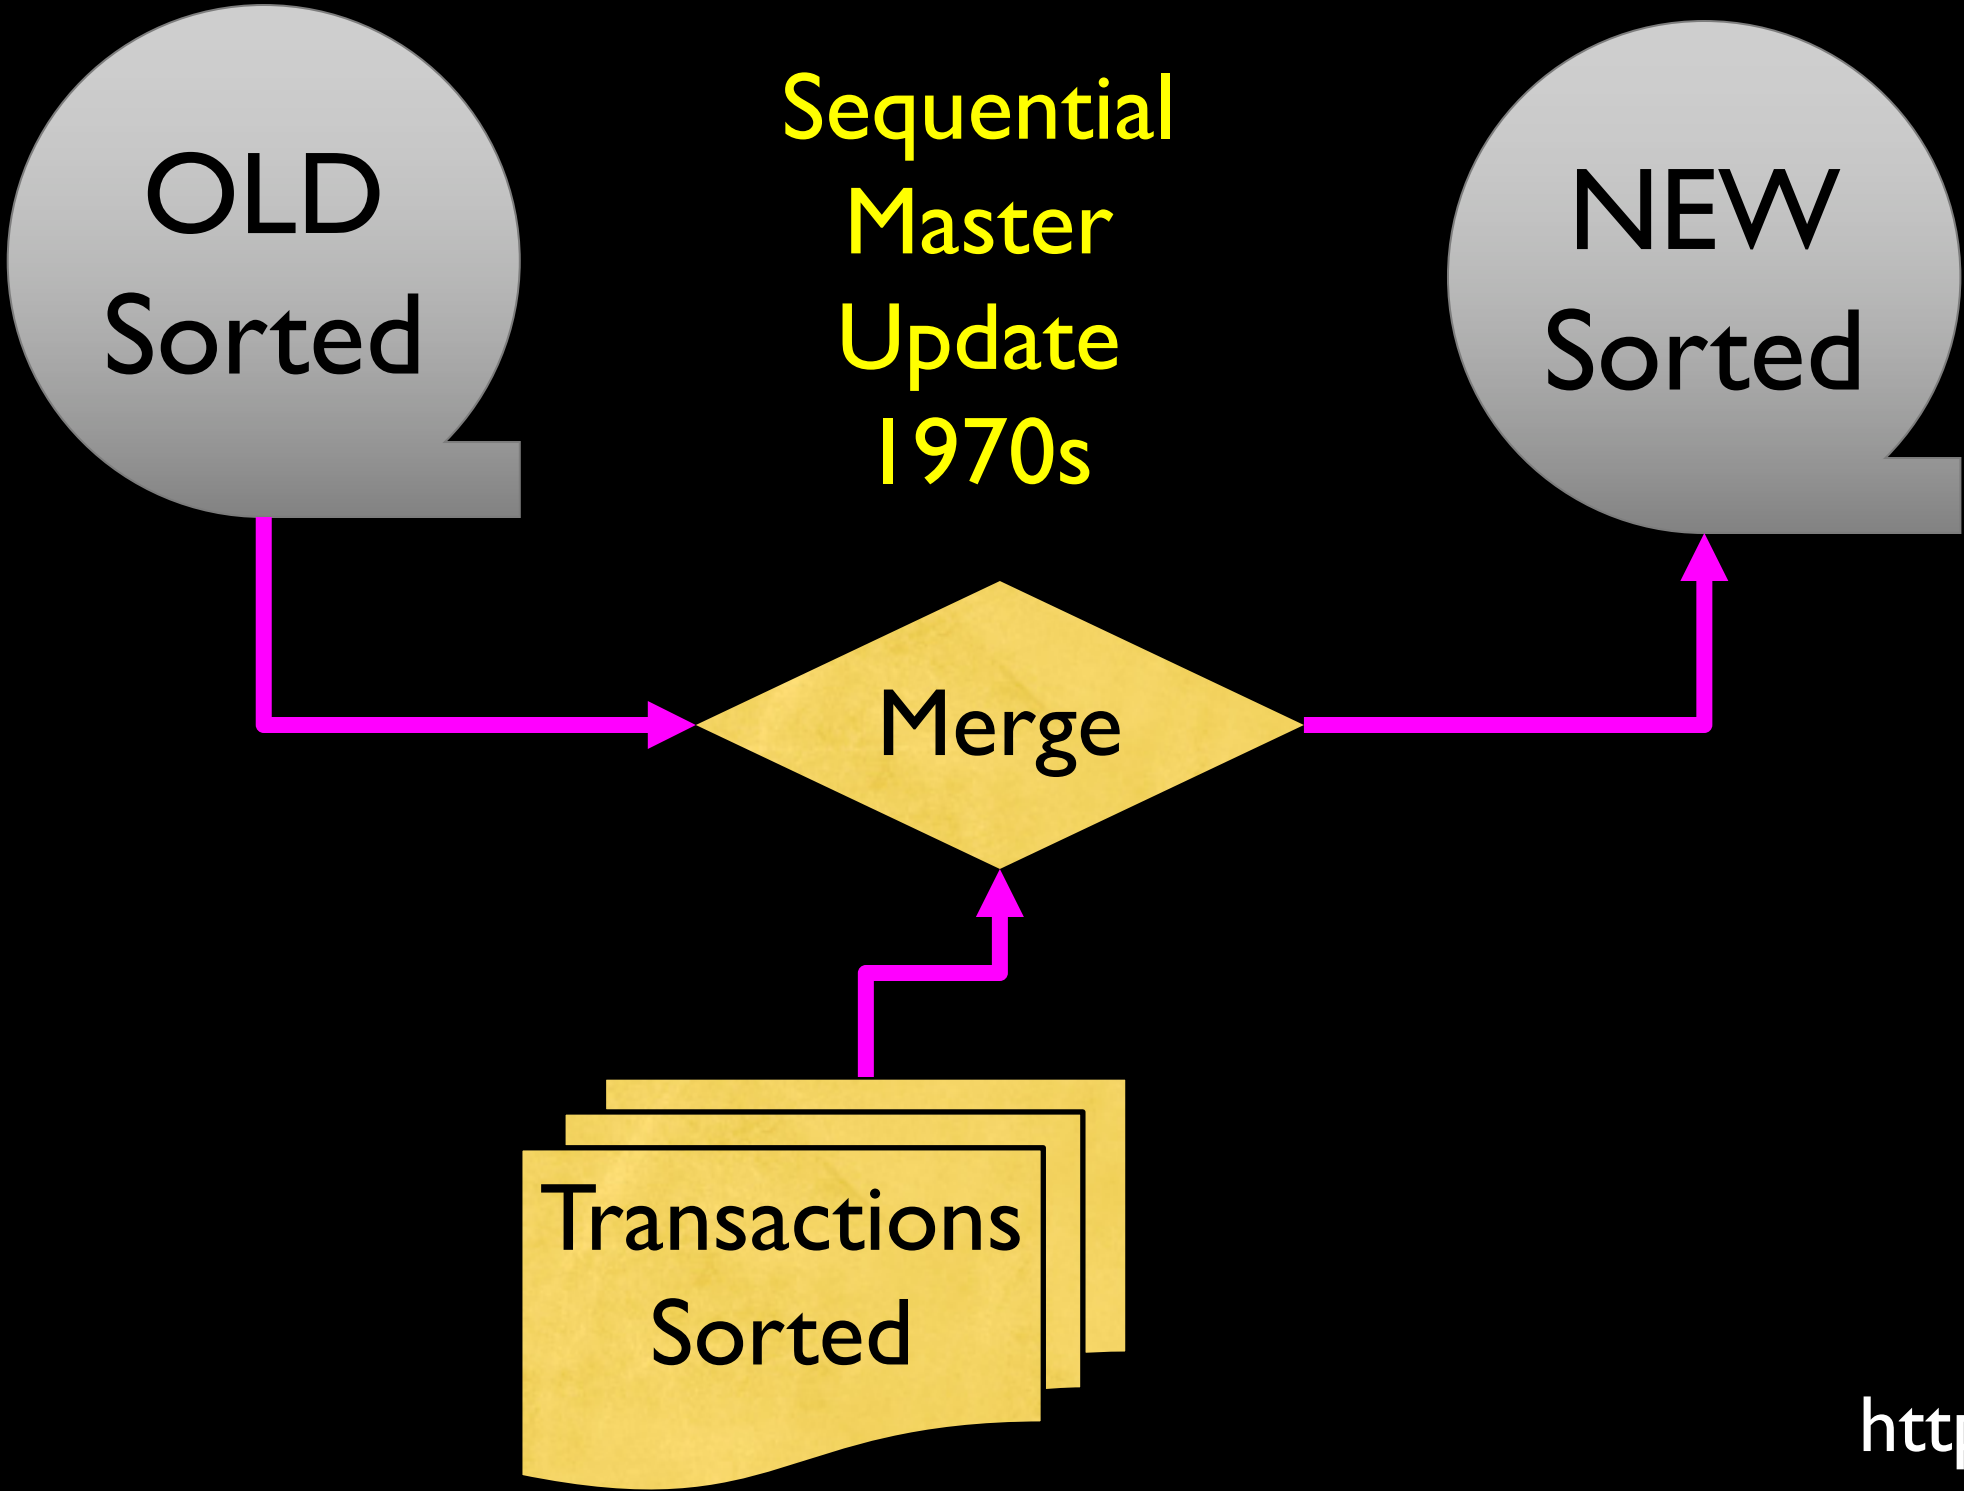


Relational Databases and MySQL

Charles Severance
www.wa4e.com



https://en.wikipedia.org/wiki/IBM_729

Random Access

- When you can randomly access data...
- How can you lay out data to be most efficient?
- Sorting might not be the best idea



https://en.wikipedia.org/wiki/Hard_disk_drive_platter

Relational Databases

Relational databases model data by storing rows and columns in tables. The power of the relational database lies in its ability to efficiently retrieve data from those tables - in particular, where the query involves multiple tables and the relationships between those tables.

http://en.wikipedia.org/wiki/Relational_database

Structured Query Language

- Structured Query Language (SQL) came out of a government / industry partnership
- National Institute of Standards and Technology (NIST)

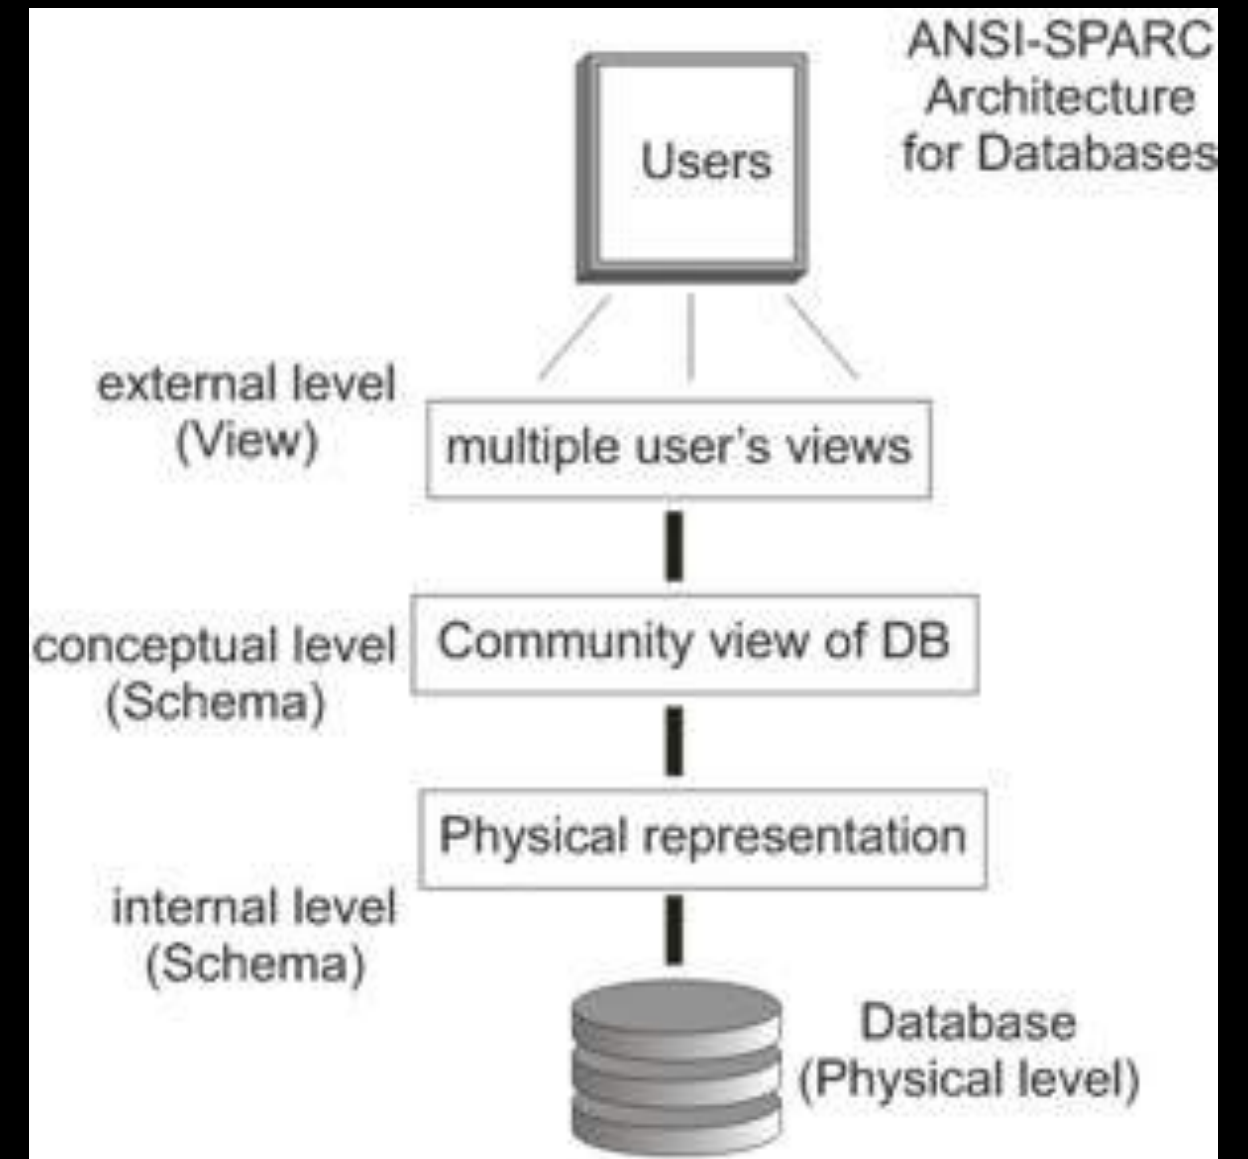


<https://youtu.be/rLUm3vst87g>

SQL

Structured Query Language is the language we use to issue commands to the database

- Create/Insert data
- Read/Select some data
- Update data
- Delete data




<http://en.wikipedia.org/wiki/SQL>

https://en.wikipedia.org/wiki/ANSI-SPARC_Architecture

Terminology

- **Database** - contains one or more tables
- **Relation (or table)** - contains tuples and attributes
- **Tuple (or row)** - a set of fields which generally represent an “object” like a person or a music track
- **Attribute (also column or field)** - one of possibly many elements of data corresponding to the object represented by the row



A **relation** is defined as a **set of tuples** that have the same **attributes**. A **tuple** usually represents an **object** and information about that object. **Objects** are typically physical objects or concepts. A **relation** is usually described as a **table**, which is organized into **rows** and **columns**. All the data referenced by an **attribute** are in the same domain and **conform to the same constraints**.

