

# DATABASES

in a few slides


# DATABASE

a collection of logically related records

UOD: Universe of Discourse  
(all the facts in the db)

# RELATIONAL DB

- all information that comprises the universe of discourse is represented as relations.
- relations can be precisely defined mathematically.
- can also be defined using tables
- relational database = a collection of tables representing a UOD






# RDBMS

- Relational database management system
  - software that enables users to define, create, and maintain a database and provide controlled access.
- RDBMS ranking
  - Oracle
  - MySQL (Oracle) - MariaDB
  - Microsoft SQL server
  - IBM DB2 (10%)

# RDBMS RANKING

1. Oracle 1467
2. MySQL 1296
3. Microsoft SQL server 1226
4. PostgreSQL 228
5. DB2 188

# DB-Engines Ranking of Relational DBMS



POSTGRESQL



HACK

# Why The Older-Than-Dirt Postgres Database Is Hot With Hipsters And Oldsters Alike

*Mainly, it just works.*



MATT ASAY · OCT 23, 2014



Boring has never looked so cool. The decades-old relational-database management system Postgres, once the forgotten older sibling to MySQL, has been on a tear the last few years. Postgres has Oracle to thank for some of its newfound sexiness, as Oracle has fumbled MySQL's community outreach at crucial moments.

But far more of Postgres' renaissance stems from the fact that it boringly, reliably works. It turns out "boring" is a critical feature in a database.

WORK

# PostgreSQL Hits 9.3, New Levels Of Popularity With The Cool Kids

*PostgreSQL has long been a great database, yet strangely ignored. Now it's rising in popularity. Why?*




MATT ASAY · SEP 10, 2013

# SQL

- S-Q-L or *sequel* (mySQL = my-S-Q-L or my-sequel)
  - ANSI standard = S-Q-L
- SQL = Structured Query Language for RDBMS
  - orig. IBM 1970s SEQUEL (**S**tructure **E**nglish **Q**uery **L**anguage)
  - the 'standard' language for RDBMS



# DB-Engines Ranking of Relational DBMS



All ways of describing stuff

relational models

SQL



8 WEEKS

All ways of describing stuff

relational models

SQL

NOSQL



240 systems in ranking, January 2015

| Rank | Last Month | DBMS                 | Database Model    | Score   | Changes |
|------|------------|----------------------|-------------------|---------|---------|
| 1.   | 1.         | Oracle               | Relational DBMS   | 1439.16 | -20.63  |
| 2.   | 2.         | MySQL                | Relational DBMS   | 1277.51 | +8.93   |
| 3.   | 3.         | Microsoft SQL Server | Relational DBMS   | 1198.61 | -1.44   |
| 4.   | 4.         | PostgreSQL           | Relational DBMS   | 254.49  | +0.48   |
| 5.   | 5.         | MongoDB              | Document store    | 250.90  | +4.38   |
| 6.   | 6.         | DB2                  | Relational DBMS   | 200.13  | -10.12  |
| 7.   | 7.         | Microsoft Access     | Relational DBMS   | 139.14  | -0.76   |
| 8.   | ↑          | 9. Cassandra         | Wide column store | 98.75   | +4.69   |
| 9.   | ↓          | 8. SQLite            | Relational DBMS   | 96.20   | +1.49   |
| 10.  |            | 10. Redis            | Key-value store   | 94.24   | +6.36   |





***cassandra***



# ALONG SIDE DB STUFF

learn about developing web apps.

OLD SKOOL: PHP

2011

## Results

| Website       | Visits      | Programming Language |
|---------------|-------------|----------------------|
| facebook.com  | 870,000,000 | PHP                  |
| youtube.com   | 790,000,000 | Python               |
| yahoo.com     | 590,000,000 | PHP*                 |
| live.com      | 540,000,000 | ASP.NET              |
| wikipedia.org | 460,000,000 | PHP                  |
| msn.com       | 450,000,000 | ASP.NET              |
| blogspot.com  | 370,000,000 | Python*              |
| baidu.com     | 310,000,000 | PHP                  |
| microsoft.com | 280,000,000 | ASP.NET              |
| qq.com        | 250,000,000 | PHP                  |
| bing.com      | 230,000,000 | ASP.NET              |
| ask.com       | 190,000,000 | SSI                  |
| taobao.com    | 170,000,000 | PHP                  |
| twitter.com   | 160,000,000 | Ruby on Rails        |



