
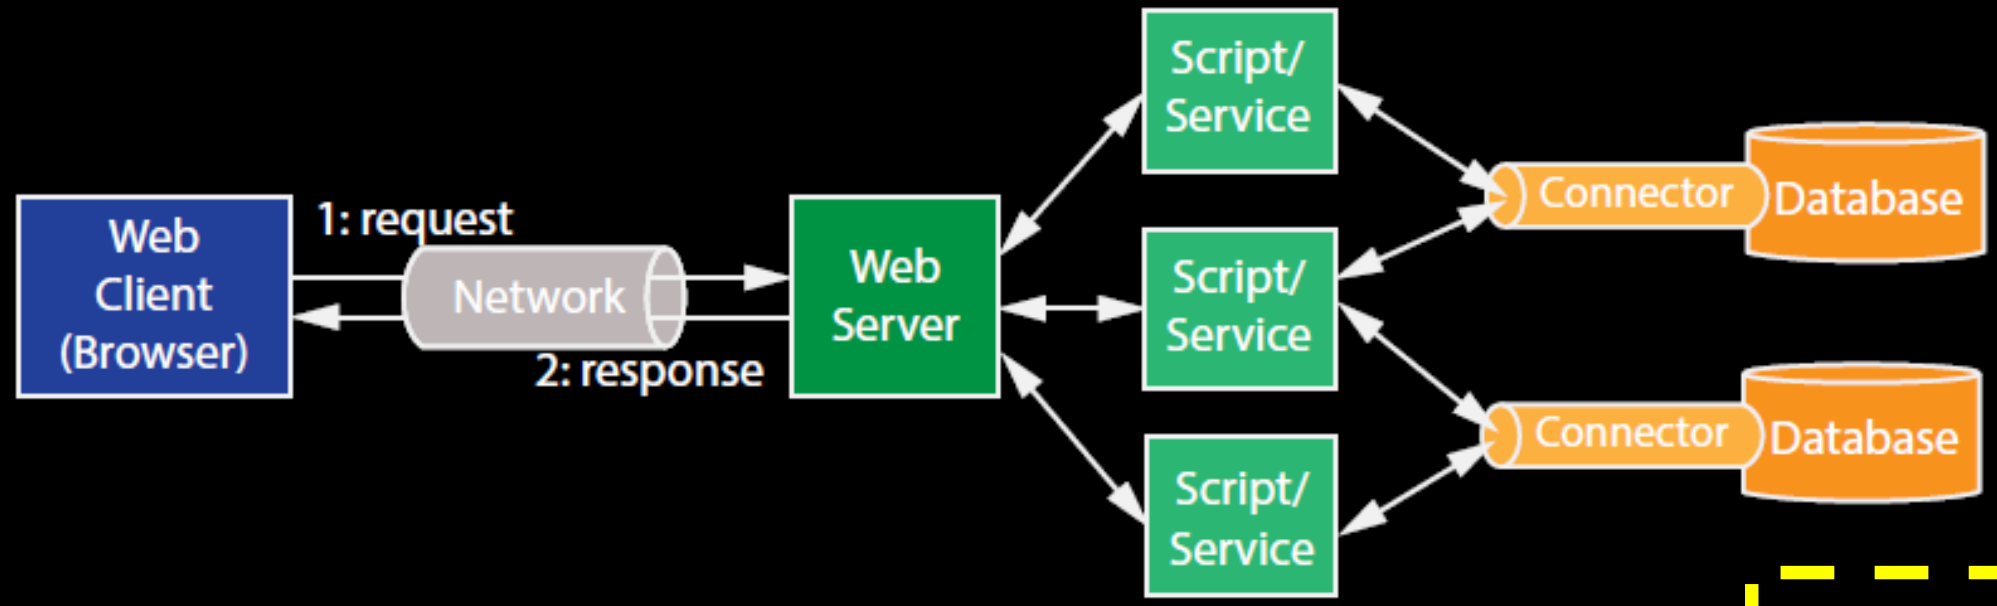


Database – Introduction & MongoDB

INF1802

Profa. Melissa Lemos





Our Focus: Database MongoDB

Database Introduction

Reference

- Jennifer Widom, Introdução a Sistemas de Banco de Dados. Coursera, Disponível online 2016, <https://pt.coursera.org/course/db>
- Ramez Elmasri and Shamkant B. Navathe, Fundamentals of Database Systems (7th Edition) Jun 18, 2015.
- Pramod J. Sadalage, Martin Fowler. “NoSQL Distilled - A Brief Guide to the Emerging World of Polyglot Persistence”, Addison-Wesley, 2013.

DB: Database

coleção de dados relacionados.

DBMS: Database Management System

coleção de programas que permitem usuários definir, construir e manipular um BD

DB System = DB + DBMS

Database Management System (DBMS) provides....

... efficient, reliable, convenient, and safe multi-user storage of and access to massive amounts of persistent data.

Database systems are extremely prevalent in the world today.

- Massive
- Persistent
- Safe
- Multi-user
- Convenient
- Efficient
- Reliable

- **Massive**

terabytes of data

much larger than can fit in the memory of a typical computing
system

- Persistent

the data in the database outlives the programs that execute on
that data

very often multiple programs will be operating on the same data.

- Safe

guarantees that the data managed by the system will stay in a consistent state, it won't be lost or overwritten when there are failures,
hardware failures, software failures, malicious users.

- Multi-user

multiple programs

many different users or applications to access the data

concurrently

- Convenient

the way that data is actually stored and laid out on disk is independent of the way that programs think about the structure of the data

high level query languages

in the query, you describe what you want out of the database but you don't need to describe the algorithm to get the data out

- Efficient

database systems have to do really thousands of queries or updates per second.

- Reliable

critically important that systems are up all the time

- Database applications may be programmed via “frameworks”
Meteor, Django, Ruby on Rails
- DBMS may run in conjunction with “middleware”
Web servers, Application servers
- Data-intensive applications may not use DBMS at all


Key concepts

- Data model
 - Description of, in general, how the data is structured
 - Relational model, XML documents, Graph data model
- Schema versus data
 - The schema sets up the structure of the database
- Data definition language (DDL)
- Data manipulation or query language (DML)

Key people

- DBMS implementer
the person who builds the system
- Database designer
the person who establishes the schema for a database
- Database application developer
the person who builds the applications or programs that are going to run on the database
- Database administrator
the person who loads the data, sort of gets the whole thing running and keeps it running smoothly

Whether you know it or not,
you're using a database every day



Evolução

tempo



**Programas e dados
na mesma memória**
Código e dados compondo um único objeto.

**Programa com
dados
armazenados**

Evolução

tempo



Sistemas de Arquivos

Surgimento do Disco.


Código e dados armazenados de forma distinta.

Dados organizados como coleções de itens relacionados, compondo arquivos.

Código com rotinas para gerência de dados (localização, recuperação e

Armazenamento de dados

Programas e dados na mesma memória



Programa com dados armazenados

Evolução


tempo

Sistemas de Bancos de Dados


Independência dos programas em relação aos dados armazenados
SGBD possui todas as funções necessárias para localização e manipulação dos dados