(wikipedia)



SI502 - Database

New Open Save Print Import Copy Paste Format Undo Redo AutoSum Sort A-Z Sort Z-A Gallery Toolbox

Sheets Charts SmartArt Graphics WordArt

A B C D

1 2 3 4 5 6 7 8

TITLE	RATING	LEN
About to Rock	3	354
Who Made Who	4	252

Tracks Albums Artists Genres +

Columns / Attributes

Rows /
Tuples

Tables / Relations

Common Database Systems

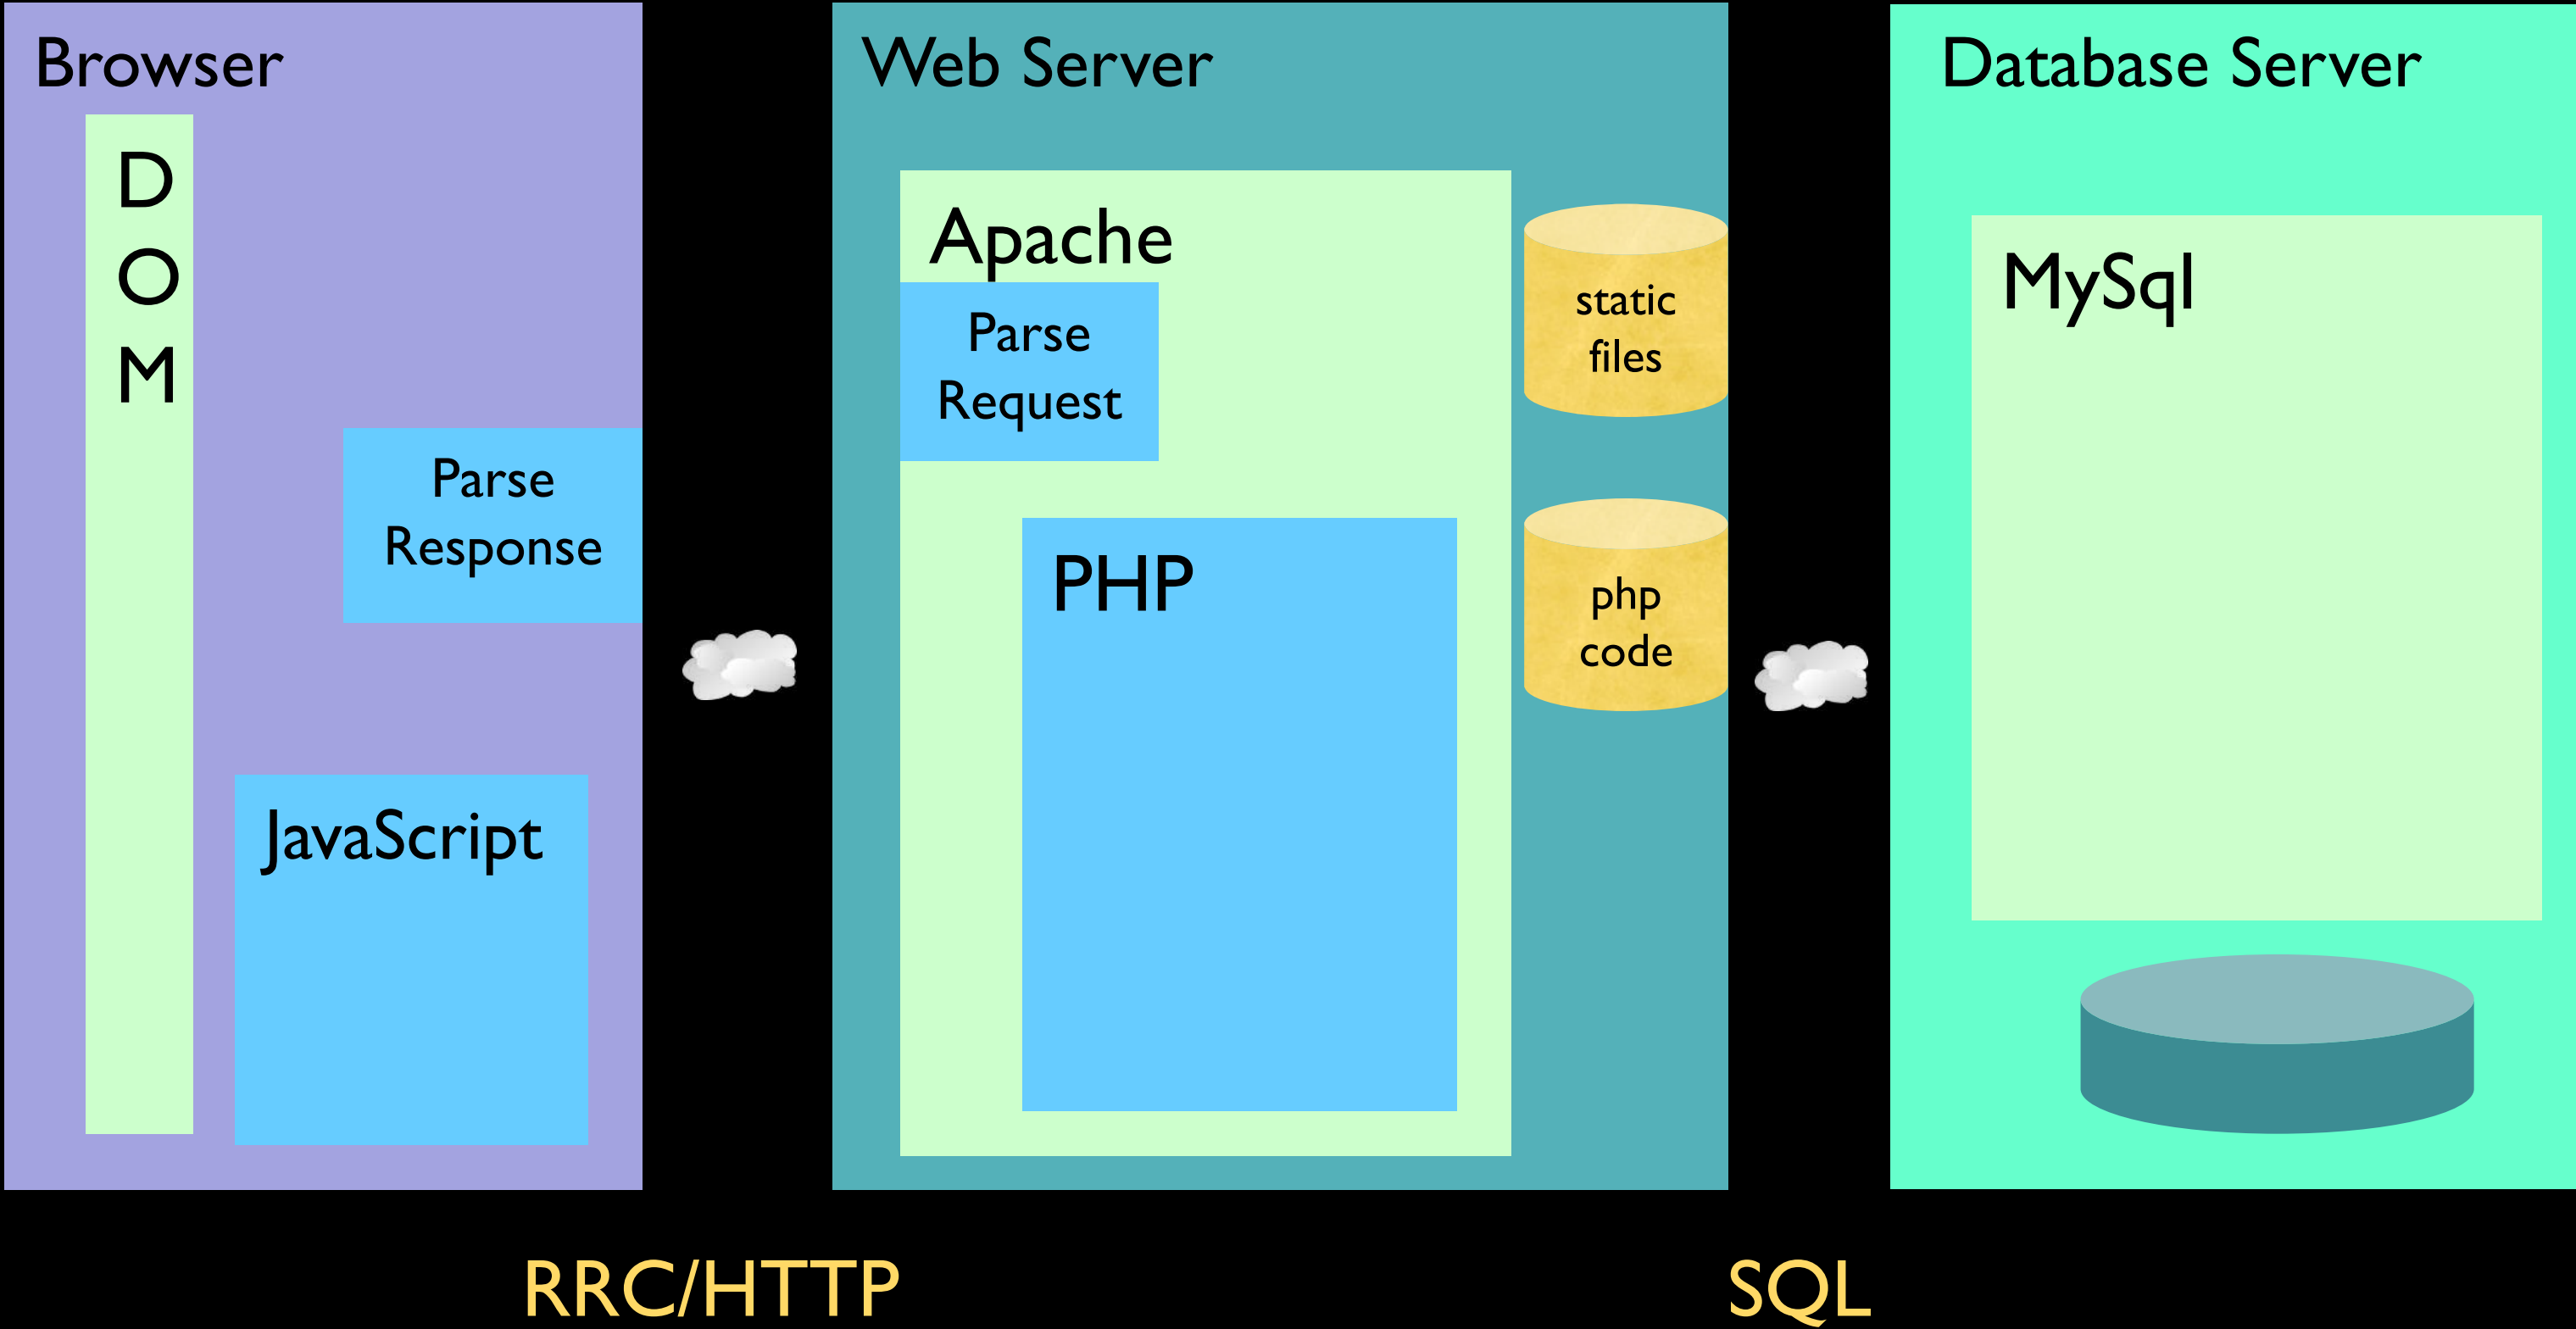
- Three major Database Management Systems in wide use
 - **Oracle** - Large, commercial, enterprise-scale, very tweakable
 - **MySQL** - Simpler but very fast and scalable - commercial open source
 - **SqlServer** - Very nice - from Microsoft (also Access)
- Many other smaller projects, free and open source
 - HSQL, SQLite, PostgreSQL ...



