



The Elastically Scalable Database™

Team

BARRY MORRIS
CEO

StreamBase Systems, IONA Technologies, Lotus Development, Leading Technologies, PROTEK



JIM STARKEY
CTO

Netfrastructure, Harbor Software, Interbase Software, Sun Microsystems, MySQL, DEC, CCA



ADAM ABBREVAYA
VP ENG

Verisign, M-Cube, Pantero, Progress, Excelon, Renaissance, MIT Lincoln Labs

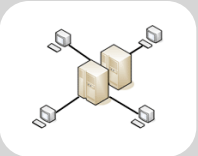
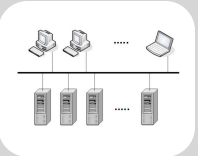
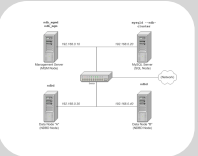
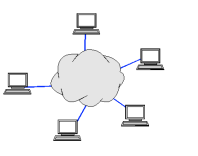


NuoDB Summary

- ▶ SQL/ACID RDBMS: Emergent Architecture
- ▶ Elastic Scale-out: Shared Nothing
- ▶ Single Logical Database: No sharding, no partitioning
- ▶ Run Anywhere: On-premises, Cloud, Hybrid
- ▶ Store Anywhere: KV-store (eg files, S3, HDFS)
- ▶ Multi-tenant: DBs share machines dynamically
- ▶ Geo-distributed: Active/active in multiple places
- ▶ Extreme Availability: Redundancy, live upgrade
- ▶ Minimal Admin: No knobs, Auto everything



Cloud Database Crisis

	Mainframe	Client-Server	Client-Cluster	Webscale
Datacenter				
Size	Megabytes	Gigabytes	Terabytes	Petabytes
Users	100's	1,000's	10,000's	100,000's
Typical TPS	100	1,000	2,000	20,000
Workload	Simple, Strings, Numbers	More Complex	Complex, Specialized	Very Complex, Hybrid Transactional, Analytical, Rich
Elasticity	Pre-provision	Pre-provision	Pre-provision	On-demand
Availability	Best Effort	Best Effort	Five 9's	100% Uptime
Multi-tenancy	Dedicated Host	Dedicated Host	Dedicated Host	Shared Resources
Multi-site	Local	Disaster Recovery	Disaster Recovery	Active/Active
Location	Datacenter	Datacenter	Datacenter	Cloud (Public, Private, Hybrid)
Developers	Hand Coding	Hand Coding, 4GLs	JDBC, ODBC	Agile, Rich Tools & Frameworks

The Cloud-SQL-ACID Test

Cloud Scalability

Elasticity
Virtualization
Extreme Availability
Geo-Distribution
Zero Admin

SQL Power

Set-based Access
Relational Algebra

NoSQL

OLD SQL



ACID Reliability

Atomicity
Consistency
Isolation
Durability



Emergent Database Architecture

"An emergent behavior can appear when a number of simple entities operate in an environment, forming more complex behaviors as a collective."

- Wikipedia



Poleposition - Single Node

Notes


MySQL 5.1

NuoDB Beta 3 - Single Node

<http://www.polepos.org>

In early tests NuoDB on a single node was 2x to 20x faster than MySQL 5.1 running the industry standard Poleposition Benchmarks.

Your mileage may vary.