Sistemas de Arquivos

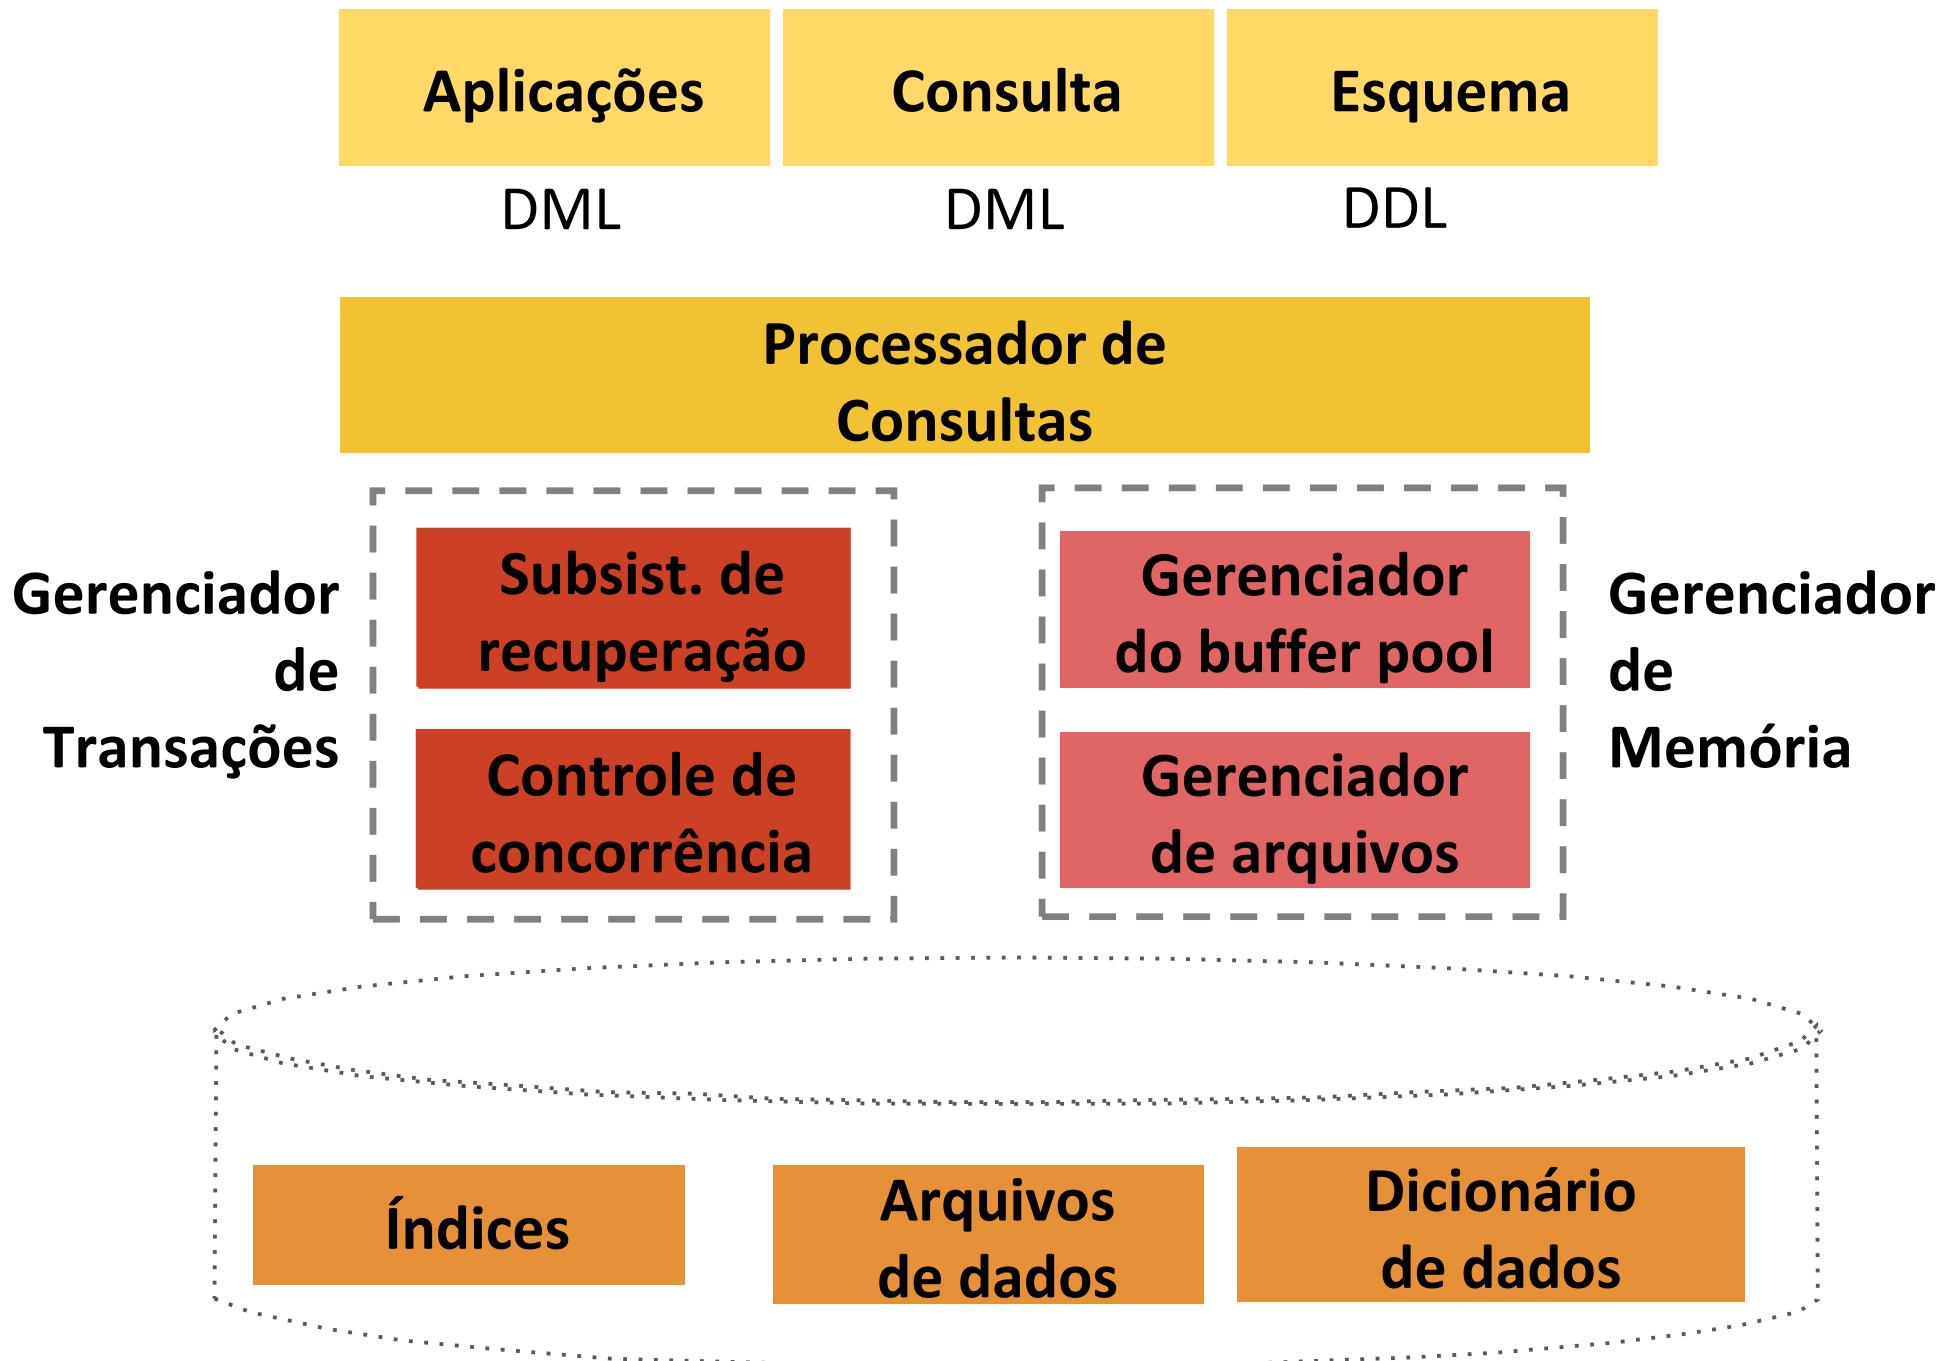


Programas e dados na mesma memória



Sistemas de BDs x Sistemas de Arquivos

- Separação entre programas e dados
- Suporte para múltiplas visões de usuário
- Compartilhamento de dados e processamento multi usuário de transações
- Armazenamento no BD da sua própria descrição ou *esquema*
- Controle para não haver redundância de dados



Arquivos de Dados

- Persistência
- Armazenamento dos objetos do banco de dados, que vão sofrendo modificações com o tempo

Dicionário de Dados

- Armazena os metadados relativos à estrutura do banco de dados.

Índices

- Proporcionam acesso rápido aos dados

Processador de Consultas

- Definição e manipulação de dados
- Funções oferecidas pelas linguagens

- DDL

- processa os comandos de descrição das estruturas do esquema, e a armazena em ***dicionário de dados***

- DML

- Realizar consultas e atualizações, inserções e remoções

- SQL é a linguagem típica de SGBD relacional que engloba uma LDD e uma LMD.

Gerenciador de Transações/ Subsistema de Recuperação

- Recuperação de dados
- Garantir que falhas durante o processamento de transações não sejam propagadas aos objetos persistentes.
 - Os objetos armazenados devem sobreviver a falhas das transações e mesmo algumas falhas de hardware
 - Mecanismo de recuperação/ *recovery*

Gerenciador de Transações/ Controle de Concorrência de Transações

- Garante que transações concorrentes serão executadas sem conflitos em seus procedimentos.
- Técnicas de controle de concorrência

Gerenciador de Memória/ Gerenciador de Arquivos

- Gerencia a alocação de espaço no armazenamento em disco e as estruturas de dados usadas para representar estas informações armazenadas em disco

Gerenciador de Memória/ Gerenciador de Buffer

- Responsável pela intermediação de dados do disco para a memória principal e pela decisão de quais dados colocar em memória cache

Gerenciador de Autorizações e Integridade

- Mantém o BD em estado consistente, satisfazendo algumas condições, chamadas de restrições de integridade.
- Implementa mecanismos de segurança de acesso para consulta, remoção, atualização e inserção de dados
 - Comandos de concessão e revogação (*grant* e *revoke*) de acesso a usuários individuais ou grupos de usuários

Desempenho

- O SGBD deve executar as funções de forma eficaz e eficiente
 - Estruturas de dados
 - Métodos de acesso
 - Técnicas de otimização

Carga, descarga, cópia, restauração

- Facilidades para
 - Carregar e descarregar o banco de dados ou parte deles
 - Fazer cópia de segurança = backup
 - Restaurar o banco de dados a partir do backup

Relational Database

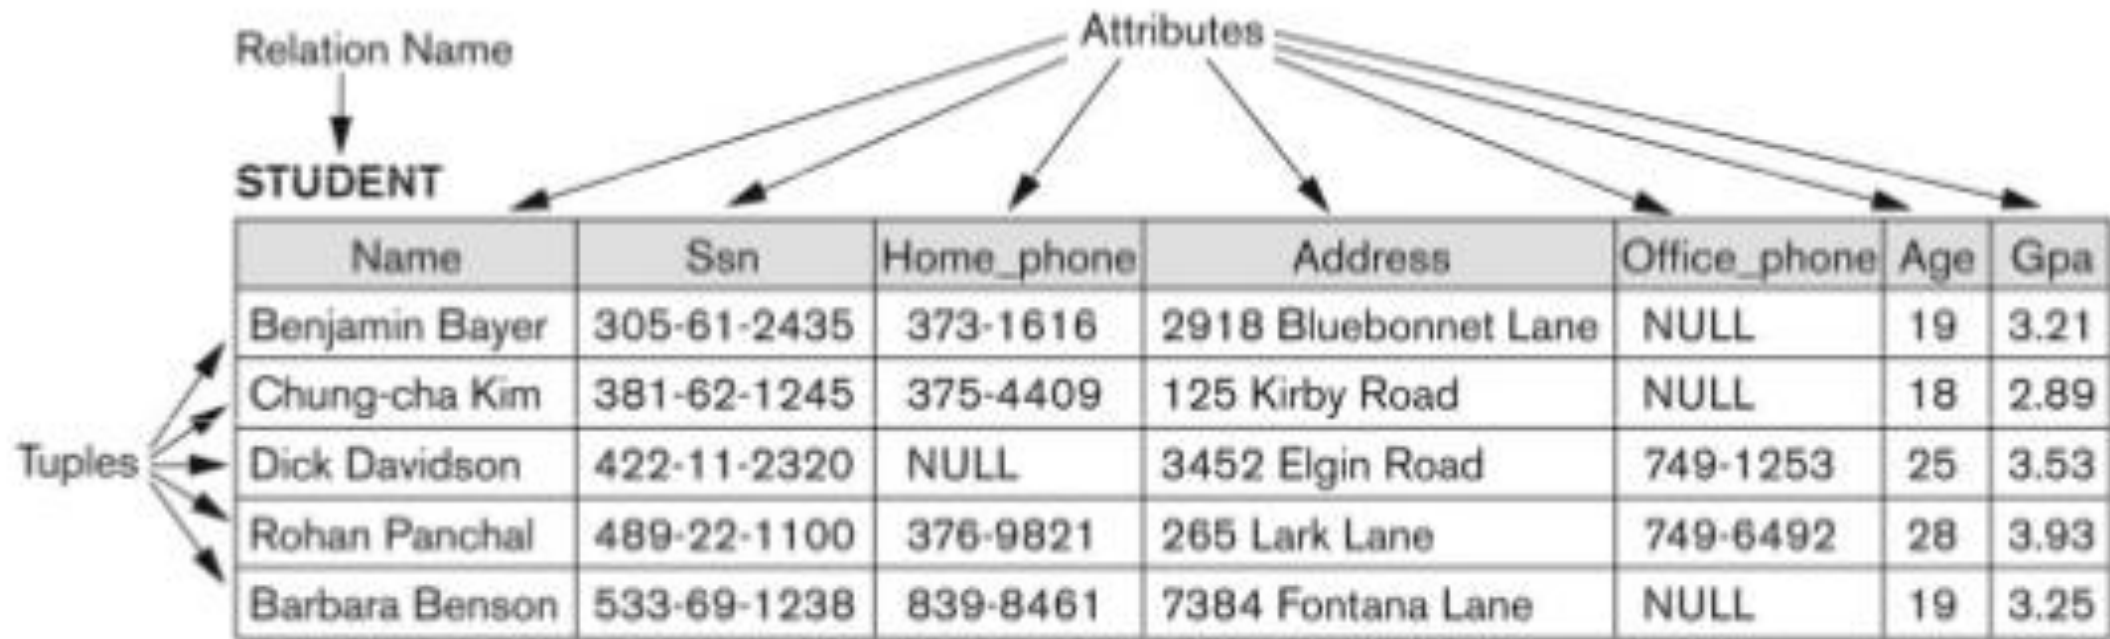


Figure 5.1

The attributes and tuples of a relation STUDENT.

STUDENT

| | | | |
|------|----------------|-------|-------|
| Name | Student_number | Class | Major |
|------|----------------|-------|-------|

COURSE

| | | | |
|-------------|---------------|--------------|------------|
| Course_name | Course_number | Credit_hours | Department |
|-------------|---------------|--------------|------------|

PREREQUISITE

| | |
|---------------|---------------------|
| Course_number | Prerequisite_number |
|---------------|---------------------|

SECTION

| | | | | |
|--------------------|---------------|----------|------|------------|
| Section_identifier | Course_number | Semester | Year | Instructor |
|--------------------|---------------|----------|------|------------|

GRADE_REPORT

| | | |
|----------------|--------------------|-------|
| Student_number | Section_identifier | Grade |
|----------------|--------------------|-------|

Figure 2.1

Schema diagram for the database in Figure 1.2.

