The NoSQL RDBMS

One of first uses of the phrase *NoSQL* is due to Carlo Strozzi, circa 1998.

NoSQL:

- A fast, portable, open-source RDBMS
- A derivative of the RDB database system (Walter Hobbs, RAND)
- Not a full-function DBMS, per se, but a shell-level tool
- User interface – Unix shell
- Based on the “operator/stream paradigm”

NoSQL Today

More recently:

- The term has taken on different meanings
- One common interpretation is “not only SQL”

Most modern NoSQL systems diverge from the relational model or standard RDBMS functionality:

The data model:	relations		documents
	tuples	vs.	graphs
	attributes		key/values
	domains		
	normalization		
The query model:	relational algebra		graph traversal
	tuple calculus	vs.	text search
			map/reduce
The implementation:	rigid schemas	vs.	flexible schemas
			(schema-less)
	ACID compliance	vs.	BASE

In that sense, NoSQL today is more commonly meant to be something like “non-relational”

NoSQL Today

(a partial, unrefined list)

Hbase Cassandra Hypertable Accumulo Amazon SimpleDB SciDB Stratosphere flare

Cloudata BigTable QD Technology SmartFocus KDI Alterian Cloudera C-Store

Vertica Qbase–MetaCarta OpenNeptune HPCC Mongo DB CouchDB Clusterpoint ServerTerrastore

Jackrabbit OrientDB Perservere CoudKit Djondb SchemaFreeDB SDB JasDB

RaptorDB ThruDB RavenDB DynamoDB Azure Table Storage Couchbase Server Riak

LevelDB Chordless GenieDB Scalaris Tokyo Kyoto Cabinet Tyrant Scalien

Berkeley DB Voldemort Dynomite KAI MemcacheDB Faircom C-Tree HamsterDB STSdb

Tarantool/Box Maxtable Pincaster RaptorDB TIBCO Active Spaces allegro-C nessDBHyperDex

Mnesia LightCloud Hibari BangDB OpenLDAP/MDB/Lightning Scality Redis

KaTree TomP2P Kumofs TreapDB NMDB luxio actord Keyspace

schema-free RAMCloud SubRecord Mo8onDb Dovetaildb JDBM Neo4 InfiniteGraph

Sones InfoGrid HyperGraphDB DEX GraphBase Trinity AllegroGraph BrightstarDB

Bigdata Meronymy OpenLink Virtuoso VertexDB FlockDB Execom IOG Java Univ Netwrk/Graph Framework

OpenRDF/Sesame Filament OWLim NetworkX iGraph Jena SPARQL OrientDb

ArangoDB AlchemyDB Soft NoSQL Systems Db4o Versant Objectivity Starcounter

ZODB Magma NEO PicoList siaqodb Sterling Morantex EyeDB

HSS Database FramerD Ninja Database Pro StupidDB KiokuDB Perl solution Durus

GigaSpaces Infinispan Queplix Hazelcast GridGain Galaxy SpaceBase JoafipCoherence

eXtremeScale MarkLogic Server EMC Documentum xDB eXist Sedna BaseX Qizx

Berkeley DB XML Xindice Tamino Globals Intersystems Cache GT.M EGTM

U2 OpenInsight Reality OpenQM ESENT jBASE MultiValue Lotus/Domino

eXtremeDB RDM Embedded ISIS Family Prevayler Yserial Vmware vFabric GemFire Btrieve

KirbyBase Tokutek Recutils FileDB Armadillo illuminate Correlation Database FluidDB

Sleet DB Twisted Storage Rindo Sherpa tin Dryad SkyNet Disco

MUMPS Adabas XAP In-Memory Grid eXtreme Scale MckoiDDB Mckoi SQL Database

Primary NoSQL Categories

- General Categories of NoSQL Systems:
  - Key/value store
  - (wide) Column store
  - Graph store
  - Document store

- Compared to the relational model:
  - Query models are not as developed.
  - Distinction between abstraction & implementation is not as clear.

Key/Value Store

"Dynamo: Amazon's Highly Available Key-value Store," DeCandia, G., et al., SOSP'07, 21st ACM Symposium on Operating Systems Principles.

