



Structured Data Storage

Xgen Congress Short Course
2010

Adam Kraut

BioTeam Inc.

- **Independent Consulting Shop:**

Vendor/technology agnostic

- Staffed by:**

- ▶ Scientists forced to learn High Performance IT to conduct research

- **Our specialty:**

Bridging the gap between Science & IT



Data Management Buzzwords

- Linked Data
- NoSQL
- Distributed Database
- Non-Relational (Schema-free)
- Document-based
- Object-based
- Key-value
- Partitioning
- Fault Tolerance

One Size Does Not Fit All


- RDBMS have become ubiquitous
 - ▶ Often synonymous with the term database
- Databases precede the implementation relational systems
- Structured storage extends far beyond the relational realm
- 90% of applications are using 10% of the features of modern RDBMS

Scaling RDBMS

- “An infinitely scalable relational database is an engineering impossibility” – Werner Vogels



Database taxonomies



Feature-first

- Oracle
 - SQL Server
 - PostgreSQL
 - MySQL
-
- Even in large enterprises, a single database instance can support the entire workload

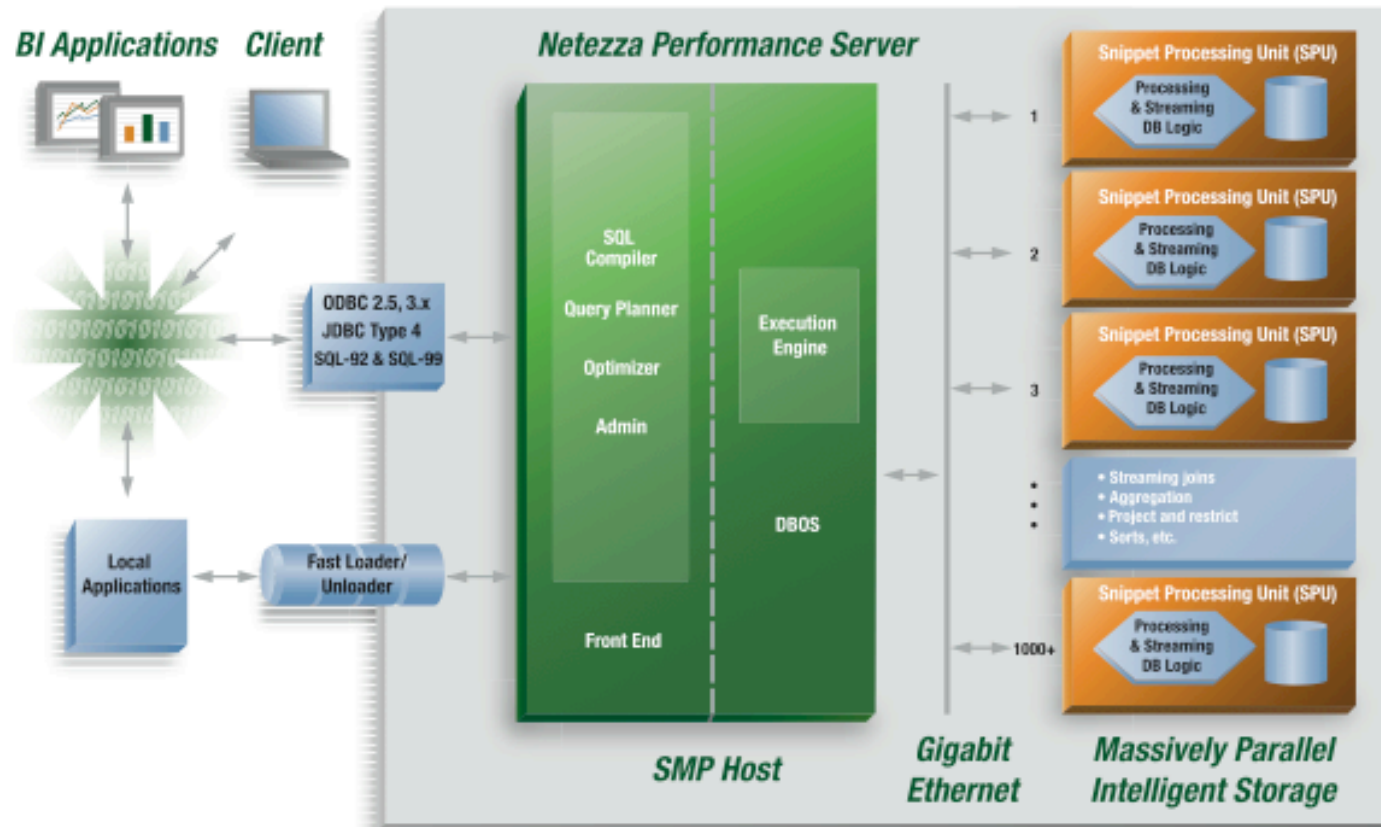
Scale-first

- Scale is more important than features
- When a single RDBMS won't cut it
 - ▶ Shard the data over a large number of systems
- Full relational model is impossible to scale
 - ▶ Cross-instance joins
 - ▶ Aggregations
 - ▶ Globally unique secondary indexes
 - ▶ Global stored procedures
- Examples
 - ▶ DB2 Parallel Edition
 - ▶ Oracle Real Application Clusters (RAC)

Purpose Optimized Storage

- Special purpose
- Often designed to beat commercial RDBMS on specific benchmarks
- Event stream processing
- Data warehousing products
- Examples
 - ▶ Aster Data, Netezza, Greenplum

High-Performance Streaming Architecture: Asymmetric Massively Parallel Processing



Simple Structured Storage

- Simple, cheap, fast
- Low operational burden
- Examples
 - ▶ BerkeleyDB
 - ▶ SimpleDB
 - ▶ MongoDB

Alternative Database Engines

- BerkeleyDB - <http://www.oracle.com/database/berkeley-db>
- memcached - <http://memcached.org>
