

NoSQL AND BIG DATA FOR THE ORACLE DBA

Iggy Fernandez

[NoCOUG Journal](#) Editor

<http://iggyfernandez.wordpress.com>

The views expressed here are my own and not necessarily those of Oracle and its affiliates



- Celebrated 25th anniversary in November 2011
- Quarterly full-day conferences around the San Francisco Bay Area. Keynote address plus three technical tracks.
- One-year membership costs \$95. First conference is free.
- NoCOUG members get printed copies of the NoCOUG Journal. Also available online for free download.
- 100th edition of the NoCOUG Journal contains a freshly typeset copy of Dr. Codd's first paper on relational theory.
http://www.nocoug.org/Journal/NoCOUG_Journal_201111.pdf
- Latest edition features interview with C. J. Date and Hugh Darwen "No! to SQL!, No! to NoSQL!"
http://www.nocoug.org/Journal/NoCOUG_Journal_Latest.pdf.



- The origins of NoSQL
- Performance, Scalability, and Availability without NoSQL
- The false premise of NoSQL
- The NoSQL product landscape
- Learning resources
- *What* makes Relational so sacred?
- The mistakes of the relational camp
- Bonus slides
 - NewSQL
 - NoSQL Buyer's Guide



- Installation and operation of NoSQL products



The Origins of NoSQL

NoSQL AND BIG DATA FOR THE ORACLE DBA



1. Extreme availability
2. Extreme performance
3. Extreme scalability



- “Customers should be able to view and add items to their shopping cart even if disks are failing, network routes are flapping, or data centers are being destroyed by tornados. Therefore, the service responsible for managing shopping carts requires that it can always write to and read from its data store, and that **its data needs to be available across multiple data centers.**” — [Dynamo: Amazon’s Highly Available Key-value Store](#)
- “Experience at Amazon has shown that **data stores that provide ACID guarantees tend to have poor availability.** ... Dynamo targets applications that operate with weaker consistency (the “C” in ACID) if this results in high availability.” — [Dynamo: Amazon’s Highly Available Key-value Store](#)



- “There are many services on Amazon’s platform that only need primary-key access to a data store. For many services, such as those that provide best seller lists, shopping carts, customer preferences, session management, sales rank, and product catalog, **the common pattern of using a relational database would lead to inefficiencies and limit scale and availability.** Dynamo provides a simple primary-key only interface to meet the requirements of these applications.” — [Dynamo: Amazon’s Highly Available Key-value Store](#)



- “Since each service uses its distinct instance of Dynamo, its initial design targets a scale of up to hundreds of storage hosts [only].”
— [Dynamo: Amazon’s Highly Available Key-value Store](#)



Dynamo Solution— Functional Segmentation

- Best seller lists, shopping carts, customer preferences, session management, sales rank, and product catalog
- Increases overall site availability by avoiding a single point of failure
- No distributed transactions
- Eventual consistency



- employee (**employee#**, name, birthdate)
- jobhistory (**employee#**, jobdate, title)
- salaryhistory (**employee#**, jobdate, salarydate, salary)
- children (**employee#**, childname, birthyear)



- Asynchronous replication
- Eventual consistency



- “Shopping carts are stored as binary objects (i.e., blobs) identified by unique keys. No operations span multiple data items and there is no need for relational schema.” — [Dynamo: Amazon’s Highly Available Key-value Store](#)
- “Using tables to store objects is like driving your car home and then disassembling it to put it in the garage. It can be assembled again in the morning, but **one eventually asks whether this is the most efficient way to park a car**” — attributed to Esther Dyson, Editor of the Release 1.0 newsletter.



- Functional segmentation
- Sharding
- No distributed transactions
- Asynchronous replication
- Eventual consistency
- BLOBs
- No SQL
- Primary-key access
- Autocommit



- Functional segmentation
 - Sharding
 - No distributed transactions
 - Asynchronous replication
 - Eventual consistency
 - BLOBs
 - No SQL
 - Primary-key access
 - Autocommit
- Functional segmentation
 - Sharding
 - No distributed transactions
 - Asynchronous replication
 - Eventual consistency
 - SQL
 - Oracle
 - Local transactions with ACID
 - Middle-tier caching
 - Middle-tier constraint checking

The False Premise of NoSQL

NoSQL AND BIG DATA FOR THE ORACLE DBA



- “Nonatomic values can be discussed within the relational framework. Thus, some domains may have relations as elements. These relations may, in turn, be defined on nonsimple domains, and so on. For example, one of the domains on which the relation employee is defined might be salary history.”—Codd, E. F. [A relational model of data for large shared data banks](#). (1970).
- **employee** (
 employee#,
 name,
 birthdate,
 jobhistory (jobdate, title, **salaryhistory** (salarydate, salary)),
 children (childname, birthyear)
)



- employee (**employee#**, name, birthdate)
- jobhistory (**employee#**, **jobdate**, title)
- salaryhistory (**employee#**, **jobdate**, **salarydate**, salary)
- children (**employee#**, **childname**, birthyear)



- Data from multiple tables stored in the same block
- Hash clusters and indexed clusters
- Clustered tables can be indexed
- No join penalty
- Can define “object-relational” views and INSTEAD OF triggers