The basic data model:

- Database is a collection of key/value pairs
- The key for each pair is unique

No requirement for normalization (and consequently dependency preservation or lossless join)

Primary operations:

- insert(key,value)
- delete(key)
- update(key,value)
- lookup(key)

Additional operations:

- variations on the above, e.g., reverse lookup
- iterators

DynamoDB
Azure Table Storage
Riak
Rdis
Aerospike
FoundationDB
LevelDB
Berkeley DB
Oracle NoSQL Database
GenieDb
BangDB
Chordless
Scalaris
Tokyo Cabinet/Tyrant
Scalien
Voldemort
Dynomite
KAI
MemcacheDB
Faircom C-Tree
LSM
KitaroDB
HamsterDB
STSdb
TarantoolBox
Maxtable
Quasardb
Pincaster
RaptorDB
TIBCO Active Spaces
Allegro-C
nessDB
HyperDex
SharedHashFile
Symas LMDB
Sophia
PickleDB
Mnesia
LightCloud
Hibari
OpenLDAP
Genomu
BinaryRage
Elliptics
Dbreeze
RocksDB
TreodeDB
(www.nosql-database.org
www.db-engines.com
www.wikipedia.com)

Wide Column Store

"Bigtable: A Distributed Storage System for Structured Data," Chang, F., et al., OSDI'06: Seventh Symposium on Operating System Design and implementation, 2006.

The basic data model:

- Database is a collection of key/value pairs
- Key consists of 3 parts – a row key, a column key, and a time-stamp (i.e., the version)
- Flexible schema - the set of columns is not fixed, and may differ from row-to-row

Warning #1!

One last column detail:

- Column key consists of two parts – a column family, and a qualifier

Accumulo
Amazon SimpleDB
BigTable
Cassandra
Cloudata
Cloudera
Druid
Flink
Hbase
Hortonworks
HPCC
Hyupertable
KAI
KDI
MapR
MonetDB
OpenNeptune
Qbase
Splice Machine
Sqrrl
(www.nosql-database.org
www.db-engines.com
www.wikipedia.com)

