


Lecture 3

DATABASES




Database



A database is an organized collection of [data](#), generally stored and accessed electronically from a computer system

Database



Edgar Frank "Ted" Codd (19 August 1923 – 18 April 2003)

while working for [IBM](#), invented the [relational model](#) for [database](#) management, initial paper was

"A Relational Model of Data for Large Shared Data Banks"

SQL

S-Q-L or “sequel” - Structured Query Language

used in programming and designed for managing data held in a [relational database management system](#) (RDBMS)

SQL

SQL was initially developed at [IBM](#) by [Donald D. Chamberlin](#) and [Raymond F. Boyce](#) after learning about the relational model from [Ted Cod](#) in the early 1970s. This version, initially called SEQUEL (Structured English Query Language), was designed to manipulate and retrieve data stored in IBM's original quasi-relational database management system, [System R](#), which a group at [IBM San Jose Research Laboratory](#) had developed during the 1970s. ^[15]

SQL clauses - w3c site

Sample:

Web SQL Database

https://www.w3schools.com/sql/trysql.asp?filename=trysql_select_all

SQL statements - w3c site

Basic commands:

CREATE TABLE, SELECT (JOINS) INSERT, UPDATE, DELETE

SQL statements - CREATE TABLE

Sample:

- `CREATE TABLE Persons`
- `(`
 - `PersonID int,`
 - `LastName varchar(255),`
 - `FirstName varchar(255),`
 - `Address varchar(255),`
 - `City varchar(255)`
- `);`

SQL statements - INSERT

Sample:

```
INSERT INTO Customers (CustomerName, ContactName, Address,  
City, PostalCode, Country)
```

```
VALUES ('Cardinal', 'Tom B. Erichsen', 'Skagen 21',  
'Stavanger', '4006', 'Norway');
```

SQL statements - SELECT

Sample:

```
SELECT column1, column2, ...
```

```
FROM table_name;
```

SQL statements - DELETE

Sample:

```
DELETE FROM table_name WHERE condition;
```

SQL statements - Other

Sample:

- `Select top,`
- `Select into,`
- `Case,`
- `Null values, order by, group, by`
- `Stored procedures`
- `Null values`

SQL Database types

Developed in 1970s to deal with first wave of data storage applications

MySQL, Postgres, Microsoft SQL Server, Oracle Database

SQL Database types

MySQL -> Postgres -> Microsoft SQL Server -> Oracle Database



MS SQL Server - Management Studio

The screenshot displays the Microsoft SQL Server Management Studio interface. The main window shows a query execution result for a table named 'dbo.Appointments'. The query executed is a 'Script for SelectTopRows command from SSMS'. The results are displayed in a table with the following columns: ResourceId, ServiceId, SourceSite, StartDate, Status, TenantId, TotalPrice, and IsNeed. The status of the appointments varies, including 'Canceled', 'Booked', and 'Paid'. A status bar at the bottom indicates 'Query executed successfully.' and '2848 rows'.