EMPLOYEE

| | | | | | | | | | |
|-------|-------|-------|------------|-------|---------|-----|--------|-----------|-----|
| Fname | Minit | Lname | <u>Ssn</u> | Bdate | Address | Sex | Salary | Super_ssn | Dno |
|-------|-------|-------|------------|-------|---------|-----|--------|-----------|-----|

DEPARTMENT

| | | | |
|-------|----------------|---------|----------------|
| Dname | <u>Dnumber</u> | Mgr_ssn | Mgr_start_date |
|-------|----------------|---------|----------------|

DEPT_LOCATIONS

| | |
|----------------|------------------|
| <u>Dnumber</u> | <u>Dlocation</u> |
|----------------|------------------|

PROJECT

| | | | |
|-------|----------------|-----------|------|
| Pname | <u>Pnumber</u> | Plocation | Dnum |
|-------|----------------|-----------|------|

WORKS_ON

| | | |
|-------------|------------|-------|
| <u>Essn</u> | <u>Pno</u> | Hours |
|-------------|------------|-------|

DEPENDENT

| | | | | |
|-------------|-----------------------|-----|-------|--------------|
| <u>Essn</u> | <u>Dependent_name</u> | Sex | Bdate | Relationship |
|-------------|-----------------------|-----|-------|--------------|

Figure 5.5
Schema diagram for
the COMPANY
relational database
schema.

| EMPLOYEE | FNAME | MINIT | LNAME | <u>SSN</u> | BDATE | ADDRESS | SEX | SALARY | SUPERSSN | DNO |
|----------|----------|-------|---------|------------|------------|--------------------------|-----|--------|-----------|-----|
| | John | B | Smith | 123456789 | 1965-01-09 | 731 Fondren, Houston, TX | M | 30000 | 333445555 | 5 |
| | Franklin | T | Wong | 333445555 | 1955-12-08 | 638 Voss, Houston, TX | M | 40000 | 888665555 | 5 |
| | Alicia | J | Zelaya | 999887777 | 1968-07-19 | 3321 Castle, Spring, TX | F | 25000 | 987654321 | 4 |
| | Jennifer | S | Wallace | 987654321 | 1941-06-20 | 291 Berry, Bellaire, TX | F | 43000 | 888665555 | 4 |
| | Ramesh | K | Narayan | 666884444 | 1962-09-15 | 975 Fire Oak, Humble, TX | M | 38000 | 333445555 | 5 |
| | Joyce | A | English | 453453453 | 1972-07-31 | 5631 Rice, Houston, TX | F | 25000 | 333445555 | 5 |
| | Ahmad | V | Jabbar | 987987987 | 1969-03-29 | 980 Dallas, Houston, TX | M | 25000 | 987654321 | 4 |
| | James | E | Borg | 888665555 | 1937-11-10 | 450 Stone, Houston, TX | M | 55000 | null | 1 |

| DEPT_LOCATIONS | <u>DNUMBER</u> | <u>DLOCATION</u> |
|----------------|----------------|------------------|
| | 1 | Houston |
| | 4 | Stafford |
| | 5 | Bellaire |
| | 5 | Sugarland |
| | 5 | Houston |

| DEPARTMENT | <u>DNAME</u> | <u>DNUMBER</u> | <u>MGRSSN</u> | <u>MGRSTARTDATE</u> |
|------------|----------------|----------------|---------------|---------------------|
| | Research | 5 | 333445555 | 1988-05-22 |
| | Administration | 4 | 987654321 | 1995-01-01 |
| | Headquarters | 1 | 888665555 | 1981-06-19 |

| WORKS_ON | <u>ESSN</u> | <u>PNO</u> | HOURS |
|----------|-------------|------------|-------|
| | 123456789 | 1 | 32.5 |
| | 123456789 | 2 | 7.5 |
| | 666884444 | 3 | 40.0 |
| | 453453453 | 1 | 20.0 |
| | 453453453 | 2 | 20.0 |
| | 333445555 | 2 | 10.0 |
| | 333445555 | 3 | 10.0 |
| | 333445555 | 10 | 10.0 |
| | 333445555 | 20 | 10.0 |
| | 999887777 | 30 | 30.0 |
| | 999887777 | 10 | 10.0 |
| | 987987987 | 10 | 35.0 |
| | 987987987 | 30 | 5.0 |
| | 987654321 | 30 | 20.0 |
| | 987654321 | 20 | 15.0 |
| | 888665555 | 20 | null |

| PROJECT | <u>PNAME</u> | <u>PNUMBER</u> | <u>PLOCATION</u> | <u>DNUM</u> |
|---------|-----------------|----------------|------------------|-------------|
| | ProductX | 1 | Bellaire | 5 |
| | ProductY | 2 | Sugarland | 5 |
| | ProductZ | 3 | Houston | 5 |
| | Computerization | 10 | Stafford | 4 |
| | Reorganization | 20 | Houston | 1 |
| | Newbenefits | 30 | Stafford | 4 |

SELECT

Query: Retrieve the birthdate and address of the employee whose name is 'John'.

```
SELECT    BDATE, ADDRESS  
FROM      EMPLOYEE  
WHERE     FNAME='John'
```

INSERT

INSERT INTO EMPLOYEE
VALUES

```
('Richard','K','Marini',  
  '653298653', '30-DEC-52',  
  '98 Oak Forest,Katy,TX',  
  'M',  
  37000,  
  '987654321',  
  4 )
```


DELETE

DELETE FROM EMPLOYEE

WHERE LNAME='Brown'

UPDATE

```
UPDATE    PROJECT
SET       PLOCATION = 'Bellaire',      DNUM = 5
WHERE     PNUMBER=10
```

1980

1990



2000


2010

**Relational
Dominance**











Google



BigTable
e

amazon



Dynamo



<http://nosql-database.org/>

noSQL Database

NoSQL is an accidental neologism.

There is no prescriptive definition—all you can make is an observation of common characteristics.

- Flexible schema
- Quicker/cheaper to set up
- Massive scalability
- Relaxed consistency → higher performance & availability
- Cluster friendly

Not every data management/analysis problem is best solved **exclusively** using a traditional relational DBMS

“NoSQL” = “Not Only SQL”

- “SQL” = Traditional Relational DBMS
- Recognition over past decade

Not every data management/analysis problem is best solved using a traditional relational DBMS

- “No SQL” = Not using traditional relational DBMS
- “No SQL” ≠ Don’t use SQL language

Scalability of NoSQL Database vs Traditional Relational Database




Image Credit: DataJobs.com

NoSQL Data Models

KEY-

VALUE



COLUMN



DOCUMENT



GRAPH



mongoDB

| Relational | mongoDB |
|------------|--------------------------------|
| Database | Database |
| Table | Collection |
| Row | Document |
| Index | Index |
| Join | Embedded Document or Reference |

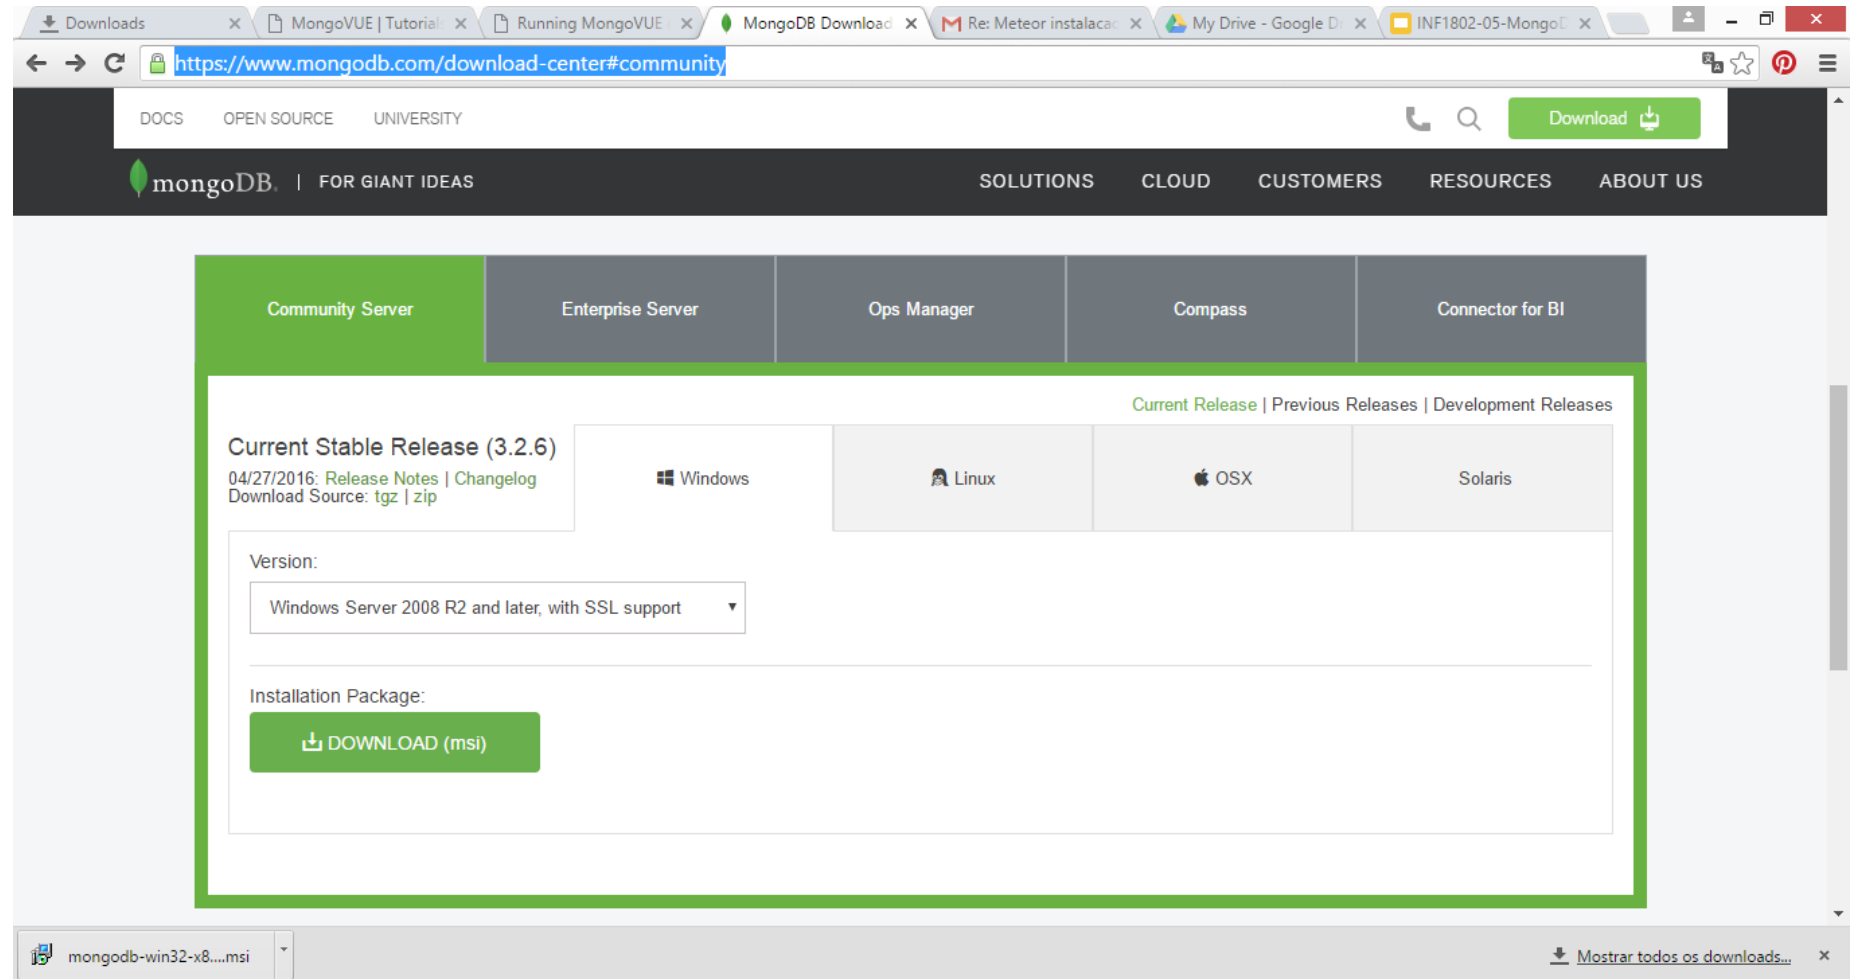
Collection

```
{    _id: <ObjectId1>,
      fname: "John",
      minit: "B",
      lname: "Smith",
      ssn: 123456789,
      address: "731 Houston, TX",
      sex: "M",
      salary: 30000,
      superssn: <ObjectId2>,
      dno: <ObjectId3>
}
```