Programming languages used in most popular websites\*

| Websites ⇅                | Popularity<br>(unique visitors) <sup>[1]</sup> ⇅ | Front-end<br>(Client-side) ⇅ | Back-end<br>(Server-side) ⇅   | Database ⇅   | Notes   |
|---------------------------|--|------------------------------|---|--|---|
| Google.com <sup>[2]</sup> | 1,100,000,000                                    | JavaScript                   | C, C++, Go, <sup>[3]</sup> Java, Python                                       | BigTable, <sup>[4]</sup><br>MariaDB <sup>[5]</sup> | The most used search engine in the world  |
| YouTube.com               | 1,000,000,000                                    | Flash,<br>JavaScript         | C/C++, Python, Java, <sup>[6]</sup> Go <sup>[7]</sup>                         | MySQL, BigTable                                    | The most visited video sharing site   |
| Facebook.com              | 900,000,000                                      | JavaScript                   | Hack, PHP, C++, Java, Python, Erlang,<br>D, <sup>[8]</sup> Xhp <sup>[9]</sup> | MySQL, <sup>[10]</sup> HBase                       | The most visited social networking site   |
| Yahoo                     | 750,000,000                                      | JavaScript                   | JavaScript, <sup>[11]</sup> PHP   | MySQL,<br>PostgreSQL <sup>[12]</sup>               | Yahoo is presently <sup>[when?]</sup> transitioning to<br>node.js <sup>[11]</sup> |
| Amazon.com                | 500,000,000                                      | JavaScript                   | Java, C++, Perl <sup>[13]</sup>   |  | Popular internet shopping site  |
| Wikipedia.org             | 475,000,000                                      | JavaScript                   | PHP   | MySQL,<br>MariaDB <sup>[14]</sup>                  | "MediaWiki" is programmed in PHP; free online<br>encyclopedia                     |
| Twitter.com               | 290,000,000                                      | JavaScript                   | C++, Java, Scala, Ruby on Rails <sup>[15]</sup>                               | MySQL <sup>[16]</sup>                              | 140 characters social network   |
| Bing                      | 285,000,000                                      | JavaScript                   | ASP.NET   | Microsoft SQL<br>Server                            |   |



### Back-end (Server-side) table in most popular websites

[illegible]



# COOL KIDS USE

- server side Javascript: node.js
- Python

# GOAL

- Learning skills that will help you get a cool job
- Learning skills that you can use on the job

# FIRST 8 WEEKS

- a solid working knowledge of how to design relational database systems for web applications
- learn the basics of web app development
- combine the 2: data driven web apps



# FRAMEWORK

- supports development of dynamic websites, web applications and web resources.
- alleviates overhead of doing it all yourself.
- microframework - core of the framework is simple. (read “easy to get started”)

# FRAMEWORKS

- support Model-view-controller pattern.
- means data model is separate from ui
- this modularity enables code reuse and reduces complexity.
- a lot of options for frameworks - Ruby on Rails, Django, Symfony

# FRAMEWORK HAZARDS

- full use of a framework would hide a lot of the database heavy lifting.
- would keep the database a mystery.
- you design it and it magically works.
- we are taking the middle road



# MODEL VIEW CONTROLLER

responsive



# Ron Zacharski

[Home](#)[Courses ▾](#)[Publications](#)[Books ▾](#)[About ▾](#)[News](#)

## Home



## Current Courses

### cs220

#### Computer Programming

This honors course uses a particular language, Java, to help you develop rudimentary programming skills.>>

### cs350

#### Databases & web apps

This course provides an introduction to databases and their application, particularly in data driven web applications.>>





server








server

url request to server



respond w/ webpage & assets



user clicks on link



respond w/ webpage & assets







Ron's Amazon.com

Today's Deals

Gift Cards

Sell

Help



New Year New You  
Sponsored by Intuit

Shop by  
Department ▾

Search

Books ▾

Go

Hello, Ron  
Your Account ▾

Your  
Prime ▾

Cart ▾

Wish  
List ▾

Books

Advanced Search

New Releases

Best Sellers

The New York Times® Best Sellers

Children's Books

Textbooks

Textbook Rentals

Sell Us Your Books

## Books at Amazon

# Save up to 90% on Textbooks

[> Learn more](#)

