

Index

■ A

Aerospike, 91, 217
Aerospike query language (AQL), 218
AJAX. *See* Asynchronous JavaScript and XML (AJAX)
Alternative persistence model, 92
Amazon
 ACID RDBMS, 46
 Dynamo, 14, 45–46
 DynamoDB, 219
 hashing, 47–48
 key-value stores, 51
 NWR notation, 49–50
 SOA, 45
Amazon Web Services (AWS), 15
Apache Cassandra, 218. *See also* Cassandra
Apache HBase, 220
Apache Kudu, 211. *See also* Hbase
Append only file (AOF), 94
Asynchronous JavaScript and XML (AJAX), 15
Atomic, Consistent, Independent, and Durable (ACID) transactions, 9–10, 128
AWS. *See* Amazon Web Services (AWS)

■ B

Berkeley analytics data stack and spark
 AMPlab, 99
 BDAS, 100
 DAG, 101
 Hadoop, 99–100
 JDBC-compliant database, 101
 MapReduce, 99
 MLBase component, 100
 RDD, 101
 spark architecture, 101
 spark processing
 elements, 102
 spark SQL, 100
 spark streaming, 100

Big Data revolution
 cloud computing, 22
 competing definitions, 21
 industrial revolution, 22
 IoT, 22
 social networks and smartphones, 22
Binary JSON (BSON), 157
Blockchain, 212
Bloom filters, 161
Boolean bit logic, 214
B-tree index structure, 158–159
Business intelligence (BI) practices, 193

■ C

Cache-less architecture, 92
CAP theorem
 partition tolerance, 44
 RAC solution, 44
Cascading Style Sheets (CSS), 54
Cassandra, 211
 cluster node, 120
 consistent hashing, 120–121
 data model, 153, 155
 gossip, 119
 node adding, 122
 order-preserving partitioners, 124
 replicas, 124–125
 snitches, 126
 virtual nodes, 122–123
Cassandra consistency
 hinted handoff, 136–137
 LWT (*see* Lightweight transactions (LWT))
 read consistency, 135
 read repair, 136–137
 replication factor, 134
 timestamps and granularity, 137
 vector clocks, 138–140
 write consistency, 134–135

Cassandra Query Language (CQL), 218
 column structures, 175
 cqlsh program, 175
 JDBC, 177

CGI. *See* Common Gateway Interface (CGI)

Cloud computing, 15

Codd's 13th rule (nonsubversion), 198

Columnar architecture and Column database architectures
 advantage, 77, 79
 aggregate operations, 78
 columnar and row-oriented storage, comparison, 77
 compression, 79
 disadvantage, 79
 data backed, 81
 data warehouse, 81
 delta store, 81
 insert, column store, 80
 IO and CPU optimizations, 78
 columnar technology, 84
 large-scale bulk sequential loads, 82
 oracle's hybrid columnar compression scheme, 84
 projections
 database table, 83
 pre-join, 83
 superprojection, 82
 RLV, 81
 Tuple Mover, 81
 vertica, 81
 write optimization, 82
 write penalty, 79
 write-optimized delta store, 81

Column family structure, 151–152

Common Gateway Interface (CGI), 40

Consistency models
 ACID, 128, 130
 Cassandra (*see* Cassandra consistency)
 HBase, 132–133
 MongoDB, 131
 MVCC, 128, 130
 transactional consistency, 130
 transaction sequence number, 130
 two-phase-commit (2PC), 130

Copenhagen interpretation, 213

Couchbase, 61, 219
 queries, 202
 N1QL, 198

CouchDB, 61

CQL. *See* Cassandra Query Language (CQL)

CRDT. *See* Convergent replicated data types (CRDT)

Cryptocurrency, 212

CSS. *See* Cascading Style Sheets (CSS)

C-store, 81

Cypher graph query language, 199, 222

■ **D**

DAG. *See* Directed acyclic graph (DAG)