<http://iggyfernandez.wordpress.com/2013/07/28/no-to-sql-and-no-to-nosql/>



Oracle Table Clusters Example

```
INSERT INTO employees VALUES (1, 'IGNATIUS', '01-JAN-1970');
```

```
INSERT INTO children VALUES (1, 'INIGA', '01-JAN-2001');
```

```
INSERT INTO children VALUES (1, 'INIGO', '01-JAN-2002');
```

```
INSERT INTO job_history VALUES (1, '01-JAN-1991', 'PROGRAMMER');
```

```
INSERT INTO job_history VALUES (1, '01-JAN-1992', 'DATABASE ADMIN');
```

```
INSERT INTO salary_history VALUES (1, '01-JAN-1991', '1-FEB-1991', 1000);
```

```
INSERT INTO salary_history VALUES (1, '01-JAN-1991', '1-MAR-1991', 1000);
```

```
INSERT INTO salary_history VALUES (1, '01-JAN-1992', '1-FEB-1992', 2000);
```

```
INSERT INTO salary_history VALUES (1, '01-JAN-1992', '1-MAR-1992', 2000);
```

```

CREATE OR REPLACE VIEW employees_view AS
SELECT
  employee#,
  name,
  birth_date,
  CAST
  (
    MULTISET
    (
      SELECT child_name, birth_date
      FROM children
      WHERE employee#=e.employee#
    )
    AS children_tab
  ) children,
  CAST
  (
    MULTISET
    (
      SELECT

```

```

  job_date,
  title,
  CAST
  (
    MULTISET
    (
      SELECT salary_date, salary
      FROM salary_history
      WHERE employee#=e.employee#
      AND job_date=jh.job_date
    )
    AS salary_history_tab
  ) salary_history
FROM job_history jh
WHERE employee#=e.employee#
)
AS job_history_tab
) job_history
FROM employees e;

```

```
SELECT * FROM employees_view WHERE employee# = 1;
```

```

EMPLOYEE# NAME          BIRTH_DAT
-----
CHILDREN(CHILD_NAME, BIRTH_DATE)
-----
JOB_HISTORY(JOB_DATE, TITLE, SALARY_HISTORY(SALARY_DATE, SALARY))
-----
1 IGNATIUS          01-JAN-70
CHILDREN_TAB(CHILDREN_REC('INIGA', '01-JAN-01'), CHILDREN_REC('INIGO', '01-JAN-02'))
JOB_HISTORY_TAB(JOB_HISTORY_REC('01-JAN-91', 'PROGRAMMER', SALARY_HISTORY_TAB(SALARY_HISTORY_REC('01-FEB-91', 1000), SALARY_HISTORY_REC('01-MAR-91', 1000))), JOB_HISTORY_REC('01-JAN-92', 'DATABASE ADMIN', SALARY_HISTORY_TAB(SALARY_HISTORY_REC('01-FEB-92', 2000), SALARY_HISTORY_REC('01-MAR-92', 2000))))

```

Plan hash value: 2117652374

Id	Operation	Name	Starts	E-Rows	E-Bytes	Cost	A-Rows	A-Time	Buffers
0	SELECT STATEMENT		1			1	1	00:00:00.01	1
* 1	TABLE ACCESS HASH	CHILDREN	1	1	32		2	00:00:00.01	1
* 2	TABLE ACCESS HASH	SALARY_HISTORY	2	1	44		4	00:00:00.01	3
* 3	TABLE ACCESS HASH	JOB_HISTORY	1	1	32		2	00:00:00.01	1
* 4	TABLE ACCESS HASH	EMPLOYEES	1	845	27040		1	00:00:00.01	1



```

INSERT INTO employees_view
VALUES
(
  2,
  'YGNACIO',
  '01-JAN-70',
  CHILDREN_TAB(CHILDREN_REC('INIGA', '01-JAN-01'), CHILDREN_REC('INIGO', '01-JAN-02')),
  JOB_HISTORY_TAB
  (
    JOB_HISTORY_REC
    (
      '01-JAN-91',
      'PROGRAMMER',
      SALARY_HISTORY_TAB(SALARY_HISTORY_REC('01-FEB-91', 1000), SALARY_HISTORY_REC('01-MAR-91', 1000))
    ),
    JOB_HISTORY_REC
    (
      '01-JAN-92',
      'DATABASE ADMIN',
      SALARY_HISTORY_TAB(SALARY_HISTORY_REC('01-FEB-92', 2000), SALARY_HISTORY_REC('01-MAR-92', 2000))
    )
  )
);


```

```
SQL> SELECT * FROM employees_view WHERE employee# = 2;
```

```

EMPLOYEE# NAME          BIRTH_DAT
-----
CHILDREN(CHILD_NAME, BIRTH_DATE)
-----
JOB_HISTORY(JOB_DATE, TITLE, SALARY_HISTORY(SALARY_DATE, SALARY))
-----
      2 YGNACIO          01-JAN-70
CHILDREN_TAB(CHILDREN_REC('INIGA', '01-JAN-01'), CHILDREN_REC('INIGO', '01-JAN-02'))
JOB_HISTORY_TAB(JOB_HISTORY_REC('01-JAN-91', 'PROGRAMMER', SALARY_HISTORY_TAB(SALARY_HISTORY_REC('01-FEB-91', 1000), SALARY_HISTORY_
REC('01-MAR-91', 1000))), JOB_HISTORY_REC('01-JAN-92', 'DATABASE ADMIN', SALARY_HISTORY_TAB(SALARY_HISTORY_REC('01-FEB-92', 2000), S
ALARY_HISTORY_REC('01-MAR-92', 2000))))