- BigTable - <http://labs.google.com/papers/bigtable.html>
- HBase - <http://hadoop.apache.org/hbase>
- CouchDB - <http://couchdb.apache.org>
- MongoDB - <http://www.mongodb.org>
- Tokyo Cabinet - <http://1978th.net/tokyocabinet>
- Redis - <http://code.google.com/p/redis>
- Riak - <http://riak.basho.com>
- Cassandra - <http://incubator.apache.org/cassandra>

MongoDB

- Document-oriented storage (JSON-like schema)
- Written in C++
- Fast, in-place updates
- Replication, and fail-over support
- Auto-sharding
- MapReduce for aggregations
 - ▶ Written in Javascript



`{name: "mongo", type: "db"}`

Redis

- Fast, in memory key-value store
- STRING, LIST, SET, and ZSET data types
- Persistence via async snapshots
- Perfect Data Structures/State/Cache Server



SimpleDB

- Hierarchical structure, not a table
- Schema-less
 - ▶ Attributes only exist when associated with a value
 - ▶ No NULL values
- Limited query capability
 - ▶ No SQL
 - ▶ No joins
- All data is stored as text
 - ▶ No data types
- Limited Attribute Sizes (1024 bytes)
- Eventual consistency model
 - ▶ Information may be slightly out of date

NoSQL Hype vs. Reality

- Schema-free
- Scalable
- Fast
- Hierarchical data structures
- No general-purpose query language
 - ▶ Yet another language to learn
- Many-to-many relationships are problematic
- Lacking tool support
- Lacking library support

BioTeam's Use of NoSQL


- Today...
- WikiLIMS
 - ▶ Semantic MediaWiki (RDF Triple store)
- Configuration Management Framework
 - ▶ Chef (CouchDB)
- Amazon Web Services Workflows
 - ▶ SimpleDB to store state

Schema-Free

ID	Category	Subcat.	Name	Color	Size	Make	Model
Item_01	Clothes	Sweater	Cathair Sweater	Siamese	Small, Medium, Large		
Item_02	Clothes	Pants	Designer Jeans	Paisley Acid Wash	30x32, 32x32, 32x34		
Item_03	Clothes	Pants	Sweatpants	Blue, Yellow, Pink	Large		
Item_04	Car Parts	Engine	Turbos			Audi	S4
Item_05	Car Parts	Emissions	02 Sensor			Audi	S4
Item_06	Motorcycle Parts	Bodywork	Fender Eliminator	Blue		Yamaha	R1
Item_07	Motorcycle Parts, Clothing	Clothing	Leather Pants	Black	Small, Medium, Large		