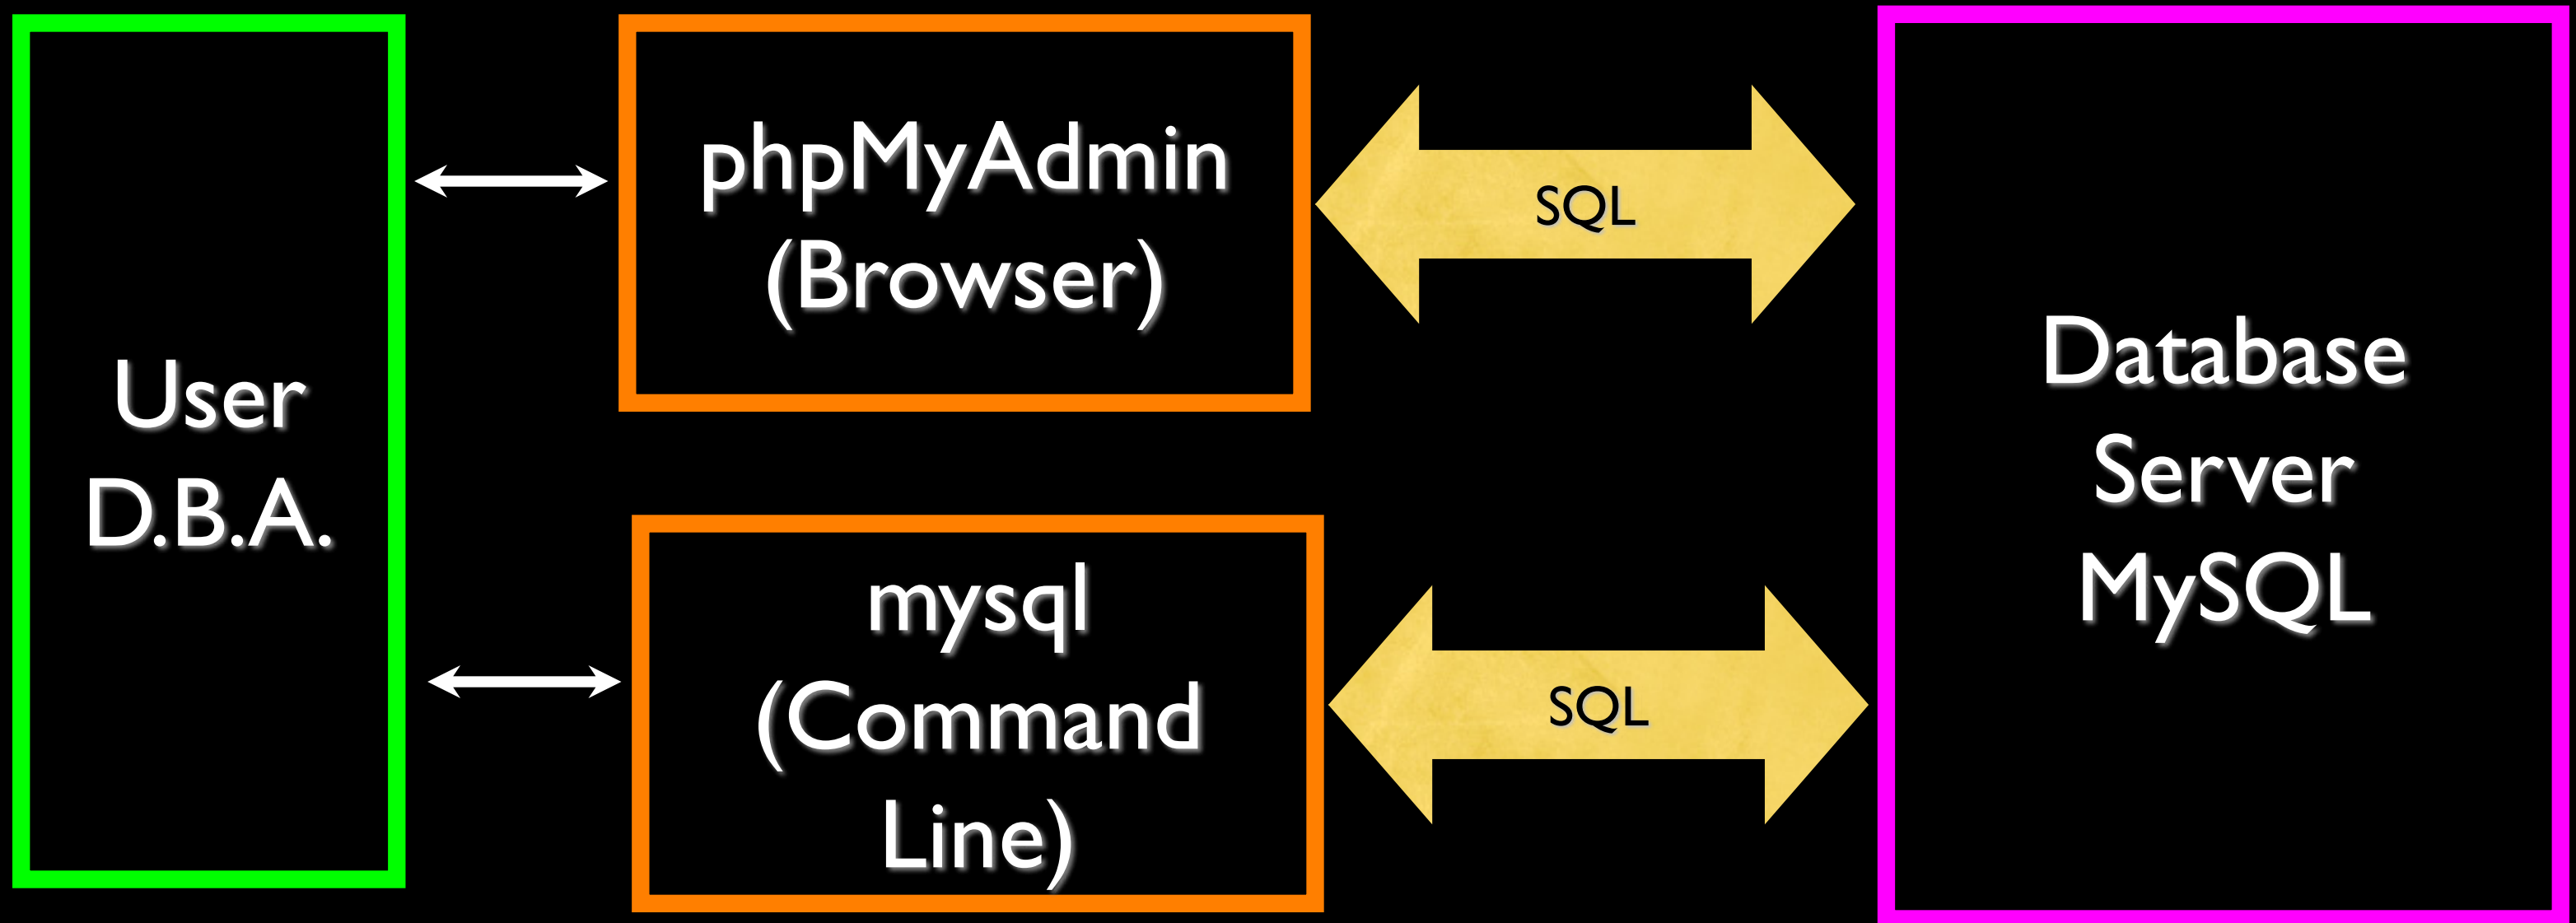
Basic SQL Operations



Time



Using SQL





The screenshot shows the XAMPP 1.7.4 installation page in Mozilla Firefox. The browser's address bar shows `http://localhost/xampp/`. The page title is "XAMPP 1.7.4". The main heading is "XAMPP for Windows" with a language selector showing "English" as the selected language. The left sidebar contains a navigation menu with the following items: Documentation, Components, PHP (with sub-items: phpinfo(), CD Collection, Biorhythm, Instant Art, Phone Book), Perl (with sub-items: perlinfo(), Guest Book), J2EE (with sub-items: Status, Tomcat examples), and Tools (with sub-items: phpMyAdmin, Webalizer, Mercury Mail, FileZilla FTP). The main content area has a yellow background and displays the following text: "Welcome to XAMPP for Windows!", "Congratulations: You have successfully installed XAMPP on this system!", "Now you can start using Apache and Co. You should first try »Status« on the left navigation to make sure everything works fine.", "For OpenSSL support please use the test certificate with <https://127.0.0.1> or <https://localhost>", and "Good luck, Kay Vogelgesang + Kai 'Oswald' Seidler". A large yellow arrow points from the "Status" link in the left sidebar to the main content area. The footer of the page shows "©2002-2010 ...APACHE". The browser's status bar at the bottom says "Done".



XAMPP 1.7.4 - Mozilla Firefox

File Edit View History Bookmarks Tools Help

http://localhost/xampp/

Most Visited Getting Started Latest Headlines

XAMPP 1.7.4

Documentation Components

PHP
phpinfo()
CD Collection
Biorhythm
Instant Art
Phone Book

Perl
perlinfo()
Guest Book

J2EE
Status
Tomcat examples

Tools
phpMyAdmin
Webalizer
Mercury Mail
FileZilla FTP

©2002-2010
...APACHE

Done

localhost / localhost | phpMyAdmin 3.3.9 - Mozilla Firefox

File Edit View History Bookmarks Tools Help

http://localhost/phpmyadmin/

Most Visited Getting Started Latest Headlines

localhost / localhost | phpMyAdmin 3....

phpMyAdmin

- cdcol (1)
- chuck
- information_schema (37)
- mysql (24)
- performance_schema (17)
- phpmyadmin (8)
- test
- webauth (1)

Please select a database

localhost

Databases

SQL

Status

Variables

Charsets

Engines

Privileges

Replication

Processes

Export

Import

Synchronize

Run SQL query/queries on server "localhost":

show databases;

Bookmark this SQL query:

☐ Let every user access this bookmark

☐ Replace existing bookmark of same name

[Delimiter ;] ☒ Show this query here again

Go

Done

XAMPP 1.7.4 - Mozilla Firefox

File Edit View History Bookmarks Tools Help

http://localhost/xampp/

Most Visited Getting Started

XAMPP 1.7.4

Documentation Components

PHP
phpinfo()
CD Collection
Biorhythm
Instant Art
Phone Book

Perl
perlinfo()
Guest Book

J2EE
Status
Tomcat examples

Tools
phpMyAdmin
Webalizer
Mercury Mail
FileZilla FTP

©2002-2010
...APACHE

Done

localhost / localhost | phpMyAdmin 3.3.9 - Mozilla Firefox

File Edit View History Bookmarks Tools Help

http://localhost/phpmyadmin/

Most Visited Getting Started

phpMyAdmin

- cdcol (1)
- chuck
- information_schema (37)
- mysql (24)
- performance_schema (17)
- phpmyadmin (8)
- test
- webauth (1)

Please select a database

Done

localhost / localhost | phpMyAdmin 3.3.9 - Mozilla Firefox

File Edit View History Bookmarks Tools Help

http://localhost/phpmyadmin/

Most Visited Getting Started Latest Headlines

phpMyAdmin

- cdcol (1)
- chuck
- information_schema (37)
- mysql (24)
- performance_schema (17)
- phpmyadmin (8)
- test
- webauth (1)

Please select a database

Done

localhost

Databases SQL Status Variables Charsets Engines Privileges

Replication Processes Export Import Synchronize


✓ Your SQL query has been executed successfully

SHOW DATABASES

☐ Profiling [Edit] [Create PHP Code] [Refresh]

+ Options

Database
information_schema
cdcol
chuck
mysql
performance_schema
phpmyadmin
test
webauth



