
TIM 50 - Business Information Systems

Lecture 16

Instructor: Terry Allen

UC Santa Cruz

11/16/2011

Outline

- Announcements
- MySQL case (cont'd)
- Databases
- Student Presentation ?

Announcements

- Reading for 11/21
 - Messerschmitt 11.2 (pp.333-335)
 - Akamai Case (reader pp.217-236)
- No class 11/23
- Homework #3 out
- Database assignment will be posted later this week

mySQL Case

mySQL

What does mySQL make?

How Successful is mySQL?

- Visibility: Fortune magazine, more mentions on www
- Reaction from giants
- Revenue growth 2001 700k, 2002 6.2m, 2003 10m
- Good performance reviews
- Recent SAP alliance
- But Market share tiny:
 - \$10 million out of \$10 billion market!

- **Why Success?**
 - Good Technology
 - Competing (large) DBMSs bloated with features most don't need
 - Innovative OSS model

MySQL

How does OSS work?

Two Types of License:

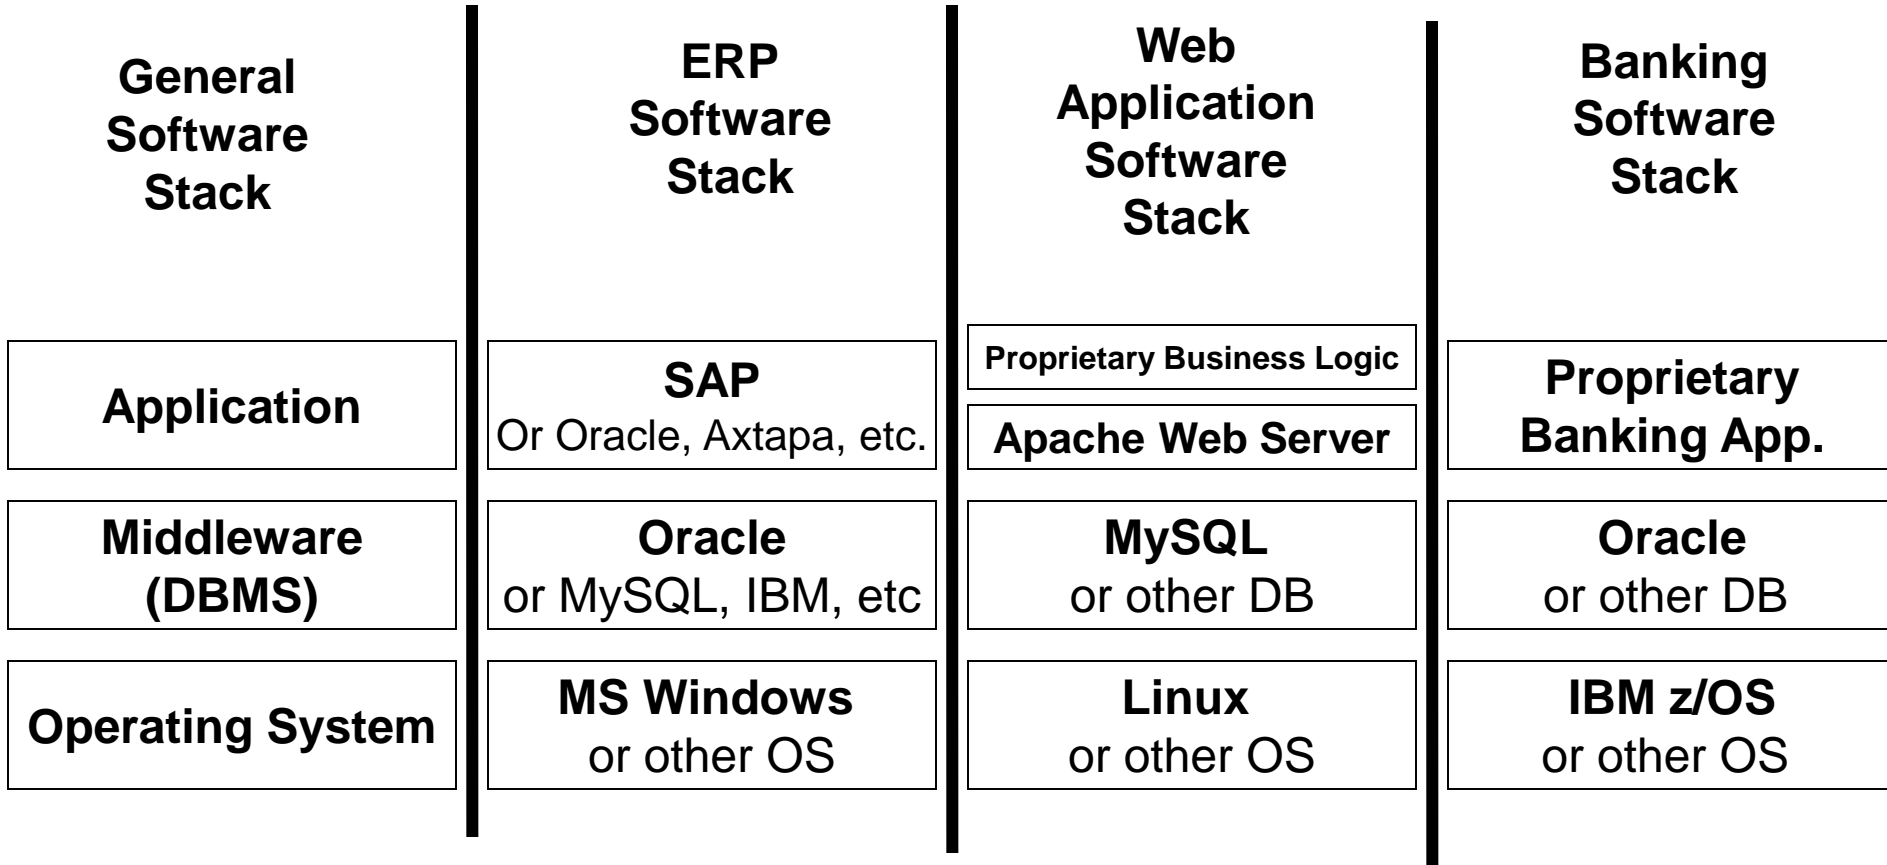
■ GPL

- Free
- No Support
- Any software that uses MySQL as a module must itself be released under GPL

■ Commercial License

- Support
- Could be distributed with non-open source software
- Not Free:
 - MySQL: Classic \$250, Pro \$495 (for ~ 50 users)
 - Compare to:
 - MSFT \$3150 single proc for 50 users
 - IBM \$33000 single proc for 50 users
 - Oracle \$40000 single proc for 50 users

Aside: DB's in different software stacks



