

Index

- , 17
- ∪, 17
- ×, 16
- π, 17
- σ, 17
- ×, 17
- 1NF, 29
- 2NF, 29
- 2PC, 83
- 2PC-UTRB, 89
- 2PL, 49
- 3NF, 29
- 3PC, 86
- 4NF, 29
- 5NF, 30
- 6NF, 30

- abort, 13
- ACA, 48
- ACID properties, 10
- aggregate operators, 21
- aggregation, 22
- aggressive, 52
- always in future, 127
- always in past, 127
- anomalies, 42
- asynchronously, 65
- atomic commitment, 71, 83
- Atomicity, 10
- attribute, 15
- attribute branch, 111
- attributes, 109
- autonomy, 61
- avoids cascading aborts, 48
- axes, 112
- axis, 112

- bag semantics, 20
- balanced, 93
- BAV, 65, 96
- BCNF, 29
- before images, 49
- begin, 10, 13
- binary, 93
- blocked, 84

- BM, 11
- both as view, 65
- both-as-view, 96
- bottom-up design process, 62
- bounded model, 124
- Boyce-Codd Normal Form, 29
- branching model, 124
- buffer manager, 11

- candidate key, 17, 29, 116
- cascading abort, 48
- CCR, 84
- CDATA section, 110
- CDM, 57
- character data, 110
- chronon, 123
- closing tag, 109
- clustered, 75
- columns, 17
- comma separated value, 8
- command language, 9
- commit, 10, 13
- commitment, concurrency, and recovery, 84
- committed projection, 44
- common concepts, 92
- common data model, 57
- communication costs, 78
- communication failure, 71
- compatible, 47
- complete, 13
- complex type, 115
- component databases, 59
- component schema, 57, 59
- concepts, 92
- conceptual level, 9
- conceptual modelling, 31
- conceptual schema, 9
- concurrency control, 10, 41
- conflict, 46, 47, 93
- conflict detection, 94
- conflict equivalent, 46
- conflict resolution, 94
- conflict serialisable, 46
- conservative, 52
- conservative 2PL, 53

- Consistency, 10
- constraint, 96
- construct, 92
- context, 112
- continuous model, 123
- coordinator, 83
- cost model, 72
- CSR, 46
- CSV, 8

- data definition language, 9
- data integration, 62, 65, 92
- data manager, 12
- data manipulation language, 9
- data model, 7
- data modelling language, 7
- data warehouse, 60, 92
- database, 7
- database administrator, 9
- database management system, 8
- database management systems, 10
- DBA, 9
- DBMS, 8, 10
- DDB, 57
- DDBMS, 57
- DDL, 9
- deadlock, 53, 54
- dependency, 28
- derived horizontal fragmentation, 62
- determinant, 29
- difference, 17
- dirty read, 48
- discrete linear bounded, 123
- discrete model, 123
- distributed data independence, 58
- distributed database, 57
- distributed DBMS, 57
- DM, 12
- DML, 9
- document type definition, 114
- double index, 75
- drill down, 22
- DTD, 114
- Durability, 10

- edge schema, 119
- election, 84
- element, 109
- element branch, 112
- element schema, 119
- end, 13
- end of file, 55
- entity reference, 110
- Entity Relationship, 31
- entity-relationship model, 91
- EOF, 55
- equivalent, 45
- equivalent schemas, 96
- ER, 31
- ER Model, 91
- export schema, 59
- extensible mark-up language, 109
- external level, 9
- external schema, 9, 59

- FD, 28
- federated database, 60
- federated multi-DBMS, 61
- federated schemas, 59
- fetch, 11
- fifth normal form, 30
- filtered, 91
- filters, 24
- first normal form, 29
- flat files, 8
- flush, 11
- FLWOR, 113
- for let where order return, 113
- foreign key, 118
- fourth normal form, 29
- fragments, 62
- fully locked, 51
- functional dependency, 28

- GAV, 65
- generators, 24
- GLAV, 65
- global as view, 65
- global local as view, 65
- global queries, 69
- global schema, 58
- global schemas, 61
- global transaction manager, 69
- global transactions, 69
- growing phase, 50
- GTM, 69

- HDB, 57
- heterogeneous database, 57
- homogeneous, 61
- homonyms, 98
- horizontal fragmentation, 62
- hybrid fragmentation, 64

- inconsistent analysis, 43
- integration step, 93
- intermediate query language, 24
- internal level, 8

- intra-query parallelism, 64
- IQL, 24
- Isolation, 10

- JD, 29
- join, 16, 72, 78
 - lossless, 64
 - natural, 16
 - semi, 16
 - since, 125
 - until, 125
- join dependency, 29

- ladder, 93
- LAV, 65
- LDM, 69
- linear model, 124
- list comprehension, 24
- livelock, 53
- livelocks, 54
- local as view, 65
- local data manager, 69
- local scheduler, 69
- local schema, 57
- local transaction manager, 69
- locks, 49
- logical distribution, 59
- look-across, 34
- look-here, 34
- loosely coupled, 61
- lost update, 42
- LS, 69
- LTM, 69

- many-many, 35
- mappings, 65
- marking-up, 109
- materialised global schema, 60
- media failures, 12
- mediator, 60
- merging strategy, 93
- migrate, 62
- missing information, 20
- multi-DBMS, 57, 61
- multi-valued dependency, 28
- multidimensional, 23
- MVD, 28