### Popular in Books

- Best Books of the Month
- New Year New You in Books
- Books Made into Movies: Read it before you see it
- 100 Books to Read in a Lifetime
- Children's Books
- Award Winners
- Best Books of 2014

Best Books  
OF THE  
Month

[> Learn more](#)

Book Club  
Picks

[> Learn more](#)

Children's  
Books

[> Shop now](#)






server

url request to server



respond w/ webpage & assets



user clicks on link



respond w/ just the requested data





Start reading *Sit Down and Shut Up* on the free [Kindle Reading App](#) or on your Kindle **in under a minute**. Don't have a Kindle? [Get your Kindle here](#).



Click to open expanded view



Listen: Play Sample

# Sit Down and Shut Up: Punk Rock Commentaries on Buddha, God, Truth, Sex, Death & Dogen's Treasury of the Right Dharma Eye [Kindle Edition]

Brad Warner (Author)  
★★★★★ (49 customer reviews)

Digital List Price: ~~\$15.95~~ What's this?  
Print List Price: ~~\$15.95~~  
Kindle Price: **\$9.99**  
You Save: **\$5.96 (37%)**

- Length: 256 pages
- Don't have a Kindle? [Get your Kindle here](#).
- Whispersync for Voice: Ready



**Free Kindle Reading App**  
Anybody can read Kindle books—even without a Kindle device—with the **FREE** Kindle app for smartphones, tablets and computers.

To get the free app, enter your email address or mobile phone number.

Enter your email or mobile phone number

Send me the link



**Audible Narration**  
[Switch back and forth](#) between reading the Kindle book and listening to the Audible narration with Whispersync for Voice. Add narration for a reduced price of **\$3.99** when you buy the Kindle book.

| Formats                           | Amazon Price        | New from | Used from |
|-----------------------------------|---------------------|----------|-----------|
| Kindle Edition                    | \$9.99              | —        | —         |
| Paperback                         | \$11.93             | \$4.98   | \$0.97    |
| Audible Audio Edition, Unabridged | \$13.96 or 1 credit |          |           |

Buy now with 1-Click®

Deliver your Kindle book to:  
Ron's 2nd Kindle

## Add Audible Narration

☐ [Sit Down and Shut Up: Punk Rock Commentaries on Buddha, God, Truth, Sex, Death, and Dogen's Treasury of the Right Dharma Eye](#)  
Narrated by Brad Warner  
~~\$19.95~~ \$3.99

Give as a Gift

Enter a promotion code or gift card

Add to Wish List

## Try it free

Sample the beginning of this book for free

Send sample now

Deliver to:

Ron's 2nd Kindle

How sampling works  
Available on your Mac

Share

## Also available on Kindle



**Forging Zero (The Le...**  
by King, Sara  
★★★★★ (822)  
Kindle Price: ~~\$3.99~~ \$0.99

Shop now

Download  
to Your Device  
[Learn more](#)



# START W/ OLD SCHOOL

quick transition to responsive (angular)

THAT'S WHAT WE ARE GOING  
TO DO



WHAT WE ARE NOT DOING

# THIS COURSE

- Applications of Databases
- ‘Real DB courses’
  - CPSC410 Database Principles and Design
  - UT CS347:Data Management

Topics related to the engineering and design of database systems, including: data models; database and schema design; schema normalization and integrity constraints; query processing; query optimization and cost estimation; transactions; recovery; concurrency control; isolation and consistency; distributed, parallel, and heterogeneous databases; adaptive databases; trigger systems; key-value stores; object-relational mappings; streaming databases; DB as a service. Lecture and readings from original research papers. 6.830 includes semester-long project and paper.



A programmer's guide to DB

How to develop data driven web apps.

# DESIGNED FOR THOSE WITH

- no database or web development experience
- some experience but have holes in knowledge

# NOT

- skills on web design
- skills on usability
- database experts



# EXAMPLE

- me

# NETEZZA



- asymmetric multiprocessing
- some apps 14,000% faster than legacy db.