Command Line

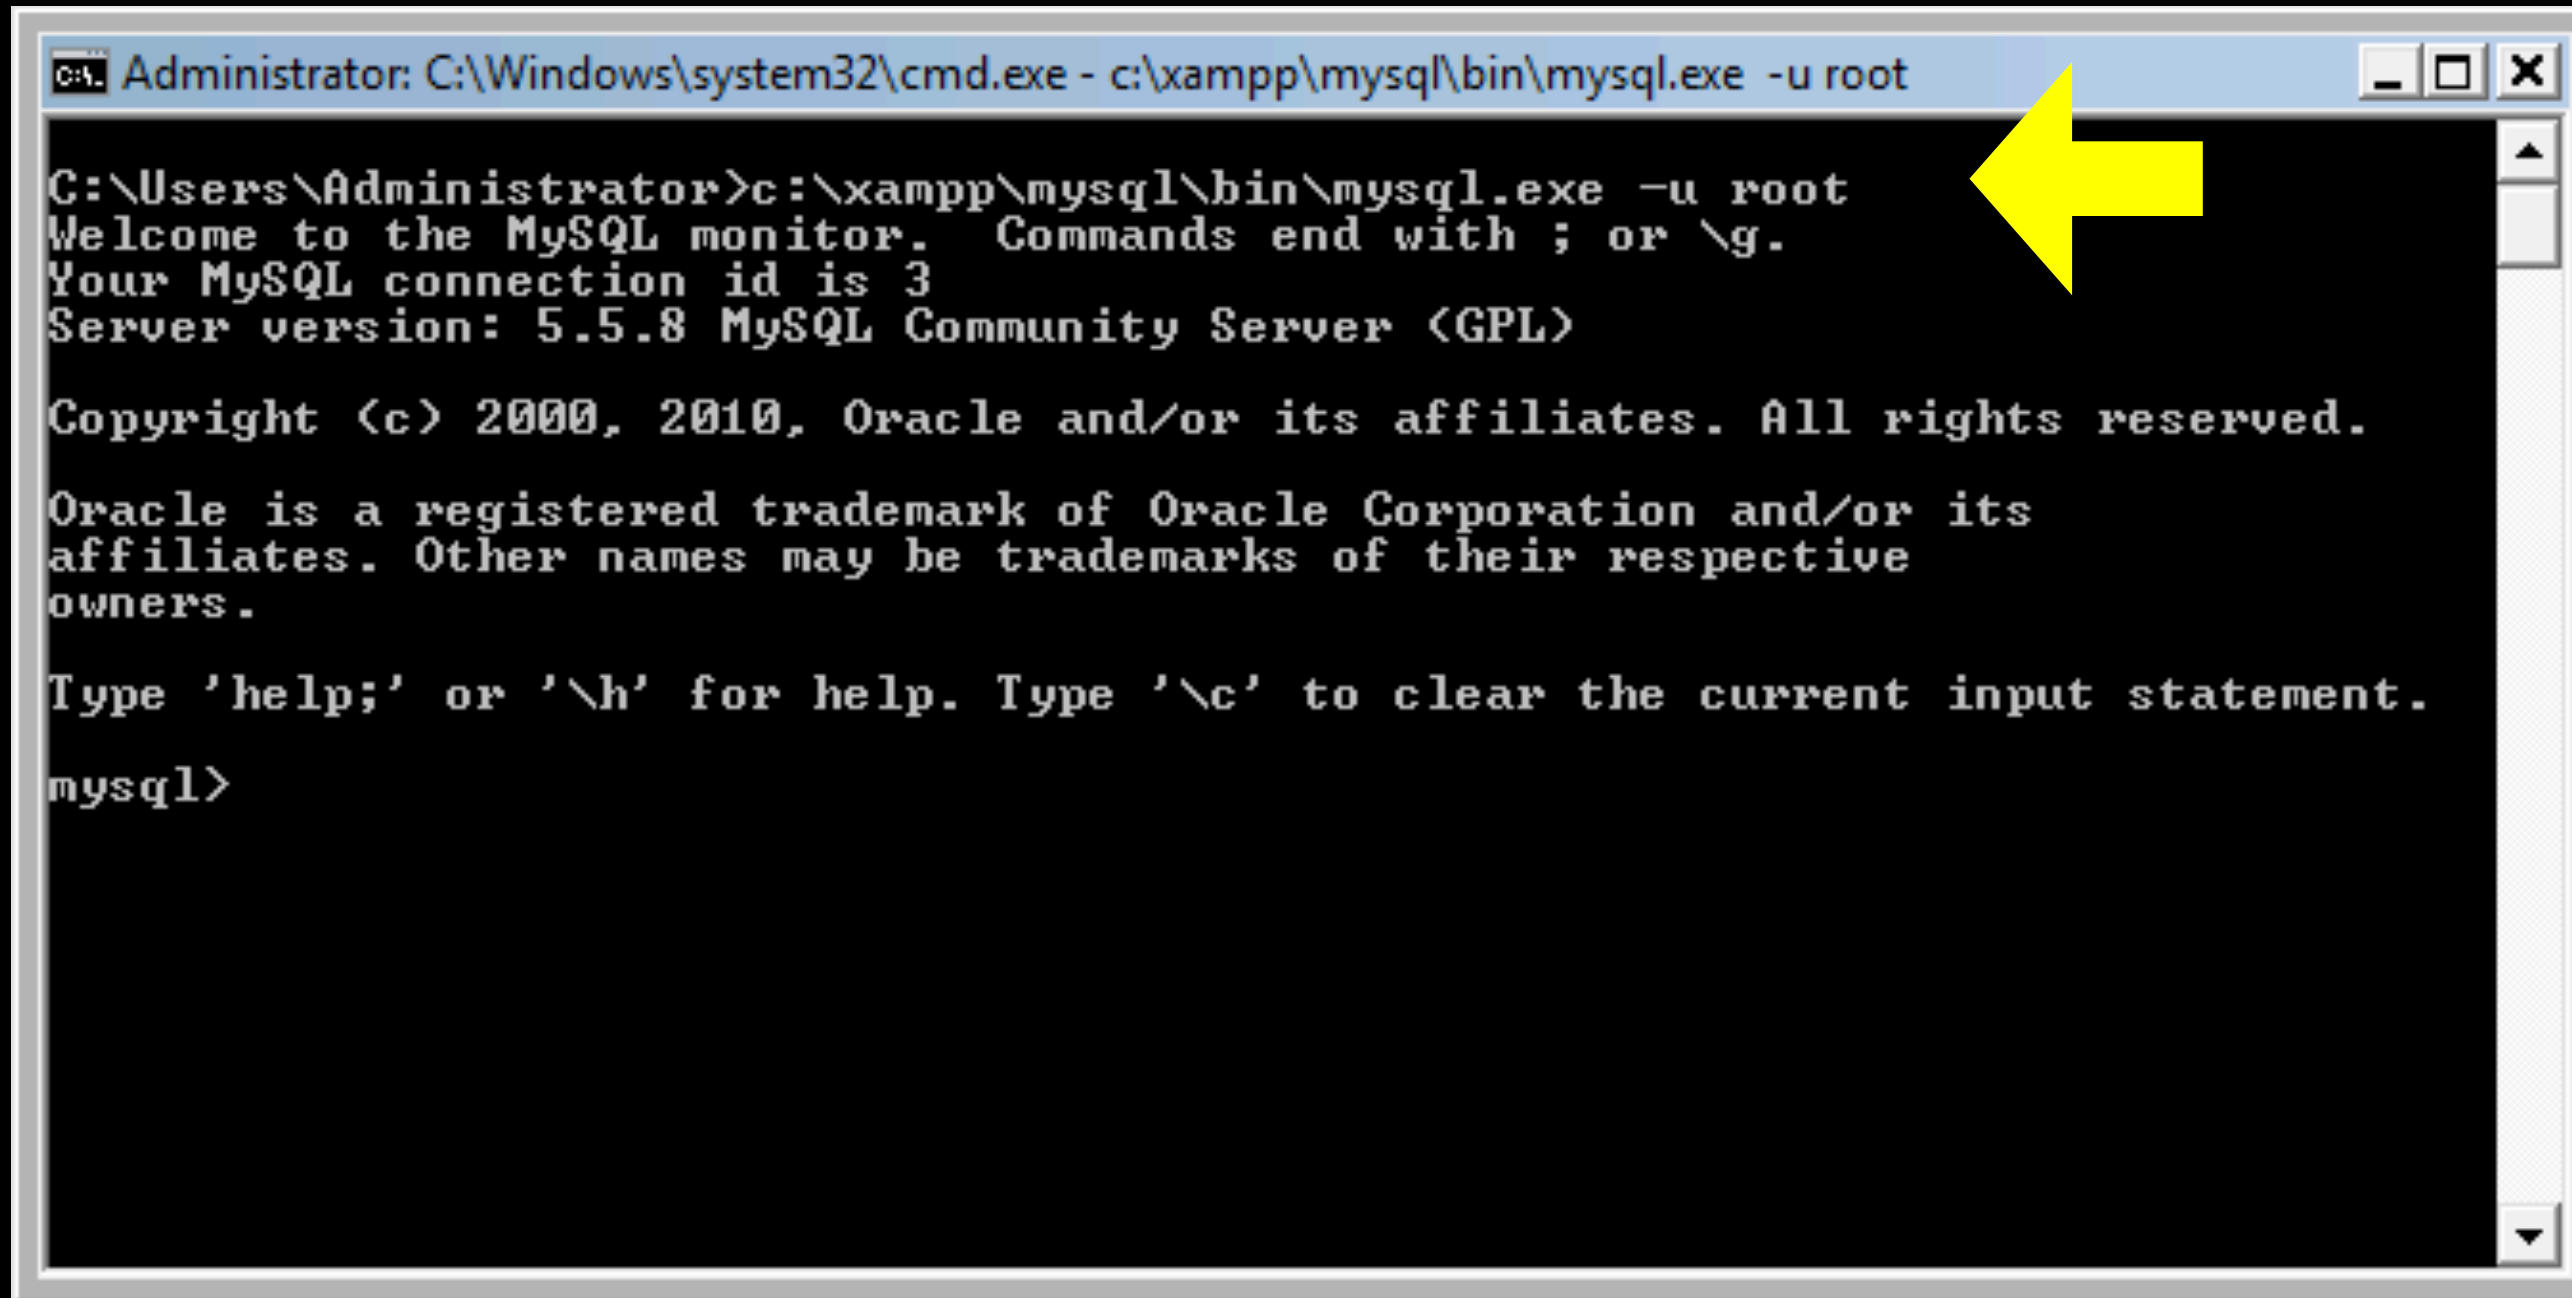
After Control Panel is running...

- Macintosh

- `/Applications/MAMP/Library/bin/mysql --host=localhost -uroot -p`
- Enter "root" when prompted for the password