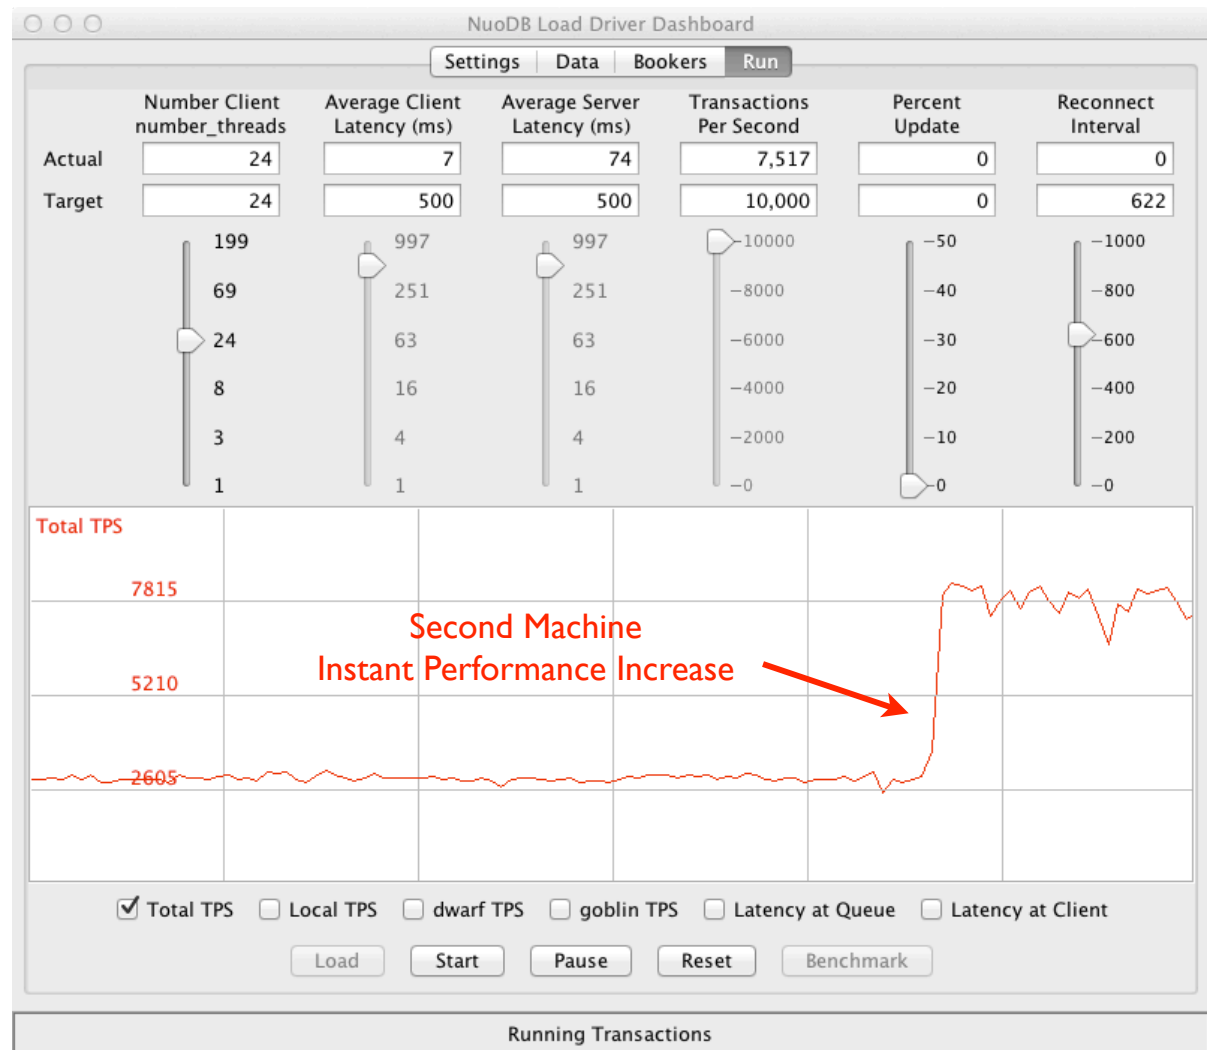


- ▶ Time taken for given benchmark, normalized to NuoDB = 1
- ▶ Less is Better



Adding a Second Machine

- Second machine typically doubles TPS
- Second machine is added to live database while it is running at 1,000's of TPS
- Performance increase is immediate
- BTW - you can take either machine away and the database keeps running without data loss



How does it work?