<http://aws.amazon.com/>

Schema-Free MySQL



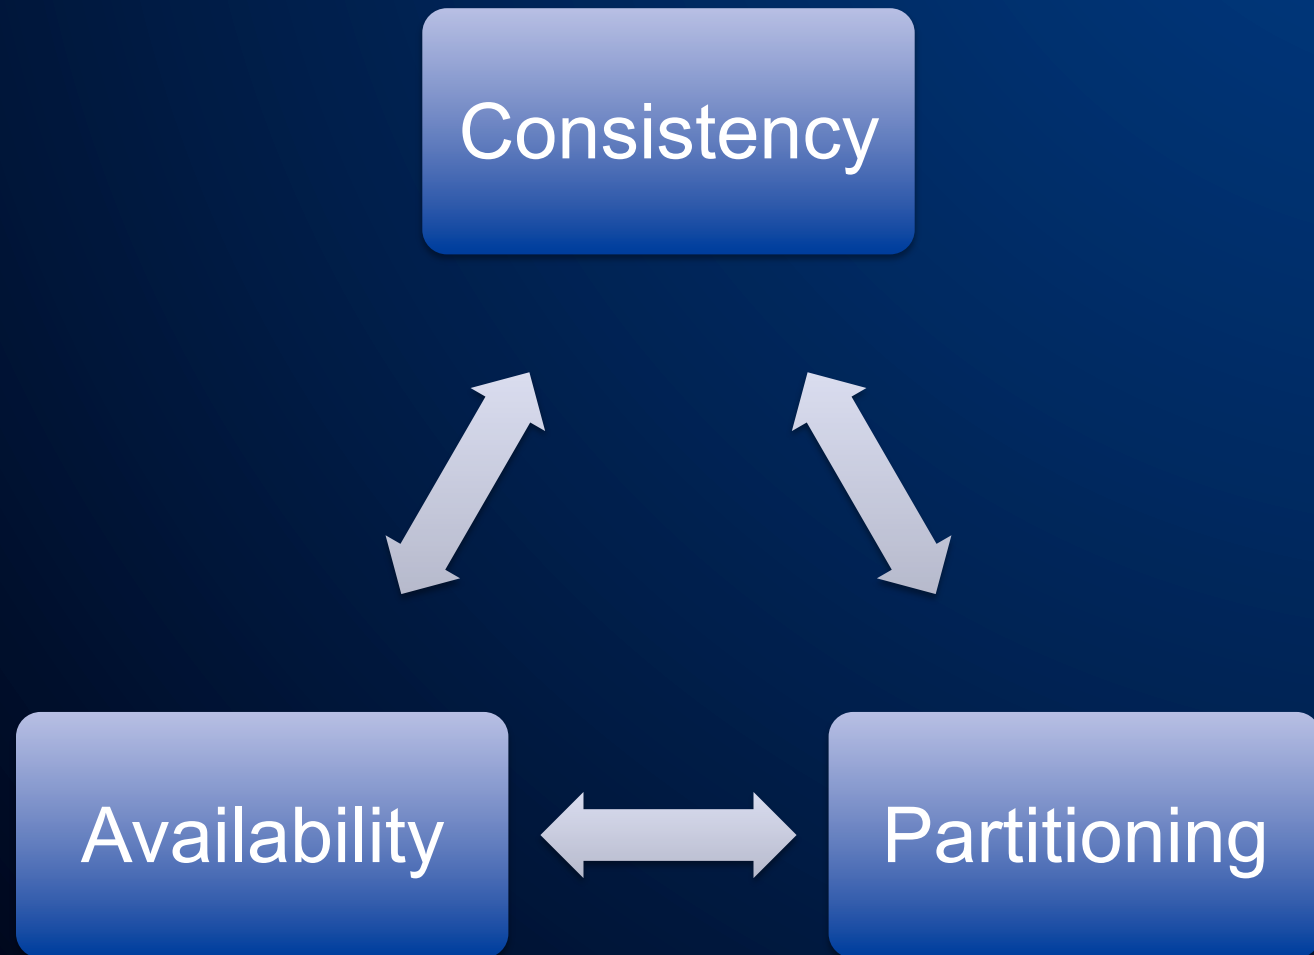
Properties of Distributed Systems

- Design for failure
 - ▶ Disks will fail
 - ▶ When is RAID6 unacceptable (2019?)
 - ▶ Nodes will fail
- Must maintain data consistency
 - ▶ Is it useful otherwise?
- Network partitioning

Eventual Consistency

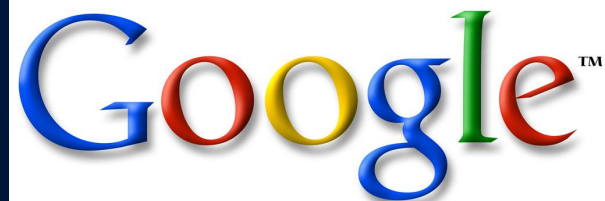
- “when no updates occur for a long period of time, eventually all updates will propagate through the system and all the replicas will be consistent”
- **Eventually** all clients will see the updates

CAP Theorem



Balancing the tradeoffs...