3.75hrs  
Netezza

72 hrs  
Oracle

# NETEZZA

## 10000 Series Specifications

| <b>Hardware</b>                                   | <b>10050</b>   | <b>10100</b> | <b>10200</b> | <b>10400</b> | <b>10600</b> | <b>10800</b> |
|---|--|--------------|--------------|--------------|--------------|--------------|
| <b>Racks</b> (Standard 19")                       | 1(50% populated)   | 1            | 2            | 4            | 6            | 8            |
| <b>Total SPUs</b><br>(Query Processing Nodes)     | 56   | 112          | 224          | 448          | 672          | 896          |
| <b>User Data</b> (max.)                           | 6.25 TB  | 12.5 TB      | 25 TB        | 50 TB        | 75 TB        | 100 TB       |
| <b>Physical Cabinet Specifications (per rack)</b> |  |              |              |              |              |              |
| <b>Height</b>                                     | 80.5" / 2050 mm  |              |              |              |              |              |
| <b>Width</b>                                      | 24" / 610 mm   |              |              |              |              |              |
| <b>Depth</b>                                      | 38" / 970 mm   |              |              |              |              |              |
| <b>Weight</b>                                     | 1,100 lbs / 520 kg (max.)  |              |              |              |              |              |
| <b>Operating Temperature</b>                      | 50F to 95F / 10C to 35C  |              |              |              |              |              |
| <b>Cooling Requirements</b>                       | 12,000 BTU/hour  |              |              |              |              |              |
| <b>Power Requirements</b>                         | Survivable A/B dual power feed, 30A fusing   |              |              |              |              |              |
| <b>Safety</b>                                     | UL/CSA/EN60950   |              |              |              |              |              |
| <b>Emissions</b>                                  | FCC Part 15, ICES-003, AUS/NZ C-Tick, VCCI and EN55022 Class A; European Immunity: EN55024 |              |              |              |              |              |



# NETEZZA

**BI Applications Client**



Local  
Applications

ODBC 2.5, 3.x  
JDBC Type 4  
SQL/92

Fast Loader/  
Unloader

**Netezza Performance Server**

SQL  
Compiler  
Query Plan  
Optimize  
Admin

Execution  
Engine

DBOS

Front End

**SMP Host**

**Gigabit  
Ethernet**

1

2

3

•

•

•

1000+

**Snippet Processing Unit (SPU)**

Processing  
& Streaming  
DB Logic

**Snippet Processing Unit (SPU)**

Processing  
& Streaming  
DB Logic

**Snippet Processing Unit (SPU)**

Processing  
& Streaming  
DB Logic

• Streaming joins  
• Aggregation  
• Sorts, etc.

**Snippet Processing Unit (SPU)**

Processing  
& Streaming  
DB Logic

**Massively Parallel  
Intelligent Storage**

# SANDIA

- building citation graphs
  - 1 million papers - 25 million citations
    - mySQL dual processor desktop - 40 hrs
    - Netezza .38 hrs. (288M citations - 50 hrs)

# NETEZZA ENGINEERS

- expert at designing db
- expert at writing efficient queries



# SUMMARY

- not expert
- sufficient expertise to create web app. startup

# ACCELERATED CLASS

- covering 15 weeks of material in 8 weeks.
- not a survey or overview class
- guarantee you that you will know this stuff backwards and forwards.
- focus on software development skills
- this is the last time the course will be 3 credits.

# WHAT DO WE DO IN THE LAST 7 WEEKS



All ways of describing stuff

relational models

SQL

All ways of describing stuff

noSQL



**Cassandra**



relational models

SQL




{name: "mongo", type: "DB"}



DIGRESSION




● A    ● A-    ● B+    ● B    ● B-    ● C    ● F




2014

● A    ● A-    ● B+    ● B    ● B-    ● C    ● F

2013




● A    ● A-    ● B+    ● B    ● B-    ● C    ● D



2012



● A    ● A-    ● B+    ● B    ● B-    ● C    ● F



2011

# GUT-FEELING

not cognitively hard.

not complex algorithms or formulas

puzzles

compare to 405

fun

# GOAL

- useful
- useful for job interviews.



back to

RELATIONAL DATABASES





# TEAM DELIVERABLES

- sent to [ron.zacharski@gmail.com](mailto:ron.zacharski@gmail.com) by tonight
- section number, team number and name of team
- team picture with names of people in correct order
- individual deliverable: 3x5 card w/ real name, avatar name, & team number