```



“There are, of course, several possible ways in which a system can detect inconsistencies and respond to them. In one approach the system checks for possible inconsistency whenever an insertion, deletion, or key update occurs. **Naturally, such checking will slow these operations down.** If an inconsistency has been generated, details are logged internally, and if it is not remedied within some reasonable time interval, either the user or someone responsible for the security and integrity of the data is notified. Another approach is to conduct consistency checking as a batch operation once a day or less frequently.” —Codd, E. F. [A relational model of data for large shared data banks.](#) (1970).



The NoSQL Landscape

NoSQL AND BIG DATA FOR THE ORACLE DBA



- Key-Value Databases—Dynamo, Riak, Oracle NoSQL
- Document Databases—Mongodb
- Column Family Databases—BigTable, HBase, Cassandra
- Graph Databases—Neo4J
- Big Data—Hadoop

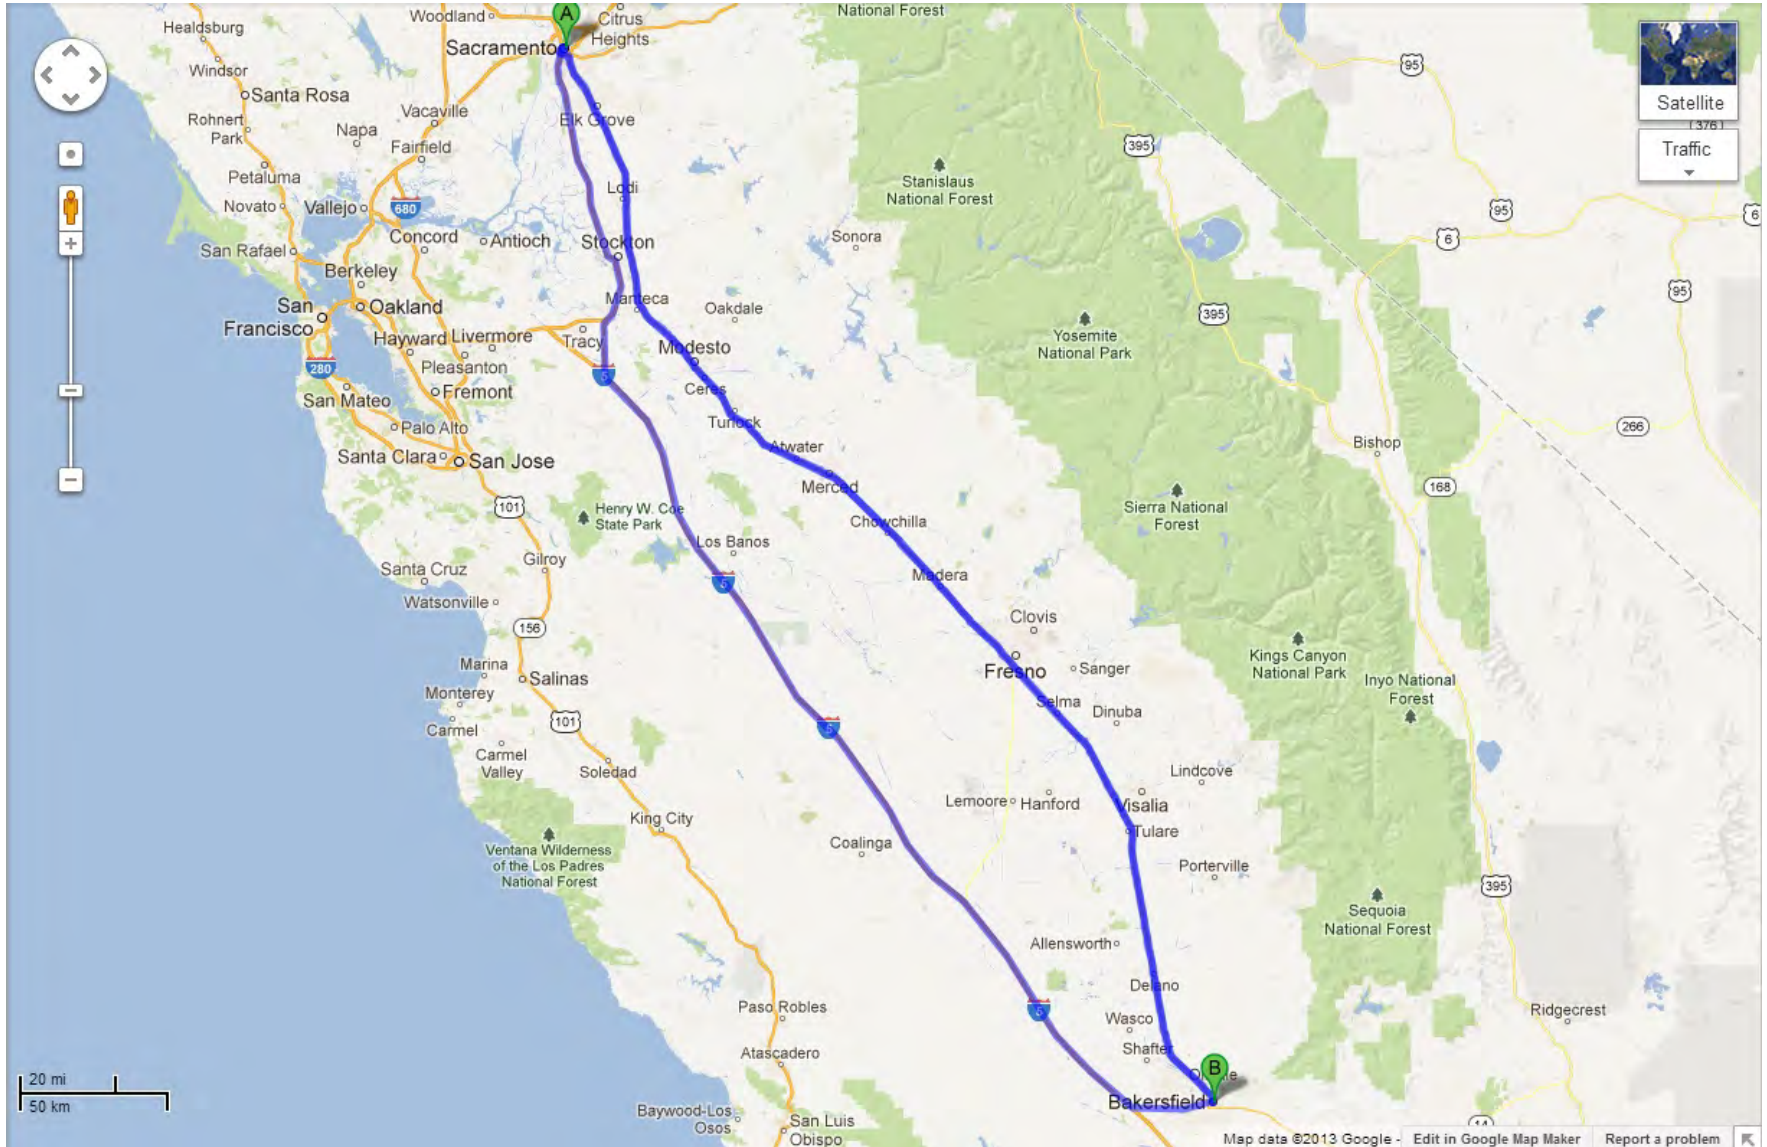


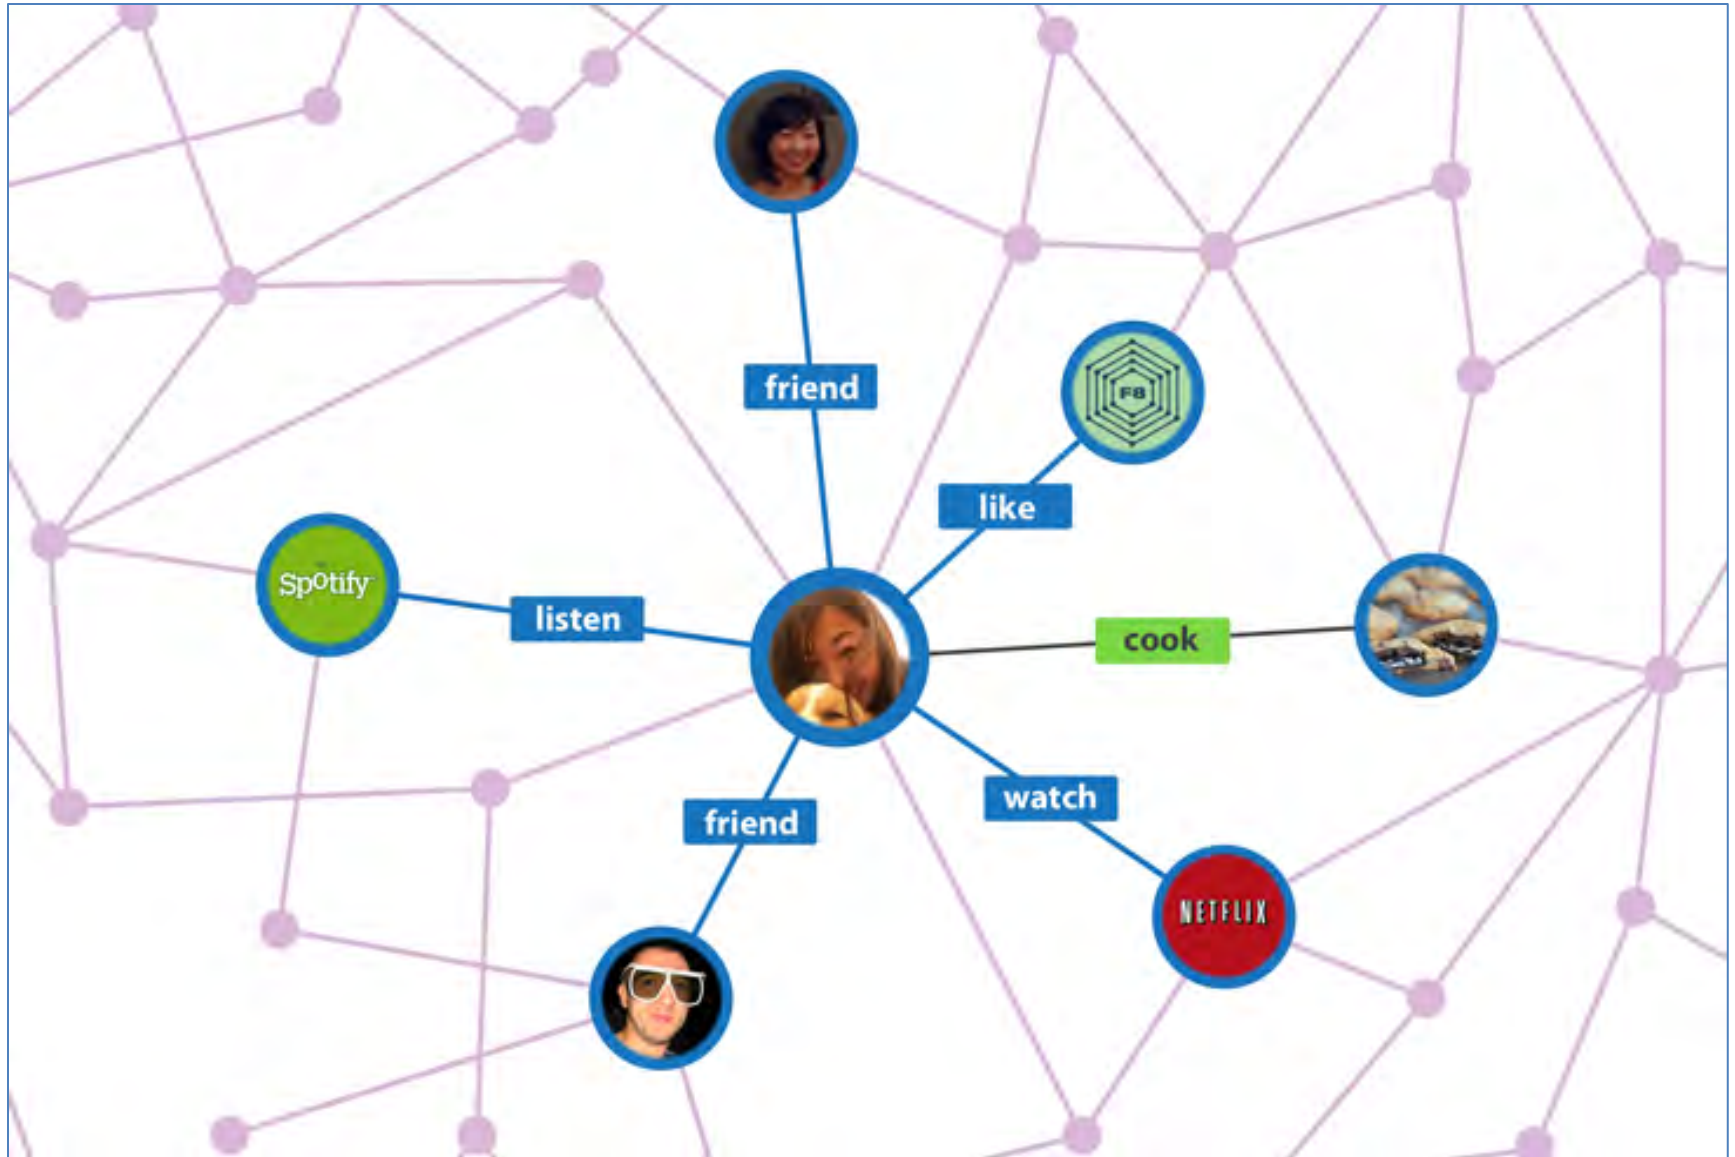
```
{
  "_id": ObjectId("4efa8d2b7d284dad101e4bc9"),
  "Last Name": "DUMONT",
  "First Name": "Jean",
  "Date of Birth": "01-22-1963"
}

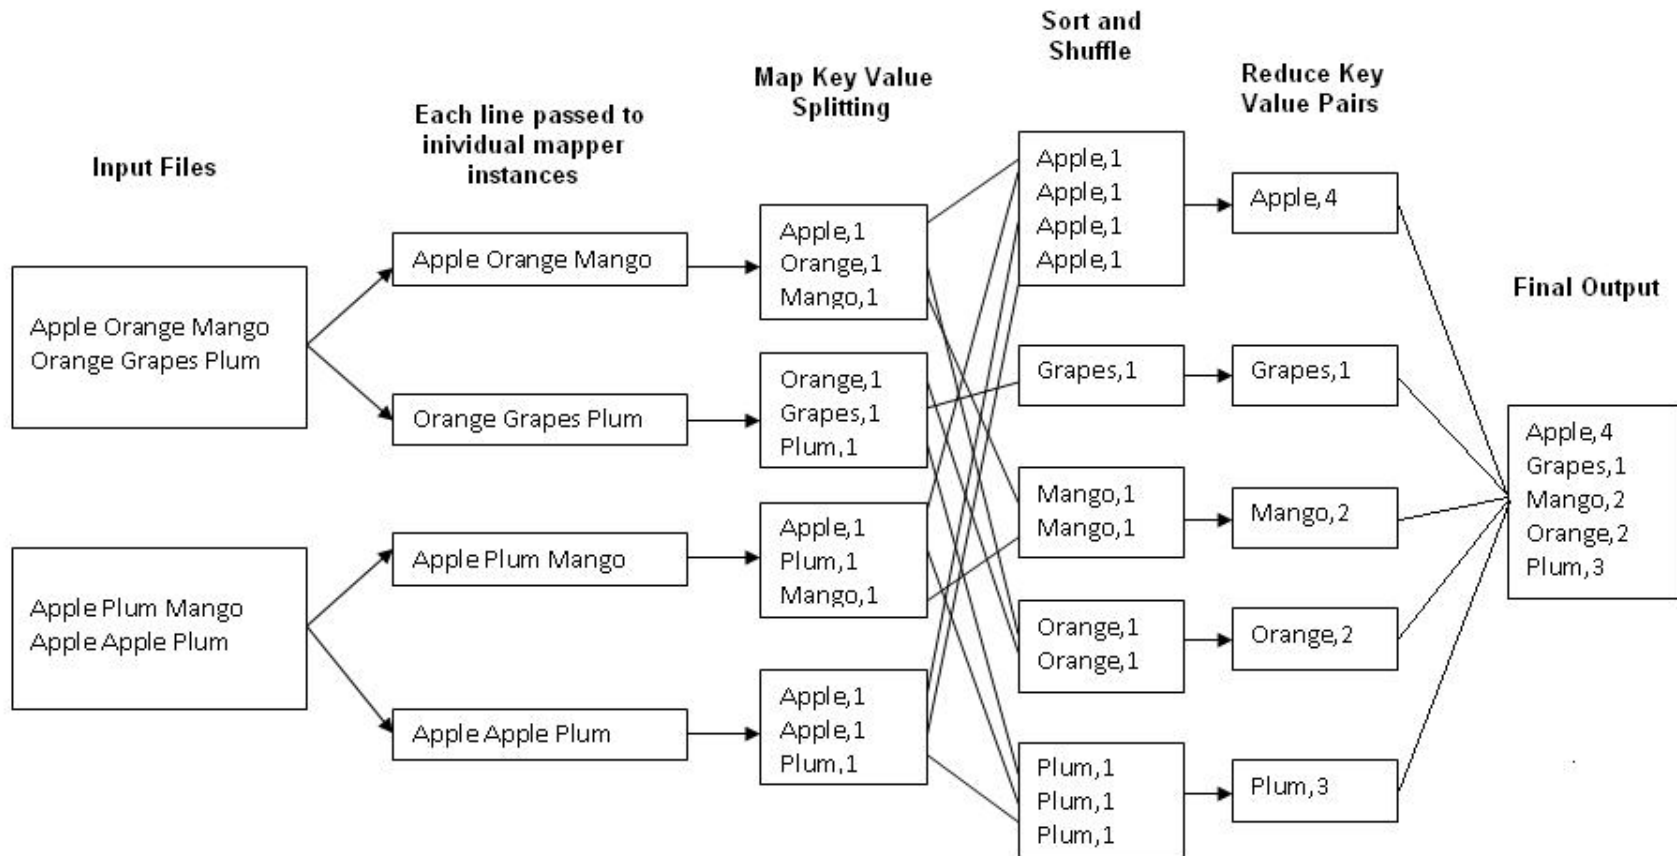
{
  "_id": ObjectId("4efa8d2b7d284dad101e4bc7"),
  "Last Name": "PELLERIN",
  "First Name": "Franck",
  "Date of Birth": "09-19-1983",
  "Address": "1 chemin des Loges",
  "City": "VERSAILLES"
}
```

```
{
  "_id": ObjectId("4efa8d2b7d284dad101e4bc7"),
  "Last Name": "PELLERIN",
  "First Name": "Franck",
  "Date of Birth": "09-19-1983",
  "phoneNumber": [
    {
      "type": "home",
      "number": "212 555-1234"
    },
    {
      "type": "fax",
      "number": "646 555-4567",
      "verified": false
    }
  ],
  "Address": {
    "Street": "1 chemin des Loges",
    "City": "VERSAILLES"
  },
  "Months at Present Address": 37
}
```