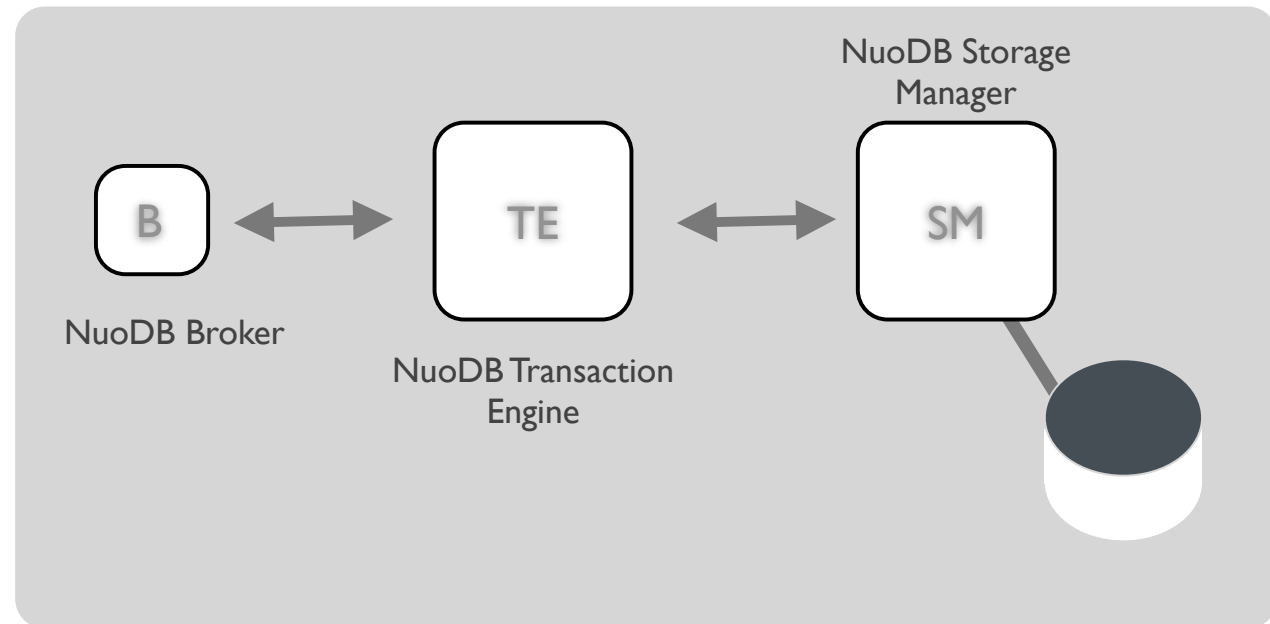

US Patent Abstract

“A multi-user, elastic, on-demand, distributed relational database management system. The database is fragmented into distributed objects called atoms. Any change to a copy of an atom at one location is replicated to all other locations containing a copy of that atom. Transactional managers operate to satisfy the properties of atomicity, consistency, isolation and durability.”

(The patent will issue in March 2012, about 12 months from filing, and without any Office Actions. The examiners found no prior art.)

Let's walk through what the description means ...