- Windows

- `c:\xampp\mysql\bin\mysql.exe -u root -p`
- Press enter when prompted for password



```
Administrator: C:\Windows\system32\cmd.exe - c:\xampp\mysql\bin\mysql.exe -u root

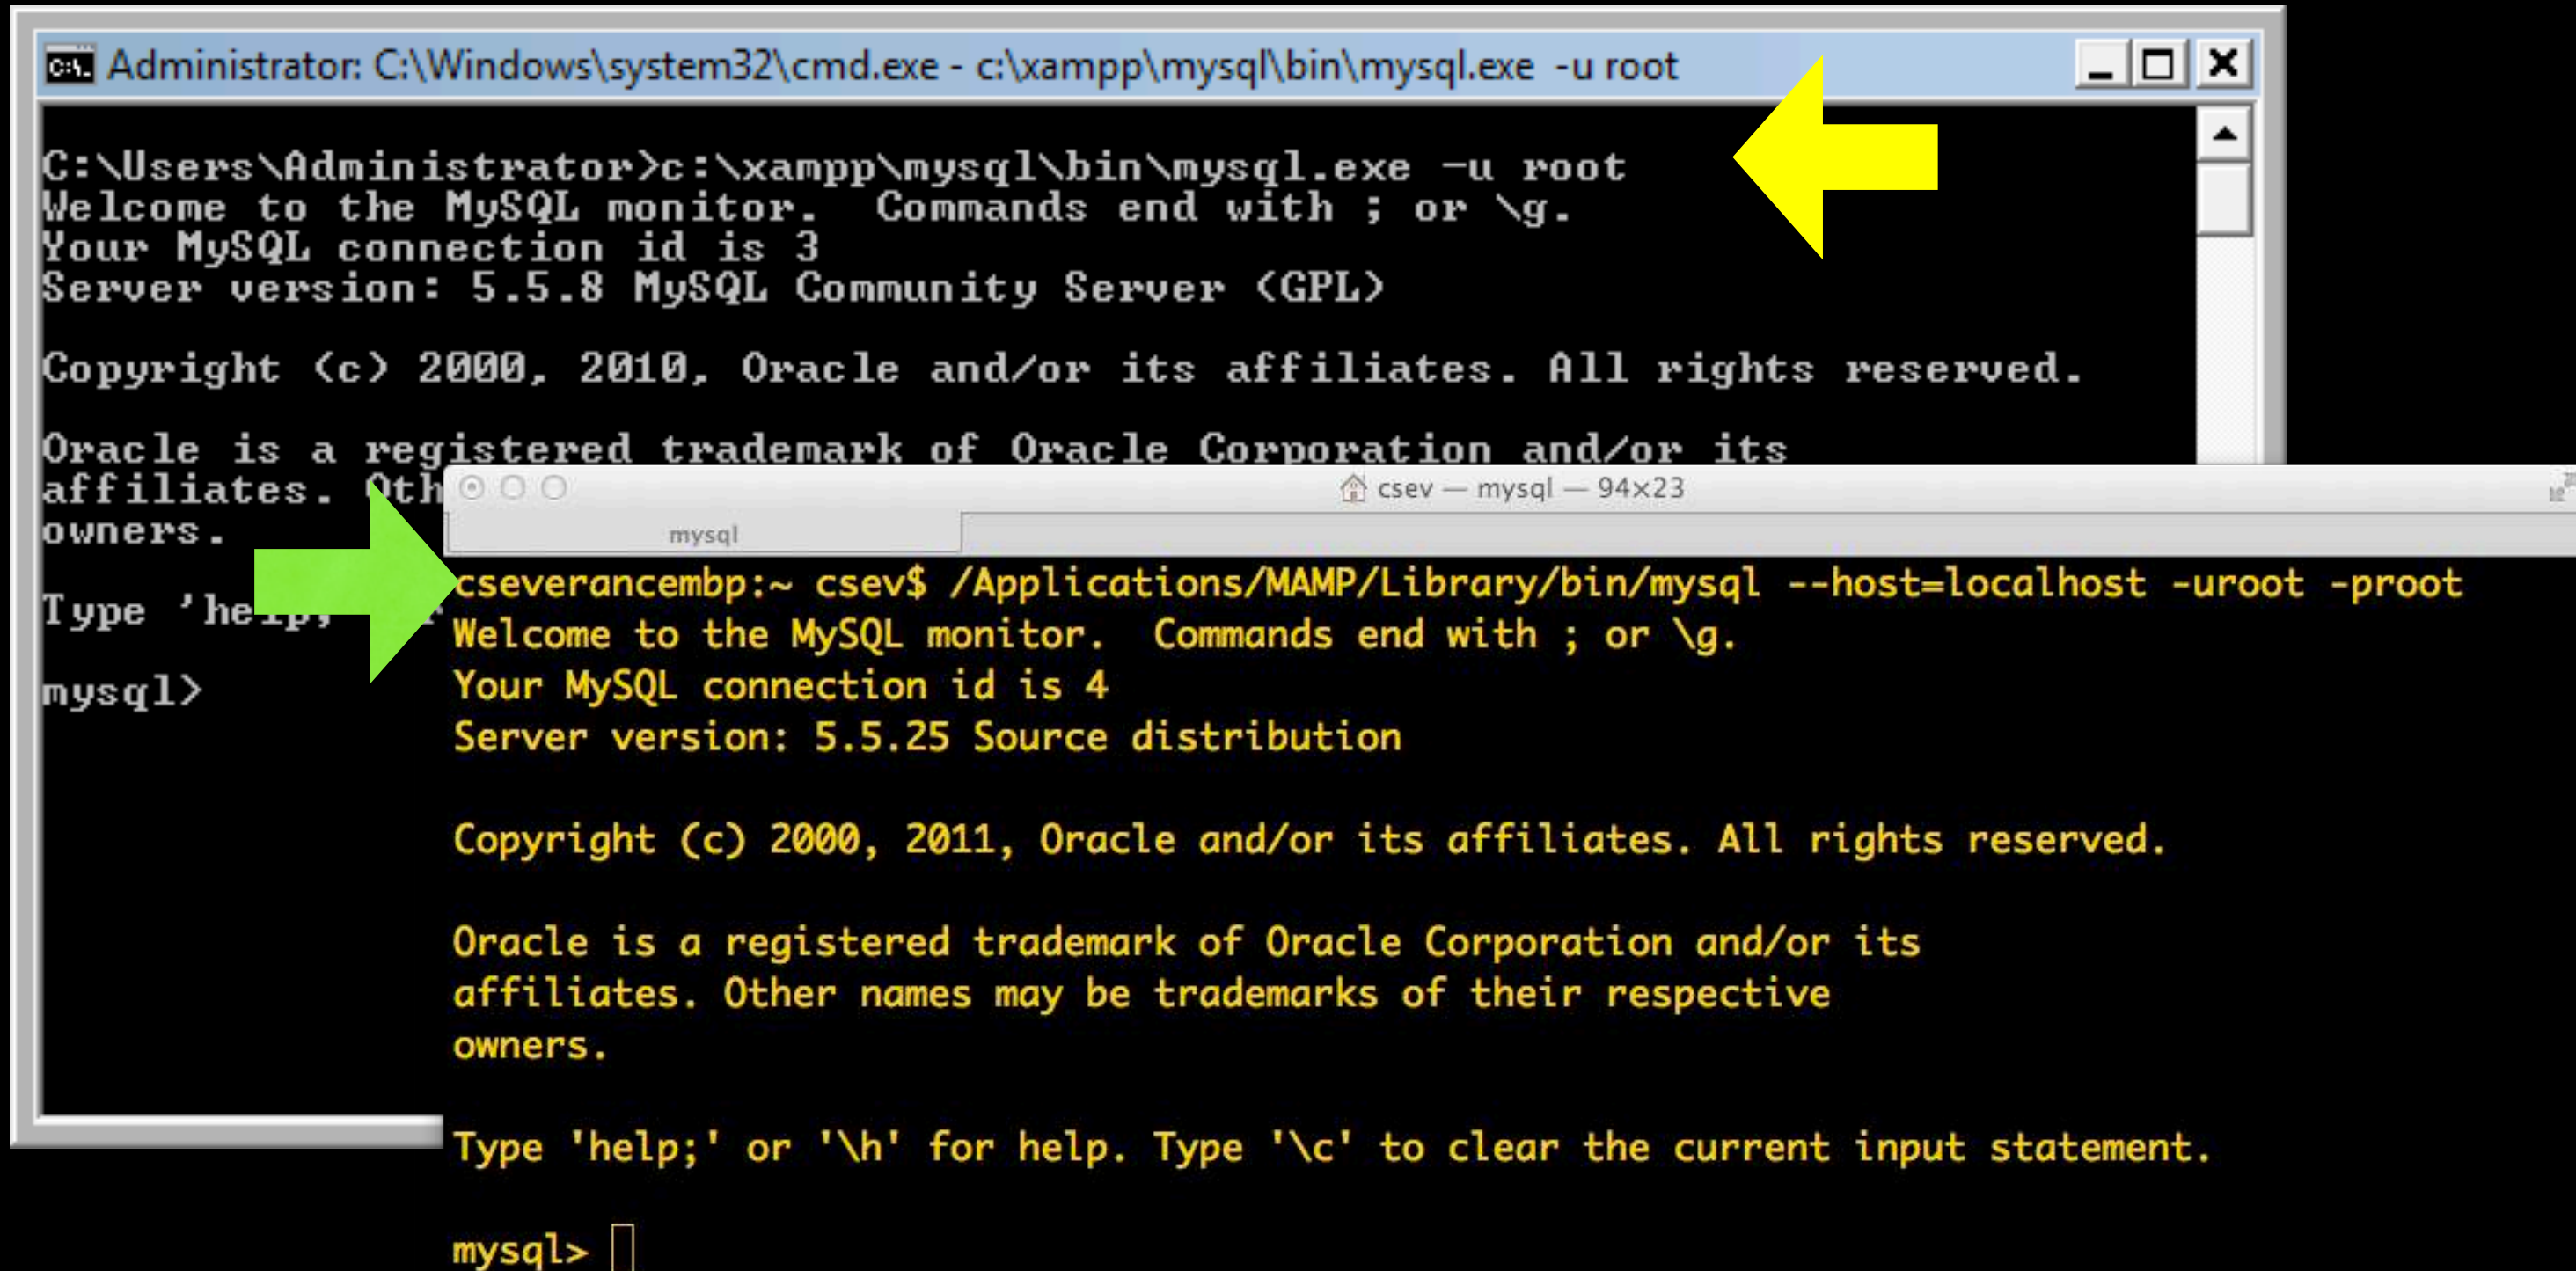
C:\Users\Administrator>c:\xampp\mysql\bin\mysql.exe -u root
Welcome to the MySQL monitor.  Commands end with ; or \g.
Your MySQL connection id is 3
Server version: 5.5.8 MySQL Community Server (GPL)

Copyright (c) 2000, 2010, Oracle and/or its affiliates. All rights reserved.

Oracle is a registered trademark of Oracle Corporation and/or its
affiliates. Other names may be trademarks of their respective
owners.

Type 'help;' or '\h' for help. Type '\c' to clear the current input statement.

mysql>
```



```
Administrator: C:\Windows\system32\cmd.exe - c:\xampp\mysql\bin\mysql.exe -u root

C:\Users\Administrator>c:\xampp\mysql\bin\mysql.exe -u root
Welcome to the MySQL monitor.  Commands end with ; or \g.
Your MySQL connection id is 3
Server version: 5.5.8 MySQL Community Server (GPL)

Copyright (c) 2000, 2010, Oracle and/or its affiliates. All rights reserved.

Oracle is a registered trademark of Oracle Corporation and/or its
affiliates. Other names may be trademarks of their respective
owners.

Type 'help;' or '\h' for help. Type '\c' to clear the current input statement.

mysql>

cseverancembp:~ csev$ /Applications/MAMP/Library/bin/mysql --host=localhost -uroot -proot
Welcome to the MySQL monitor.  Commands end with ; or \g.
Your MySQL connection id is 4
Server version: 5.5.25 Source distribution

Copyright (c) 2000, 2011, Oracle and/or its affiliates. All rights reserved.

Oracle is a registered trademark of Oracle Corporation and/or its
affiliates. Other names may be trademarks of their respective
owners.

Type 'help;' or '\h' for help. Type '\c' to clear the current input statement.

mysql>
```

Your First MySQL Command

`show databases`

Kind of like `print('hello world')`


```
mysql> show databases;
+-----+
| Database |
+-----+
| information_schema |
| cdcol |
| chuck |
| mysql |
| sakai |
| test |
+-----+
6 rows in set (0.06 sec)
```

```
mysql> 
```

If this does not work, stop
and figure out why.

Some of these are part of
MySQL and store internal
data - don't mess with them.

Creating a Database

Command Line:

```
CREATE DATABASE People;  
USE People;
```

MAMP

localhost:8888/MAMP/?language=English

StartphpInfoXCachephpMyAdminSQLiteManagerFAQTry MAMP PRO

MAMP & MAMP PRO powered by appsolute GmbH

phpMyAdmin

localhost

DatabasesSQLStatusUsersExportImportSettingsSynchronizeMore

1 databases have been dropped successfully.

DROP DATABASE `Users` ;

[Inline] [Edit] [Create PHP Code]

(Recent tables)

- csonline
- information_schem
- misc
- moodle
- mysql
- performance_sche

Databases

Create database ?

Peopleutf8_general_ciCreate

	Database	
<input type="checkbox"/>	csonline	
<input type="checkbox"/>	information_schema	
<input type="checkbox"/>	misc	
<input type="checkbox"/>	moodle	
<input type="checkbox"/>	mysql	
<input type="checkbox"/>	performance_schema	
Total: 6		

Check All / Uncheck All With selected:

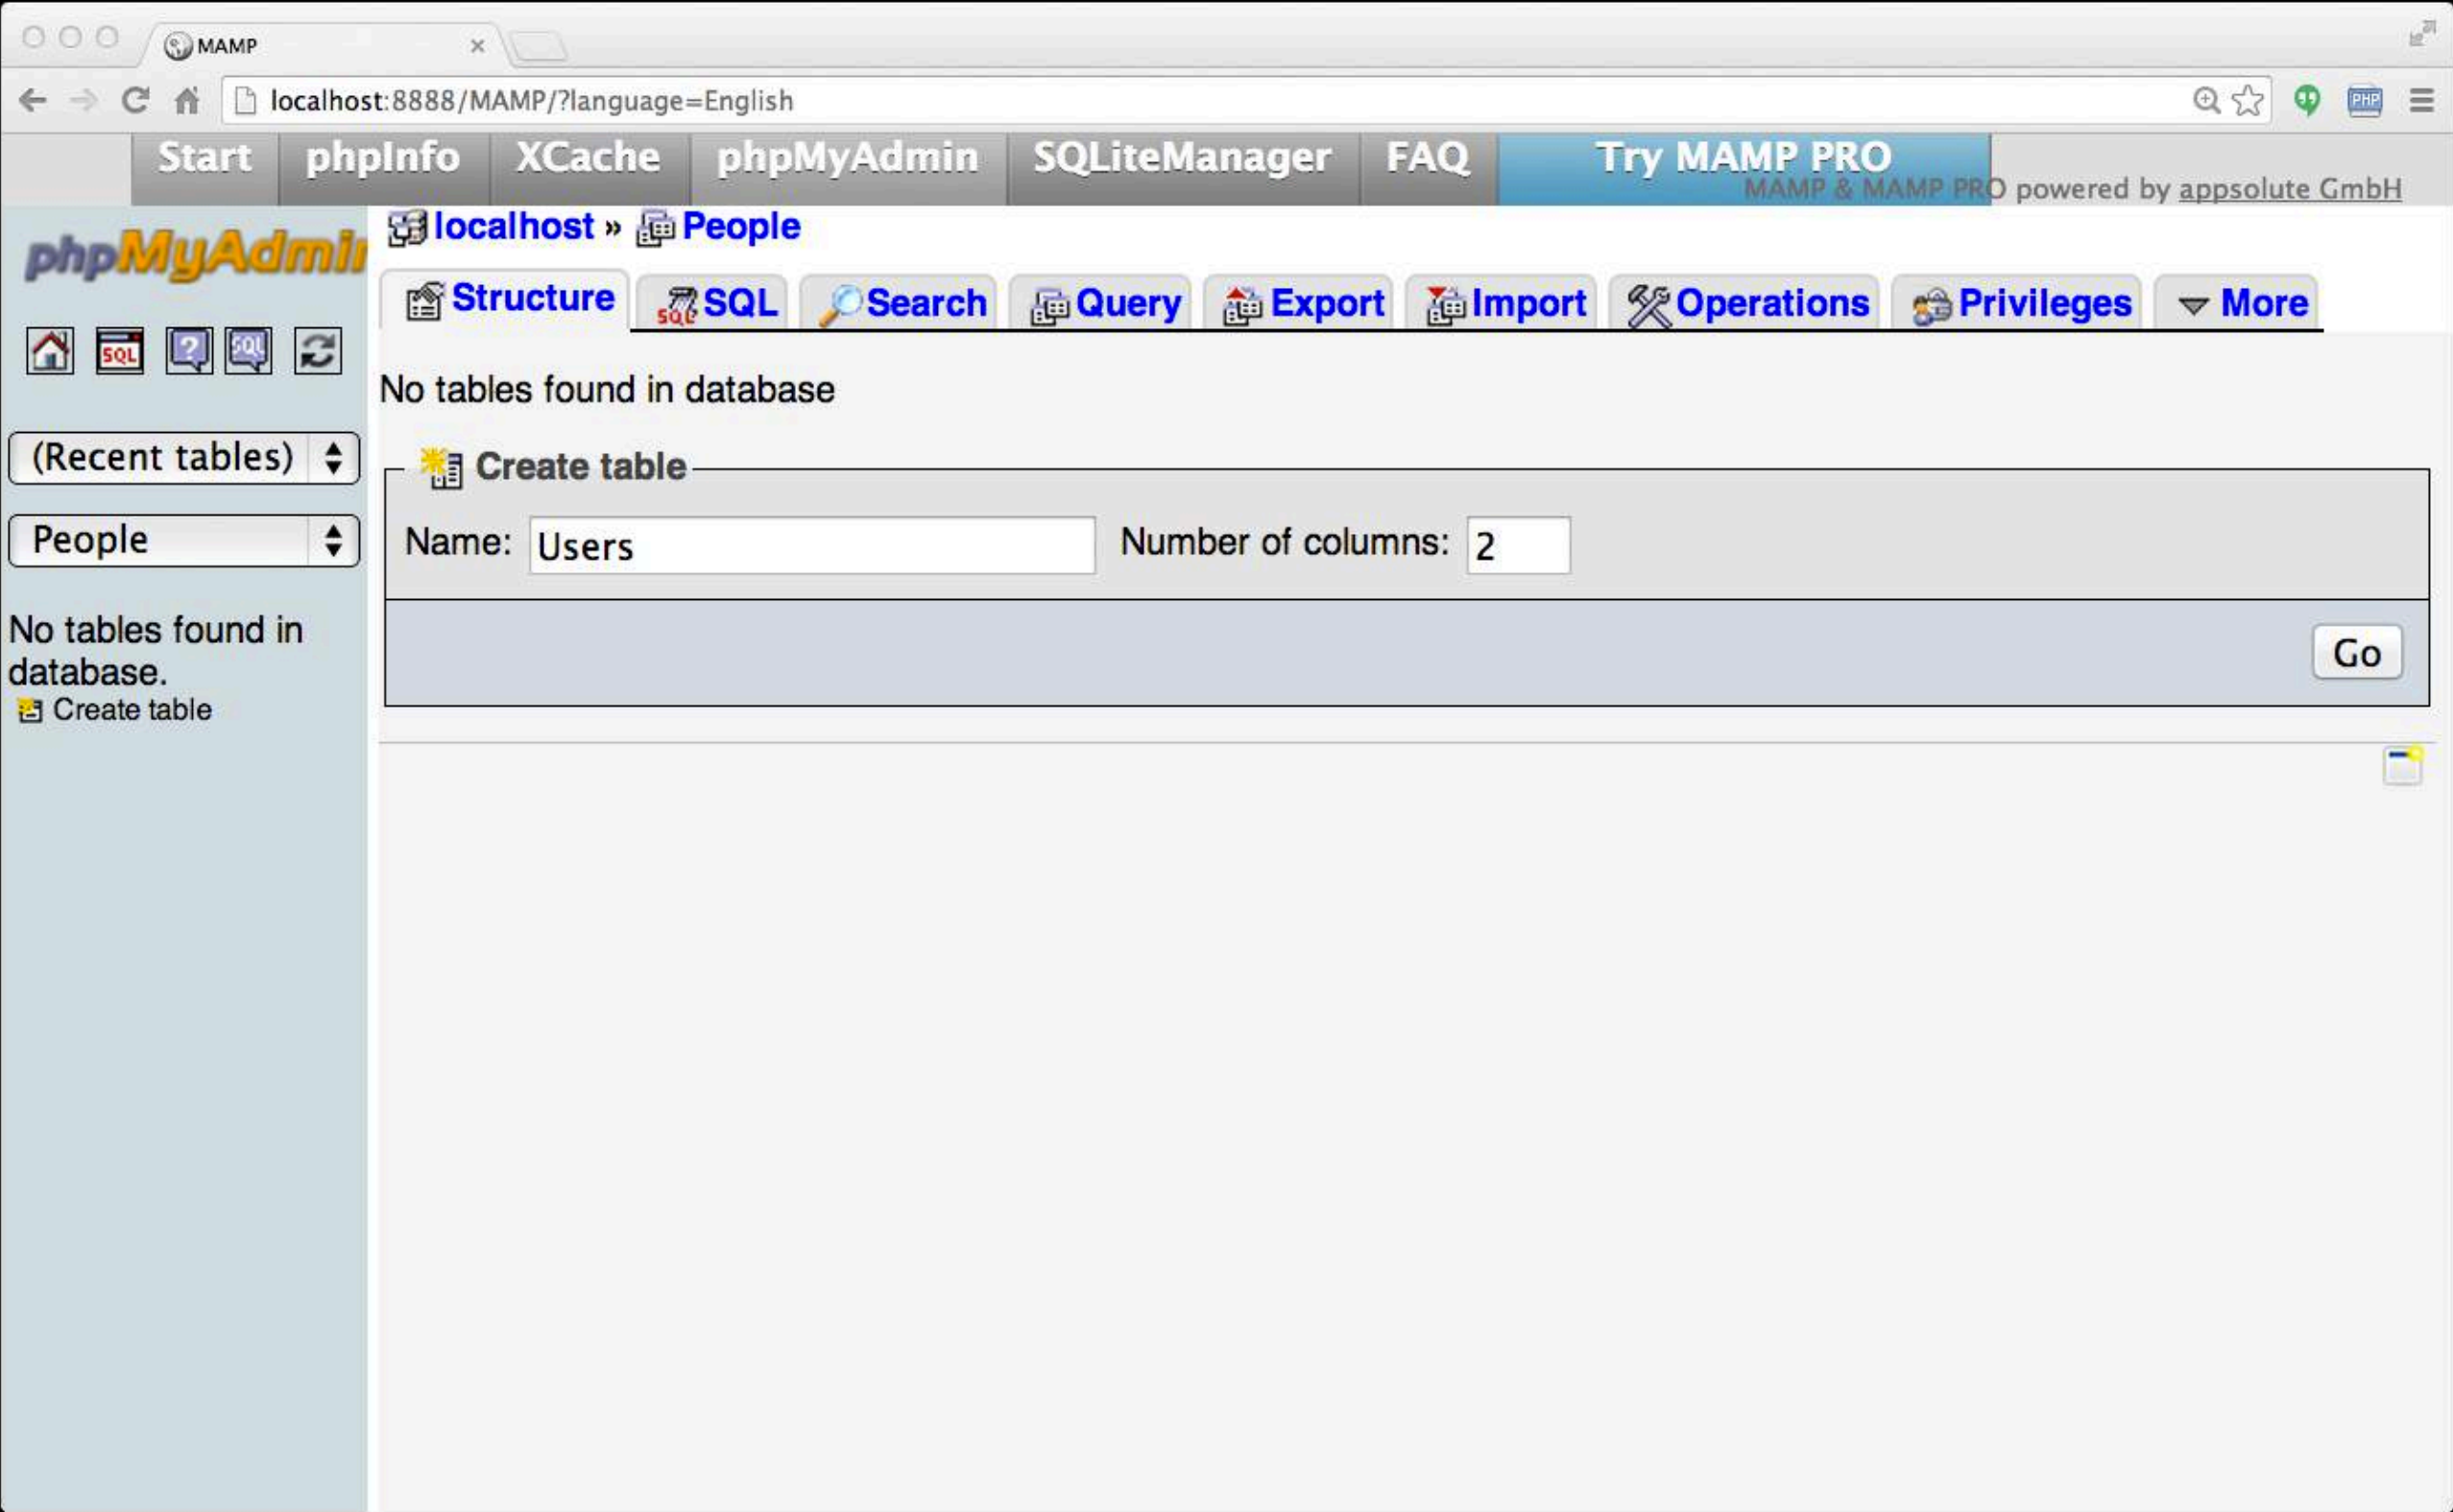
Enable Statistics

Start Simple - A Single Table

- Let's make a table of Users in our People database
- Two columns - name and email

```
CREATE TABLE Users(  
    name VARCHAR(128),  
    email VARCHAR(128)  
);
```

```
DESCRIBE Users;
```

MAMP

localhost:8888/MAMP/?language=English

StartphpInfoXCachephpMyAdminSQLiteManagerFAQTry MAMP PRO

MAMP & MAMP PRO powered by appsolute GmbH

phpMyAdmin

Home

SQL

Help

SQL

Refresh

(Recent tables) ▾

People ▾

No tables found in database.
Create table

Table name: UsersAdd 1 column(s)Go

Structure ?

Name	Type ?	Length/Values ?	Default ?	Collation
name	VARCHAR ▾	128	None ▾	▾
email	VARCHAR ▾	128	None ▾	▾

Table comments:

Storage Engine: InnoDB ▾

Collation: ▾

PARTITION definition: ?

SaveCancel

MAMP

localhost:8888/MAMP/?language=English

StartphpInfoXCachephpMyAdminSQLiteManagerFAQTry MAMP PRO

MAMP & MAMP PRO powered by appsolute GmbH

phpMyAdmin















localhost » People » Users







BrowseStructureSQLSearchInsertExportImportOperationsTriggers

MySQL returned an empty result set (i.e. zero rows). (Query took 0.0003 sec)

```
SELECT *
FROM `Users`
LIMIT 0 , 30
```

☐ Profiling [Inline] [Edit] [Explain SQL] [Create PHP Code] [Refresh]

#	Name	Type	Collation	Attributes	Null	Default	Extra	Action
<input type="checkbox"/>	1	name	varchar(128)	utf8_general_ci		No	None	      
<input type="checkbox"/>	2	email	varchar(128)	utf8_general_ci		No	None	      

☐ Check All / Uncheck All With selected: 