Graph Database—Shortest Path Problem







Map

Map

```
(  
  String key,  
  String value  
):  
// key: document name  
// value: document contents  
for each word w in value:  
  EmitIntermediate(w, "1");
```

Reduce

Reduce

```
(  
  String key,  
  Iterator values  
):  
// key: a word  
// values: a list of counts  
int result = 0;  
for each v in values:  
  result += ParseInt(v);  
Emit(AsString(result));
```



- “Dirty secret of Big Data is you can not able deploy if you not have SQL expert on staff.” — DevOps Borat
- Pig, Hive, Tenzing, Impala



- [*Seven Databases in Seven Weeks*](#) by Eric Redmond and Jim R. Wilson
- [*NoSQL Distilled*](#) by Pramod Sadalage and Martin Fowler
- [KVLite](#) for Windows (single-process version of Oracle NoSQL Database)
- Oracle NoSQL hands-on [workshop](#) by Anuj Sahni
- Oracle NoSQL [presentation](#) by Anuj Sahni
- [Cloudera QuickStart VM](#) (available in VMware, VirtualBox and KVM flavors)
- [Benchmarking Cloud Serving Systems With YCSB](#) by B. F. Cooper, A. Silberstein, E. Tam, R. Ramakrishnan, and R. Sears



What makes relational so sacred?