Wide Column Store

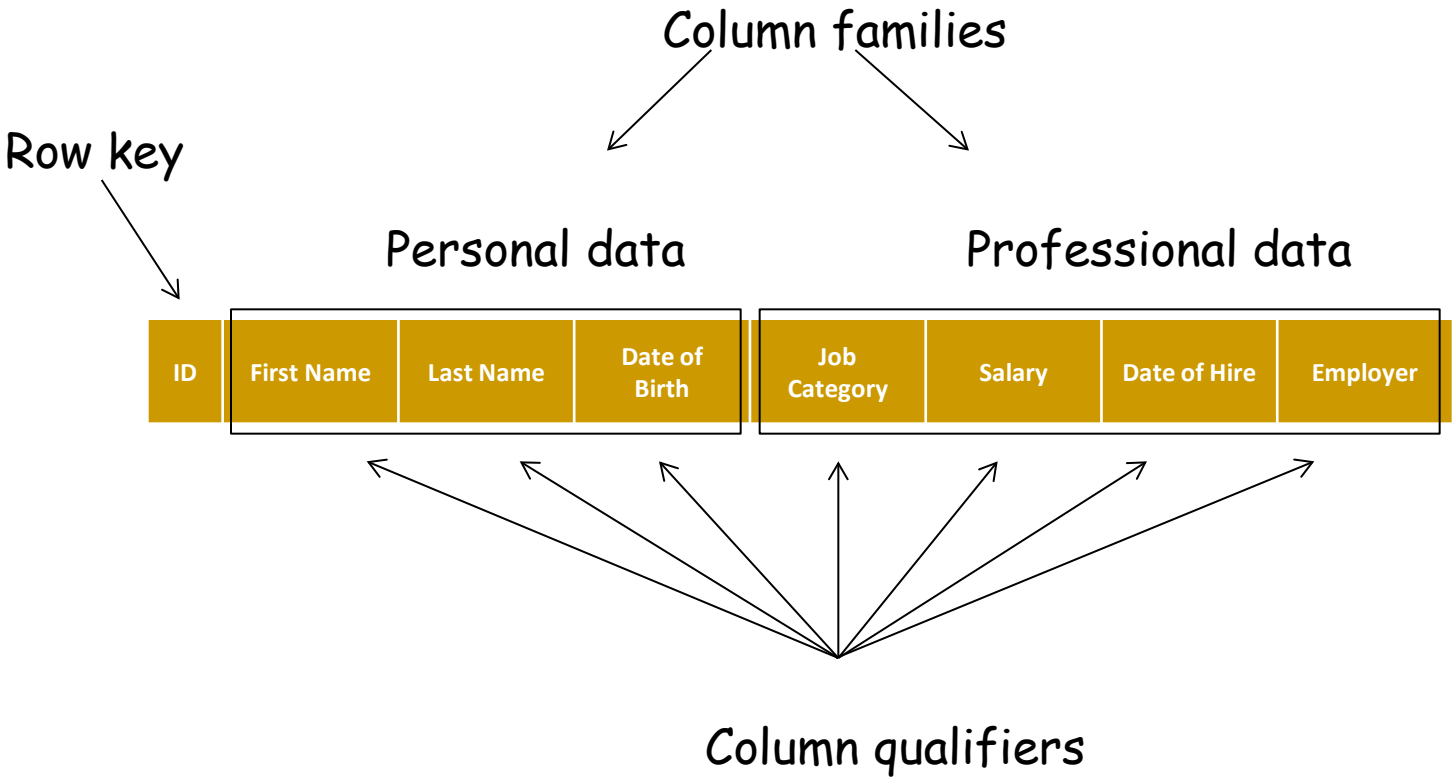