Database Installation

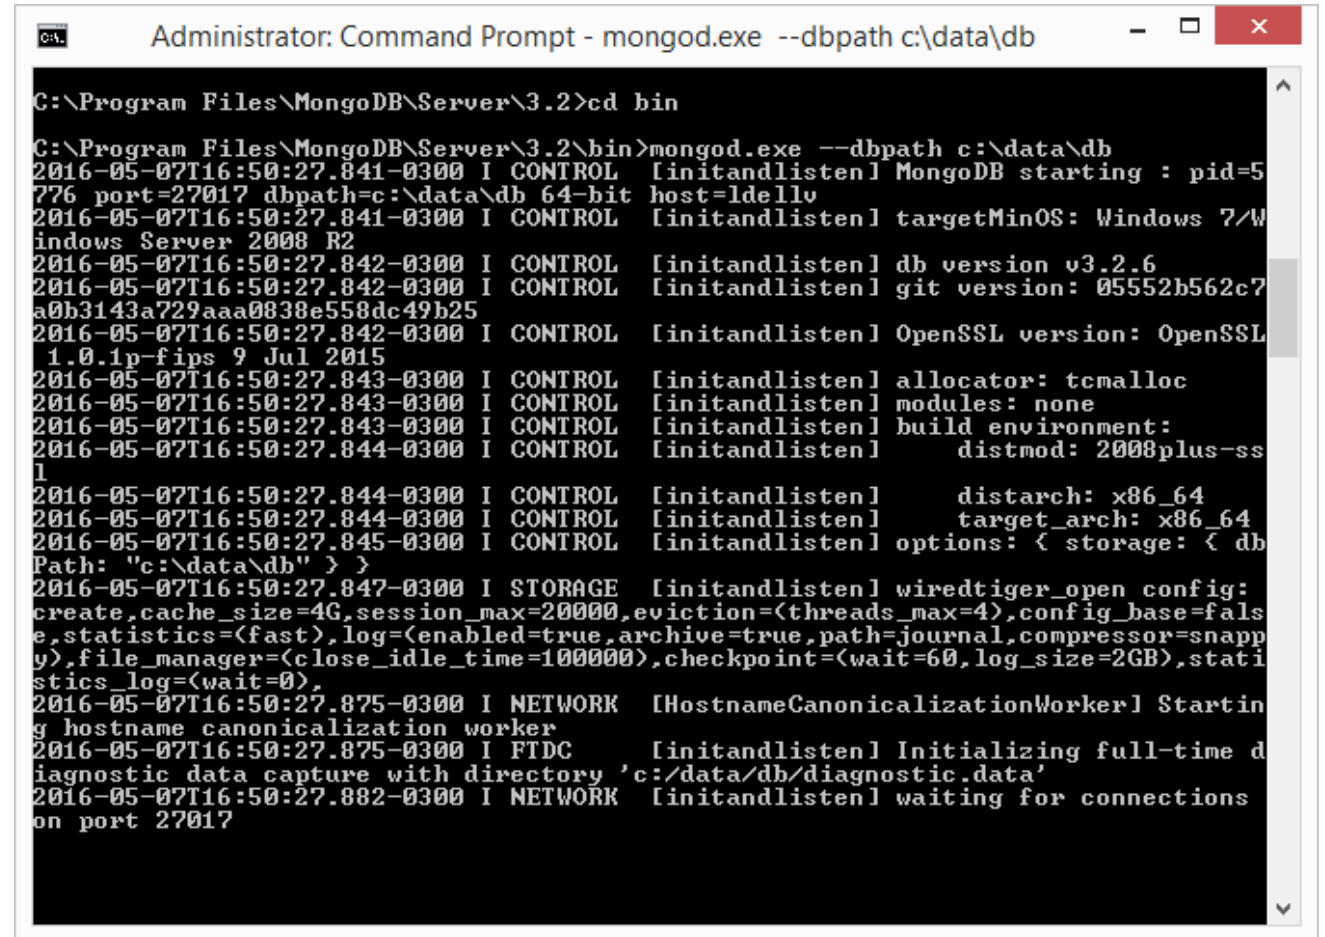
Step 1 - MongoDB Download and Installation (msi file, double click to install)

<https://www.mongodb.com/download-center#community>



Step 2 - Run the MongoDB server

```
C:\mongodb\bin\mongod.exe --dbpath c:\data\db
```



```
Administrator: Command Prompt - mongod.exe --dbpath c:\data\db
C:\Program Files\MongoDB\Server\3.2>cd bin
C:\Program Files\MongoDB\Server\3.2\bin>mongod.exe --dbpath c:\data\db
2016-05-07T16:50:27.841-0300 I CONTROL [initandlisten] MongoDB starting : pid=5776 port=27017 dbpath=c:\data\db 64-bit host=ldellv
2016-05-07T16:50:27.841-0300 I CONTROL [initandlisten] targetMinOS: Windows 7/Windows Server 2008 R2
2016-05-07T16:50:27.842-0300 I CONTROL [initandlisten] db version v3.2.6
2016-05-07T16:50:27.842-0300 I CONTROL [initandlisten] git version: 05552b562c7a0b3143a729aaa0838e558dc49b25
2016-05-07T16:50:27.842-0300 I CONTROL [initandlisten] OpenSSL version: OpenSSL 1.0.1p-fips 9 Jul 2015
2016-05-07T16:50:27.843-0300 I CONTROL [initandlisten] allocator: tcmalloc
2016-05-07T16:50:27.843-0300 I CONTROL [initandlisten] modules: none
2016-05-07T16:50:27.843-0300 I CONTROL [initandlisten] build environment:
2016-05-07T16:50:27.844-0300 I CONTROL [initandlisten] distmod: 2008plus-ssl
2016-05-07T16:50:27.844-0300 I CONTROL [initandlisten] distarch: x86_64
2016-05-07T16:50:27.844-0300 I CONTROL [initandlisten] target_arch: x86_64
2016-05-07T16:50:27.845-0300 I CONTROL [initandlisten] options: { storage: { dbPath: "c:\data\db" } }
2016-05-07T16:50:27.847-0300 I STORAGE [initandlisten] wiredtiger_open config: create,cache_size=4G,session_max=20000,eviction=(threads_max=4),config_base=false,statistics=(fast),log=(enabled=true,archive=true,path=journal,compressor=snappy),file_manager=(close_idle_time=100000),checkpoint=(wait=60,log_size=2GB),statistics_log=(wait=0),
2016-05-07T16:50:27.875-0300 I NETWORK [HostnameCanonicalizationWorker] Starting hostname canonicalization worker
2016-05-07T16:50:27.875-0300 I FTDC [initandlisten] Initializing full-time diagnostic data capture with directory 'c:\data\db\diagnostic.data'
2016-05-07T16:50:27.882-0300 I NETWORK [initandlisten] waiting for connections on port 27017
```

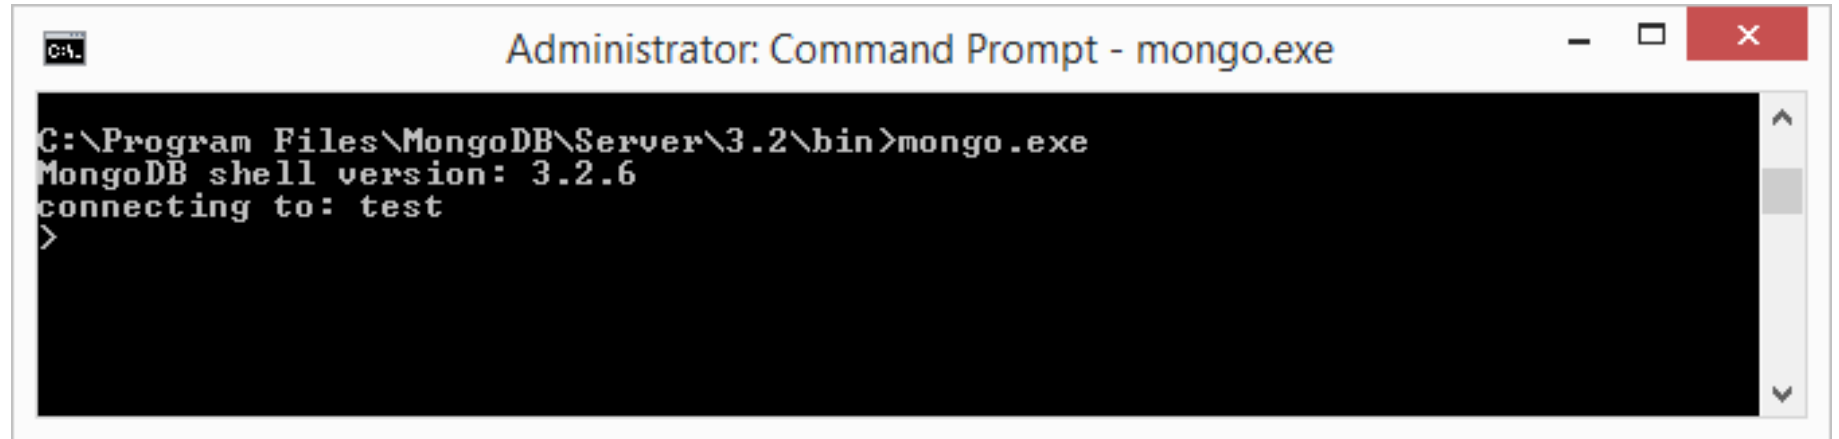
From

<https://docs.mongodb.com/manual/tutorial/install-mongodb-enterprise-on-windows/>

<http://www.mkyong.com/mongodb/how-to-install-mongodb-on-windows/>

Step 3 - Connect to MongoDB Server
through the mongo.exe shell, open another Command Prompt.

C:\mongodb\bin\mongo.exe



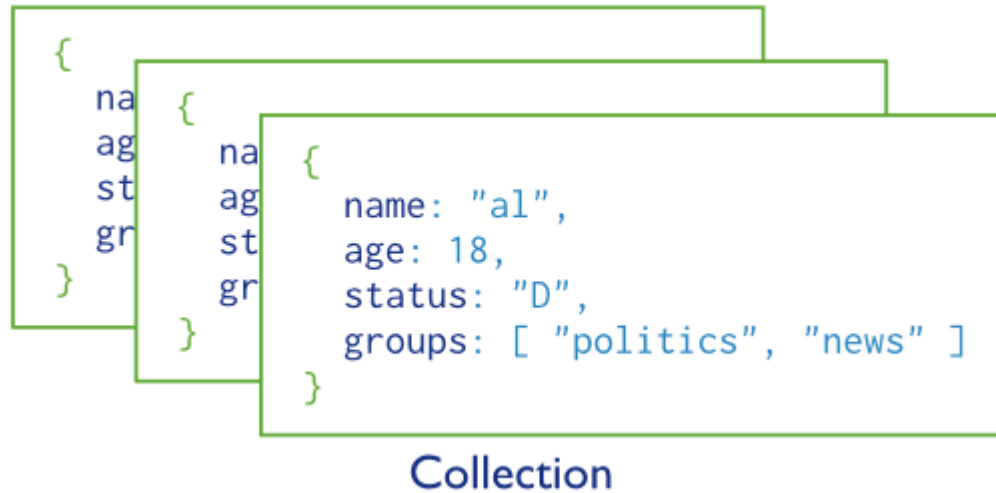
```
C:\Program Files\MongoDB\Server\3.2\bin>mongo.exe
MongoDB shell version: 3.2.6
connecting to: test
>
```

Data Model

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

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

MongoDB stores data in the form of *documents*, which are JSON-like field and value pairs. Documents are analogous to structures in programming languages that associate keys with values (e.g. dictionaries, hashes, maps, and associative arrays).