- Which companies are competitors?
- Which are complimentary to each other?
- Which are both?

mySQL

- Which segments of market is mySQL strong in?
 - Large Companies or Small Companies?
 - Web applications or Critical Enterprise data?
- Why would a major enterprise want to pay so much more for an Oracle or IBM DB?

My SQL: market

	Small 20%	Medium 30%	Large 50%
Enterprise wide data 90%	Microsoft		Oracle IBM Reliability Scalability Support Longevity
Web Sites 10%	My SQL Cost		


How should mySQL grow in order to meet it's stated goal of getting to \$100 million In revenue?

My SQL: Growth Strategy

	Small 20%	Medium 30%	Large 50%
Enterprise wide data 90%	Microsoft		Oracle IBM Reliability Scalability Support Longevity
Web Sites 10%	My SQL Cost	Maybe?	

- + Many of these customers already using MySQL with websites
- + Less emphasis on global organization
- + Leverage SAP alliance
- - Up against Microsoft.

My SQL: Growth Strategy

	Small 20%	Medium 30%	Large 50%
Enterprise wide data 90%	Microsoft		Oracle IBM Reliability Scalability Support Longevity
Web Sites 10%	My SQL Cost		Maybe?

- + builds on existing brand and strengths
- - Market not so big

Database Management

Databases & MySQL Quiz

- 1) In a relational table, each record is represented by
 - a) a row
 - b) a column
 - c) either a row or a column

- 2) Who are the top three RDBMS competitors of MySQL?