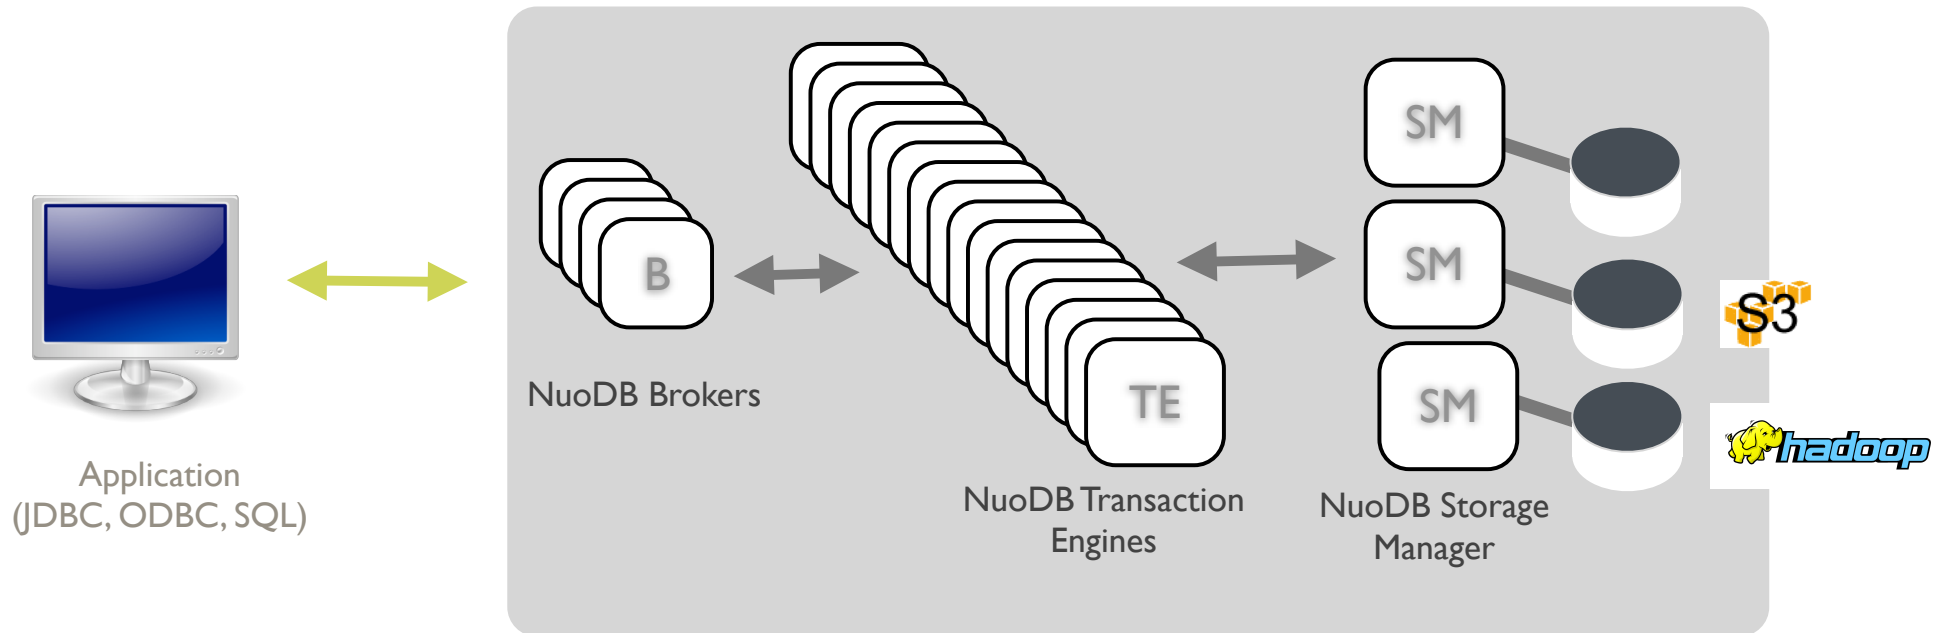
Baseline System



ATOM	Everything is an ATOM
Storage Manager	Key-value Storage
Transaction Engine	Diskless Node Performs ACID Transactions
Broker	Client Load-balancing



How NuoDB Works



Tx Scalability	Add as many of TE, SM or B as you like
I/O Scalability	Distributed IO, eg Hadoop HDFS
100% Uptime	No single point of failure, geo-distributed
Optimal Utilization	Multiple Databases on finite resources
Low People Costs	Single Console Management



NuoDB Summary

- ▶ SQL/ACID RDBMS: Emergent Architecture
- ▶ Elastic Scale-out: Shared Nothing
- ▶ Single Logical Database: No sharding, no partitioning
- ▶ Run Anywhere: On-premises, Cloud, Hybrid
- ▶ Store Anywhere: KV-store (eg files, S3, HDFS)
- ▶ Multi-tenant: DBs share machines dynamically
- ▶ Geo-distributed: Active/active in multiple places
- ▶ Extreme Availability: Redundancy, live upgrade
- ▶ Minimal Admin: No knobs, Auto everything





nuodb

The Elastically Scalable Database™



Copyright © NuoDB 2012