MongoDB stores all documents in [collections](#). A collection is a group of related documents that have a set of shared common indexes. Collections are analogous to a table in relational databases.

Database Operations

<https://docs.mongodb.com/manual/core/crud-introduction/>

Query

Query

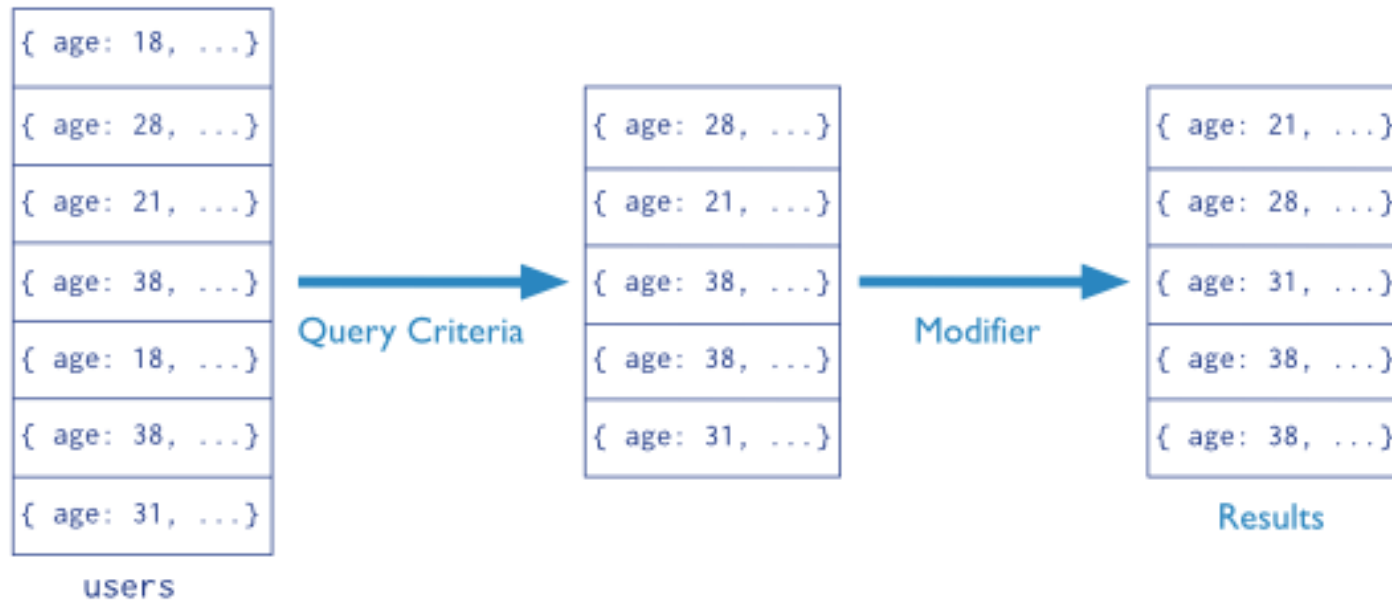
In MongoDB a query targets a specific collection of documents. Queries specify criteria, or conditions, that identify the documents that MongoDB returns to the clients.

```
db.users.find()
```

```
select *  
from users
```

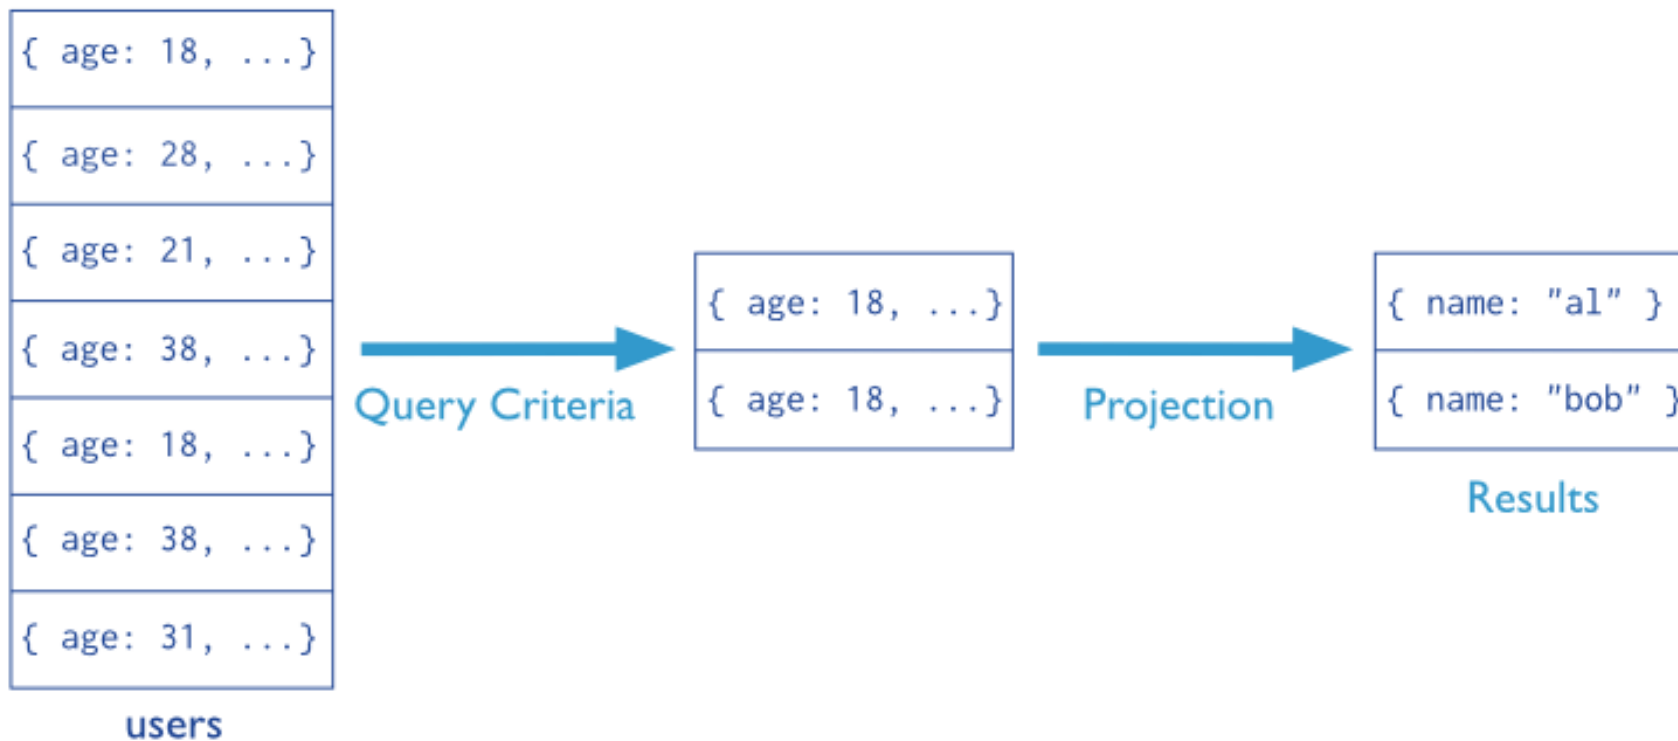
Query - You can optionally modify queries to impose limits, skips, and sort orders.

Collection Query Criteria Modifier
`db.users.find({ age: { $gt: 18 } }).sort({age: 1 })`



Query - A query may include a *projection* that specifies the fields from the matching documents to return.

Collection Query Criteria Projection
`db.users.find({ age: 18 }, { name: 1, _id: 0 })`



```
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

Data Modification

Operations that Create, Update or Delete Data

<https://docs.mongodb.com/manual/core/write-operations-introduction/>


```
Collection
↓
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, ... } |

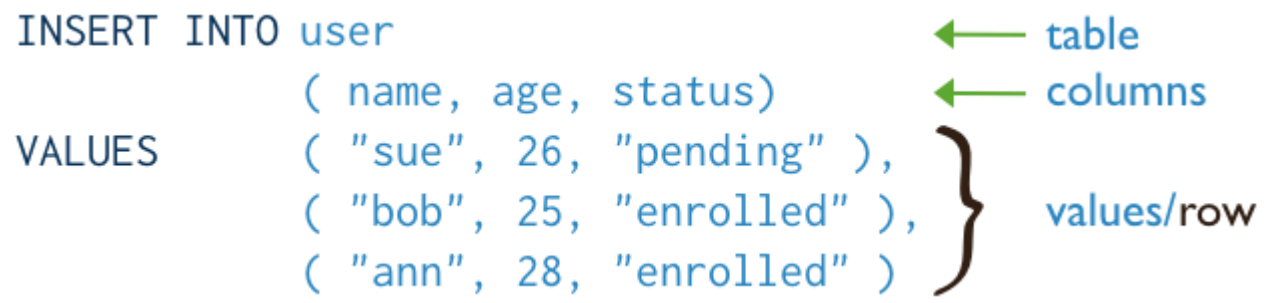
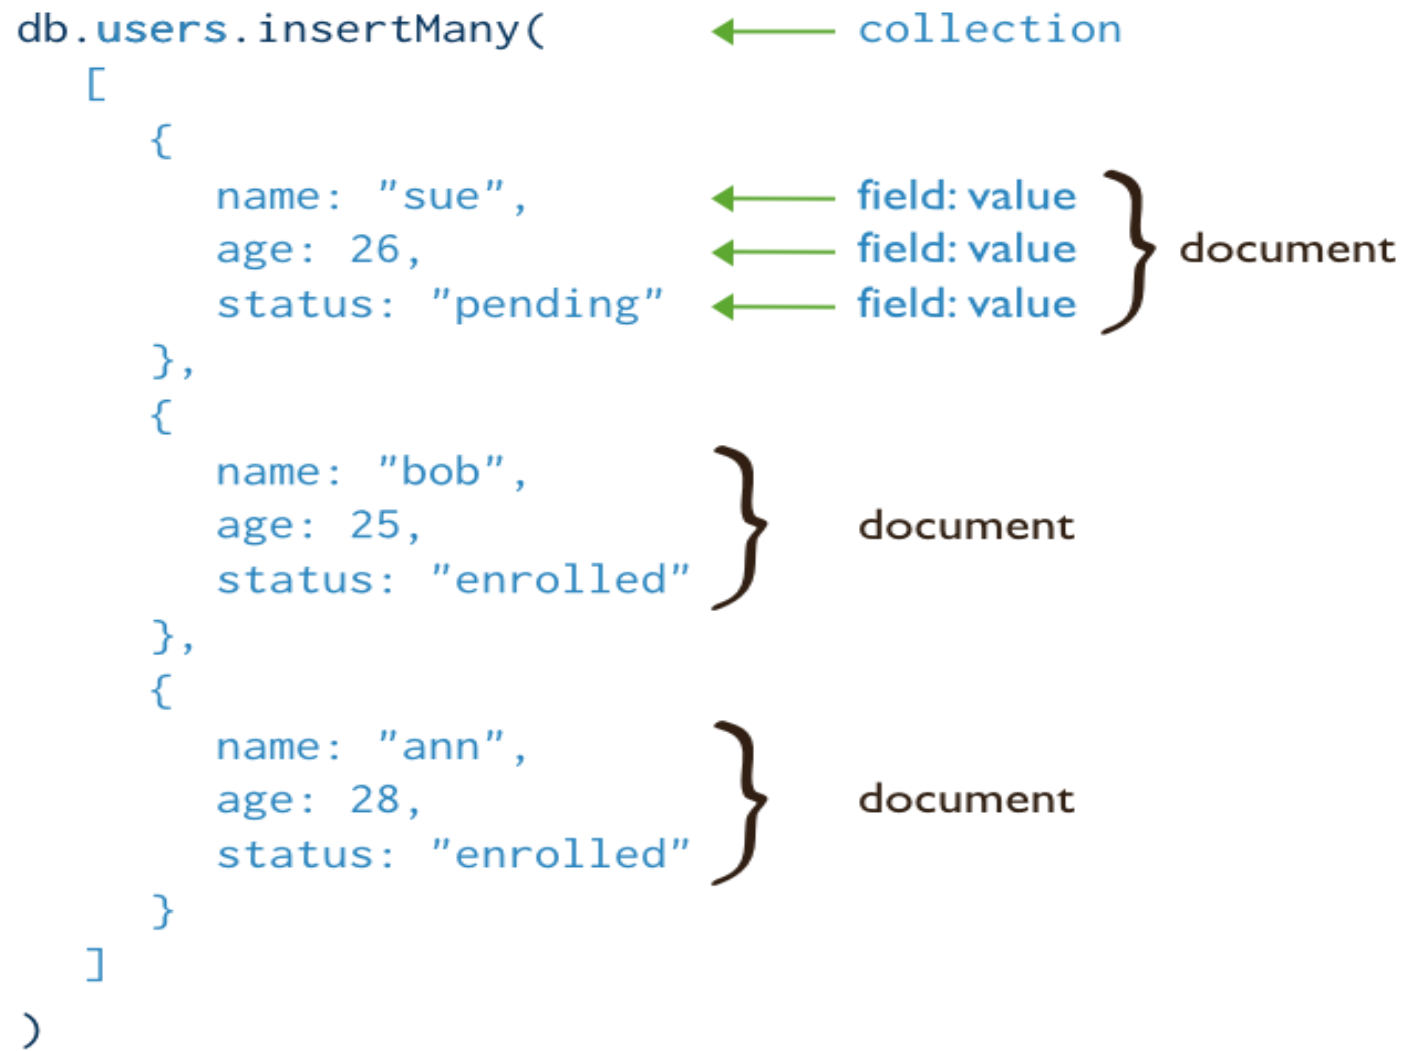
users

```
INSERT INTO users
      ( name, age, status )
VALUES ( "sue", 26, "A" )
```