- 3) Linux is
 - a) a proprietary operating system
 - b) a proprietary DBMS
 - c) an open source operating system
 - d) an open source DBMS

Relational table

Table

Employee		
Name	Address	Dept
Record		

Field/Attribute

Remember: Data properties

Field name

Table name

EMPLOYERS

EMPL_ID	EMPL_NAME	EMPL_POSITION	DEPT.ID
100	Alice	Manager	1
101	Bob	Programmer	1
102	Chris	Manager	2
103	David	Accountant	2

Field Value

Field Type
e.g. EMPL_ID is INTEGER
EMPL_NAME is STRING

The Relational Model

NAME	POSITION	DEPARTMENT	DEPT. ADDRESS
Alice	Manager	IT	San Jose
Bob	Programmer	IT	San Jose
Chris	Manager	Finance	New York
David	Accountant	Finance	New York

- Not all data are stored in the same table
 - Avoid data replication/redundancy
 - Enable faster/efficient computations over data

The Relational Model

EMPLOYERS

EMPL_ID	EMPL_NAME	EMPL_POSITION	DEPT.ID
100	Alice	Manager	1
101	Bob	Programmer	1
102	Chris	Manager	2
103	David	Accountant	2

DEPARTMENTS

DEPT. ID	DEPARTMENT	DEPT. ADDRESS
1	IT	San Jose
2	Finance	New York

The Relational Model

Tables are connected using these **KEYS**

EMPLOYERS

EMPL_ID	EMPL_NAME	EMPL_POSITION	DEPT.ID
100	Alice	Manager	1
101	Bob	Programmer	1
102	Chris	Manager	2
103	David	Accountant	2

Each table should have a unique identifier for each record: **KEY**

DEPARTMENTS

DEPT. ID	DEPARTMENT	DEPT. ADDRESS
1	IT	San Jose
2	Finance	New York

The Relational Model



Secondary Key



EMPLOYERS

EMPL_ID	EMPL_NAME	EMPL_POSITION	DEPT.ID
100	Alice	Manager	1
101	Bob	Programmer	1
102	Chris	Manager	2
103	David	Accountant	2

Primary Keys



DEPARTMENTS

DEPT. ID	DEPARTMENT	DEPT. ADDRESS
1	IT	San Jose
2	Finance	New York


The Relational Model

- **Relational Model**

- Based on relational algebra
- Specifies a suite of *operations/operators* that can be performed on tables
- These operations are expressed by **SQL commands**

SQL interface

- SQL (Structured Query Language)
- Forms an "interface" between an application and the DBMS
 - For manipulating, and extracting data from database
- Standardized, not vendor specific
- Encapsulates various internal details
 - Data partitioning and replication
 - Host mapping
 - File representation
 - etc.



Database Operations

EMPLOYERS

EMPL_ID	EMPL_NAME	EMPL_POSITION	DEPT_ID
100	Alice	Manager	1
101	Bob	Programmer	1
102	Chris	Manager	2
103	David	Accountant	2