ResourceId	ServiceId	SourceSite	StartDate	Status	TenantId	TotalPrice	IsNeed		
2...	zero: 0.012 - 0.0...	120	NULL	NULL	2019-09-17 13:00:00.0000000	Canceled	25	0.02	1
2...	zero: 0.012 - 0.0...	120	NULL	NULL	2019-09-17 13:00:00.0000000	Canceled	25	0.02	1
2...	Boere: 0.012 - 1...	120	NULL	NULL	2019-09-17 13:00:00.0000000	Canceled	25	0.02	1
2...	zero: 0.012 - 0.0...	120	NULL	NULL	2019-09-18 13:00:00.0000000	Canceled	25	0.02	1
2...	ero: 200.004 - 1...	40	NULL	NULL	2019-09-28 13:00:00.0000000	Booked	3	200.00	0
2...	zero: 0.012 - 0.0...	120	NULL	NULL	2019-09-18 13:00:00.0000000	Canceled	25	0.02	1
2...	ero: 55.200 - 46...	26	NULL	NULL	2019-09-22 13:00:00.0000000	Paid	3	110.40	1
2...	ero: 51.000 - 42...	9	NULL	NULL	2019-09-28 13:00:00.0000000	Canceled	3	520.20	1
2...	zero: 55.200 - 46...	5	NULL	http://localhost:3001/book/	2019-09-26 13:00:00.0000000	Canceled	3	124.20	1
2...	ero: 69.000 - 57...	5	NULL	http://localhost:3001/book/	2019-10-05 13:00:00.0000000	Canceled	3	124.20	1
2...	ero: 69.000 - 57...	5	NULL	http://localhost:3001/book/	2019-10-12 13:00:00.0000000	Canceled	3	124.20	1
2...	ero: 69.000 - 57...	5	NULL	http://localhost:3001/book/	2019-10-26 13:00:00.0000000	Canceled	3	124.20	1
2...	ero: 55.200 - 46...	5	NULL	http://localhost:3001/book/	2019-10-23 13:00:00.0000000	Canceled	3	110.40	1
2...	zero: 55.200 - 46...	5	NULL	http://localhost:3000/book/	2019-10-16 13:00:00.0000000	Canceled	3	110.40	1
2...	zero: 55.200 - 46...	5	NULL	http://search.zapytai.by/book/	2019-10-07 13:00:00.0000000	Canceled	3	110.40	1
2...	zero: 55.200 - 46...	5	NULL	http://localhost:3000/book/	2019-09-29 13:00:00.0000000	Canceled	3	110.40	1

NoSQL Database types

- Document
- Graph
- Key-value .
- Wide-column

NoSQL considerations

- Large volumes of rapidly changing structured, semi-structured, and unstructured data
- Agile sprints, quick schema iteration, and frequent code pushes
- Object-oriented programming that is easy to use and flexible
- Geographically distributed scale-out architecture instead of expensive, monolithic architecture

NoSQL considerations

Since its first MongoDB project in 2012, Baidu has grown its cluster to 600 nodes storing 200 billion documents and 1PB of data, powering over 100 apps.



NoSQL Database types - Document

- Document databases pair each key with a complex data structure known as a document. Documents can contain many different key-value pairs, or key-array pairs, or even nested documents.

NoSQL Database types - Graph stores

Graph stores are used to store information about networks of data, such as social connections. Graph stores include Neo4J and Giraph.

NoSQL Database types - Wide-column

Wide-column stores such as Cassandra and HBase are optimized for queries over large datasets, and store columns of data together, instead of rows

NoSQL Database types - Key-value

Key-value stores are the simplest NoSQL databases. Every single item in the database is stored as an attribute name (or 'key'), together with its value. Examples of key-value stores are Riak and Berkeley DB. Some key-value stores, such as Redis, allow each value to have a type, such as 'integer', which adds functionality.

NoSQL Mongo DB



mongoDB®

MongoDB uses [JSON](#)-like documents with [schema](#).

NoSQL Mongo DB - Robomongo

The screenshot shows the Robo 3T - 1.3 interface. The left sidebar displays a tree view of the database structure under 'Development (1)'. The main window shows a query execution result for the 'Positions' collection. The query is `db.getCollection('Positions').find({})`. The result is a table with 19 rows, each representing a document in the collection. Each document has a unique ObjectId and contains 27 to 31 fields.