← table
← columns
← values/row

```
db.users.insertOne( ← collection
  {
    name: "sue", ← field: value
    age: 26, ← field: value
    status: "pending" ← field: value } document
  }
)
```

```
INSERT INTO users ← table
      ( name, age, status ) ← columns
VALUES ( "sue", 26, "A" ) ← values/row
```



```
db.users.update(
  { age: { $gt: 18 } },
  { $set: { status: "A" } },
  { multi: true }
)
```

← collection
← update criteria
← update action
← update option

```
UPDATE users
SET status = 'A'
WHERE age > 18
```

← table
← update action
← update criteria

```
db.users.updateOne(           ← collection
  { age : { $lt : 18 } } ,    ← update filter
  { $set: { status : "reject" } ← update action
)
```

```
UPDATE users                 ← table
SET   status = 'reject'     ← update action
WHERE age < 18              ← update filter
LIMIT 1                     ← update limit
```

```
db.users.updateMany(
  { age: { $lt: 18 } },
  { $set: { status: "reject" } }
)
```

← collection
← update filter
← update action

```
UPDATE users
SET    status = 'reject'
WHERE age < 18
```

← table
← update action
← update filter

```
db.users.remove(      ← collection  
  { status: "D" }    ← remove criteria  
)
```

```
DELETE FROM users    ← table  
WHERE status = 'D'  ← delete criteria
```

```
db.users.deleteOne(
  { status: "Rejected" }
)
```

← collection
← delete filter

```
DELETE FROM users
WHERE      status = 'reject'
LIMIT     1
```

← table
← delete filter
← delete limit

```
db.users.deleteMany(
  { status: "Rejected" }
)
```

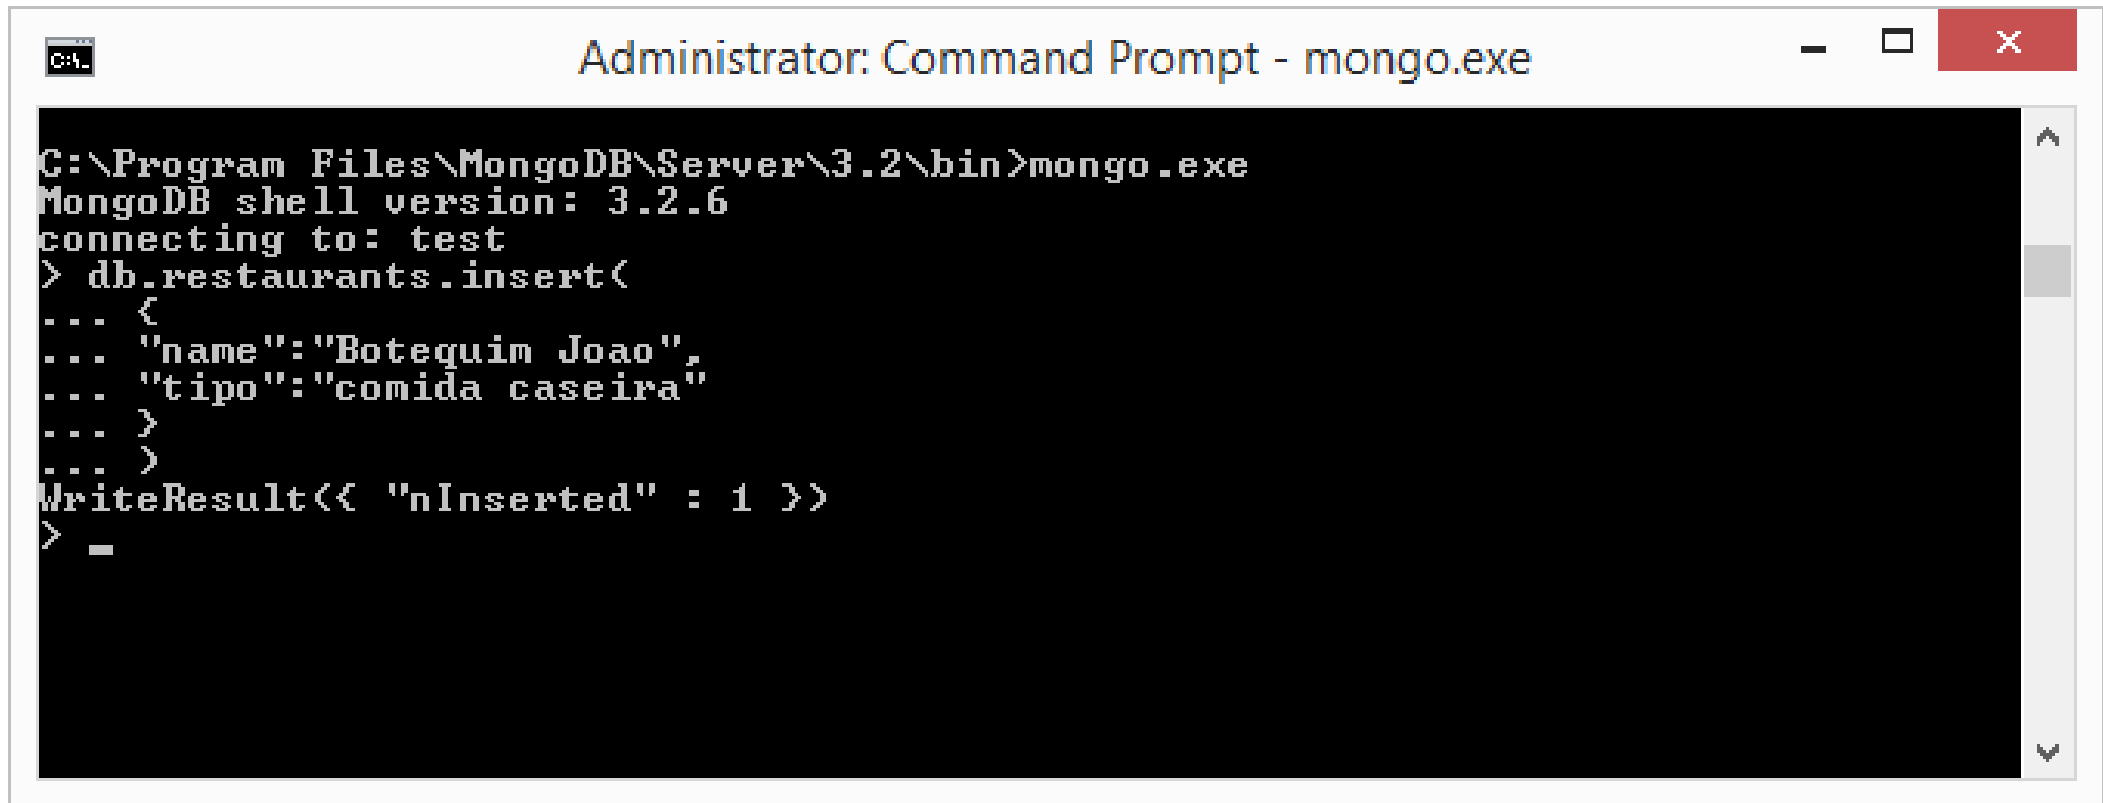
← collection
← delete filter

```
DELETE FROM users
WHERE      status = 'reject'
```

← table
← delete filter

Step 4 - Insert a document into a collection named restaurants. The operation will create the collection if the collection does not currently exist.

```
db.restaurants.insert
```



```
C:\Program Files\MongoDB\Server\3.2\bin>mongo.exe
MongoDB shell version: 3.2.6
connecting to: test
> db.restaurants.insert<
... <
...   "name": "Botequim Joao",
...   "tipo": "comida caseira"
... >
... >
WriteResult<< "nInserted" : 1 >>
> -
```

Step 5 - Query for All Documents in a Collection. To return all documents in a collection, call the find() method without a criteria document.

```
db.restaurants.find()
```



```
Administrator: Command Prompt - mongo.exe
> db.restaurants.find()
{ "_id" : ObjectId("572e4c431d569630bfd9890e"), "name" : "Botequim Joao", "tipo"
: "comida caseira" }
> -
```

<https://docs.mongodb.com/getting-started/shell/query/>

Step 6 - Specify Equality Conditions.

```
db.restaurants.find({<field1>: <value1>, <field2>: <value2>, ...})
```

```
> db.restaurants.insert<
... {
...   "name": "Dominos",
...   "tipo": "pizzaria"
... }
... >
WriteResult<< "nInserted" : 1 >>
>
>
> db.restaurants.insert<
... {
...   "name": "Gula Gula",
...   "tipo": "variada"
... }
... >
WriteResult<< "nInserted" : 1 >>
>
>
> db.restaurants.find<
< "_id" : ObjectId("572e4c431d569630bfd9890e"), "name" : "Botequim Joao", "tipo"
: "comida caseira" >
< "_id" : ObjectId("572e50401d569630bfd9890f"), "name" : "Dominos", "tipo" : "pi
zzaria" >
< "_id" : ObjectId("572e506c1d569630bfd98910"), "name" : "Gula Gula", "tipo" : "
variada" >
>
>
> db.restaurants.find<<"name": "Dominos">>
< "_id" : ObjectId("572e50401d569630bfd9890f"), "name" : "Dominos", "tipo" : "pi
zzaria" >
```

Relationships, Joins

<https://docs.mongodb.com/manual/core/write-operations-introduction/>

No Joins

The first and most fundamental difference that you'll need to get comfortable with is MongoDB's lack of joins.