DEPARTMENTS

DEPT_ID	DEPARTMENT	DEPT_ADDRESS
1	IT	San Jose
2	Finance	New York

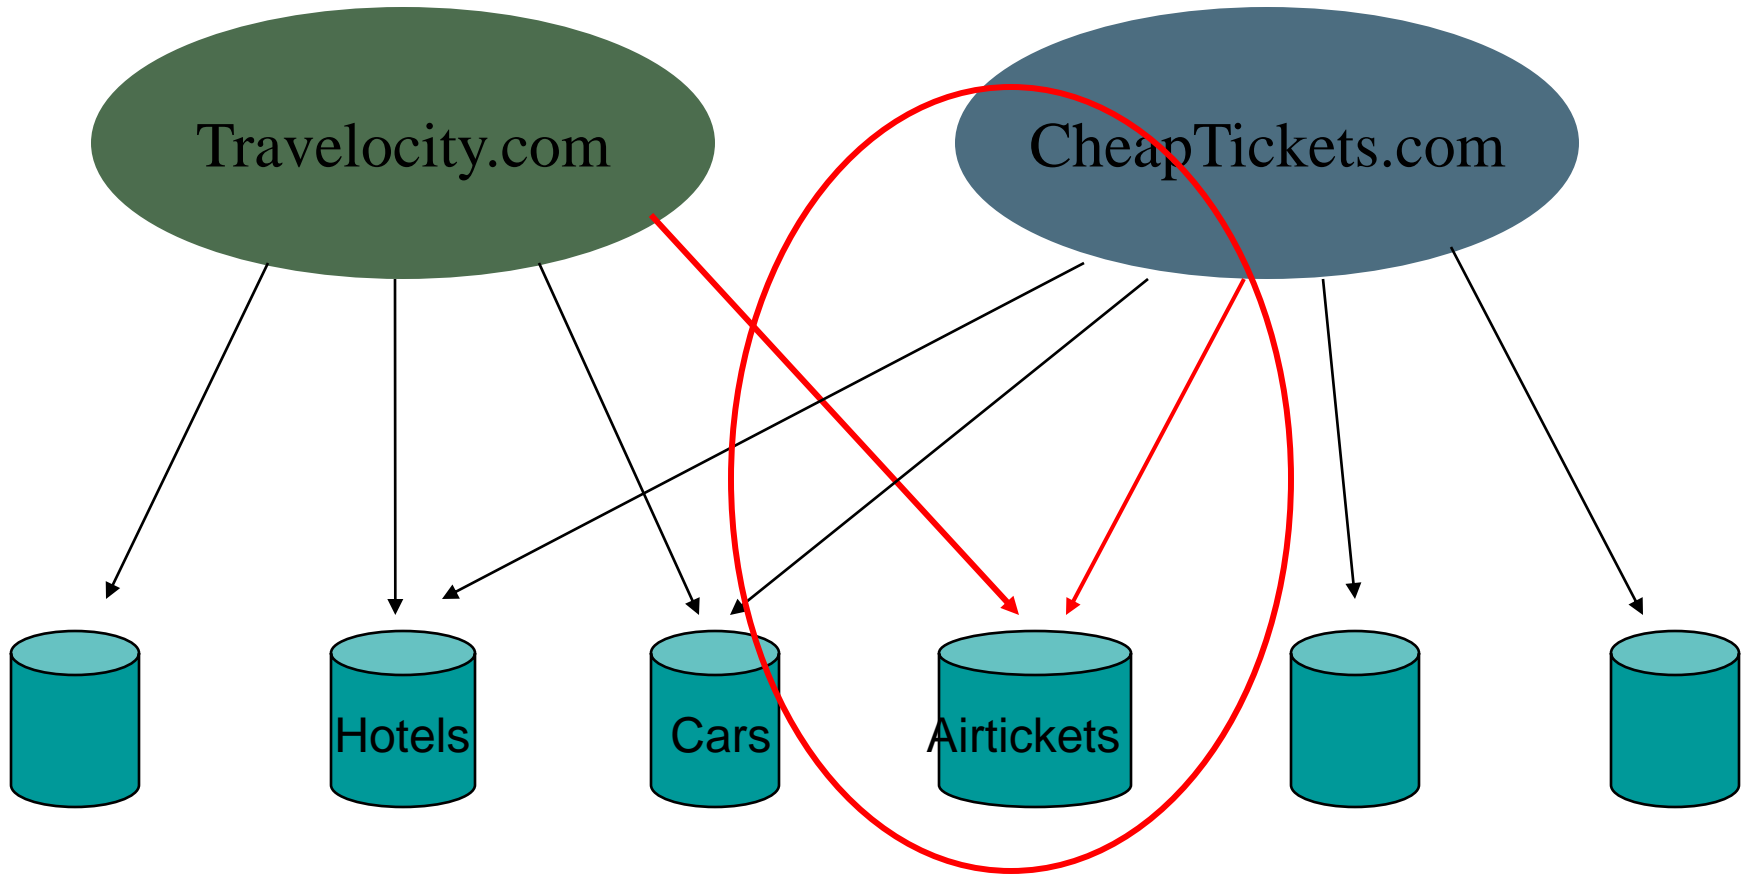
JOIN EMPL_DEPT



EMPL_ID	EMPL_NAME	EMPL_POSITION	DEPT.	ADDRESS
100	Alice	Manager	IT	San Jose
101	Bob	Programmer	IT	San Jose
102	Chris	Manager	Finance	New York
103	David	Accountant	Finance	New York

mySQL Case

Example - Travel Agency



Two applications might try to access & update the same resource concurrently