- Three properties of shared-data systems
 - ▶ Consistency of the data
 - ▶ Availability of the system
 - ▶ Partition tolerance
- Only two can be achieved at any given time
- Network partitions are a given
- See also: Project Triangle
 - ▶ Good, Fast, Cheap... pick two.

The Google logo, featuring the word "Google" in its signature multi-colored font (blue, red, yellow, blue, green, red) with a trademark symbol (TM) to the upper right. It is set against a white rectangular background.


Google™

The Amazon.com logo, featuring the text "amazon.com" in a bold, black, sans-serif font. Below the text is a yellow curved arrow pointing from the 'a' to the 'z'. To the right of the arrow is the tagline "and you're done." in a smaller, black, sans-serif font, followed by a trademark symbol (TM). The entire logo is set against a white rectangular background.

amazon.com[®]
and you're done.[™]

The Facebook logo, featuring the word "facebook" in a white, lowercase, sans-serif font. It is set against a blue rectangular background, which is itself enclosed within a white border.

facebook[®]



Map Reduce

BigTable

Dynamo



Thrift

Cassandra

The Facebook logo, consisting of the word "facebook" in a white, lowercase, sans-serif font, set against a blue rectangular background with a white border.

Things to consider...

- Nested data structures
- Document-model
- BLOBS
- Natural partitions
- Client access patterns
- Eventual Consistency

NoSQL in Practice

- Choose the right storage system for your data
- De-normalize your data
- No ACID guarantees
- Do JOINS in your application code
- Less well suited for
 - ▶ Highly-transactional systems
 - ▶ Traditional BI systems
 - ▶ Problems that require SQL