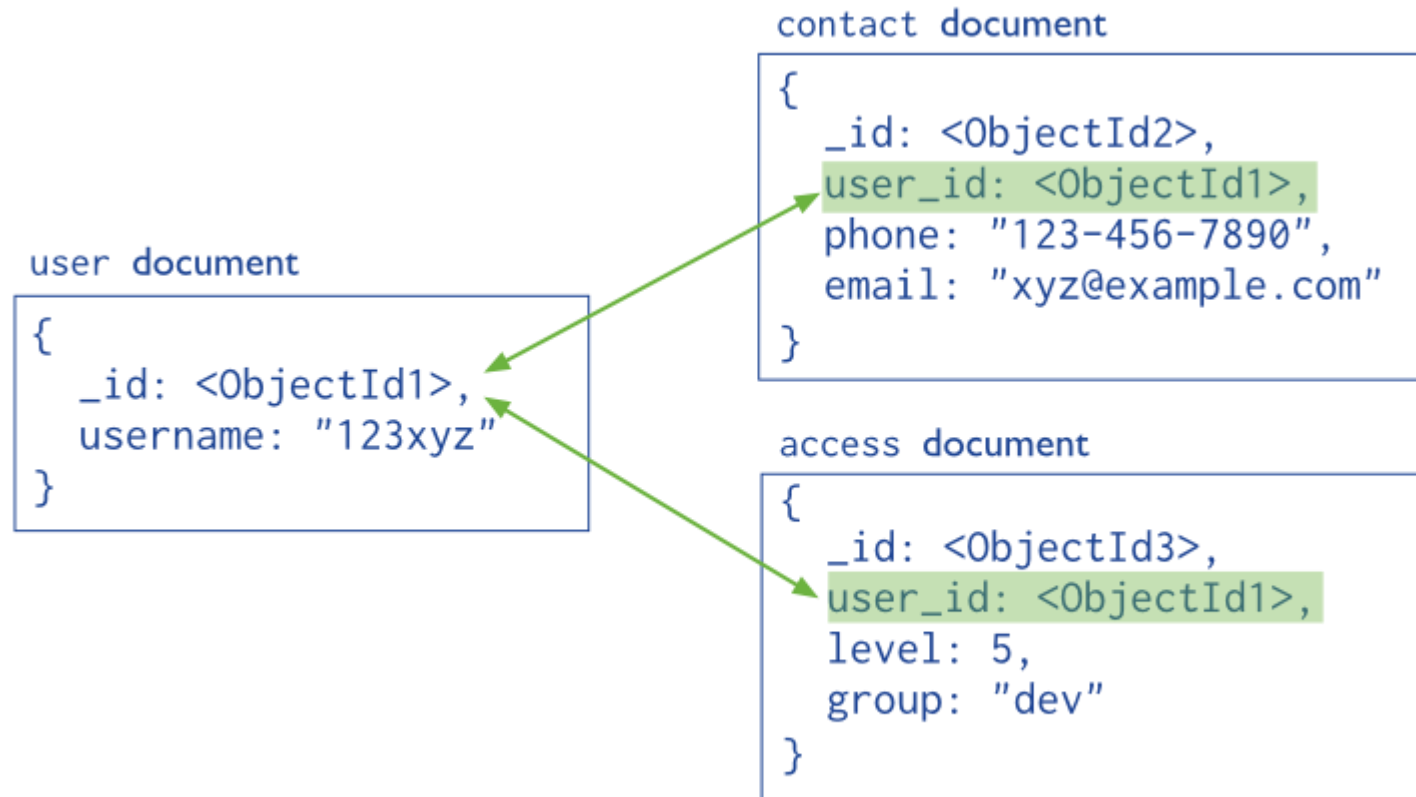
Normalized Data Models

In general, use normalized data models:

when embedding would result in duplication of data but would not provide sufficient read performance advantages to outweigh the implications of the duplication.

to represent more complex many-to-many relationships.

to model large hierarchical data sets.



Embedded Document - Relationship 1:1



Embedded Document - Relationship 1:n

```
{
  _id: "joe",
  name: "Joe Bookreader",
  addresses: [
    {
      street: "123 Fake Street",
      city: "Faketon",
      state: "MA",
      zip: "12345"
    },
    {
      street: "1 Some Other Street",
      city: "Boston",
      state: "MA",
      zip: "12345"
    }
  ]
}
```


Document Reference - Relationship 1:n

```
{
  title: "MongoDB: The Definitive Guide",
  author: [ "Kristina Chodorow", "Mike Dirolf" ],
  published_date: ISODate("2010-09-24"),
  pages: 216,
  language: "English",
  publisher: {
    name: "O'Reilly Media",
    founded: 1980,
    location: "CA"
  }
}
```

```
{
  name: "O'Reilly Media",
  founded: 1980,
  location: "CA",
  books: [12346789, 234567890, ...]
}

{
  _id: 123456789,
  title: "MongoDB: The Definitive Guide",
  author: [ "Kristina Chodorow", "Mike Dirolf" ],
  published_date: ISODate("2010-09-24"),
  pages: 216,
  language: "English"
}

{
  _id: 234567890,
  title: "50 Tips and Tricks for MongoDB Developer",
  author: "Kristina Chodorow",
  published_date: ISODate("2011-05-06"),
  pages: 68,
  language: "English"
}
```

Model One-to-One Relationships with Embedded Documents

<https://docs.mongodb.com/manual/tutorial/model-embedded-one-to-one-relationships-between-documents/#data-modeling-example-one-to-one>

Model One-to-Many Relationships with Embedded Documents

<https://docs.mongodb.com/manual/tutorial/model-embedded-one-to-many-relationships-between-documents/#data-modeling-example-one-to-many>

Embedding

- For 1:1 or 1:many
- Document limit to 16MB, consider document growth

Referencing

- `_id` field is referenced in the related document
- Application runs 2nd query to retrieve the data
- Data duplication vs performance gain
- Object referenced by many different sources
- Models complex Many:Many & hierarchical structures

Using a mongoDB Client

Suggestion: <http://www.mongoclient.com/>

MongoClient

Step 1 - Download zip file

from <https://github.com/rsercano/mongoclient>

The screenshot shows a web browser window with the GitHub repository page for `rsercano/mongoclient`. The browser's address bar shows the URL `https://github.com/rsercano/mongoclient`. The page header includes the GitHub logo, navigation links (Personal, Open source, Business, Explore), and a search bar. The repository name `rsercano / mongoclient` is displayed, along with statistics for Watch (25), Star (258), and Fork (29).


Below the repository name, there are buttons for "New pull request", "New file", "Find file", and "Download ZIP". The "Download ZIP" button is highlighted. Below these buttons, a commit message is shown: "rsercano added notifications for #73" with the latest commit hash `8db4933` and the time "19 hours ago".

A table of files and folders is displayed, showing the repository structure:

| File/Folder | Description | Time |
|-------------------------|-----------------------------|--------------|
| <code>.electrify</code> | Create command | 7 days ago |
| <code>.meteor</code> | electron staff | 7 days ago |
| <code>.scripts</code> | chimp args | 2 months ago |
| <code>client</code> | added notifications for #73 | 19 hours ago |
| <code>lib</code> | resolves #53 | a day ago |
| <code>packages</code> | under heavy development #53 | 2 days ago |
| <code>public</code> | electron staff | 7 days ago |

At the bottom of the browser window, a download bar is visible, showing the downloaded file `mongoclient-master.zip`. Other files in the download bar include `windows-portable-x6...zip` and `mongodb-win32-x8...msi`.

Step 2 - Unzip it. Example
C:\mongoclient-master.



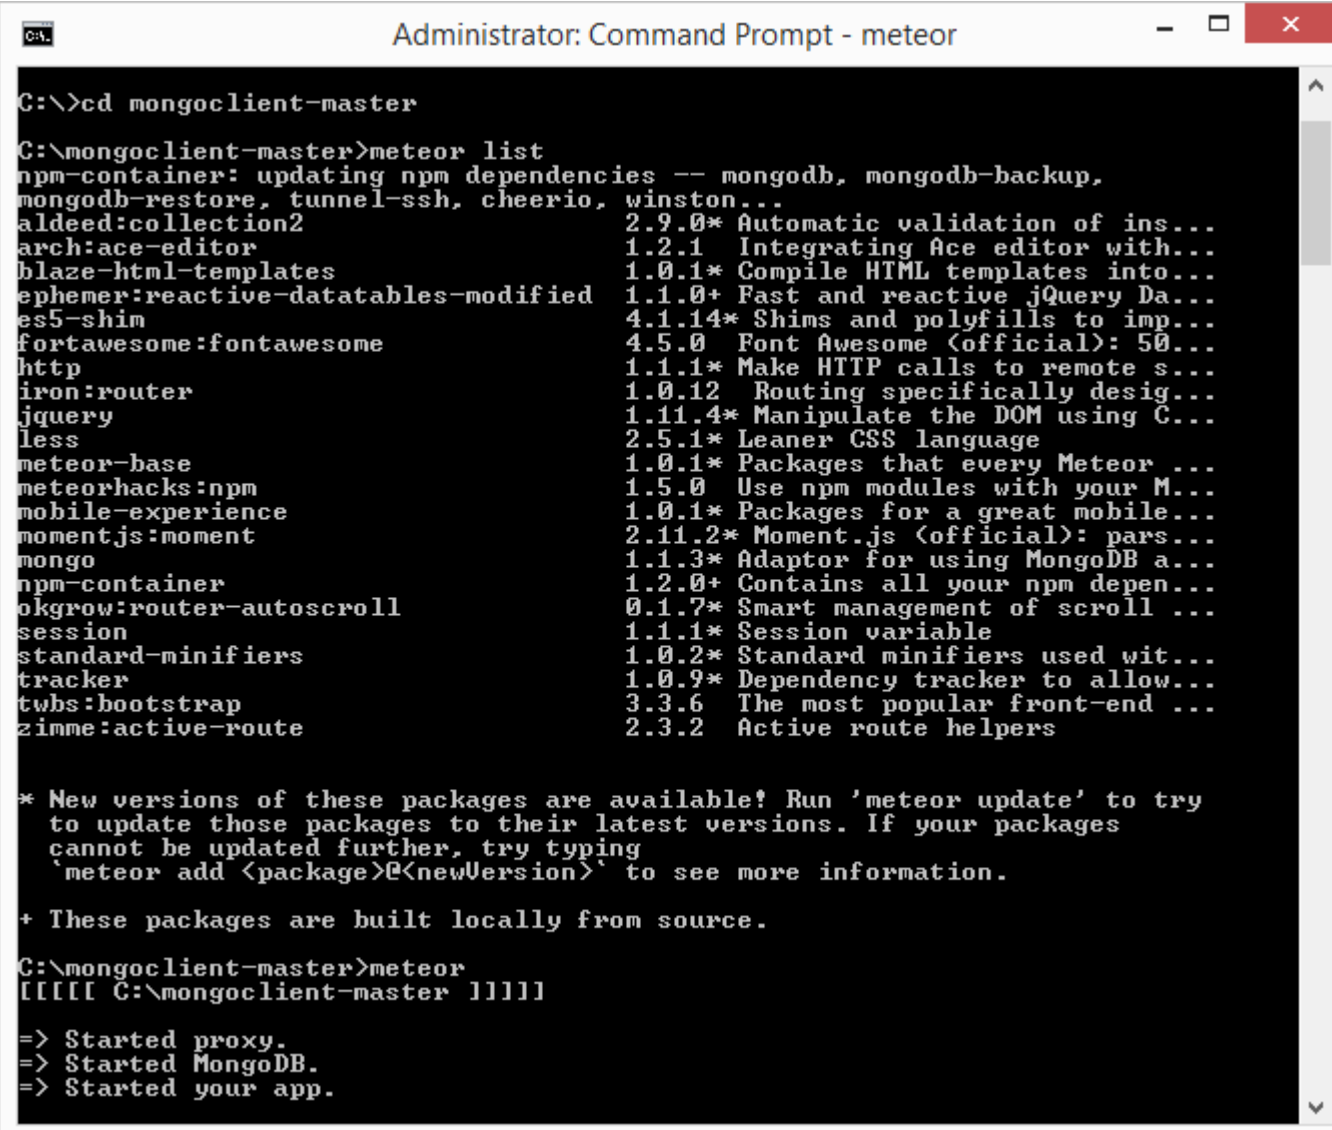
Step 3 - Pre-requisite:
Meteor installed.