NoSQL AND BIG DATA FOR THE ORACLE DBA



“Future users of large data banks must be protected from having to know how the data is organized in the machine (the internal representation).” CODD, E. F. [A relational model of data for large shared data banks](#). (1970).



“Surely, in the choice of logical data structures that a system is to support, **there is one consideration of absolutely paramount importance - and that is the convenience of the majority of users.** [emphasis added] ... To make formatted data bases readily accessible to users (especially casual users) who have little or no training in programming we must provide the simplest possible data structures and almost natural language. ... What could be a simpler, more universally needed, and more universally understood data structure than a table? Why not permit such users to view all the data in a data base in a tabular way?” —Codd, E. F. Normalized Data Base Structure: A Brief Tutorial. (1971).



“There is a large class of users who, while they are not computer specialists, would be willing to learn to interact with a computer in a reasonably high-level, non-procedural query language. Examples of such users are **accountants, engineers, architects, and urban planners**. [emphasis added] It is for *this* [emphasis added] class of users that SEQUEL is intended.” —CHAMBERLIN, D. AND BOYCE, R. [SEQUEL: A Structured English Query Language](#) (1974).



“Codd gave a seminar and a lot of us went to listen to him. This was as I say a revelation for me because Codd had a bunch of queries that were fairly complicated queries and since I’d been studying CODASYL, I could imagine how those queries would have been represented in CODASYL by programs that were five pages long that would navigate through this labyrinth of pointers and stuff. Codd would sort of write them down as one-liners. These would be queries like, “Find the employees who earn more than their managers.” **He just whacked them out and you could sort of read them, and they weren’t complicated at all, and I said, “Wow.”** [emphasis added] This was kind of a conversion experience for me, that I understood what the relational thing was about after that. —Donald Chamberlin, the creator of the SQL language in [The 1995 SQL Reunion: People, Projects, and Politics](#)



Derivability, Redundancy and Consistency of Relations

“The large, integrated data banks of the future will contain many relations of various degrees in stored form. It will not be unusual for this set of stored relations to be redundant. Two types of redundancy are defined and discussed. One type may be employed to improve accessibility of certain kinds of information which happen to be in great demand. When either type of redundancy exists, those responsible for control of the data bank should know about it and have some means of detecting any “logical” inconsistencies in the total set of stored relations. Consistency checking might be helpful in tracking down unauthorized (and possibly fraudulent) changes in the data bank contents.” —Codd, E. F. Derivability, Redundancy and Consistency of Relations Stored in Large Data Banks. (1969).



You are allowed to teach a certain course only if:

1. You have been employed for at least one year, or

2. You have attended that course first and the trainer of that course offering attends your first teach as participant

—Example from [Applied Mathematics for Database Professionals](#) by Lex de Haan and Toon Koppelaars




```
CREATE ASSERTION employees_a1  
AS CHECK  
(  
  (SELECT COUNT(*) FROM employees) >= 50  
)
```



```
CREATE ASSERTION employees_departments_fk
AS CHECK
(
  NOT EXISTS
  (
    SELECT * FROM employees e
    WHERE NOT EXISTS
    (
      SELECT * FROM departments d
      WHERE d.department_id = e.department_id
    )
  )
)
```



The relational model is sacred because it gives application software developers the ability to assert and enforce consistency of data in databases.



The mistakes of the relational camp

NoSQL AND BIG DATA FOR THE ORACLE DBA



Mistake #1—De-emphasizing physical database design

“Future users of large data banks must be protected from having to know how the data is organized in the machine (the internal representation).” —Codd, E. F. [A relational model of data for large shared data banks](#). (1970).



- DBAs bear chief responsibility for the performance of SQL statements.
- Applications should be designed without reference to the way data is stored, e.g., index organized tables, hash clusters, partitions, etc.
- Application programmers should not tailor their SQL statements to make use of existing indexes. DBAs should instead create traps to catch badly performing SQL at runtime and create new indexes as necessary to make them perform better.
- It is not necessary to review the Query Execution Plan of an SQL statement before releasing it into a production environment. It is further not necessary to freeze the Query Execution Plan of an SQL statement before releasing it into a production environment. It is desirable that Query Execution Plans change in response to changes in the statistical information that the query optimizer relies upon. Such changes are always for the better.
- The most common cause of poorly performing SQL is the failure of the DBA to collect statistical information on the distribution of data for the use of the query optimizer. This statistical information should be refreshed frequently.



“There is a large class of users who, while they are not computer specialists, would be willing to learn to interact with a computer in a reasonably high-level, non-procedural query language. Examples of such users are **accountants, engineers, architects, and urban planners**. It is for this class of users that SEQUEL is intended.” CHAMBERLIN, D. AND BOYCE, R. [SEQUEL: A Structured English Query Language](#) (1974).