Print view

Relation view

Propose table structure

Add column(s) ☒ At End of Table ☐ At Beginning of Table ☐ After

name

Go

+ Indexes

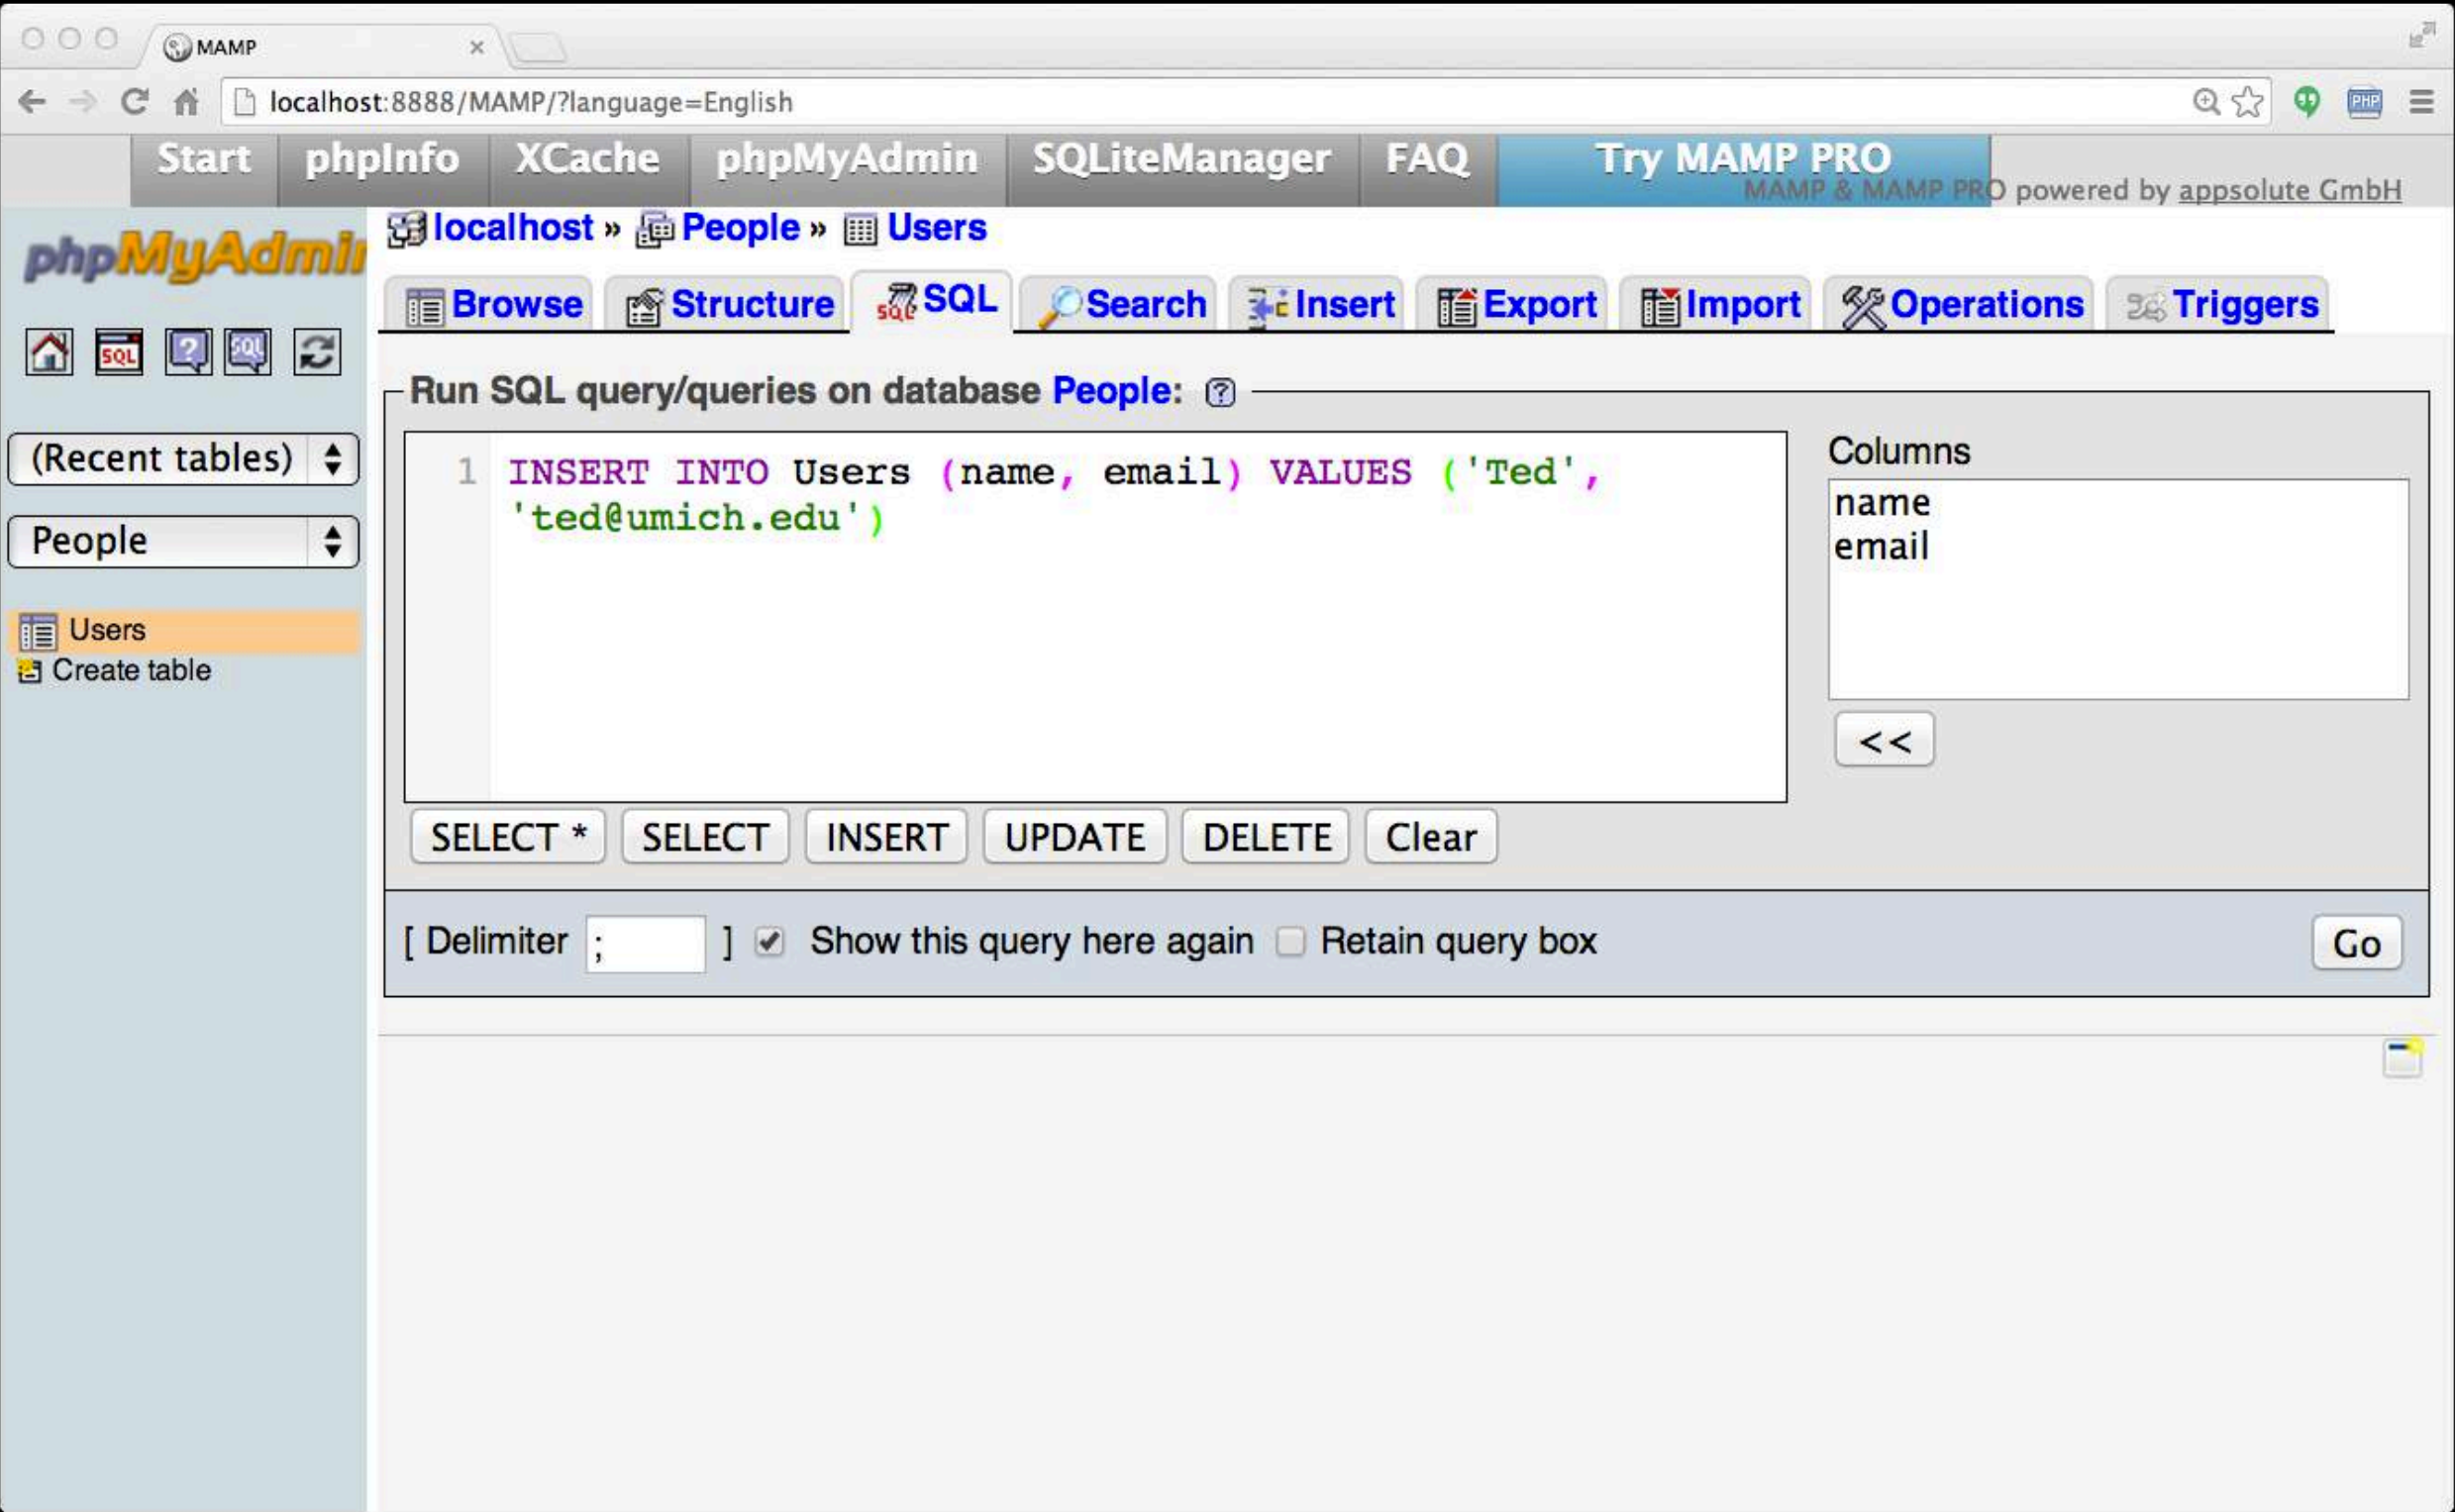
Information

Space usage			Row Statistics	
Type	Usage		Statements	Value
Data	16	KiB	Format	Compact
Index	0	B	Collation	utf8_general_ci
Total	16	KiB	Creation	Jan 26, 2014 at 04:57 PM

SQL: Insert

The **INSERT** statement inserts a row into a table

```
INSERT INTO Users (name, email) VALUES ('Chuck', 'csev@umich.edu') ;
INSERT INTO Users (name, email) VALUES ('Somesh', 'somesh@umich.edu') ;
INSERT INTO Users (name, email) VALUES ('Caitlin', 'cait@umich.edu') ;
INSERT INTO Users (name, email) VALUES ('Ted', 'ted@umich.edu') ;
INSERT INTO Users (name, email) VALUES ('Sally', 'sally@umich.edu') ;
```



MAMP

localhost:8888/MAMP/?language=English

StartphpInfoXCachephpMyAdminSQLiteManagerFAQTry MAMP PRO

MAMP & MAMP PRO powered by appsolute GmbH

phpMyAdmin

localhost » People » Users

BrowseStructureSQLSearchInsertExportImportOperationsTriggers

Showing rows 0 - 4 (~5 total , Query took 0.0004 sec)

SELECT *
FROM `Users`
LIMIT 0 , 30

Profiling [Inline] [Edit] [Explain SQL] [Create PHP Code] [Refresh]

Show : Start row: 0 Number of rows: 30 Headers every 100 rows

+ Options

				name	email
<input type="checkbox"/>				Chuck	csev@umich.edu
<input type="checkbox"/>				Sally	sally@umich.edu
<input type="checkbox"/>				Somesh	somesh@umich.edu
<input type="checkbox"/>				Caitlin	cait@umich.edu
<input type="checkbox"/>				Ted	ted@umich.edu

Check All / Uncheck All With selected:

Show : Start row: 0 Number of rows: 30 Headers every 100 rows

Query results operations

Print viewPrint view (with full texts)ExportDisplay chartCreate view

SQL: Delete

Deletes a row in a table based on selection criteria

```
DELETE FROM Users WHERE email='ted@umich.edu'
```

MAMP

localhost:8888/MAMP/?language=English

Start phpInfo XCache phpMyAdmin SQLiteManager FAQ Try MAMP PRO

MAMP & MAMP PRO powered by appsolute GmbH

localhost » Users » People

Browse Structure SQL Search Insert Export Import Operations Triggers

Run SQL query/queries on database Users: ?

1 DELETE FROM Users WHERE email='ted@umich.edu'

Columns
name
email

SELECT * SELECT INSERT UPDATE DELETE Clear

[Delimiter ;] ☒ Show this query here again ☐ Retain query box

Go

SQL: Update

Allows the updating of a field with a **WHERE** clause

```
UPDATE Users SET name='Charles' WHERE email='csev@umich.edu'
```


StartphpInfoXCachephpMyAdminSQLiteManagerFAQTry MAMP PRO

localhost » People » Users

BrowseStructureSQLSearchInsertExportImportOperationsTriggers

Showing rows 0 - 3 (~4 total , Query took 0.0003 sec)

```
SELECT *
FROM `Users`
LIMIT 0 , 30
```

Profiling [Inline] [Edit] [Explain SQL] [Create PHP Code] [Refresh]

Show : Start row: 0 Number of rows: 30 Headers every 100 rows

+ Options


				name	email
<input type="checkbox"/>				Charles	csev@umich.edu
<input type="checkbox"/>				Sally	sally@umich.edu
<input type="checkbox"/>				Somesh	somesh@umich.edu
<input type="checkbox"/>				Caitlin	cait@umich.edu

Check All / Uncheck All With selected:

Show : Start row: 0 Number of rows: 30 Headers every 100 rows

Query results operations

Print view Print view (with full texts) Export Display chart Create view



Retrieving Records: **Select**

Retrieves a group of records - you can either retrieve all the records or a subset of the records with a **WHERE** clause

```
SELECT * FROM Users
```

```
SELECT * FROM Users WHERE email='csev@umich.edu'
```

MAMP

localhost:8888/MAMP/?language=English

StartphpInfoXCachephpMyAdminSQLiteManagerFAQTry MAMP PRO

MAMP & MAMP PRO powered by appsolute GmbH

localhost » People » Users

Loading

BrowseStructureSQLSearchInsertExportImportOperationsTriggers

Showing rows 0 - 3 (4 total, Query took 0.0002 sec)

SELECT *
FROM Users
LIMIT 0 , 30

Profiling [Edit] [Explain SQL] [Create PHP Code] [Refresh]

Show : Start row: 0 Number of rows: 30 Headers every 100 rows

				name	email
<input type="checkbox"/>				Charles	csev@umich.edu
<input type="checkbox"/>				Sally	sally@umich.edu
<input type="checkbox"/>				Somesh	somesh@umich.edu
<input type="checkbox"/>				Caitlin	cait@umich.edu

Check All / Uncheck All With selected:

Show : Start row: 0 Number of rows: 30 Headers every 100 rows

Query results operations

Print view Print view (with full texts) Export Display chart Create view

StartphpInfoXCachephpMyAdminSQLiteManagerFAQTry MAMP PRO

MAMP & MAMP PRO powered by appsolute GmbH

localhost » People » Users

Loading

BrowseStructureSQLSearchInsertExportImportOperationsTriggers

[Delimiter ;] [Show this query here again] [Retain query box] Go

Showing rows 0 - 0 (1 total, Query took 0.0002 sec)

SELECT *
FROM Users
WHERE email = 'csev@umich.edu'
LIMIT 0 , 30

Profiling [Edit] [Explain SQL] [Create PHP Code] [Refresh]

Show : Start row: 0 Number of rows: 30 Headers every 100 rows

nameemail

Charlescsev@umich.edu

Check All / Uncheck All With selected:

Show : Start row: 0 Number of rows: 30 Headers every 100 rows

Query results operations

Print viewPrint view (with full texts)ExportDisplay chartCreate view

Sorting with ORDER BY

You can add an **ORDER BY** clause to **SELECT** statements to get the results sorted in ascending or descending order

```
SELECT * FROM Users ORDER BY email
```

MAMP

localhost:8888/MAMP/?language=English

StartphpInfoXCachephpMyAdminSQLiteManagerFAQTry MAMP PRO

MAMP & MAMP PRO powered by appsolute GmbH

localhost » People » Users

Loading

BrowseStructureSQLSearchInsertExportImportOperationsTriggers

phpMyAdmin

(Recent tables)

People

Users

Create table

SELECT *
FROM Users
ORDER BY email
LIMIT 0 , 30

☐ Profiling [Edit] [Explain SQL] [Create PHP Code] [Refresh]