Wide Column Store

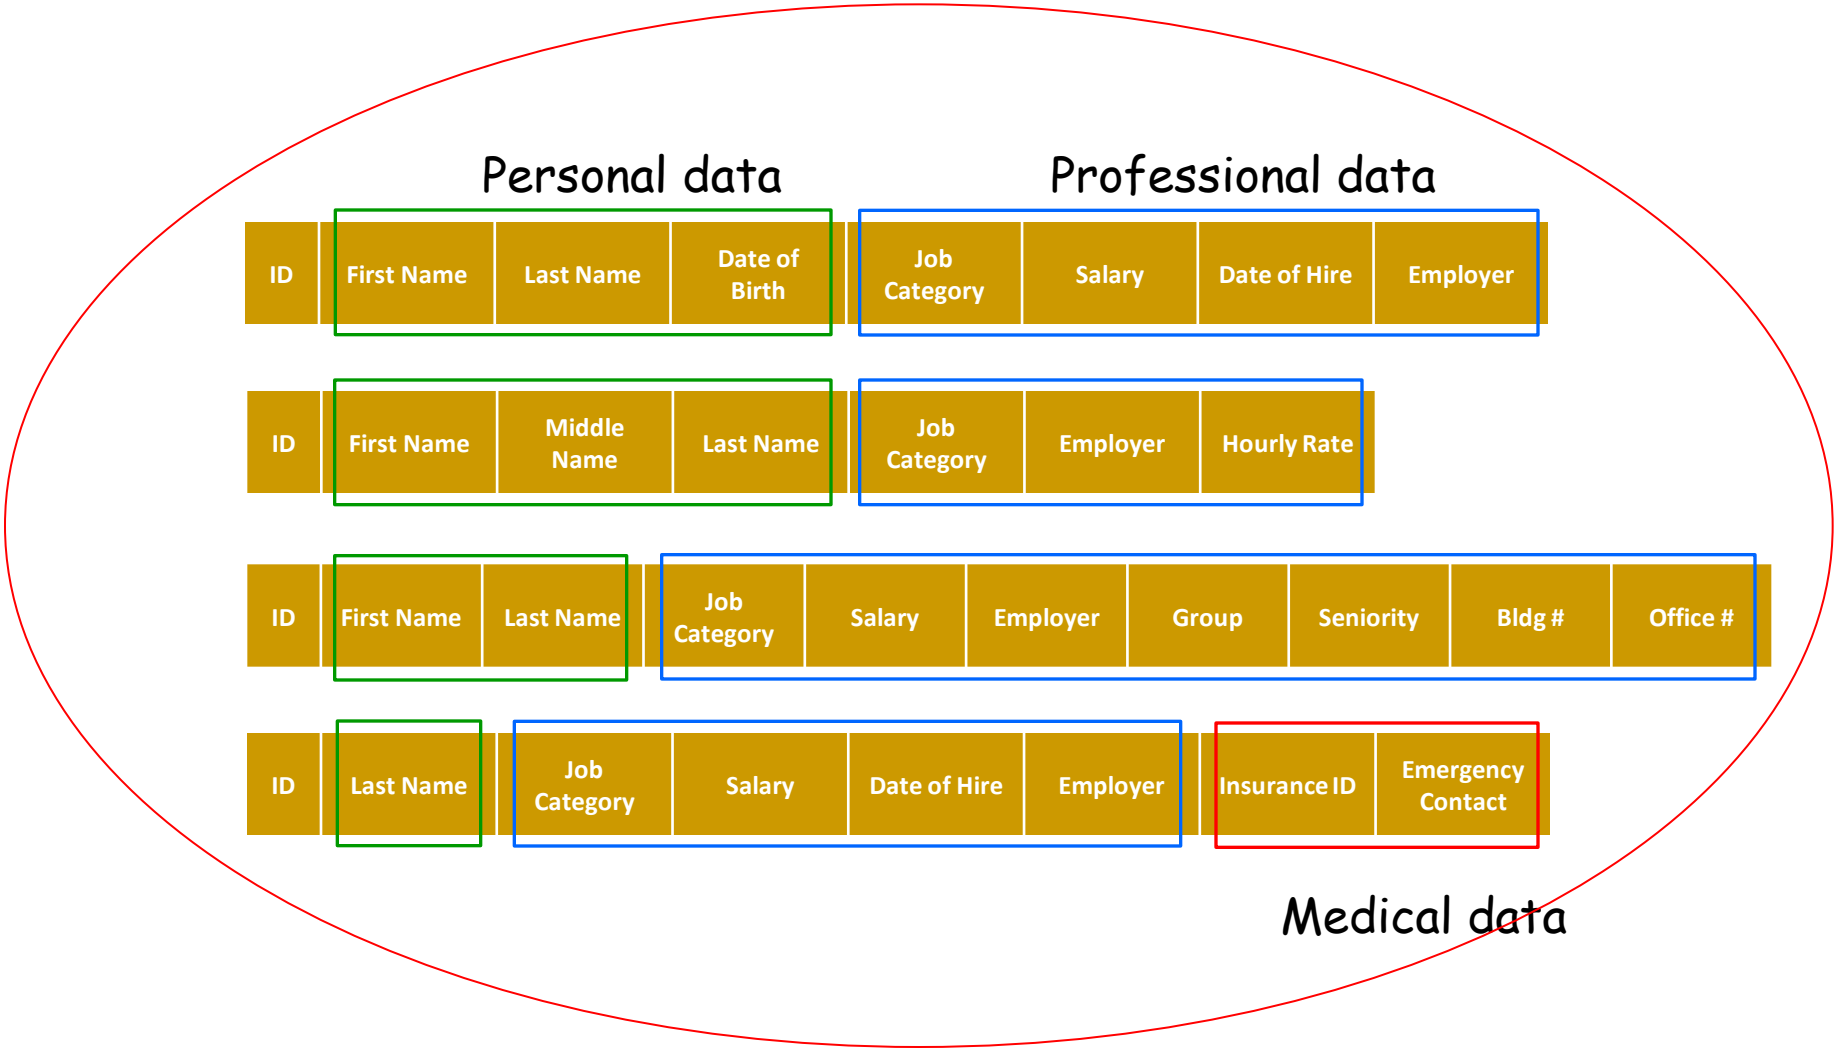

One "table"