Normalized set

- **employee** (
 employee#,
 name,
 birthdate,
 jobhistory (
 jobdate,
 title,
 salaryhistory(salarydate, salary)
)
 children (childname, birthyear)
)

Normalized set

- **employee** (**man#**, **name**,
 birthdate, jobhistory,
 children)
- **jobhistory** (**man#**, **jobdate**,
 title, salaryhistory)
- **salaryhistory** (**man#**,
 jobdate, **salarydate**, salary)
- **children** (**man#**, **childname**,
 birthyear)



Dr. Codd's Justifications for First Normal Form

- “A relation whose domains are all simple can be represented in storage by a two-dimensional column-homogeneous array ... Some more complicated data structure is necessary for a relation with one or more nonsimple domains.”
- “The simplicity of the array representation which becomes feasible when all relations are cast in normal form is not only an advantage for storage purposes but also for communication of bulk data between systems which use widely different representations of the data.”
- “The second-order predicate calculus (rather than first-order) is needed because the domains on which relations are defined may themselves have relations as elements.”



“Using [flat] tables to store objects is like driving your car home and then disassembling it to put it in the garage. It can be assembled again in the morning, but one eventually asks whether this is the most efficient way to park a car.” —DYSON, E. *Review 1.0*. (September 1988).



- ... each stored table should occupy one physical file
- ... data should be stored in row-major order
- ... stored tables have only one storage representation each
- ... data should be stored in normalized form only
- ... a single data block only contain data from a single table
- ... data should not be *stored* in compact forms



“Clearly, the majority of users should not have to learn either the relational calculus or algebra in order to interact with data bases. However, requesting data by its properties is far more natural than devising a particular algorithm or sequence of operations for its retrieval. Thus, a calculus-oriented language provides a good target language for a more user-oriented source language.” —Codd, E. F. Relational Completeness of Data Base Sublanguages. (1972)



Relational Calculus

```
SELECT
  first_name,
  last_name
FROM employees mgr
WHERE EXISTS (
  SELECT *
  FROM employees emp
  WHERE emp.manager_id =
mgr.employee_id
  AND emp.salary > mgr.salary
)
```

Relational Algebra

SELECT DISTINCT

```
  mgr.first_name,
  mgr.last_name
FROM employees mgr
INNER JOIN employees emp
ON emp.manager_id =
mgr.employee_id
WHERE emp.salary >
mgr.salary
```



“I became interested in the CBO’s selectivity calculations trying to understand why it comes up with some of the ridiculously low cardinality estimates (like 1 when in reality there are 80,000+) which then lead to disastrous access plans that take hours, provided they finish at all, instead of minutes or seconds.” —Wolfgang Breitling, author of Tuning by Cardinality Feedback

(http://asktom.oracle.com/pls/apex/f?p=100:11:0:::~:P11_QUESTION_ID:4344365159075#6261591022323)



“Normalization is a step-by-step reversible process of replacing a given collection of relations by successive collections in which the relations have a progressively simpler and more regular structure. The objectives of normalization are ... To free the collection of relations from undesirable insertion, update and deletion dependencies” —Codd, E. F. Normalized Data Base Structure: A Brief Tutorial. (1971).



- Unavailable in Oracle Database, SQL Server, DB2, MySQL, and PostgreSQL
- Limited support in [Oracle Rdb](#)
 - The predicate in a CHECK table constraint can refer directly to any column in the table and can refer to columns in other tables in the database through column select expressions in the predicate.



“We don't use databases. We don't use indexes. We store all our data in compressed text files. Each compressed text file contains one year of data for one location. There is a separate subdirectory for each year. We have a terabyte of data going back to 1901 so we currently have 113 subdirectories. The performance is just fine, thank you.”

<http://iggyfernandez.wordpress.com/2013/01/22/we-dont-use-databases-we-dont-use-indexes/>



- Limited subset of SQL
- SQL procedures compiled and linked into the execution engine
- Sharding
- Replication
- Only one transaction operates in a shard at a time
- No locking, no latching, no concurrency, no redo



Bonus Slide— NoSQL Buyer's Guide

- The first factor is the buyer's performance requirements ... **Only if the performance requirements are extremely severe should buyers rule out present relational DBMS products on this basis.** [emphasis added] Even then buyers should design performance tests of their own, rather than rely on vendor-designed tests or vendor-declared strategies.
- The second factor is reduced costs for developing new databases and new application programs ...
- The third factor is protecting future investments in application programs by acquiring a DBMS with a solid theoretical foundation

CODD, E. F. An Evaluation Scheme for Database Management Systems that are claimed to be Relational. (1985)



- Amazon requirements: Extreme performance, extreme scalability, and extreme availability
- Amazon solution: Functional segmentation, Sharding, Multi-master replication, BLOBs
- eBay has the same goals as Amazon but uses Oracle and SQL for its e-commerce platform



- Data from multiple tables stored in the same block
- Hash clusters and indexed clusters
- Clustered tables can be indexed
- No join penalty
- Can define “object-relational” views and INSTEAD OF triggers

- De-emphasizing physical database design
- Discarding nested relations
- Favoring relational calculus over relational algebra
- Equating the normalized set with the stored set
- Marrying relational theory to ACID DBMS
- Ignoring SQL-92 CREATE ASSERTION



The relational model is sacred because it gives application software developers the ability to assert and enforce consistency of data in databases.



THANK YOU!



Q & A