Key	Value	Type
(1)	ObjectId("5c2de661b166fe00165eed...") { 27 fields }	Object
(2)	ObjectId("5c667db3f9267f00165110...") { 26 fields }	Object
(3)	ObjectId("59108954b7965e001b817...") { 31 fields }	Object
(4)	ObjectId("5911958902c967001bf9fc...") { 31 fields }	Object
(5)	ObjectId("5911958902c967001bf9fc...") { 31 fields }	Object
(6)	ObjectId("5912d46c5e41a3001b0c9...") { 31 fields }	Object
(7)	ObjectId("5912d7bedda927001b8f76...") { 31 fields }	Object
(8)	ObjectId("5912dc74dda927001b8f76...") { 31 fields }	Object
(9)	ObjectId("5912dd2cdda927001b8f76...") { 31 fields }	Object
(10)	ObjectId("5912dec3dda927001b8f7...") { 31 fields }	Object
(11)	ObjectId("5912ded2dda927001b8f7...") { 31 fields }	Object
(12)	ObjectId("5912ded2dda927001b8f7...") { 31 fields }	Object
(13)	ObjectId("5912ded3dda927001b8f7...") { 31 fields }	Object
(14)	ObjectId("5912ded3dda927001b8f7...") { 31 fields }	Object
(15)	ObjectId("5912df48dda927001b8f7...") { 30 fields }	Object
(16)	ObjectId("5912e15cdda927001b8f7...") { 30 fields }	Object
(17)	ObjectId("5912ee02dda927001b8f7...") { 30 fields }	Object
(18)	ObjectId("5912f068dda927001b8f7...") { 31 fields }	Object
(19)	ObjectId("5912f0a1dda927001b8f76...") { 31 fields }	Object

NoSQL Mongo DB - Robomongo

The screenshot shows the Robomongo interface with a query executed in the 'Positions' collection. The query is:

```
db.getCollection('Positions').find({'_id': ObjectId("5c49d938e13960001678b3ba")})
```

The results are displayed in a table with columns 'Key', 'Value', and 'Type'. The 'selected' field is expanded to show its nested structure.

Key	Value	Type
milestones	{ 5 fields }	Object
placementFee	{ 2 fields }	Object
resumeReceivingMethod	{ 1 field }	Object
status	{ 3 fields }	Object
remainingPositions	1	Int32
created	2019-01-24 12:02:00.685Z	Date
lastUpdate	2019-09-24 08:12:49.123Z	Date
remainingRecruiters	0	Int32
employerId	ObjectId("5aa156159e649e0013c0c796")	ObjectId
recruiterId	ObjectId("590f2a2cf36d281c3b96969")	ObjectId
jobTitle	TFT-2118-GROUP-SKILLS-TEST	String
location	{ 10 fields }	Object
employmentType	Full-time	String
numOfPositions	1	Int32
functionalRole	Quality Assurance	String
desiredCandidateProfile	{ 5 fields }	Object
searchCountries	[3 elements]	Array
searchLanguages	[1 element]	Array
selected	[4 elements]	Array
[0]	{ 2 fields }	Object
[1]	{ 2 fields }	Object
[2]	{ 2 fields }	Object
field	skills	String
items	[9 elements]	Array
[0]	{ 2 fields }	Object
required	mandatory	String
group	[4 elements]	Array
[0]	{ 2 fields }	Object
value	Mentor Graphics	String
uid	0	String
[1]	{ 2 fields }	Object
[2]	{ 2 fields }	Object
[3]	{ 2 fields }	Object
[3]	{ 2 fields }	Object
[1]	{ 2 fields }	Object
[2]	{ 3 fields }	Object
[3]	{ 3 fields }	Object
[4]	{ 3 fields }	Object

NoSQL Mongo DB



mongoDB®

MongoDB uses [JSON](#)-like documents with [schema](#).

NoSQL Mongo DB



mongoDB®

MongoDB uses [JSON](#)-like documents with [schema](#).

```
{
  name "sue",
  age: 26,
  status "A",
  groups [ "news", "sports" ]
}
```

← field: value
← field: value
← field: value
← field: value

NoSQL Mongo DB



mongoDB®

MongoDB uses [JSON](#)-like documents with [schema](#).

```
{
  name "al",
  age 18,
  status "D",
  groups ["politics", "news"]
}
```

Collection

Insert data

```
Collection      Document
  ↓             ↓
db.users.insert(
  {
    name: "sue",
    age: 26,
    status: "A",
    groups: [ "news", "sports" ]
  }
)
```

