices"





In the data center

On the administrator console