Open Command Prompt.

In the mongoclient-master
folder:

1.Run the command
meteor list to install the
application mongoclient.

2.Run the command
meteor to start the
application



```
C:\>cd mongoclient-master

C:\mongoclient-master>meteor list
npm-container: updating npm dependencies -- mongodb, mongodb-backup,
mongodb-restore, tunnel-ssh, cheerio, winston...
aldeed:collection2          2.9.0* Automatic validation of ins...
arch:ace-editor             1.2.1 Integrating Ace editor with...
blaze-html-templates        1.0.1* Compile HTML templates into...
ephemer:reactive-datatables-modified 1.1.0+ Fast and reactive jQuery Da...
es5-shim                    4.1.14* Shims and polyfills to imp...
fontawesome:fontawesome    4.5.0 Font Awesome (official): 50...
http                        1.1.1* Make HTTP calls to remote s...
iron:router                 1.0.12 Routing specifically desig...
jquery                      1.11.4* Manipulate the DOM using C...
less                        2.5.1* Leaner CSS language
meteor-base                 1.0.1* Packages that every Meteor ...
meteorhacks:npm            1.5.0 Use npm modules with your M...
mobile-experience          1.0.1* Packages for a great mobile...
momentjs:moment            2.11.2* Moment.js (official): pars...
mongo                      1.1.3* Adaptor for using MongoDB a...
npm-container              1.2.0+ Contains all your npm depen...
okgrow:router-autoscroll   0.1.7* Smart management of scroll ...
session                    1.1.1* Session variable
standard-minifiers         1.0.2* Standard minifiers used wit...
tracker                    1.0.9* Dependency tracker to allow...
twbs:bootstrap             3.3.6 The most popular front-end ...
zimme:active-route         2.3.2 Active route helpers

* New versions of these packages are available! Run 'meteor update' to try
to update those packages to their latest versions. If your packages
cannot be updated further, try typing
'meteor add <package>@<newVersion>' to see more information.

+ These packages are built locally from source.

C:\mongoclient-master>meteor
[[[[[ C:\mongoclient-master ]]]]]

=> Started proxy.
=> Started MongoDB.
=> Started your app.
```

Step 4 - Open the browser.

<http://localhost:3000/>

1. Click Connection

2. Add Connection

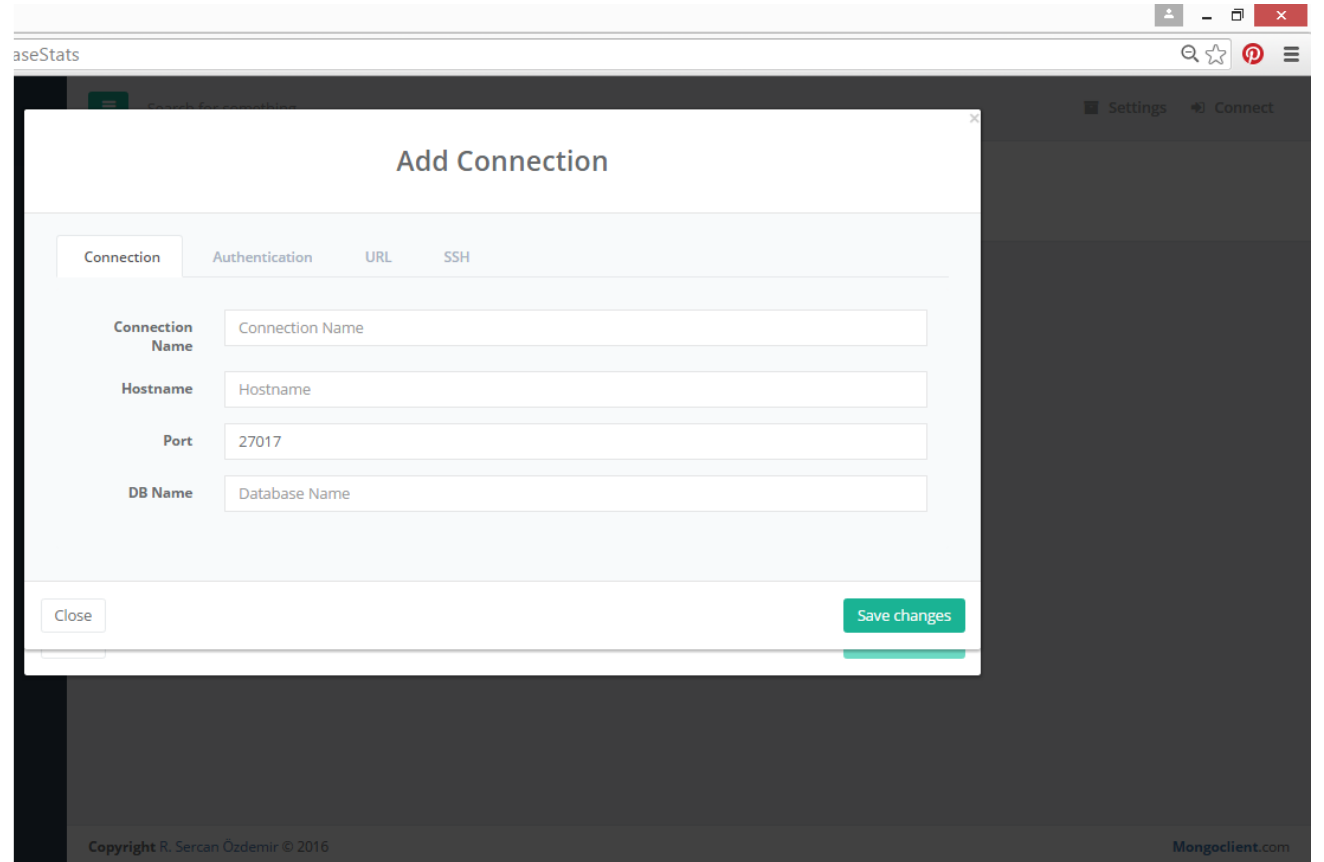
Connection Name: inf1802
(you can choose)

Hostname: localhost

Port: 27017

DB Name: test

1. Connect



INF1802
localhost : 27017
test ▾

- Database Stats
- Management <
- Collections **Add** ▾
 - restaurants **Drop**

feedback

Browse

Home / Browse

Idellv

The server is up for 1103 seconds

The performances which are shown here are coming from **serverStatus** and **stats** queries. Therefore, before having a look at below graphs consider **mongoclient** is already executing **serverStatus** at regular intervals. See [here](#) for more info

Mongodb version is **3.2.6**
Process id is **2036**
Process name is **mongod.exe**

| | |
|--------------------|-----------------------|
| 1 Collections | 3 Objects |
| 0.00 MBs Data Size | 0.04 MBs Storage Size |
| 1 Indexes | 0.04 MBs Index Size |
| 0.00 MBs File Size | 0 Extents |

Memory Usage



The screenshot shows a web browser window with the URL `localhost:3000/browseCollection`. The page title is "MongoClient | Featured M". The browser's address bar contains the URL. The page has a navigation bar with a menu icon, "Refresh Collections", "Settings", and "Disconnect". The main content area displays "restaurants" and "Home / restaurants". On the right side, there are statistics for the collection: "Count: 3", "Index Count: 1", "Size: 0.00 MBs", and "Total Index Size: 0.04 MBs". A vertical scrollbar is visible on the right. In the bottom left corner, there is a "feedback" link.

MongoClient | Featured M

localhost:3000/browseCollection

Refresh Collections Settings Disconnect

restaurants

Home / **restaurants**

Count:
3

Index Count:
1

Size:
0.00 MBs

Total Index Size:
0.04 MBs

feedback

MongoClient | Featured M x

localhost:3000/browseCollection

Query

find

Selector

1

By default, valid **ObjectID** and **ISODate (YYYY-MM-DD HH:mm:ss)** strings are being converted into MongoDB objects, **unchecked** right upper corner checkboxes to change this behaviour

Options

Choose one or more options..

Execute

feedback

POR 21:16

MongoClient | Featured M X

localhost:3000/browseCollection

Execute

find - restaurants

Tree

```
array [3]
  0 {3}
    _id : 572e4c431d569630bfd9890e
    name : Botequim Joao
    tipo : comida caseira
  1 {3}
    _id : 572e50401d569630bfd9890f
    name : Dominos
    tipo : pizzeria
  2 {3}
```

feedback

Create Collection

INF1802

Name

cliente

Is Capped

Size (If it's capped)

Max Documents (If it's capped)

Auto Index ID

Close

Create

Disconnect

3 Objects

0.04 MBs
Storage Size

0.04 MBs
Index Size

0 Extents

MongoClient | Featured M X

localhost:3000/databaseStats

INF1802
localhost : 27017
test

Database Stats

Management

Collections Add

- client Drop
- restaurants Drop

feedback

Search for something... Refresh Collections Settings Disconnect

Browse

Home / Browse

Idellv

The server is up for 1540 seconds

The performances which are shown here are coming from **serverStatus** and **stats** queries. Therefore, before having a look at below graphs consider **mongoclient** is already executing **serverStatus** at regular intervals. See [here](#) for more info

Mongodb version is **3.2.6**
Process id is **2036**
Process name is **mongod.exe**

| | |
|--------------------|-----------------------|
| 1 Collections | 3 Objects |
| 0.00 MBs Data Size | 0.04 MBs Storage Size |
| 1 Indexes | 0.04 MBs Index Size |
| 0.00 MBs File Size | 0 Extents |

Memory Usage

500 MB

false

Query You can either use mongodb or sql queries



Query

insertMany

Docs

```
1 {  
2   "name" : "John",  
3   "email" : "john@gmail.com",  
4   "tags" : ["sports","politics"],  
5   "grade" : "A",  
6   "score" : 11  
7 }
```

By default, valid **ObjectID** and **ISODate (YYYY-MM-DD HH:mm:ss)** strings are being converted into MongoDB objects, **unchecked** right upper corner checkboxes to change this behaviour
Furthermore, you can provide an **array** to insert **multiple** docs, or an **object** for **single** one

Execute

feedback

Execute

find - restaurants x find - client x

```
Tree
```

```
array [1]
  0 {6}
    _id : 572e895f3f2252100b708bac
    name : John
    email : john@gmail.com
    tags [2]
      0 : sports
      1 : politics
    grade : A
    score : 11
```

feedback