Document

```
{
  name: "sue",
  age: 26,
  status: "A",
  groups: [ "news", "sports" ]
}
```

insert



Collection

{ name: "al", age: 18, ... }
{ name: "lee", age: 28, ... }
{ name: "jan", age: 21, ... }
{ name: "kai", age: 38, ... }
{ name: "sam", age: 18, ... }
{ name: "mel", age: 38, ... }
{ name: "ryan", age: 31, ... }
{ name: "sue", age: 26, ... }

users

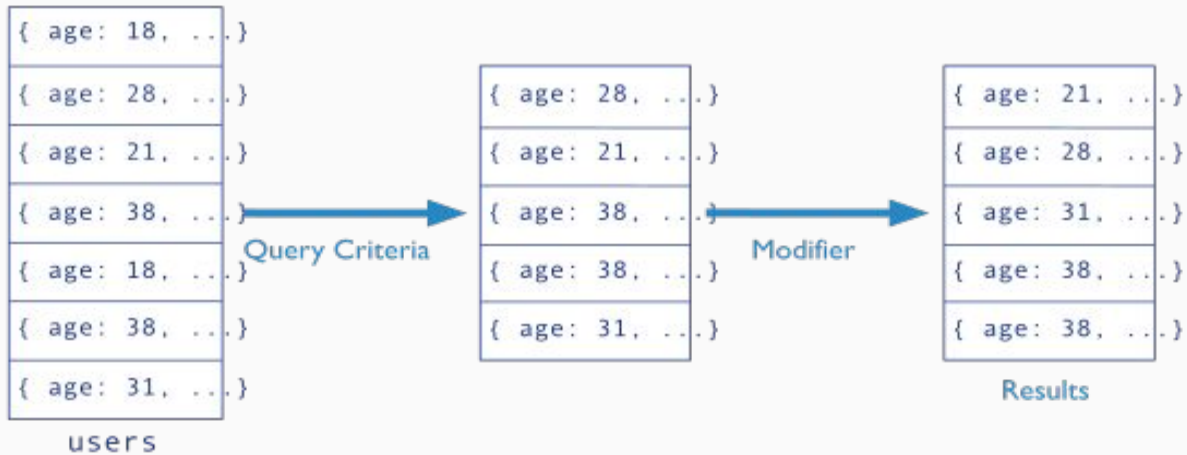
NoSQL Mongo DB



mongoDB®

Query

Collection Query Criteria Modifier
`db.users.find({ age: { $gt: 18 } }).sort({age`



NoSQL Mongo DB



mongoDB®

```
db.users.find(  
  { age: { $gt: 18 } },  
  { name: 1, address: 1 }  
) .limit(5)
```

← collection
← query criteria
← projection
← cursor modifier

```
SELECT _id, name, address  
FROM users  
WHERE age > 18  
LIMIT 5
```