Wide Column Store

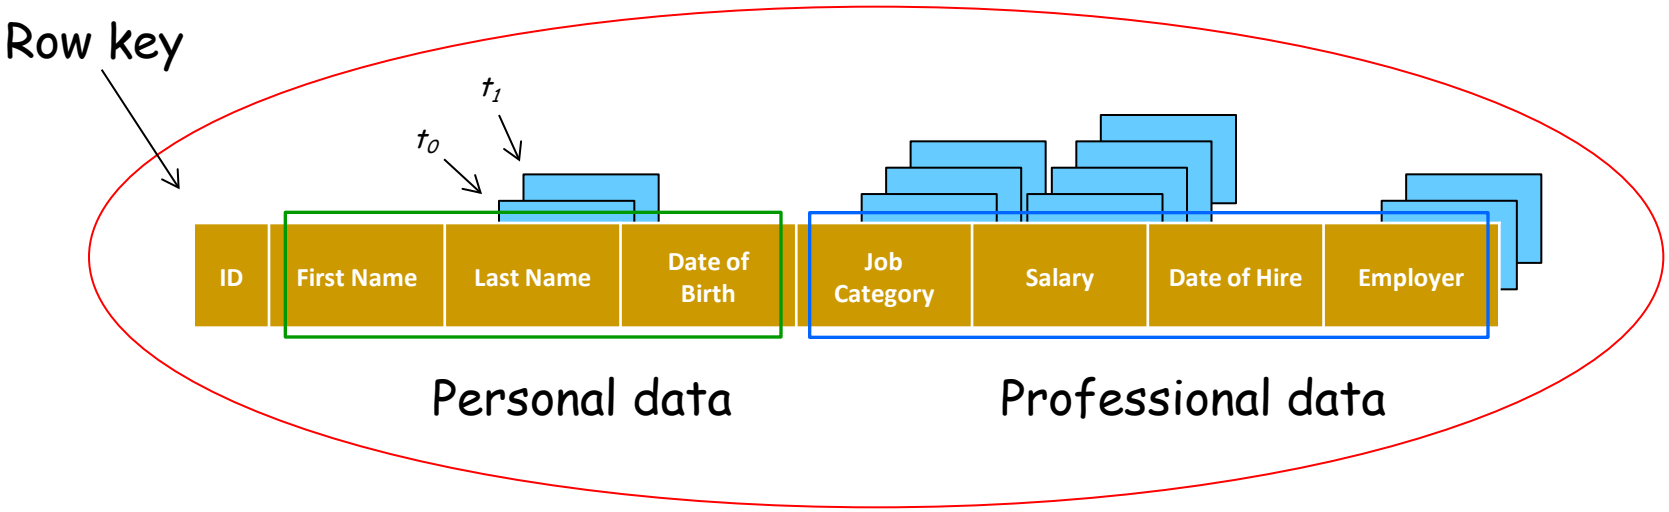

One "row"

One "row" in a wide-column NoSQL database table
=
Many rows in several relations/tables in a relational database

Graph Store

Neo4j - “The Neo Database – A Technology Introduction,” 2006.

The basic data model:

- Directed graphs
- Nodes & edges, with properties, i.e., “labels”