- n-ary, 93
- next, 127
- node test, 112
- non-blocking, 86
- non-recoverable execution, 48
- normalisation, 29

- null, 20
- null value, 20

- object role modelling, 31
- objects, 13
- OLAP, 21
 - relational, 21
- OLTP, 21
- one-many, 35
- one-one, 35
- online analytical processing, 21
- online transaction processing, 21
- opening tag, 109
- operational distribution, 59
- operations, 13
- optimistic concurrency control, 49
- ORM, 31

- pages, 11
- partial virtual global schema, 60
- path expression, 112
- pathway, 96
- phantom read, 43
- physical heterogeneity, 57
- pivot, 24
- plans, 12
- predicate, 55
- predicate locking, 55
- predicates, 24
- previous, 127
- primary copies, 65
- primary key, 29
- primary-backup, 80
- primitive transformation, 96
- processing instructions, 111
- product, 17
 - since, 125
 - until, 125
- project, 17

- QP, 71
- query optimisation, 12, 72
- query plan, 76
- query processing, 65
- query processor, 71

- RA, 15
- RC, 48
- RDBMS, 15, 17
- read locks, 50
- read-write conflicts, 80, 81
- recoverable, 47, 48
- recovery, 11
- recovery manager, 12

- redundancy, 28
- related concepts, 93
- relation, 15
- relational algebra, 15
- relational DBMS, 15, 17
- relational OLAP, 21
- relationally complete, 16, 24
- replicated, 62
- resource managers, 83
- RM, 12
- ROLAP, 21
- rollback, 13
- rolldown, 22
- rollup, 22

- scheduler, 12
- schema, 8, 15
- schema conforming, 94
- schema improvement, 94
- schema integration, 58, 62, 65, 92
- schema merging, 94
- schema transformation, 95
- scheme, 24, 96
- second normal form, 29
- select, 17
- select project join, 19
- semantic heterogeneity, 57
- semi-join, 16
- semi-join method, 78
- semi-structured data, 8
- Sequel, 17
- serial, 44
- serialisable, 45
- server, 83
- shrinking phase, 51
- simple type, 115
- since, 124
- single index, 75
- single-DBMS, 61
- sites, 57
- sixth normal form, 30
- snapshot, 123
- sometime in future, 127
- sometime in past, 127
- SPJ, 19
- SQL, 15, 17
 - AVG, 21
 - COUNT, 21
 - CREATE, 17
 - CREATE UNIQUE INDEX, 17
 - CROSS JOIN, 19
 - CUBE, 23
 - DATENAME, 24
 - DECLARE, 25
 - DISTINCT, 20, 22
 - EXCEPT, 17, 20
 - FALSE, 20
 - GROUP BY, 21
 - IS, 20
 - IS NOT, 20
 - IS NOT NULL, 21
 - IS NULL, 21
 - JOIN, 18
 - MAX, 21
 - MIN, 21
 - NATURAL JOIN, 19
 - PRIMARY KEY, 29
 - ROLLUP, 23
 - SELECT, 17
 - STDDEV, 21
 - SUM, 21
 - TRUE, 20
 - UNION, 17, 20
 - UNION ALL, 20
 - UNION DISTINCT, 20
 - UNIQUE, 29
 - UNKNOWN, 20
 - VARIANCE, 21
- SQL-Server, 2
- ST, 49
- stable storage, 11
- state-machine, 80
- step expression, 112
- strict, 49
- strict 2PL, 53
- structured data, 8
- structured query language, 15, 17
- stylesheets, 111
- subset entity, 32
- superset entity, 32
- Sybase, 2
- synchronously, 65
- synonyms, 98
- system failures, 12

- tables, 17
- tags, 109
- temporal database, 123
- temporal normal form, 127
- temporal relational algebra, 125
- temporary tables, 24
- theta join, 16
- third normal form, 29
- three-phase commit, 86
- tightly coupled, 61
- time-stamping, 49
- TM, 12
- TNF, 127

- top-down design process, 62
- TRA, 125
- Transact-SQL, 2
- transaction, 9, 10
- transaction manager, 12, 69
- transaction time, 123
- transformed, 91
- trivial MVD, 29
- two-phase commit, 83
- two-phase locking, 49

- UML, 31
- unbalanced, 93
- unbounded model, 124
- unfederated multi-DBMS, 61
- unified modelling language, 31
- uniform timed reliable broadcast, 89
- union, 17
 - compatible, 17
- union compatible, 62
- universe of discourse, 7
- unknown, 20
- unstructured-data, 8
- until, 124
- UoD, 7
- US logic, 124
- UTRB, 89

- valid time, 123
- vertical fragmentation, 64
- view, 65
- view equivalent, 45
- view integration, 92
- view serialisable, 46
- virtual global schema, 60
- volatile storage, 11
- VSR, 46

- waits-for graph, 54
- WFG, 54
- workflow, 60
- wrapper, 57
- write locks, 50
- write-locks-all, 80
- write-write conflicts, 80, 81

- XML, 109
- XML declaration, 111
- XML document, 109
- XML Schema, 115
- XPath, 111
- XPath tree, 111
- XQuery, 113