← projection
← table
← select criteria
← cursor modifier

NoSQL Mongo DB

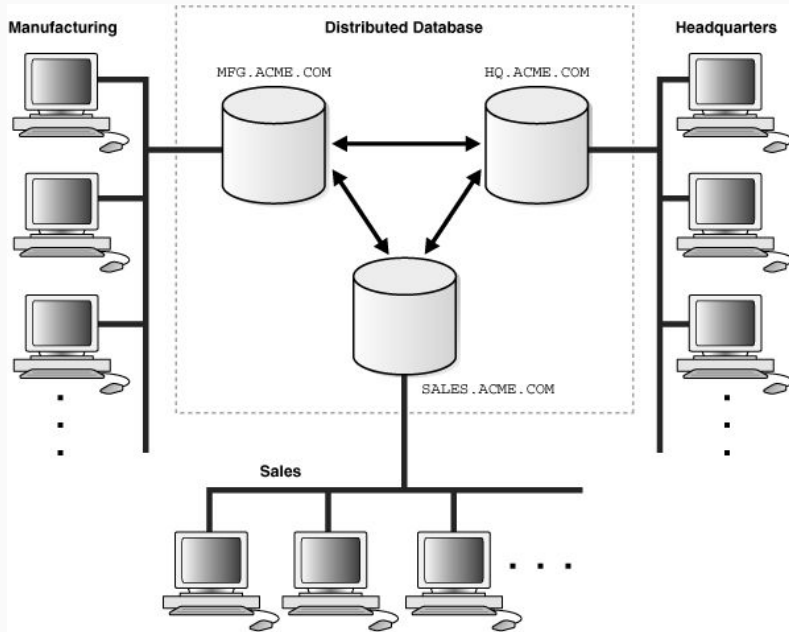


mongoDB®

MongoDB stores data in flexible, JSON-like documents, meaning fields can vary from document to document and data structure can be changed over time.

MongoDB does support a rich, ad-hoc query language of its own.