Database
 bewildering array, 215
 BI frameworks, 197
 blockchain, 212
 Cambrian explosion, 214
 Cloudera distribution of Hadoop, 201
 consistency models, 195–196
 convergent, 210
 Criticisms of next generation
 business intelligence, 193
 compromises, 193
 decision points, 194
 de-normalization, 193
 Edgar Codd's key critiques, 193
 high-level logical model, 194
 IDMS and IMS, 193
 inconsistent, 193
 navigational model, 193
 nonrelational systems, 192–193
 RDBMS, 194
 unambiguous and nonredundant view, 194
 disruptive database technologies, 211
 Dynamo-style eventual consistency, 210
 graph compute engine, 202
 hybrid capabilities, 197
 in ACID RDBMS systems, 210
 incompatible technologies, 195
 JSON embedded in RDBMS, 201
 JSON via Oracle REST, 204–206
 languages, 198–199
 modern RDBMS, 214
 MongoDB users, 210
 Oracle Big Data SQL, 201
 Oracle graph, 207
 Oracle JSON support, 202–203
 Oracle sharding, 208–210
 Oracle's RAC clustered, 210
 Oracle tables, 206–207
 ORDS, 201
 possible convergence, schema models, 198
 quantum computing, 213–214
 RDBMS incumbents, 195
 relational model, 197
 revolution
 competitive challenges, 192
 graph databases, 192
 Hadoop and Spark, 192
 Internet of Things (IoT), 191
 nonrelational operational databases, 192
 predominant drivers, 191

- relational model, 192
- SQL, 192
- transactions, 192
- sharded distributed database, 202
- storage, 199–201
- storage technologies, 211–212
- strict multi-record ACID transactions, 195

Database Management System (DBMS), 6

Database survey

- Aerospike, 217
- Cassandra, 218
- CouchBase, 219
- DB-Engines site, 217
- DynamoDB, 219
- HBase, 220
- MarkLogic, 221
- MongoDB, 221
- Neo4J, 222
- non-trivial score, 217
- NuoDB, 223
- Oracle RDBMS, 223
- Redis, 224
- “revolutionary”, 217
- Riak, 225
- SAP Hana, 225
- TimesTen, 226
- Vertica, 227
- VoltDB, 227

Data models

- BigTable and HBase, 145, 151–152
- Cassandra, 153–156
- document databases, 146
- graph databases, 146
- JSON, 156–157
- key-value, 145
- key-value stores
 - conflict resolution, 148
 - CRDT, 148–150
 - data-type agnostic, 148
 - Riak, 148, 150
 - secondary indexes, 148
- relational, 145–147

Data warehousing schemas

- CPU and IO intensive, 76
- CRUD operations, 75
- OLAP, 75
- OLTP system, 75
- snowflake schema, 75
- star schemas, 75–76

Directed acyclic graph (DAG), 101, 198

- Apache Tez *project*, 181
- MapReduce paradigm, 181

Distributed relational databases

- client-server, 106
- mainframe, 105
- monolithic database server, 106

- MPP databases, 108
- RAC cluster database, 109
- replication
 - approach, 107
 - log-based, 107
 - standby database, 107
 - transaction log, 107

- shared disk, 107, 109

- shared-nothing, 108
- web architectures, 106

Document databases, 15

- JSON, 57
- nonrelational database, 53
- XML, 54–57

Durable distributed cache (DDC) architecture, 223

■ E

Early database systems

- definition, 4
- electronic computers, 5
- human civilization and technology, 4
- indexing methods, 5
- tabulating machines and punched cards, 5

EC2. *See* Elastic Compute Cloud (EC2)

EHCC. *See* Enhanced Hybrid Columnar

Compression (EHCC)

Elastic Compute Cloud (EC2), 15

Enhanced Hybrid Columnar Compression (EHCC), 84

eXtensible Markup Language (XML)

- CSS, 54
- database architecture, 56
- relational systems, 57
- tools and standards, 54
- XQuery statement, 54

■ F

Facebook, 14

Fast projection index, 84

First database revolution

- data handling code, 6
- DBMS, 6
- network and hierarchical model, 6–7

■ G

Google

- hardware platform, 23
- MapReduce, 26–27
- modular data center, 24
- PageRank, 23
- software, 25

Google Cloud BigTable, 151

Google Modular Data Center, 24

Graph databases, 192
 definition, 66
 graph compute engines, 73
 Gremlin, 71, 73
 index-free adjacency, 73
 internal storage, 73
 Neo4j, 69–71
 property, 69–71
 RDBMS patterns, 67–68
 RDF, 68–69
 SPARQL, 68–69

■ **H**

Hadoop
 analytic processing, 28
 architecture, 29–30
 ecosystem, 37
 HBase, 32–34
 Hive, 34–35
 Nutch, 28
 open-source project, 28
 Pig, 36
 Hadoop Distributed File System (HDFS), 29
 HANA architecture, 95
 HBase
 architecture, 115, 117
 caching and data locality, 117–118
 catalog tables, 116
 DataNode, 118
 Hadoop HDFS file system, 115
 HDFS, 32
 HDFS DataNodes, 115
 implementation, 115
 master node, 119
 master server, 116
 OpenTSDB, 118
 random access database services, 115
 real-time random access database, 115
 region replicas, 119
 RegionServer, 116, 118–119
 vs. relational model, 33
 rowkey ordering, 118
 short-circuit reads, 117
 tables, 116
 Zookeeper service, 116
 HDFS. *See* Hadoop Distributed File System (HDFS)
 Hierarchical model, 6–7
 Hive
 architecture, 35
 Impala, 35
 SQL processing layer, 34
 Hive Query Language (HQL), 34