Show : Start row: 0 Number of rows: 30 Headers every 100 rows

				name	email
<input type="checkbox"/>				Caitlin	cait@umich.edu
<input type="checkbox"/>				Charles	csev@umich.edu
<input type="checkbox"/>				Sally	sally@umich.edu
<input type="checkbox"/>				Somesh	somesh@umich.edu

☐ Check All / Uncheck All With selected:

Show : Start row: 0 Number of rows: 30 Headers every 100 rows

Query results operations

Print view Print view (with full texts) Export Display chart Create view

The LIKE Clause

We can do wildcard matching in a **WHERE** clause
using the **LIKE** operator

```
SELECT * FROM Users WHERE name LIKE '%e%'
```

MAMP

localhost:8888/MAMP/?language=English

StartphpInfoXCachephpMyAdminSQLiteManagerFAQTry MAMP PRO

MAMP & MAMP PRO powered by appabsolute GmbH

localhost » People

Loading

StructureSQLSearchQueryExportImportOperationsPrivilegesMore

Showing rows 0 - 1 (2 total, Query took 0.0003 sec)

SELECT *
FROM Users
WHERE name LIKE 'e%'
LIMIT 0 , 30

☐ Profiling [Edit] [Explain SQL] [Create PHP Code] [Refresh]

Show : Start row: 0

Number of rows: 30

Headers every 100

rows

				name	email
<input type="checkbox"/>				Charles	csev@umich.edu
<input type="checkbox"/>				Somesh	somesh@umich.edu

☐ Check All / ☐ Uncheck All With selected:

Show : Start row: 0

Number of rows: 30

Headers every 100

rows

Query results operations

Print view Print view (with full texts) Export Display chart Create view

The LIMIT Clause

- The **LIMIT** clause can request the first "n" rows, or the first "n" rows after some starting row. Note: the first row is zero, not one.
- WHERE and ORDER BY clauses happen **before** the LIMIT is applied.
- The limit can be a count or a starting row and count (starts from 0).

```
SELECT * FROM Users ORDER BY email DESC LIMIT 2;  
SELECT * FROM Users ORDER BY email LIMIT 1,2;
```


StartphpInfoXCachephpMyAdminSQLiteManagerFAQTry MAMP PRO

localhost » People

StructureSQLSearchQueryExportImportOperationsPrivilegesMore

[Delimiter ;] [x] Show this query here again [] Retain query box [Go]

Showing rows 2 - 1 (2 total, Query took 0.0003 sec) [email: SOMESH@UMICH.EDU - SALLY@UMICH.EDU]

SELECT *
FROM Users
ORDER BY email DESC
LIMIT 2

[] Profiling [Edit] [Explain SQL] [Create PHP Code] [Refresh]

				name	email
<input type="checkbox"/>				Somesh	somesh@umich.edu
<input type="checkbox"/>				Sally	sally@umich.edu

Check All / Uncheck All With selected:

Query results operations

Print view Print view (with full texts) Export Display chart Create view

MAMP

localhost:8888/MAMP/?language=English

StartphpInfoXCachephpMyAdminSQLiteManagerFAQTry MAMP PRO

MAMP & MAMP PRO powered by appabsolute GmbH

localhost » People

StructureSQLSearchQueryExportImportOperationsPrivilegesMore

[Delimiter ;] ☒ Show this query here again ☐ Retain query box Go

Showing rows 1 - 2 (2 total, Query took 0.0002 sec) [email: CSEV@UMICH.EDU - SALLY@UMICH.EDU]

SELECT *
FROM Users
ORDER BY email
LIMIT 1 , 2

☐ Profiling [Edit] [Explain SQL] [Create PHP Code] [Refresh]

name

email

☐

Charles

csev@umich.edu

☐

Sally

sally@umich.edu

Check All / Uncheck All

With selected:

Query results operations

Print view

Print view (with full texts)

Export

Display chart

Create view

Counting Rows with SELECT

You can request to receive the **count** of the rows that would be retrieved instead of the rows

```
SELECT COUNT(*) FROM Users;
```

```
SELECT COUNT(*) FROM Users WHERE email='csev@umich.edu'
```

SQL Summary

```
INSERT INTO Users (name, email) VALUES ('Ted', 'ted@umich.edu')
```

```
DELETE FROM Users WHERE email='ted@umich.edu'
```

```
UPDATE Users SET name='Charles' WHERE email='csev@umich.edu'
```

```
SELECT * FROM Users WHERE email='csev@umich.edu'
```

```
SELECT * FROM Users ORDER BY email
```

```
SELECT * FROM Users WHERE name LIKE '%e%'
```

```
SELECT * FROM Users ORDER BY email LIMIT 1,2;
```

```
SELECT COUNT(*) FROM Users WHERE email='csev@umich.edu'
```

This is not too exciting (so far)

- Tables pretty much look like big, fast programmable spreadsheets with rows, columns, and commands.
- The power comes when we have more than one table and we can exploit the relationships between the tables.



Data Types in SQL

Looking at Data Types

- Text fields (small and large)
- Binary fields (small and large)
- Numeric fields
- AUTO_INCREMENT fields

String Fields

- Understand character sets and are indexable for searching
- **CHAR** allocates the entire space (faster for small strings where length is known)
- **VARCHAR** allocates a variable amount of space depending on the data length (less space)

Text Fields

- Have a character set - paragraphs or HTML pages
 - **TINYTEXT** up to 255 characters
 - **TEXT** up to 65K
 - **MEDIUMTEXT** up to 16M
 - **LONGTEXT** up to 4G
- Generally not used with indexing or sorting - and only then limited to a prefix

Binary Types (rarely used)

- Character = 8 - 32 bits of information depending on character set
- Byte = 8 bits of information
 - **BYTE**(n) up to 255 bytes
 - **VARBINARY**(n) up to 65K bytes
- Small Images - data
- Not indexed or sorted

Binary Large Object (BLOB)

- Large raw data, files, images, word documents, PDFs, movies, etc.
- No translation, indexing, or character set
 - **TINYBLOB**(n) - up to 255
 - **BLOB**(n) - up to 65K
 - **MEDIUMBLOB**(n) - up to 16M
 - **LOB**(n) - up to 4G



Integer Numbers

Integer numbers are very efficient, take little storage, and are easy to process because CPUs can often compare them with a single instruction.

- **TINYINT** (-128, 128)
- **SMALLINT** (-32768, +32768)
- **INT** or **INTEGER** (2 Billion)
- **BIGINT** - (10^{18} ish)

Floating Point Numbers

Floating point numbers can represent a wide range of values, but accuracy is limited.

- **FLOAT** (32-bit) 10^{38} with seven digits of accuracy
- **DOUBLE** (64-bit) 10^{308} with 14 digits of accuracy

Dates

- **TIMESTAMP** - 'YYYY-MM-DD HH:MM:SS' (1970, 2037)
- **DATETIME** - 'YYYY-MM-DD HH:MM:SS'
- **DATE** - 'YYYY-MM-DD'
- **TIME** - 'HH:MM:SS'
- Built-in MySQL function NOW()



Database Keys and Indexes

AUTO_INCREMENT

Often as we make multiple tables and need to JOIN them together we need an integer primary key for each row so we can efficiently add a reference to a row in some other table as a foreign key.

```
DROP TABLE Users;
```

```
CREATE TABLE Users (  
    user_id INT UNSIGNED NOT NULL  
        AUTO_INCREMENT,  
    name VARCHAR(128),  
    email VARCHAR(128),  
    PRIMARY KEY(user_id),  
    INDEX(email)  
)
```

MySQL Functions

Many operations in MySQL need to use the built-in functions (like `NOW()` for dates).

- <http://dev.mysql.com/doc/refman/5.0/en/string-functions.html>
- <http://dev.mysql.com/doc/refman/5.0/en/date-and-time-functions.html>

Indexes


- As a table gets large (they always do), scanning all the data to find a single row becomes very costly
- When drchuck@gmail.com logs into FaceBook, they must find my password amongst 500 million users
- There are techniques to greatly shorten the scan as long as you create data structures and maintain those structures - like shortcuts
- Hashes or Trees

MySQL Index Types

- **PRIMARY KEY** - Very little space, exact match, requires no duplicates, extremely fast for integer fields
- **INDEX** - Good for individual row lookup and sorting / grouping results - works best with exact matches or prefix lookups - can suggest HASH or BTREE

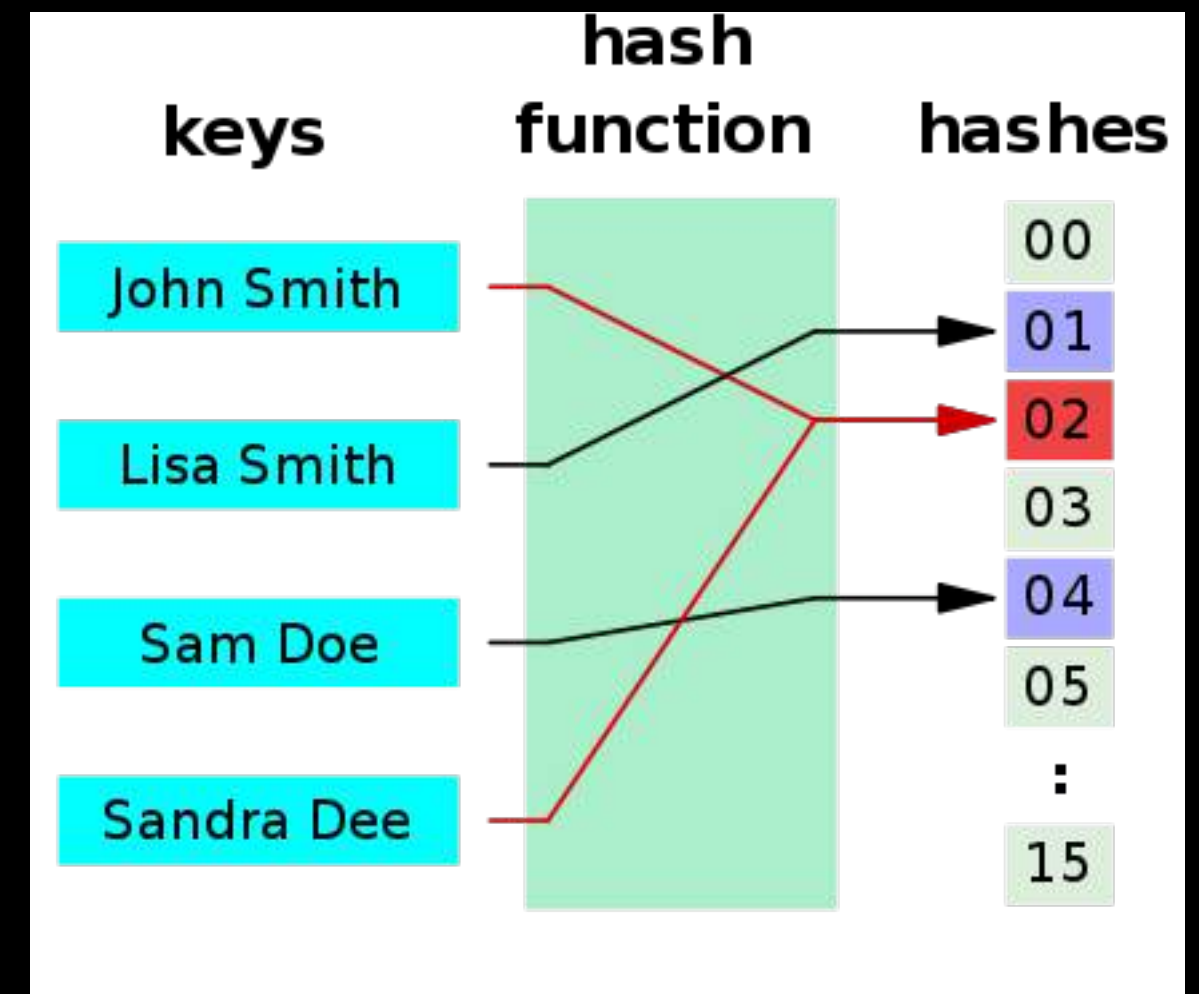
B-Trees

A B-tree is a tree data structure that keeps data sorted and allows searches, sequential access, insertions, and deletions in logarithmic amortized time. The B-tree is optimized for systems that read and write large blocks of data. It is commonly used in databases and file systems.



Hashes

A hash function is any algorithm or subroutine that maps large data sets to smaller data sets, called keys. For example, a single integer can serve as an index to an array (cf. associative array). The values returned by a hash function are called hash values, hash codes, hash sums, checksums, or simply hashes. Hash functions are mostly used to accelerate table lookup or data comparison tasks such as finding items in a database...



Specifying Indexes

```
DROP TABLE Users;
```

```
CREATE TABLE Users (  
    user_id INT UNSIGNED NOT NULL  
        AUTO_INCREMENT,  
    name VARCHAR(128),  
    email VARCHAR(128),  
    PRIMARY KEY(user_id),  
    INDEX(email)  
)
```

```
ALTER TABLE Users ADD INDEX ( email ) USING BTREE
```


Summary

- SQL allows us to describe the shape of data to be stored and give many hints to the database engine as to how we will be accessing or using the data.
- SQL is a language that provides us operations to Create, Read, Update, and Delete (CRUD) our data in a database.

Acknowledgements / Contributions



These slides are Copyright 2010- Charles R. Severance (www.dr-chuck.com) as part of www.wa4e.com and made available under a Creative Commons Attribution 4.0 License. Please maintain this last slide in all copies of the document to comply with the attribution requirements of the license. If you make a change, feel free to add your name and organization to the list of contributors on this page as you republish the materials.

Initial Development: Charles Severance, University of Michigan
School of Information

Insert new Contributors and Translators here including names and dates

Continue new Contributors and Translators here