Distributed databases



Next generation of database evolution

CAP Theorem - Distributed databases



Eric Allen Brewer

The [CAP theorem](#) about distributed network applications in the late 1990s

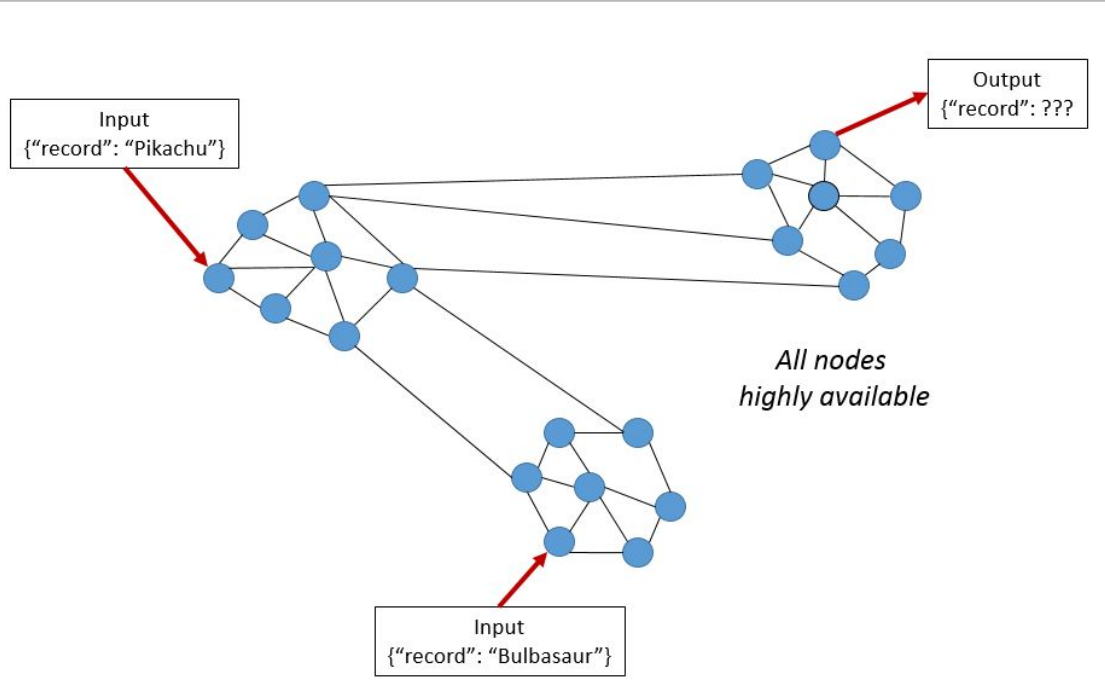
CAP Theorem - Distributed databases

Consistency: Every read receives the most recent write or an error

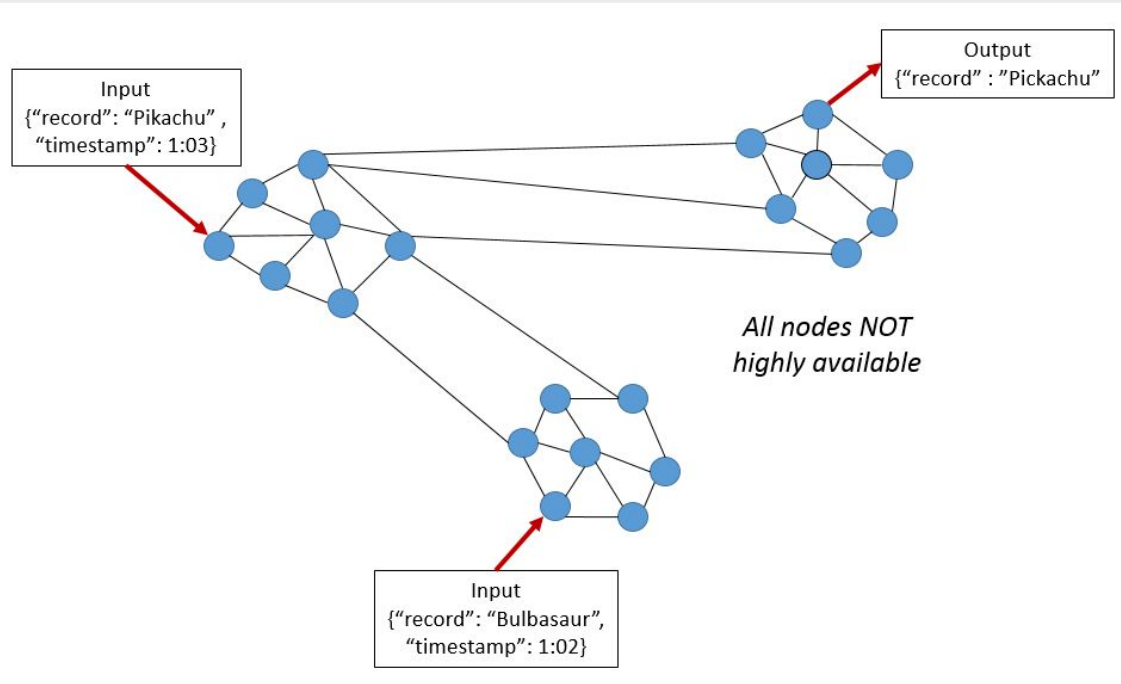
Availability: Every request receives a (non-error) response, without the guarantee that it contains the most recent write

Partition tolerance: The system continues to operate despite an arbitrary number of messages being dropped (or delayed) by the network between nodes