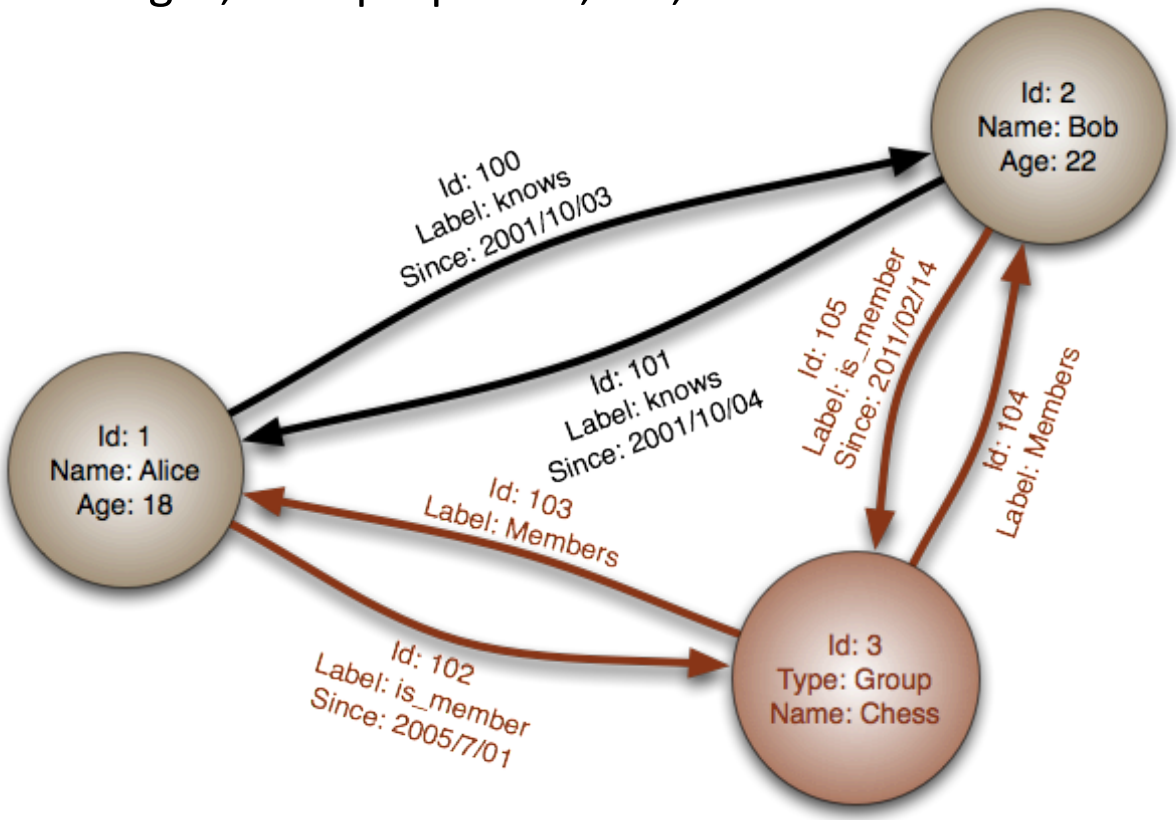

AllegroGraph
ArangoDB
Bigdata
Bitsy
BrightstarDB
DEX/Sparksee
Execom IOG
Fallen *
Filament
FlockDB
GraphBase
Graphd
Horton
HyperGraphDB
IBM System G Native Store
InfiniteGraph
InfoGrid
jCoreDB Graph
MapGraph
Meronymy
Neo4j
Orly
OpenLink virtuoso
Oracle Spatial and Graph
Oracle NoSQL Datbase
OrientDB
OQGraph
Ontotext OWLIM
R2DF
ROIS
Sones GraphDB
SPARQLCity
Sqrrl Enterprise
Stardog
Teradata Aster
Titan
Trinity
TripleBit
VelocityGraph
VertexDB
WhiteDB
(www.nosql-database.org
www.db-engines.com
www.wikipedia.com)

Document Store

MongoDB - “How a Database Can Make Your Organization Faster, Better, Leaner,” February 2015.

The basic data model:

- The general notion of a document – words, phrases, sentences, paragraphs, sections, subsections, footnotes, etc.
- Flexible schema – subcomponent structure may be nested, and vary from document-to-document.
- Metadata – title, author, date, embedded tags, etc.
- Key/identifier.

One implementation detail:

- Formats vary greatly – PDF, XML, JSON, BSON, plain text, various binary, scanned image.

AmisaDB
ArangoDB
BaseX
Cassandra
Cloudant
Clusterpoint
Couchbase
CouchDB
Densodb
Djondb
EJDB
Elasticsearch
eXist
FleetDB
iBoxDB
Inquire
JasDB
MarkLogic
MongoDB
MUMPS
NeDB
NoSQL embedded db
OrientDB
RaptorDB
RavenDB
RethinkDB
SDB
SisoDB
Terrastore
ThruDB
(www.nosql-database.org
www.db-engines.com
www.wikipedia.com)

ACID vs. BASE

Database systems traditionally support ACID requirements:

Atomicity, Consistency, Isolation, Durability

In a distributed web applications the focus shifts to:

Consistency, Availability, Partition tolerance

CAP theorem - At most two of the above can be enforced at any given time.

Reducing consistency, at least temporarily, maintains the other two.

Thus, distributed NoSQL systems are typically said to support some form of BASE:

Basic Availability

Soft state

Eventual consistency*