Example: SNPs in SimpleDB

```
~/Projects/AWS-Training/SimpleDB-Demo-1 % ruby find_snps.rb
New Aws::SdbInterface using per_request-connection mode
Opening new HTTPS connection to sdb.amazonaws.com:443
Closing HTTPS connection to sdb.amazonaws.com:443
RightAws::ActiveSdb::Base
{"name"=>["rs2228714"], "seq5"=>["ATGAAACGAGTAGCGGCAGCAATAGTTCTCGAGAGGATATATCAAGTAGA"], "id"=>"2228714", "seq3"=>["GTTAGTAGTAGTATCGTTATACTA
ATTATGTTTTGGATCCCCAACTTA"], "observed"=>["A/G"]}
{"name"=>["rs2228728"], "seq5"=>["TAGTCTCAAAGGTCACAGCAGCAGCAGCAACAATCAGAGCAACAATGGAG"], "id"=>"2228728", "seq3"=>["TCCCGATAGCTACCAAAATATCTT
GATGATGCGTCACGGTGATCGCATTG"], "observed"=>["A/G"]}
RightAws::ActiveSdb::Base
{"name"=>["rs2228751"], "seq5"=>["GATTTTCTTTATTAGCTTTTGAATTTCTTTTATTTTTTAATCAACT"], "id"=>"2228751", "seq3"=>["GTATATCACATTTTCGGTGACATT
CACTGATGGCAGTCCAGATCAGTTTT"], "observed"=>["A/G"]}
RightAws::ActiveSdb::Base
{"name"=>["rs2228769"], "seq5"=>["TAAATGACTAAAACATAATTAATAGCTGCAACTCCAAATGCGTGACCAA"], "id"=>"2228769", "seq3"=>["CTTGATGATCAAAAGAAAACAC
ATGAGAGAAGATAAAGATTTAAGCG"], "observed"=>["A/G"]}
{"name"=>["rs2228789"], "seq5"=>["AATACTATTAACATAATTTCAAAAATTATTAGTCCCTTTGAAGACCTCCT"], "id"=>"2228789", "seq3"=>["AATTATTTTTCCAAAACCGGAAAT
AATTTACGCCGATAACTTAAAAATC"], "observed"=>["A/G"]}
RightAws::ActiveSdb::Base
{"name"=>["rs2228805"], "seq5"=>["AGAAAAGAGAGAGAGTCAAAGCAGTAATGAATTACCTGTGTATGAAATGT"], "id"=>"2228805", "seq3"=>["CCTGTGACATTTTATGAGGATATG
TGGCTATTATCCAGCGGTACCACGAA"], "observed"=>["A/G"]}
RightAws::ActiveSdb::Base
{"name"=>["rs2228811"], "seq5"=>["AATCTTGTTGTTTGTCAATTTGTCTCGATATGTTTTTTTATGGAGATT"], "id"=>"2228811", "seq3"=>["TCAACTTCAGAGATGGGAAAAGCT
AGCCATTAGTTTCCGATTACGCTGA"], "observed"=>["A/G"]}
{"name"=>["rs2228822"], "seq5"=>["AGCTAGACCTTTGAGCACGGCGAGCTTGGTGAATTTGGCGATCAGTCATA"], "id"=>"2228822", "seq3"=>["TGCAACCTCTTCGTTATCAATAGA
GGATCTTCTAGCCCCGAGATGAACAT"], "observed"=>["A/G"]}
{"name"=>["rs2228823"], "seq5"=>["AGCTGTCCATCGGAATATGTCGAGAGCTAGATCTGGATCGGATTGAGCTC"], "id"=>"2228823", "seq3"=>["TCGGAGTTTGTGAATCCTGGTTTT
AGATTCTGAATCCATGTTTCGAATTG"], "observed"=>["A/G"]}
{"name"=>["rs2228824"], "seq5"=>["CTGGTTTATGATTCTGAATCCATGTTTCGAATTGAGTTTCGAGTGGTGTT"], "id"=>"2228824", "seq3"=>["AGGGGAAGTGTGTGATTCCGGTGA
AGATGAGACGGAAGAGAATGATAGAGA"], "observed"=>["A/G"]}
{"name"=>["rs2228862"], "seq5"=>["GTAAAGTGTTTGGAGGTATACCTACTTTCCGTTAGTGAATGAAAATAAA"], "id"=>"2228862", "seq3"=>["ATTACTGGTTCCTAAACACCTA
ATCCATGCACCAACAACCAAGTGAAT"], "observed"=>["A/G"]}
{"name"=>["rs2228870"], "seq5"=>["AAAACTCCAGGAAATCCCGCAATTTCTGTAGACTTGATGAACCTCTCTGT"], "id"=>"2228870", "seq3"=>["GTTGGTCGAGTAAATGATGATT
AAAAGTTTTGGGTTTTAAACATCAC"], "observed"=>["A/G"]}
{"name"=>["rs2228881"], "seq5"=>["GGGAGTTGGAATCAGTTCGAGCCGATGAAGTTAAGTACCCAGAGAAGGCT"], "id"=>"2228881", "seq3"=>["CCAACCTGACCTGGTTGAACTCAC
TTCTCTCCTGGACGACACCCCTCCTA"], "observed"=>["A/G"]}
{"name"=>["rs2228891"], "seq5"=>["ATATGAAACGAGTAGCGGCAGCAATAGTTCTCGAGAGGATATATCAAGTAGA"], "id"=>"2228891", "seq3"=>["GTTAGTAGTAGTATCGTTATAC
TAATTATGTTTTGGATCCCCAACTTA"], "observed"=>["A/G"]}
~/Projects/AWS-Training/SimpleDB-Demo-1 %
```

Example: Storing short-reads

- 2.8 million unique reads loaded
- 0.5 million reads retrieved
- Single node system

Database	Load time	Retrieval time	File size
Tokyo Cabinet/Tyrant	12 minutes	3 1/2 minutes	24MB
CouchDB	5 1/2 minutes	14 1/2 minutes	236MB
MongoDB	3 minutes	4 minutes	192-960MB

NoSQL use cases

- Real-time analytics
 - ▶ Fast real-time inserts, updates, and queries
- Problems requiring high scalability
 - ▶ Tens or hundreds of servers
 - ▶ Replication/sharding built-in
- Persistent object store
 - ▶ Think persistent memcached
- Document or key-value oriented schemas
 - ▶ JSON-like data schemas

That's it

- Thanks!
- kraut@bioteam.net
- www.twitter.com/adamkraut
- www.friendfeed.com/adamk
- blog.bleedingedgebiotech.com