CAP Theorem - High Availability



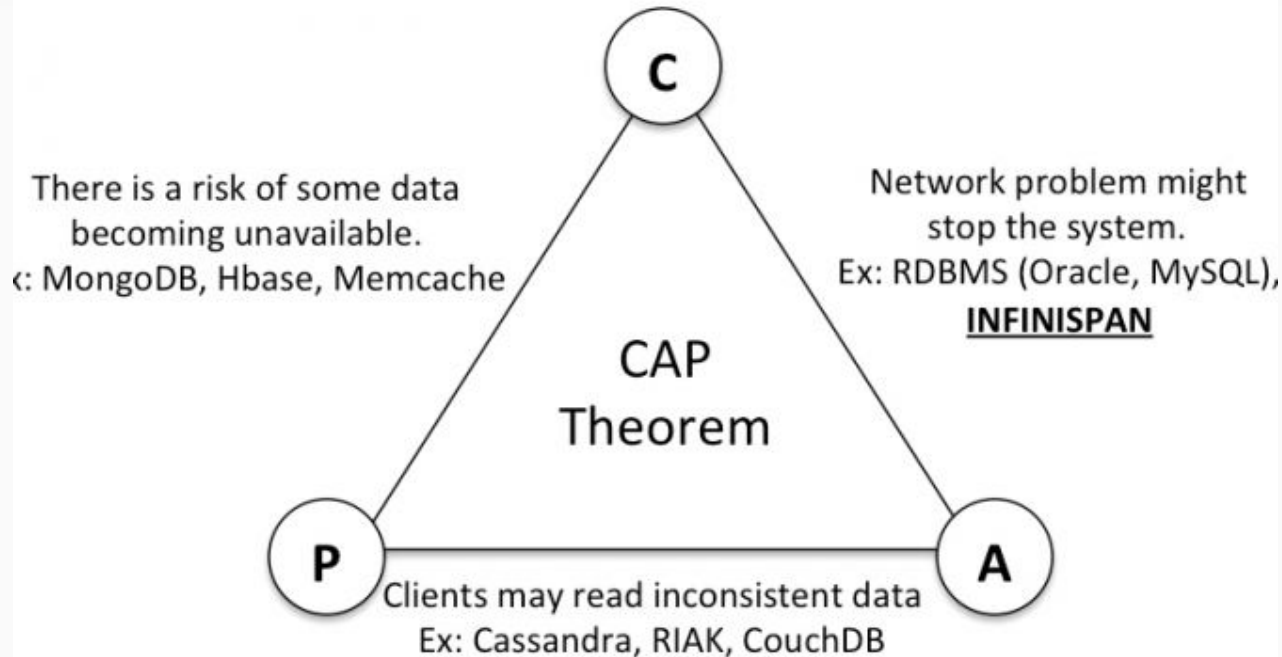
CAP Theorem - High Consistency



CAP Theorem - Partition tolerance



CAP Theorem



Cloud databases

- Azure SQL Database
- Aws
- (Amazon Aurora, Amazon Relational Database Service)
- Google cloud SQL
- Etc.

Data modeling

1. The data contained in the database
2. The relationships between data items
3. The constraints on data

Data modeling - 1

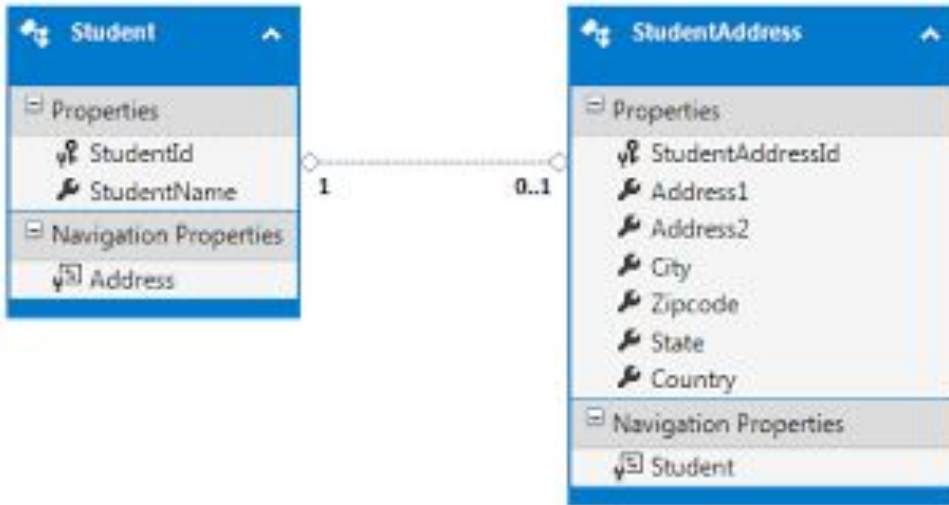
- Numbers,
- Text,
- Images,
- Binary,
- Geo etc.