■ **I**

IaaS. *See* Infrastructure as a Service (IaaS)
 IBM, 7
 Inconsistent, 193
 Index Sequential Access Method (ISAM), 5
 Infrastructure as a Service (IaaS), 15
 In-memory databases
 alternative persistence model, 92
 Big Data phenomenon, 92
 Cache-less architecture, 92
 COMMIT operations, 92
 memory cost and capacity, 91
 Oracle 12c, 98–99
 Redis, 93–94
 SAP HANA, 95
 TimesTen, 92–93
 traditional database architecture, 92
 VoltDB, 97–98
 Internet of Things (IoT), 191
 ISAM. *See* Index Sequential Access Method (ISAM)

■ **J, K**

JavaScript object notation (JSON), 15, 156
 AJAX, 57, 58
 content-management systems, 57
 CouchBase, 61
 CouchDB, 61
 databases, 58–59
 document embedding, 60
 MemBase, 61
 MongoDB, 61
 JSON. *See* JavaScript object notation (JSON)
 JSON embedded in RDBMS, 201

■ **L**

Lightweight transactions (LWT)
 compare-and-set (CAS) *pattern*, 141
 lockless architecture, 140
 optimistic locking *pattern*, 142
 Paxos protocol, 142
 processing, 142–143
 Log-based replication, 107
 Log-structured merge (LSM) trees, 1117
 architecture, 90
 bloom filters, 161
 Cassandra terminology, 160–161
 CommitLog, 160
 compaction, 162
 in-memory tree, 160

- on-disk trees, 160
- SSTables, 160–161
- Tombstones, 162
- WAL, 160

LSM. *See* Log-structured merge tree (LSM)

■ M

- Magnetic disk device, 87
- MarkLogic, 221
- Massively parallel processing (MPP)*, 108
- Membase
 - Memcached technology, 61
 - nonrelational system, 61
- Memristors, 212
- MemTable, 160
- Mesos, 100
- MongoDB, 221
 - cluster balancing, 113
 - JavaScript query and SQL, 173
 - JSON, 61
 - MySQL, 62
 - replica set and primary
 - failover, 113–114
 - replication, 113
 - sharding
 - architecture, 110
 - mechanisms, 111, 113
 - range and hash, 112
 - shard key, 111
 - tag-aware, 113
 - write concern and read preference, 115
- MongoDB query, 192–193, 199, 202, 205, 208
- Multi-level cell (MLC)*, 88
- Multi-version concurrency
 - control (MVCC)
 - advantage, 130
 - patterns, 193
 - snapshot construction, 128

■ N

- NIQL
 - analytic systems, 185
 - UNNEST command, 187
- NIQL. *See* Non-first Normal Form Query Language (NIQL)
- Neo4j, 222
 - cypher and *cypher query*, 69–70, 71
- Network model, 6–7
- Network topology aware replication
 - strategy, 124
- NewSQL
 - H-Store* and *C-Store*, 16
 - RDBMS, 16

- Non-first Normal Form Query Language (NIQL), 61
- Nonrelational distributed databases
 - ACID compliance, 110
 - balancing availability and consistency, 110
 - consistent hashing model, 110
 - hardware economics, 110
 - omniscient master, 110
 - traditional sharding architecture, 110
- Nonrelational operational databases, 192
- NoSQL
 - APIs, 169
 - cascading, 181
 - CQL, 175–177
 - DAG, 181
 - Hbase, 171–172
 - MapReduce, 177–179
 - MongoDB, 173–175
 - Pig, 179–180
 - Riak, 169–171
 - Spark *project*, 181–182
- NuoDB, 210, 223
- Nutch, 28