Data modeling - 2

1. ONE - ONE RELATIONS
2. ONE-MANY RELATIONS
3. MANY-TO-MANY RELATIONS

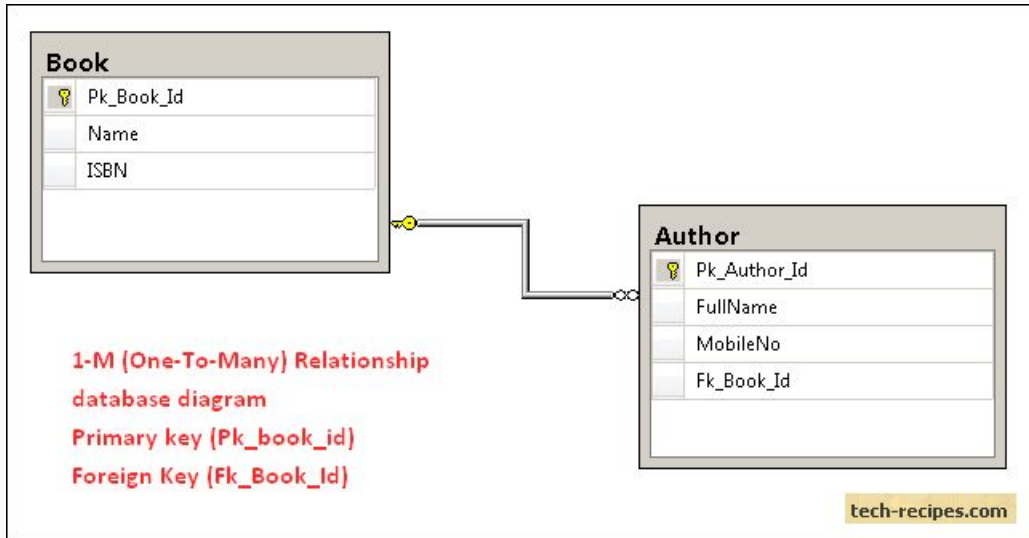
Data modeling - 2

1. ONE - ONE RELATIONS



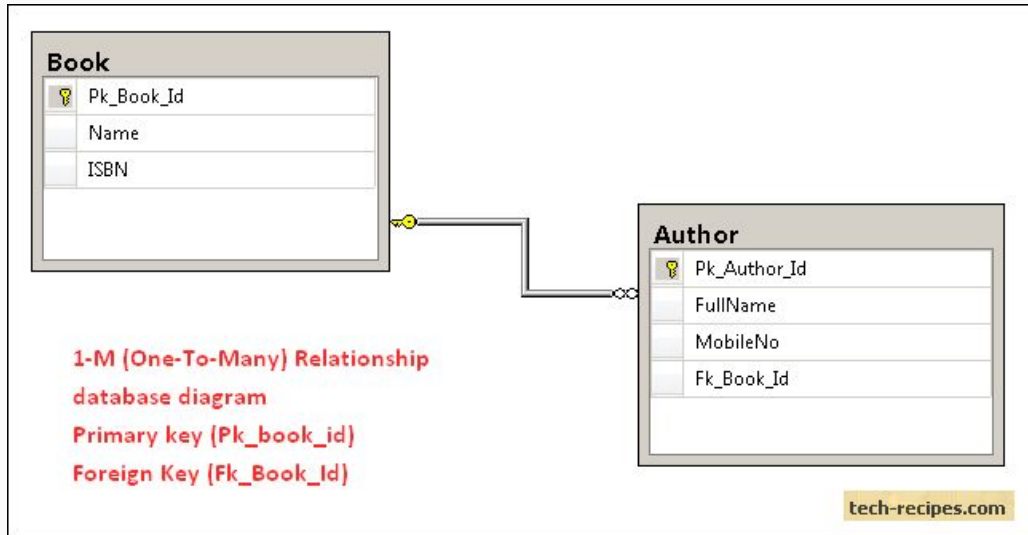
Data modeling - 2

2. ONE - MANY RELATIONS



Data modeling - 2

3. MANY- TO - MANY RELATIONS



Data modeling - 3

The constraints on data

- **not null** - each value in a column must not be NULL
- **unique** - value(s) in specified column(s) must be unique for each row in a table
- **primary key** - value(s) in specified column(s) must be unique for each row in a table and not be NULL; normally each table in a database should have a primary key - it is used to identify individual records
- **foreign key** - value(s) in specified column(s) must reference an existing record in another table (via it's primary key or some other unique constraint)
- **check**

The end