■ O

- Object Oriented Database Management System (OODBMS), 11, 13
- Object-oriented programming (OOP)
 - encapsulation, 12
 - inheritance, 12
 - RDBMS, 12
- Object-Relational Mapping (ORM), 13
- OLTP. *See* On-line Transaction Processing (OLTP)
- Online Analytic Processing, 75
- On-line Transaction Processing (OLTP), 5
- OODBMS. *See* Object Oriented Database Management System (OODBMS)
- OOP. *See* Object-oriented programming (OOP)
- openCypher graph language, 207
- Oplog, 113
- Oracle Big Data Appliance, 201
- Oracle Big Data Hadoop system, 207
- Oracle database in-memory, 98–99
- Oracle JSON support, 202–203
- Oracle Parallel Server, 208
- Oracle RDBMS, 223
- Oracle Real Application Clusters (RAC), 208
- Oracle REST Data Services (ORDS), 201
- Oracle REST interface, 207
- Oracle REST JSON query, 204
- Oracle REST query, 206
- Oracle sharding architecture, 208–209
- Oracle TimesTen In-Memory Database (TimesTen), 226
- ORM. *See* Object-Relational Mapping (ORM)

P

Pig Latin, 36
PropertyFileSnitch, 126

Q

Quantum query language (QQL), 214
Quantum search, 213
Quantum transactions, 213
QUEL, 10

R

RAC. *See* Real Application Clusters (RAC)
RackInferringSnitch, 126
RDD. *See* Resilient distributed datasets (RDD)
RDF. *See* Resource Description Framework (RDF)
Real Application Clusters (RAC), 43, 109
Redis. *See* Remote dictionary server (Redis)
Relational storage model
 B-tree index structure, 158–159
 Couchbase's *HB+-Trie*, 160
 database architecture, 158
 index blocks, 159
 RDBMS architectural pattern, 158
 Tokutek's *fractal tree index*, 160
Relational theory
 concepts, 8
 normalized and un-normalized data, 8–9
Remote dictionary server (Redis), 224
 AOF, 94
 architectural components, 94
 architecture, 94
 EMC, 93
 key-value store, 95
 key-value store architecture, 93
 memory database system, 95
 MongoDB, 95
 snapshot, 93
 virtual memory system, 94
Replica sets, 113–114
Replication factor, 124
Resilient distributed datasets (RDD), 101
Resource Description Framework (RDF), 68–69
Riak, 225
Row Level Versioned, 81

S

SAP Hana, 95–96, 225
SCN. *See* System change number (SCN)
Second database revolution
 client-server computing, 11
 IBM, 7

INGRES, 10
mainframe computer, 10
OODBMS, 11–13
OOP, 11–13
QUEL, 10
relational database model (*see* Relational theory)
SQL/DS, 10
SQL language, 10
transaction models, 9–10
Secondary index
 B-tree indexes, 163
 DIY, 163
 global and local, 165
 implementations, 166
 nonrelational operational database systems, 163
Service-oriented architecture (SOA), 45
Set-based query language (SQL)
 advantages, 183
 ANSI and ISO standard, 168
 Apache Drill *framework*, 188–190
 Hive, 183–184
 Impala, 184
 N1QL, 185–187
 NoSQL, 190
 spark, 185
 types, 168
Shard chunk, 113
Sharding
 ACID transactions, 14
 drawbacks, 43
 Facebook, 14
 memcached/replication architecture, 42
Shared-disk database architecture, 109
Simple Oracle Data Access (SODA), 204
Simple Oracle Document Access (SODA), 201
Single-level cell (SLC), 88
SOA. *See* Service-oriented architecture (SOA)
SODA REST query, 205
Solid state disk (SSD)
 Aerospike, 91
 algorithms, 89
 battery-backed RAM device, 88
 DDR RAM, 88
 economics, 89–90
 enabled databases, 90
 NAND flash, 88
 performance characteristics, 89
 SLC and MLC, 88
 write amplification, 89
SPARQL Protocol, 69
Splice Machine layers, 210
SQL/DS, 10
SSTables, 160

Star schemas, 75–76, 147
Superprojection, 82
 Sybase IQ, 81
 System change number (SCN), 130

■ **T, U**

Tabulating machines and punched cards, 5
 Tachyon, 100
 Third database revolution
 cloud computing, 15
 document databases, 15
 Google, 14
 Hadoop, 14
 NewSQL, 16
 TimesTen, 92–93, 226
 Time to live (TTL), 152
 Transaction models, ACID transaction, 9
 Transactions, 192
 Tunable consistency model, 201

■ **V**

Vertica, 81, 83, 227
 VoltDB, 97–98, 227

■ **W**

Web 2.0
 CAP theorem, 43–44
 CGI-based approaches, 40
 e-commerce, 40
 eventual consistency, 45
 open-source solution
 memcached servers and
 replication, 41
 MySQL, 41
 scale-up *solution*, 40
 sharding, 41–43
Write-Ahead Log (WAL), 117, 160

■ **X**

XML. *See* eXtensible Markup Language (XML)

■ **Y, Z**

Yet Another Resource
 Negotiator (YARN)
 Application Manager, 30
 Resource Manager, 30

